

Natural-Language Multi-Agent Simulations of Argumentative Opinion Dynamics

Gregor Betz¹

¹Karlsruhe Institute of Technology, DebateLab, Department of Philosophy Douglasstr. 24
Karlsruhe 76135 Germany

Correspondence should be addressed to gregor.betz@kit.edu

Journal of Artificial Societies and Social Simulation 25(1) 2, 2022

Doi: 10.18564/jasss.4725 Url: <http://jasss.soc.surrey.ac.uk/25/1/2.html>

Received: 06-05-2021 Accepted: 08-11-2021 Published: 31-01-2022

Abstract: This paper develops a natural-language agent-based model of argumentation (ABMA). Its artificial deliberative agents (ADAs) are constructed with the help of so-called neural language models recently developed in AI and computational linguistics. ADAs are equipped with a minimalist belief system and may generate and submit novel contributions to a conversation. The natural-language ABMA allows us to simulate collective deliberation in English, i.e. with arguments, reasons, and claims themselves—rather than with their mathematical representations (as in symbolic models). This paper uses the natural-language ABMA to test the robustness of symbolic reason-balancing models of argumentation (Mäs & Flache 2013; Singer et al. 2019): First of all, as long as ADAs remain passive, confirmation bias and homophily updating trigger polarization, which is consistent with results from symbolic models. However, once ADAs start to actively generate new contributions, the evolution of a conversation is dominated by properties of the agents *as authors*. This suggests that the creation of new arguments, reasons, and claims critically affects a conversation and is of pivotal importance for understanding the dynamics of collective deliberation. The paper closes by pointing out further fruitful applications of the model and challenges for future research.

Keywords: Opinion Dynamics, Argumentation, Natural Language Processing, Language Model

● Introduction

- 1.1 During the last decade, a variety of computational models of argumentative opinion dynamics have been developed and studied (such as in Betz 2012; Mäs & Flache 2013; Olsson 2013; Borg et al. 2017; Singer et al. 2019; Banisch & Olbrich 2021). These agent-based models of argumentation (ABMAS) have been put to different scientific purposes: to study polarization, consensus formation, or the veritistic value of argumentation; to understand the effects of different argumentation strategies, ascriptions of trustworthiness, or social networks; to provide empirically adequate, or epistemically ideal descriptions of joint deliberation.
- 1.2 Moreover, ABMAS differ radically in terms of *how* they represent argumentation, ranging from complex dialectical networks of internally structured arguments (e.g. Betz 2015; Beisbart et al. 2021), to abstract argumentation graphs (e.g. Sešelj & Straßer 2013; Taillandier et al. 2021), to flat pro/con lists (e.g. Mäs & Flache 2013), to deflationary accounts that equate arguments with evidence (e.g. Hahn & Oaksford 2007). However, all these models are symbolic in the sense that they are built with and process mathematical representations of natural language arguments, reasons, claims, etc.—rather than these natural language entities themselves.
- 1.3 This paper presents a computational model of argumentation that is decidedly not symbolic in the following sense: It is not built from abstract representations of arguments, but from the very natural language arguments and claims (which are merely *represented* in symbolic models) themselves. Our natural-language ABMA directly processes and runs on English sentences. A key component of our natural-language ABMA is what we call the artificial deliberative agent (ADA), which we construct with the help of game-changing NLP technology recently

developed in AI and computational linguistics (Vaswani et al. 2017; Devlin et al. 2019; Radford et al. 2019; Brown et al. 2020). Our design and study of ADAS bears similarities to research on dialogue systems (Zhang et al. 2020; Bao et al. 2020) and chatbots (Adiwardana et al. 2020) powered with neural language models; however, unlike these neural dialogue systems, ADAS are equipped with additional cognitive architecture, in particular a minimalist belief system. As illustrated in Figure 1, ADAS have a limited (and changing) perspective of a conversation, which determines their opinion vis-à-vis the central claims of the debate. In addition, ADAS may contribute to a conversation by generating novel posts conditional on their current perspective.

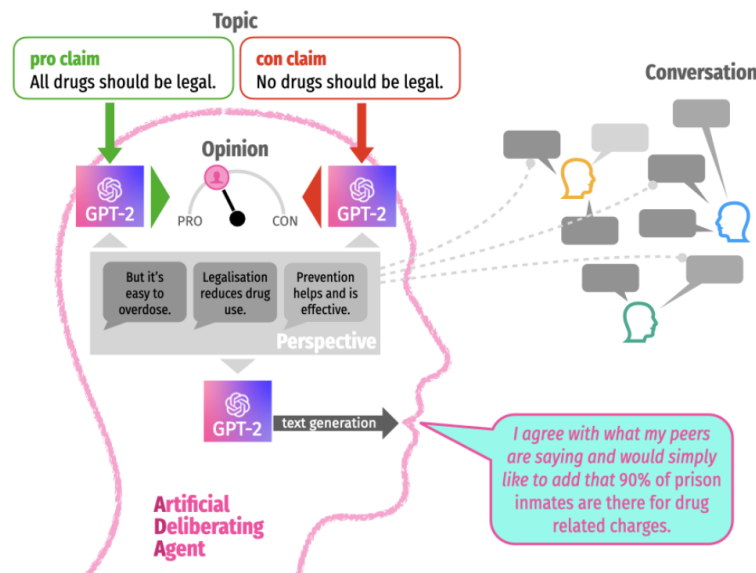


Figure 1: Basic design of artificial deliberative agents (ADAS), which we use to power natural-language agent-based models of argumentation.

- 1.4 Now, what is the motivation for developing ADAS and natural-language models of argumentative opinion dynamics in the first place?
- 1.5 A first motive for studying natural-language ABMAS is to de-idealize symbolic models and to test their results' structural robustness. If, for example, groups with over-confident agents typically bi-polarize in symbolic models but not in their natural-language counterparts, the original result is not robust and ought to be treated with care.
- 1.6 A second motive is to "reclaim new territory" by computationally investigating novel phenomena that have not been (and possibly cannot be) represented by symbolic models. Metaphorical language (Hesse 1988), slurs (Rappaport 2019), framing effects (Grüne-Yanoff 2016), or the invention of entirely new arguments (Walton & Gordon 2019) is difficult to represent in symbolic models, but relatively easy in natural-language ones. Actually, normative and epistemic biases in NLP systems are currently studied (cf. Kassner & Schütze 2020; Blodgett et al. 2020; Nadeem et al. 2020).
- 1.7 A third motive is to create computational models with implicit, natural semantics. Symbolic models of deliberation cannot escape assumptions about the "semantics" and "logic" of argument, specifically the evaluation of complex argumentation. These assumptions concern, for instance, whether individual reasons accrue by addition, whether the strength of a collection of reasons is merely determined by its weakest link, whether undefended arguments are universally untenable, whether every argument can be represented by a deductive inference, or whether our non-deductive reasoning practice is governed by probabilistic degrees of beliefs. In other words, symbolic models of argumentative opinion dynamics inevitably rest on highly contested normative theories. With natural-language ABMAS, however, there is no need to take an explicit stance regarding these theoretical issues, because the neural language model, which underlies the ADA, comes with an implicit semantics of argument and takes care of argument evaluation itself. That's why natural-language ABMAS may turn out to be neutral ground, or at least a common point of reference for symbolic models from rivaling theoretical paradigms.
- 1.8 A fourth motive is to close the gap between computational simulations on the one side and the vast amount of linguistic data about real conversations on the other side. As natural-language ABMAS do not process mathematical representations of text, but text itself, it is much more straightforward to apply and test these models on text corpora (we'll come back to this in the concluding section).

- 1.9** A fifth and final motive for studying natural-language ABMAS is to explore the capabilities of neural language models. It seems there is currently no clear scientific consensus on what to make of these AI systems (Askill 2020; Dale 2021, cf.). On the one hand, performance metrics for NLP benchmark tasks (translation, text summarization, natural language inference, reading comprehension, etc.) went literally off the chart with the advent of neural language models (see paperswithcode.com). On the other hand, some performance gains have been shown to be spurious, as they were just triggered by statistical cues in the data (e.g. Gururangan et al. 2019); what's more, the suspicion that large neural language models have simply memorized sufficiently many tasks from the Internet looms large. In this context, the ability of neural language models (mediated via ADAS) to engage in and carry on a self-sustaining, sensible conversation about a topic, while allowing ADAS to reasonably adjust their opinions in its course, may provide further evidence for the cognitive potential, if not capacities of neural language models.
- 1.10** These reasons motivate the study of natural-language ABMAS – *besides* and *in addition to* symbolic models. All of the above is *not* to say that symbolic models should be replaced by natural-language ones.
- 1.11** The paper is organized as follows. Section Model presents, step-by-step, the outline of our natural-language ABMA, including key features of ADAS, i.e.: the way an agent's opinion is elicited given her perspective, the way an agent chooses peers in a conversation, the way an agent updates her perspective, and the way an agent generates a new contribution. Note that the guiding design principle in setting up the natural-language ABMA is to rebuild symbolic reason-balancing models of argumentation (Mäs & Flache 2013; Singer et al. 2019) – which stand in the tradition of Axelrod's model of cultural dissemination (Axelrod 1997) – as faithfully as possible and to deviate from these models only where the natural language processing technology requires us to do so.
- 1.12** As further detailed in Section Experiments, we run various simulation experiments with the model to test the effects of (i) different updating strategies and (ii) active contributions to a debate. A closer look at an illustrative simulation run (Subsection An illustrative case) suggests that our model gives rise to meaningful natural language conversations and that, in particular, ADAs respond to changes in their perspective in a sensible way (as regards both opinion revision and active generation of further posts). Our main findings are reported in Subsection Global consensus and polarization effects: First of all, the natural-language ABMA with passive agents qualitatively reproduces the results of symbolic reason-balancing models regarding the effects of updating strategies on group polarization and divergence. This establishes the robustness of the originally observed effects. Secondly, active generation of novel posts heavily influences the collective opinion dynamics—to the extent that properties of the agents *qua* authors totally dominate the evolution of the conversation. So, the natural-language ABMA identifies a mechanism which is not covered in symbolic models, but which is potentially of pivotal importance for understanding the dynamics of collective deliberation.
- 1.13** We close by arguing that there are further fruitful applications of the model, which can be naturally extended to account for phenomena such as multi-dimensional opinion spaces, topic mixing, topic changes, digression, framing effects, social networks, or background beliefs (Section Discussion and Future Work). Although we report results of a preliminary sensitivity analysis in Subsection Sensitivity analysis (suggesting our findings are robust), a systematic exploration of the entire parameter space as well as of alternative initial and boundary conditions appears to be a prime desideratum of future research.

● Model

Basic design and terminology

- 2.1** A *conversation* evolves around a topic, where a *topic* is defined by a pair of central claims that characterize the opposing poles of the conversation, $\langle \text{PROCLAIMS}, \text{CONCLAIMS} \rangle$. For example, the claims {"All drugs should be legal.", "Decriminalize drugs!"} on the one side and {"No drugs should be legal.", "Drugs should be illegal."} on the opposite side may define the topic "legalization of drugs."
- 2.2** A *post* is a small (<70 words) natural language message that can be submitted to a conversation (let \mathbb{S} denote the set of all such word sequences). The conversation's timeline contains all posts that actually have been contributed to the debate ($\text{POSTS} \subseteq \mathbb{S}$), including their submission date and author. Let POSTS_t , $\text{POSTS}_{\leq t}$ refer to all posts submitted at, respectively at or before, step t .
- 2.3** Agents participate, actively or passively, in a conversation ($\text{AGENTS} = \{a_1, \dots, a_n\}$). Every agent a_i adopts a specific *perspective* on the conversation, i.e., she selects and retains a limited number of posts which have been contributed to the conversation. Formally, $\text{PERSP}_t^i = \langle p_1, p_2, \dots, p_k \rangle$ with $p_j \in \text{POSTS}_{\leq t}$ for $j = 1 \dots k$.

- 2.4 An agent's perspective fully determines her *opinion* at a given point in time, $OPIN_t^i \in [0, 1]$ (see also Subsection Opinion elicitation). In consequence, ADA's violate the total evidence requirement of rational belief formation (Suppes 1970) and our ABMA is a *bounded rationality* model (Simon 1982) of argumentative opinion dynamics.
- 2.5 Every agent a_i has a (possibly dynamically evolving) *peer group*, formally: $PEERS_t^i \subseteq AGENTS$ (cf. Subsection Peer selection). As agents update their perspective, they exchange their points of view with peers only (cf. Subsection Perspective updating).

The main loop

- 2.6 The following pseudo-code describes the main loop of the simulation of a single conversation.

Algorithm 1: Main Loop of the Simulation

```

for  $t$  in  $[1..t_{max}]$  do
  for  $i$  in AGENTS do
    determine the peers of agent  $i$  ( $\rightarrow PEERS_t^i$ );
    update the perspective of agent  $i$  ( $\rightarrow PERSP_t^i$ );
    if agent  $i$  is contributing at  $t$  then
      generate and submit a new post;
    end
    elicit the opinion of agent  $i$  ( $\rightarrow OPIN_t^i$ );
  end
end

```

Opinion elicitation

- 2.7 The opinion of agent a_i at step t is a function of a_i 's perspective at step t . We define a universal elicitation function \mathbf{O} to determine an agent's opinion:

$$OPIN_t^i = \mathbf{O}(PERSP_t^i) \quad (1)$$

$$\mathbf{O} : \mathcal{P}(\mathbb{S}) \rightarrow [0, 1] \quad (2)$$

where $PERSP_t^i$ is a sequence of posts.

- 2.8 Function \mathbf{O} is implemented with the help of neural language modeling technology (and in particular GPT-2, see Appendix A; such so-called causal language models are essentially next-word prediction machines: given a sequence of words, they estimate the probability that the next word is w , for all words w in the vocabulary). First, we transform the posts in the perspective into a single word sequence (basically by concatenating and templating, as described in Appendix B), which yields a natural language query $Q_{elic}(PERSP_t^i)$, for example: "Let's discuss legalization of drugs! Prevention helps and is effective. Legalisation reduces drug use. But it's easy to overdose. I more or less agree with what my peers are saying here. And therefore, all in all,". We elicit the opinion of the agent regarding the conversation's central claims by assessing, roughly speaking, the probability by which the agent expects a pro-claim rather than a con-claim given her perspective. More specifically, we deploy the neural language model to calculate the so-called *conditional perplexity* of the pro claims / con claims given the query previously constructed. Generally, the conditional perplexity of a sequence of words $\mathbf{w} = \langle w_1 \dots w_l \rangle$ given a sequence of words $\mathbf{v} = \langle v_1 \dots v_k \rangle$ is defined as the inverse geometric mean of conditional probabilities (Russell & Norvig 2016, 864):

$$PPL(\mathbf{w}|\mathbf{v}) = \sqrt[l]{\prod \frac{1}{p_i}}, \quad (3)$$

where $p_i := \text{Prob}(w_i | \mathbf{v} w_1 \dots w_{i-1})$ denotes the conditional probability of word w_i in sequence \mathbf{w} being the next word given all previous words of the sequence. Intuitively, if a sequence's (conditional) perplexity equals k , then, on average, every word in the sequence is as surprising as a single item being chosen out of a totality of k items (cf. Manning & Schütze 1999, 78).

- 2.9 A causal language model (LM) predicts next-word probabilities and can hence be used to calculate the conditional perplexity of a word sequence (text) given another word sequence (text). For example, the neural language model may compute $PPL(\text{"All drugs should be legal"} | \text{"Let's discuss legalization of drugs! Prevention helps and is effective. Legalisation reduces drug use. But it's easy to overdose. I more or less agree with what my peers are$

saying here. And therefore, all in all,"). Note that perplexity corresponds to inverse probability, so the higher a sequence's perplexity, the lower its overall likelihood, as assessed by the language model.

- 2.10 Furthermore, let $MPPL_{\text{PRO}}(\text{PERSP}_t^i)$ and $MPPL_{\text{CON}}(\text{PERSP}_t^i)$ denote the mean conditional perplexity averaged over all pro claims, resp. con claims, conditional on the perspective of agent i at step t , i.e.

$$MPPL_{\text{PRO}}(\text{PERSP}_t^i) = \text{mean}_{c \in \text{PROCLAIMS}} (\text{PPL}(c | \text{PERSP}_t^i)), \quad (4)$$

and similarly for $MPPL_{\text{CON}}(\text{PERSP}_t^i)$. The opinion of agent a_i at step t is then given by

$$\mathbf{O}(\text{PERSP}_t^i) = \frac{MPPL_{\text{CON}}(\text{PERSP}_t^i)}{MPPL_{\text{PRO}}(\text{PERSP}_t^i) + MPPL_{\text{CON}}(\text{PERSP}_t^i)}. \quad (5)$$

- 2.11 Function \mathbf{O} measures the extent to which an agent leans towards the pro rather than the con side in a conversation, as defined by its central claims (recall: low perplexity \sim high probability). It is a polarity measure of an agent's opinion, and we alternatively refer to the opinion thus elicited as an agent's "polarity."
- 2.12 The mean perplexities ($MPPL_{\text{PRO}}(\text{PERSP}_t^i)$, $MPPL_{\text{CON}}(\text{PERSP}_t^i)$), however, reveal more than an agent's tendency towards the pro side or the con side in a conversation. If, e.g., both $MPPL_{\text{PRO}}(\text{PERSP}_t^i)$ and $MPPL_{\text{CON}}(\text{PERSP}_t^i)$ are very large, the agent's perspective is off-topic with respect to the central claims. We define an agent's *pertinence* as

$$\mathbf{P}(\text{PERSP}_t^i) = 0.5(MPPL_{\text{PRO}}(\text{PERSP}_t^i) + MPPL_{\text{CON}}(\text{PERSP}_t^i)). \quad (6)$$

Measure \mathbf{P} allows us to track whether agents change the topic in the course of a conversation.

Peer selection

- 2.13 We explore two peer selection procedures: a simple baseline, and a bounded confidence mechanism inspired by the work of Hegselmann & Krause (2002).
- 2.14 **Universal (baseline).** Every agent is a peer of every other agent at any point in time. Formally, $\text{PEERS}_t^i = \text{AGENTS}$ for all i, t .
- 2.15 **Bounded confidence.** An agent a_j is a peer of another agent a_i at some point in time if and only if their absolute difference in opinion is smaller than a given parameter ϵ . Formally, $\text{PEERS}_t^i = \{a_j \in \text{AGENTS} : |\text{OPIN}_{t-1}^i - \text{OPIN}_{t-1}^j| < \epsilon\}$.

Perspective updating

- 2.16 Agents update their perspectives in two steps (contraction, expansion), while posts that an agent has contributed in the previous time step are always added to her perspective:

Algorithm 2: Perspective Updating, Overview

def *perspective Updating*(i, t):

```

    retrieve the old perspective,  $\text{PERSP\_NEW} = \text{PERSP}_{t-1}^i$ ;
    randomly drop  $k$  posts from  $\text{PERSP\_NEW}$  according to how long each post has been included in the
    perspective;
    if agent  $i$  has been contributing at  $t-1$  then
        add the post generated by agent  $i$  at  $t-1$  to  $\text{PERSP\_NEW}$ ;
         $k = k-1$ 
    end
    add  $k$  further posts to  $\text{PERSP\_NEW}$  according to perspective_expansion_method;
    set new perspective,  $\text{PERSP}_t^i = \text{PERSP\_NEW}$ ;

```

- 2.17 The contracted perspective, PERSP_NEW , of agent a_i at step t is expanded with posts from the perspectives of all peers, avoiding duplicates, i.e., the agent selects – according to her specific updating method – k eligible posts, with $\text{POSTS}_{\text{el}} = \bigcup_{j \in \text{PEERS}_{t-1}^i} \text{PERSP}_{t-1}^j \setminus \text{PERSP_NEW} \subseteq \text{POSTS}_{<t}$. This kind of perspective expansion is governed by one of the following methods:
- 2.18 **Random (baseline).** Randomly choose and add k eligible posts ($\subseteq \text{POSTS}_{\text{el}}$) to the perspective.

- 2.19 Confirmation bias (lazy).** First, randomly draw k posts from POSTS_{e1} ; if all chosen posts confirm the agent's opinion (given the contracted perspective PERSP_NEW), then add the k posts; else, draw another k posts from POSTS_{e1} and add the k best-confirming ones from the entire sample (of size $2k$) to the perspective.
- 2.20 Homophily (ACTB).** Choose a peer $a_j \in \text{PEERS}_t^i$ in function of the similarity between the agent's and the peer's opinion; randomly choose k posts from the perspective of peer a_j , PERSP_{t-1}^j .
- 2.21** Note that *homophily (ACTB)*, which mimics the ACTB model by Mäs & Flache (2013), evaluates the eligible posts *ad hominem*, namely, based on the opinion of the corresponding peer only, while a post's semantic content is ignored. In contrast, *confirmation bias (lazy)*, which implements 'coherence-minded' updating from the model by Singer et al. (2019), only assesses the eligible posts' argumentative role, irrespective of who actually holds the post (see also Banisch & Shamon 2021). Moreover, we have implemented a "lazy" version of *confirmation bias*, as described above, for computational reasons: a confirmation-wise assessment of *all* eligible posts is practically not feasible.
- 2.22** A full and more precise description of the perspective expansion methods is given in Appendix D.

Text generation

- 2.23** Causal language models like GPT-2 are essentially probabilistic next-word prediction machines, as noted before. Given an input sequence of words $x_1 \dots x_k$, the language model predicts—for all words w_i in the vocabulary—the probability that w_i is the next word in the sequence, $\Pr(x_{k+1} = w_i | x_1 \dots x_k)$. It is obvious that such conditional probabilistic predictions can be used to generate a text word-by-word, and there exist various ways for doing so. This kind of text generation with statistical language models is commonly referred to as *decoding*, and it represents a research field in NLP in its own (c.f. Holtzman et al. 2019; Welleck et al. 2020). Pre-studies have suggested to use randomized beam search (with nucleus sampling) as decoding algorithm (see also Appendix C). The key parameters we use to control decoding (relying on default implementations by Wolf et al. (2019)) are
- *temperature*, which rescales the predicted probabilities over the vocabulary (increasing low and decreasing high probabilities if temperature is greater than 1);
 - *top_p*, which restricts the set of eligible words by truncating the rank-ordered vocabulary (let the vocabulary be sorted by decreasing probability, and let r be the greatest rank such that the probability that the next word will be w_1 or ... or w_r is still below *top_p*, then only $w_1 \dots w_r$ are eligible for being inserted).
- 2.24** In the experiments, we explore the following two decoding profiles:

profile	temperature	top_p
narrow	1.0	0.5
creative	1.4	0.95

- 2.25** Metaphorically speaking, the narrow profile emulates a conservative, narrow-minded author who's sticking with the most obvious, likely, common, and usual options when writing a text. The creative profile, in contrast, characterizes an author who is much more willing to take surprising turns, to use unlikely phrases and unexpected sentences, who is easily carried away, prone to digress, and much harder to predict.
- 2.26** Pre-studies show that conversations are extremely noisy if each agent generates and submits a post at every time step; in the following, the probability that an agent is contributing a novel post at a given time step is set to 0.2.

Experiments

Initialisation

- 3.1** To run simulations with our natural-language ABMA, the initial perspectives of the agents ($\text{PERSP}_0^i, i = 1 \dots n$) have to contain meaningful posts that fall within the conversation's topic. Additionally, it seems desirable that the initial perspectives give rise, group-wise, to a sufficiently broad initial opinion spectrum.

- 3.2** To meet these requirements, we define topics that correspond to specific online debates on the debating platform kialo.com, from where we crawl and post-process potential posts. Post-processing involves filtering (maximum length equals 70 words) and conclusion explication. As we crawl posts from a nested pro-con hierarchy, the argumentative relation of each post to the central pro / con claims (root) can be inferred, which allows us to add, to each post, an appropriate pro / con claim as concluding statement. For example, the post "*If drugs being illegal prevented addiction, there would be no drug addicted person. Thus, there is no prevention by just keeping drugs illegal.*" is expanded by "*So, legalization of drugs is a pretty good idea.*" In order to increase diversity of posts, we expand only half of all the posts retrieved by such a conclusion statement.
- 3.3** The experiments described below are run on the topic of the legalization of drugs, initial perspectives are sampled from 660 posts, of which 442 justify or defend the pro claim (legalization).

Scenarios

- 3.4** We organize our experiments along two main dimensions, namely (i) peer & perspective updating, and (ii) agent type.
- 3.5** Regarding peer & perspective **updating**, we explore four parameter combinations (see Subsections Peer selection and Perspective updating):
- *random*: baseline update rules for peers (universal) and perspective (random);
 - *bounded confidence*: bounded confidence peer selection and random perspective updating;
 - *confirmation bias*: universal peers and lazy confirmation bias (for perspective updating);
 - *homophily*: universal peers and homophily (for perspective updating).
- 3.6** Regarding **agent type**, we distinguish passive, and two types of active (i.e., generating) agents (see Subsection Text generation):
- *listening*: agents are not generating, they only forget, share and adopt posts that have been initially provided;
 - *generating narrow*: agents can generate posts, text generation is controlled by the narrow decoding profile;
 - *generating creative*: agents can generate posts, text generation is controlled by the creative decoding profile.
- 3.7** So, all in all, the simulations are grouped in 4×3 scenarios. For each scenario, we run an *ensemble* of 150 individual simulations. (The ensemble results are published and can be inspected with this online app.)

Results

An illustrative case

- 4.1** In this subsection, we present an illustrative simulation run and follow a single ADA during a brief episode of the conversation. The case study is not intended to be representative. Its purpose is two-fold: (i) to illustrate what exactly is going on in the simulations, and (ii) to demonstrate that the model is not just producing non-sense, by showing that we can interpret the ADA's opinion trajectory as an episode of reasonable belief revision.
- 4.2** Figure 2 plots the opinion trajectories (polarity) of 20 agents over the entire simulation run. Agents select peers in accordance with the bounded confidence mechanism ($\epsilon = 0.04$). After the initialisation phase (steps 0–4), the collective opinion spectrum ranges from 0.4 to 0.8. Opinion diversity stays high for the next 20 steps, after which the collective profile starts to collapse and more and more agents settle on an opinion around 0.55. From step 65 onwards, a noisy equilibrium state has seemingly been reached.
- 4.3** Figure 2 highlights the opinion trajectory of agent a_8 . One of its outstanding features is that agent a_8 holds an initial perspective that induces a strong opinion pro legalization ($\text{OPIN}_5^8 = 0.8$). In steps 12–20, however, agent

α_8 completely reverses her opinion. We will now try to make sense of this drastic opinion reversal in terms of changes in the agent's perspective. We limit our discussion to steps 17–19, during which the opinion falls from 0.64 to 0.54 and further to 0.49 (see also Figure 2b).

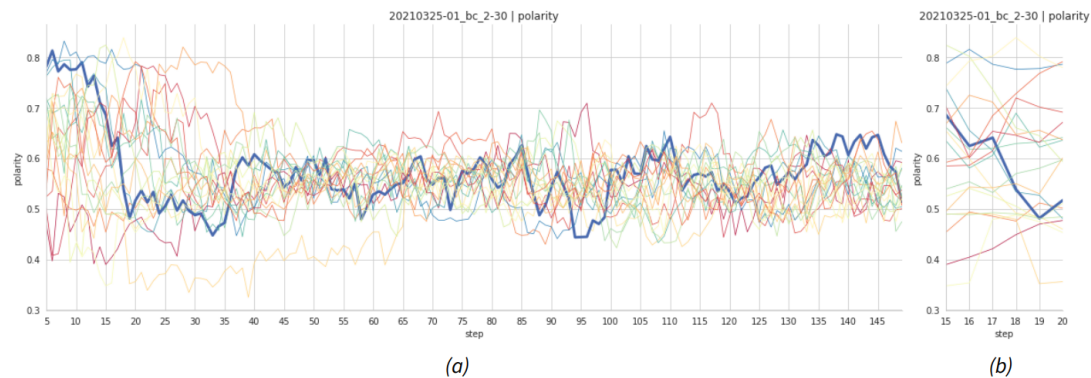


Figure 2: Opinion dynamics (polarity) in an illustrative simulation run (20 agents, context size 8, *bounded confidence / generating creative scenario*), (a): full range, (b): zoom into steps 15-20.

- 4.4 In the following display of the agent's perspectives (starting with step 17), posts are highlighted according to whether they **have been generated by the agent**, **are newly added to the perspective at this step**, **have been generated by the agent at the previous step and are hence newly added**, or **will be removed from the perspective at the next step**.
- 4.5 At step 17, the agent's perspective contains a majority of posts plainly pro legalization (17-1, 17-3, 17-5, 17-6, 17-7, 17-8) and no clear reasons against legalization (and that's how it should be, reckoning that the agent comes from the extreme pro side and has – thanks to bounded confidence – exchanged posts only with like-minded, pro-legalization peers):
- 4.6 [17-1] Professional addiction treatment is usually [expensive](<https://www.addictioncenter.com/rehab-questions/cost-of-drug-and-alcohol-treatment/>). I believe all drugs should be legal. [17-2] The term 'increase public health' is subjective. What does that mean? [17-3] Marijuana use in the Netherlands has [not increased](<https://www.opensocietyfoundations.org/voices/safe-and-effective-drug-policy-look-dutch/>) following decriminalisation; in fact, cannabis consumption is lower compared to countries with stricter legislation such as the UK. [17-4] **It might be a good idea to limit the sale of drugs to adults over the age of 18, and to state clearly that the possession and use of alcohol and cannabis by minors is prohibited.** [17-5] Legalising drugs related to date rape could bring the issue further into the public eye, allowing for more widespread education on the topic. [17-6] **The current system is not working. It's absurd to lock people up for using drugs that they choose to make themselves. If they wanted to get high, they'd do it somewhere else.** [17-7] If someone wants to go to the supermarket and pick up a few cakes, why shouldn't they? Why shouldn't they be allowed to do so? [17-8] **People should be able to ingest anything they want without getting in any trouble for it.**
- 4.7 The newly added post 17-8 represents a reason pro legalization, which might explain the slight increase of polarity compared to to step 16. Marked for removal in the next step 18 are: 17-1, an explicit pro reason, and 17-4, a rather nuanced statement which advocates a differentiated policy. Here is how these posts are replaced (cf. 18-7 and 18-8):
- 4.8 [18-1] The term 'increase public health' is subjective. What does that mean? [18-2] Marijuana use in the Netherlands has [not increased](<https://www.opensocietyfoundations.org/voices/safe-and-effective-drug-policy-look-dutch/>) following decriminalisation; in fact, cannabis consumption is lower compared to countries with stricter legislation such as the UK. [18-3] **Legalising drugs related to date rape could bring the issue further into the public eye, allowing for more widespread education on the topic.** [18-4] **The current system is not working. It's absurd to lock people up for using drugs that they choose to make themselves. If they wanted to get high, they'd do it somewhere else.** [18-5] If someone wants to go to the supermarket and pick up a few cakes, why shouldn't they? Why shouldn't they be allowed to do so? [18-6] **People should be able to ingest anything they want without getting in any trouble for it.** [18-7] **When you legalize drugs, you're going to have a lot of people who have personal vendettas against certain substances. In this case, the vendettas will probably manifest themselves into violent crime.** [18-8]

- 4.9 The post 18-7, which has just been *generated* by the agent, paints a gloomy (despite somewhat awkward) picture and predicts bad consequences of the legalization of drugs. Post 18-8, which had been previously submitted by agent a_8 and then forgotten, is now taken from another peer's perspective and re-adopted by agent a_8 . It coincidentally picks up the crime trope from 18-7, claiming that a large proportion of prison inmates have committed drug-related crimes. While 18-8 is, per se, an argumentatively ambivalent statement which can be used to argue both for and against legalization, its main effect, in this particular context, is apparently to amplify the gloomy outlook cast in preceding 18-7; it hence further strengthens the case against legalization. Given this change in perspective from step 17 to step 18, it makes perfectly sense that the agent's opinion has shifted towards the con side.
- 4.10 Moreover, note that two clear-cut reasons pro legalization are marked for removal (18-3, 18-6), which paves the way for further opinion change towards the con-side.
- 4.11 [19-1] The term 'increase public health' is subjective. What does that mean? [19-2] Marijuana use in the Netherlands has [not increased](<https://www.opensocietyfoundations.org/voices/safe-and-effective-drug-policy-look-dutch/>) following decriminalisation; in fact, cannabis consumption is lower compared to countries with stricter legislation such as the UK. [19-3] **The current system is not working. It's absurd to lock people up for using drugs that they choose to make themselves. If they wanted to get high, they'd do it somewhere else.** [19-4] If someone wants to go to the supermarket and pick up a few cakes, why shouldn't they? Why shouldn't they be allowed to do so? [19-5] **When you legalize drugs, you're going to have a lot of people who have personal vendettas against certain substances. In this case, the vendettas will probably manifest themselves into violent crime.** [19-6] **According to the Department of Justice, 75% of the federal prison population is serving time for nonviolent drug crimes. Nearly 90% of inmates in federal prisons are there for drug crimes.** [19-7] Cocaine is [highly addictive](https://en.wikipedia.org/wiki/Cocaine_dependence) and easy to become dependent on. I believe legalization of drugs is a really bad idea. [19-8] It's very easy to overdose on psychoactive substances. It's very difficult to overdose on non-psychoactive substances.
- 4.12 The perspective in step 19 newly embraces two posts against legalization, adopted from peers. Post 19-7, in particular, is an explicit con reason, post 19-8 draws attention towards overdosing and hence towards the negative effects of drug use. So, four posts in the perspective now speak against legalization – compared to 6 pro reasons and no con reason in step 17. Plus, the four con reasons are also the most recent posts (recall that order matters when prompting a language model) and, in a sense, "overwrite" the three previously stated pro claims (19-2 to 19-4). In sum, this explains the sharp opinion change from step 17 to step 19.

Global consensus and polarization effects

- 4.13 In this subsection, we characterize and compare the simulated opinion dynamics across our 12 experiments (see Subsection Scenarios), and provide results averaged over the corresponding simulation ensembles (variance within each ensemble is documented by Figures 6–8 in the Data Appendix). In line with this paper's overall argument, which pertains to the robustness of symbolic models of opinion dynamics, we will in particular compare passive and active (i.e., generating) agents. We will also discuss the result's plausibility, but let me add as caveat upfront: a comprehensive *explanation* of the findings requires a closer analysis which would go beyond the scope of this paper.
- 4.14 Based on a cluster analysis (see Appendix E for details), we measure the degree of polarization (in terms of clustering coverage and number of clusters), the frequency of bipolarization, and the frequency of full consensus in the simulated conversations. Moreover, we report opinion variance and min-max spread as divergence measures, plus average squared opinion difference – $(\text{OPIN}_t - \text{OPIN}_{t-1})^2$ – as volatility measure. Conversations are evaluated at $t = 150$.
- 4.15 Figure 3 shows how clustering metrics vary across the 4×3 scenarios. Consider, firstly, **passive agents** who are only sharing but not generating novel posts (column *listening*). With the baseline update mechanism (row *random*), 6% of the agents fall within an opinion cluster (Figure 3a), and there exists, on average, one cluster in one out of three conversations (Figure 3b). Clustering is much more pronounced for the alternative update mechanisms, with *homophily* in particular attaining more than 70% coverage and 1.65 clusters per conversation. Let us say that a conversation is in a state of *bipolarization* (*full consensus*) if and only if clustering coverage is greater than 0.9 and there exist two clusters (exists a single cluster). We observe, accordingly, no instances of bipolarization or consensus in the baseline scenario, very few instances for *bounded confidence* and *confirmation bias*, and a significant frequency of bipolarization and full consensus for *homophily* (Figure 3c,d). Consistent with the clustering coverage reported in Figure 3a, the absence of bipolarization and full consensus in the case of

bounded confidence and *confirmation bias* updating points to the fact that clustering effects, in these scenarios, are not all-encompassing and a substantial number of agents remain outside the emerging opinion groups.

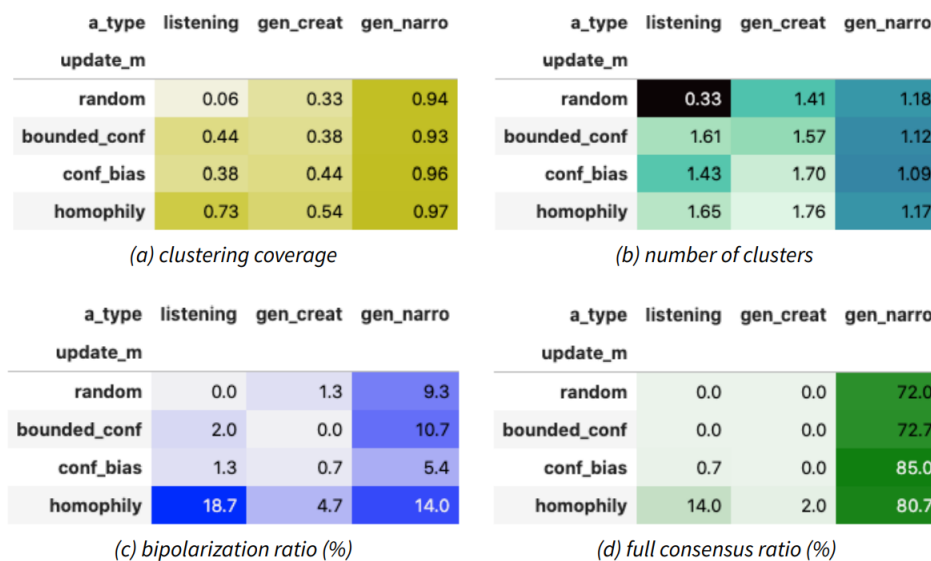


Figure 3: Clustering metrics built on automatic density-based cluster detection: clustering coverage (i.e., ratio of opinions that belong to a detected cluster), number of detected clusters, frequency of bipolarization (i.e., ratio of cases with two clusters and clustering coverage > 0.9, in percent), frequency of full consensus (i.e., ratio of cases with one cluster and clustering coverage > 0.9, in percent). See also Figure 6.

4.16 This global picture changes entirely as we turn, secondly, to **active agents** that are generating posts in line with one of the two decoding profiles (cf. Subsection Text generation). Regarding creative authors (column *gen_creat*) and irrespective of the particular update mechanism, clustering coverage lies between 0.3 and 0.6, there exist approximately 1.5–2 clusters, we observe bipolarization in up to 5% and full consensus in less than 2% of all conversations. So we find, compared to passive agents, much stronger clustering in the baseline scenario but significantly less clustering in the *homophily* scenario. Both results seem plausible: On the one hand, generation introduces a potential positive feedback, as agents may construct (viral) posts that fit to their perspective and are adopted by peers, which in turn further aligns the peers' perspectives; on the other hand, clusters of like-minded agents risk to fall apart as long as group members "think out of the box," and creatively produce posts that may run counter to the group consensus – especially so if epistemic trust is a function of the agents' opinions (as in *bounded confidence* and *homophily*) and perspective updating does not assess each candidate post individually before adopting it. Regarding narrow-minded authors (column *gen_narro*), however, clustering coverage is greater than 0.9, there exists roughly a single cluster per conversation, bipolarization is frequent, and more than 70% of the debates reach full consensus. In this case, there is no thinking "out of the box", no disruption, and the positive feedback introduced by generation seems to be the dominating effect irrespective of the particular update mechanism.

4.17 Figure 4 describes the extent to which opinions diverge in the 12 simulation experiments. Concerning **passive agents**, disagreement is most pronounced (both in terms of opinion variance and max–min spread) with *bounded confidence* updating, closely followed by the baseline scenario. We may also note that *confirmation bias*, which hardly results in bipolarization or full consensus (Figure 3c,d), brings about the strongest overall agreement nonetheless. Picking posts based on their coherence with one's perspective, and independently of who has produced or is holding them, seems to allow agents to foster mutual agreement without fully endorsing each other's perspectives, i.e., without polarization. Conversations with **active agents**, in contrast, give rise to much lower levels of disagreement, while narrow-minded authoring is even more agreement-conducive than creative generation, reducing divergence by an entire order of magnitude compared to passive agents. Again, this might be due to the positive feedback introduced with active generation.

a_type	listening	gen_creat	gen_narro	a_type	listening	gen_creat	gen_narro
update_m				update_m			
random	0.0105	0.0019	0.0003	random	0.35	0.16	0.04
bounded_conf	0.0110	0.0024	0.0014	bounded_conf	0.41	0.18	0.09
conf_bias	0.0042	0.0015	0.0002	conf_bias	0.21	0.14	0.03
homophily	0.0077	0.0016	0.0017	homophily	0.24	0.15	0.05

(a) opinion variance

(b) max-min spread

Figure 4: Divergence metrics: opinion variance, max-min opinion spread. See also Figure 7

- 4.18 As Figure 5a shows, there exist huge differences in terms of per-agent opinion volatility, which is most pronounced for **passive agents** that randomly adopt novel posts. Structured updating procedures (*bounded confidence*, *confirmation bias*, and *homophily*) have, ceteris paribus, a stabilizing effect and significantly reduce volatility. **Creative generation** produces mixed effects on volatility (depending on the update procedure), while narrow-minded agents possess maximally stable opinions (consistent with a strong positive feedback of narrow generation).
- 4.19 Finally, Figure 5b reports mean pertinence values for the different scenarios. Recall that pertinence measures the relevance of a perspective for a given pair of central pro/con claims (cf. Subsection Opinion elicitation): the lower the pertinence value, the more relevant the perspective. Accordingly, agents retain the most relevant perspectives in the baseline scenario. As soon as agents start to generate their own posts, pertinence value increases. However, the conversations stay, on average, faithful to the initial topic (that wouldn't be the case for perplexities above 20, though). Mean pertinence value is, somewhat surprisingly, even slightly lower for creative than for narrow-minded agents.

a_type	listening	gen_creat	gen_narro	a_type	listening	gen_creat	gen_narro
update_m				update_m			
random	0.0051	0.0017	0.0002	random	4.48	6.81	7.65
bounded_conf	0.0007	0.0011	0.0001	bounded_conf	5.83	6.97	7.61
conf_bias	0.0014	0.0014	0.0001	conf_bias	5.71	6.84	7.71
homophily	0.0004	0.0008	0.0001	homophily	6.10	7.08	7.74

(a) volatility

(b) pertinence

Figure 5: Per agent opinion volatility and pertinence. See also Figure 8.

Sensitivity analysis

- 4.20 By varying the number of agents per conversation, the maximum number of posts an agent can hold in her perspective, as well as update-specific parameters (epsilon interval, homophily strength) in the simulation experiments, we obtain a preliminary understanding of the model's sensitivity. Yet, these experiments fall short of a systematic exploration of the entire parameter space, which has not been carried out due to its computational costs, and is certainly a desideratum for future research.
- 4.21 In general, the model seems to yield *qualitatively* similar results when varying key parameters: In particular, structured updating (*bounded confidence*, *confirmation bias*, and *homophily*) with passive agents gives rise to polarization, and once agents get active and start to generate posts, the collective opinion evolution is dominated by decoding parameters – in the latter case, changes in community or perspective size have quantitatively very little effect.
- 4.22 Regarding *homophily* and *confirmation bias* updating with passive agents, increasing the number of agents per conversation results in more full consensus and less bipolarization. With more agents covering the ground, it seems to be more difficult for subgroups to isolate from other agents and to build a shared, sufficiently distinct perspective. Moreover, increasing the perspective size decreases the frequency of both full consensus and bipolarization, and weakens clustering in general. With more posts in a perspective it takes – ceteris paribus – longer for an agent to entirely change her perspective; characteristic collective opinion profiles hence build up

more slowly and we observe, at a fixed time step, lower polarization and consensus. Plus, it's clear from a look at the time-series that the conversations with *homophily* and *confirmation bias* have not reached a (possibly noisy) equilibrium state at $t = 150$, yet. It's therefore a desideratum for future research to run the simulations for much longer time spans.

- 4.23** Furthermore, the parameters epsilon and homophily exponent specifically control *bounded confidence*, respectively *homophily* updating. The model reacts to changes in these parameters as to be expected: As we increase the epsilon interval, we obtain (with *bounded confidence* updating / passive agents) more clustering and more agreement. Increasing the homophily exponent results (with *homophily* updating / passive agents) in stronger clustering (more consensus, more bipolarization) and greater disagreement.

● Discussion and Future Work

Structural robustness of symbolic models

- 5.1** As regards passive agents, our natural-language ABMA reproduces qualitative results obtained with symbolic reason-balancing models: *homophily* and *confirmation bias* updating lead to bipolarization, in line with the findings of Mäs & Flache (2013) and Singer et al. (2019). *Bounded confidence* updating increases polarization and disagreement, consistent with Hegselmann & Krause (2002). Due to requirements imposed by language modeling technology, the natural-language ABMA is structurally similar to, albeit not identical with the corresponding symbolic models (e.g., *confirmation bias* implements, for computational reasons, local rather than global search). In addition, the context sensitive, holistic processing of reasons in the natural-language ABMA departs from the strictly monotonic and additive reason aggregation mechanism built into the symbolic models. All these structural dissimilarities, however, further strengthen the robustness of the findings concerning passive agents.

Limitations of symbolic models

- 5.2** We have observed that, once agents start to generate and submit their own posts, the evolution of the collective opinion profile is dominated by decoding parameters (i.e., properties of the agents as authors). With active agents, we obtain entirely different results for polarization, consensus and divergence than in the experiments with passive agents. In symbolic reason balancing models, however, agents cannot generate new reasons (or rephrase, summarize, mix, and merge previous ones). So, the natural-language ABMA identifies a potentially pivotal mechanism that's currently ignored by symbolic models, whose explanatory and predictive scope seem, accordingly, to be limited to conversations and collective deliberations with a fixed set of reasons to share.

Sensitivity analysis

- 5.3** A systematic sensitivity analysis of the natural-language model seems urgent and should go beyond an exploration of the entire parameter space and longer simulation runs. First, some implementation details are worth varying (e.g., the generation of prompts used to query the language model, the post-processing of generated posts, the functional form of the confirmation measure, local search). Second, the initial conditions should be altered, too; in particular, conversations should be simulated on (and be initialized with) different topics.

Empirical applications

- 5.4** There are multiple routes for empirically applying and testing natural-language ABMAs, which are closed for purely symbolic models and which might be explored in future work – as the following suggestions are supposed to illustrate. On the one hand, one can derive and empirically check *macro* predictions of the model: (a) One may test whether groups of conservative authors are more likely to reach full consensus than groups of creative authors. (b) One might try to explain statistical properties of an observed opinion distribution in a debate by initializing the model with posts from that debate and running an entire (perturbed-physics style) ensemble of simulations. (c) Or one might check whether the macro patterns of semantic similarity (Reimers & Gurevych 2019) within a simulated conversation correspond to those in empirical discourse. On the other hand, one can

test the *micro* dynamics built into the natural language model: (a) One might verify whether deliberating persons respond to reasons and aggregate reasons in the way the ADA does. (b) Alternatively, one might try to infer – by means of the natural-language ABMA – (unobserved, evolving) agent perspectives from (observed) agent contributions so as to account for the agents' (observed) final verdicts on a topic. (c) Or one may trace how frequently posts "go viral" in the simulations and what their specific influence on the overall opinion dynamics is.

Model extensions

- 5.5 The natural-language ABMA is extremely flexible, as its agents (ADAs) understand and speak English. This allows us to address further linguistic phenomena (slurs, thick concepts) and cognitive phenomena (fallacious reasoning, framing effects) with relative ease, e.g., by systematically changing the prompts used to query the agents, by intervening in a simulated conversation and inserting targeted posts at a given time step, or by controlling for these phenomena during the initialisation. Likewise, taking opinion pertinence (in addition to opinion polarity) into account in the updating process, eliciting multi-dimensional opinions (with additional – possibly even endogenously generated – pairs of pro-con claims), and mixing multiple topics in one and the same conversation are further straight-forward and easy-to-implement extensions of the model. Obviously, it's also possible to define a neighborhood relation and simulate conversations on social networks. A further set of model extensions regards the heterogeneity of agents: As the model presented in this paper contains (except for their initial condition) identical agents, a first way to increase diversity is to allow for agent-specific (updating and decoding) parameters. Furthermore, we can model background beliefs by fixing immutable posts in an agent's perspective. Finally, there's no reason (besides a computational, practical one) to use one and the same language model to power ADAs; in principle, each agent might be simulated by a specific instance of a language model (with particular properties due to its size, pre-training and fine-tuning) – plus, these language models might actually be trained and hence *evolve* in the course of a simulated conversation. The practical effect of this last modification is that agents would display different initial (empty perspective) positions and that an agent might have different opinions at two points in time although she holds one and the same perspective.

Lessons for AI

- 5.6 This paper has adopted recent technology from *AI and NLP* to advance computational models of argumentative opinion dynamics in the fields of *formal and social epistemology* and *computational social science*. Now, this *might* in turn have repercussions for AI: The fact that we've been able to simulate self-sustainable rational argumentative opinion dynamics suggests that the language model we're using to power agents possesses minimal argumentative capabilities and is, in particular, able to process and respond to reasons in a sensible way. Otherwise, the successful simulation of collective deliberation would be quite miraculous and unexplainable. Plus, our experiments can be interpreted – inside out – as a single agent's attempt to think through a topic by consecutively adopting alternative perspectives (and hence mimicking a deliberation); which suggests that language models are capable of sensible self-talk, consistent with Shwartz et al. (2020) and Betz et al. (2021). Finally, such *argumentative* multi-agent systems might be a fruitful design pattern to address tasks in AI and NLP that are difficult to solve with standard systems built around a single agent / a single language model.

● Model Documentation

The source code (Python) for running the simulations is published at <https://github.com/debateLab/ada-simulation>. The simulation results used in this paper are published in a separate repository, which includes a link to an interactive online data-explorer app. The repository can be found at: <https://github.com/debateLab/ada-inspect>.

● Technical Appendix

Appendix A: Language model

In opinion elicitation (Subsection Opinion elicitation) and text generation (Subsection Text generation), we rely on the pretrained autoregressive language model GPT-2 (Radford et al. 2019) as implemented in the *Transformers* Python package by Wolf et al. (2019).

Appendix B: Prompt generation

Let $\text{PERSP}_t^i = \langle p_1 \dots p_k \rangle$ be the perspective of agent a_i at t . To elicit the opinion of agent a_i at $t + 1$ (Subsection Opinion elicitation), we prompt the language model with the following query:

Let's discuss legalization of drugs!

p_1

...

p_k

I more or less agree with what my peers are saying here. And therefore, all in all,

When generating a new post at $t + 1$ (Subsection Text generation), the model is prompted with

Let's discuss legalization of drugs!

p_1

...

p_k

I more or less agree with what my peers are saying here. Regarding the legalization of drugs, I'd just add the following thought:

Appendix C: Parameters

Global parameters of the simulation runs are:

number of agents	20
perspective size	8
maximum steps	150
relevance deprecation	.9
memory loss (passive)	1
memory loss (active)	2
homophily bias exponent	50
epsilon	0.04

Parameters that specifically control decoding are:

number of beams	5
repetition penalty	1.2
sampling	True

Appendix D: Perspective Updating Methods

With **homophily** updating, agent a_i chooses a peer $a_j \in \text{PEERS}^i$ (we drop time indices for convenience) from whom new posts are adopted in function of the similarity in opinion,

$$\text{sim}(i, j) = 1 - |\text{OPIN}^i - \text{OPIN}^j| \quad (7)$$

The weight agent a_i assigns to peer a_j in randomly choosing her communication partner is further determined by the homophily exponent, hpe:

$$\text{weight}(i, j) = \text{sim}(i, j)^{\text{hpe}} \quad (8)$$

With **confirmation bias** updating, agent a_i evaluates eligible posts in terms of their argumentative function. This is modeled by one-sided relevance confirmation, which measures the degree to which a post p confirms the opinion which corresponds to a given perspective **PERSP** for an agent a_i at step t :

$$\text{conf}(p) = \begin{cases} |\mathbf{O}(\text{PERSP} + p) - \text{OPIN}_0^i| & \text{if } (\mathbf{O}(\text{PERSP} + p) > \text{OPIN}_0^i) \Leftrightarrow (\text{OPIN}_{t-1}^i > \text{OPIN}_0^i) \\ 0 & \text{otherwise.} \end{cases} \quad (9)$$

Appendix E: Cluster Analysis

We carry out the cluster analysis with the help of density based clustering (Ester et al. 1996) as implemented in the Python package SciKit learn (setting $\text{eps}=0.03$ and $\text{min_samples}=3$). As the opinion trajectories are – depending on the experiment – very noisy, a clustering algorithm risks to detect merely coincidental clusters that have emerged by chance if it is applied to single data points. In order to identify stable clusters, we therefore apply the clustering algorithm to short opinion trajectories; more specifically, we cluster opinion triples $\langle \text{OPIN}_{t-2}^i, \text{OPIN}_{t-1}^i, \text{OPIN}_t^i \rangle$.

● Data Appendix

Histograms in Figures 6–8 show how the key metrics reported in the main text vary across the individual simulations runs in each of the 12 different ensembles ($N=150$).

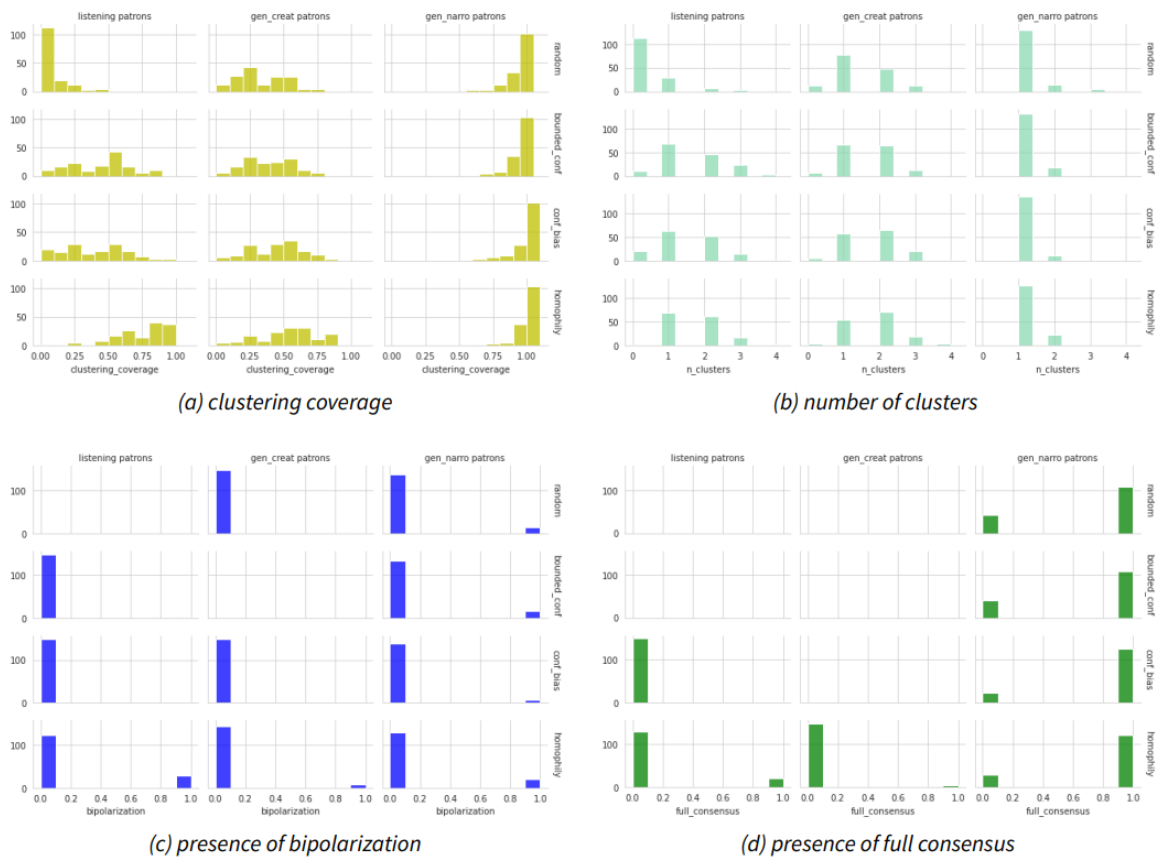


Figure 6: Histograms corresponding to Figure 3 and showing the distribution of clustering metrics within the 12 different ensembles: clustering coverage (i.e., ratio of opinions that belong to a detected cluster), number of detected clusters, presence of bipolarization (i.e., cases with two clusters and clustering coverage > 0.9, in percent), presence of full consensus (i.e., cases with one cluster and clustering coverage > 0.9, in percent).

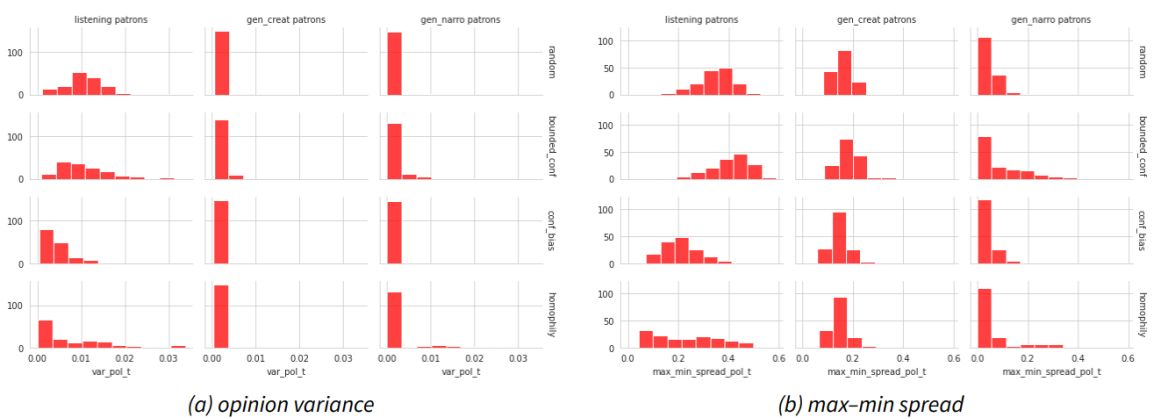


Figure 7: Histograms corresponding to Figure 4 and showing the distribution of divergence metrics within the 12 different ensembles: opinion variance, max-min opinion spread.




Figure 8: Histograms corresponding to Figure 5 and showing the distribution of per agent opinion volatility and pertinence within the 12 different ensembles.

References

- Adiwardana, D., Luong, M. T., So, D., Hall, J., Fiedel, N., Thoppilan, R., Yang, Z., Kulshreshtha, A., Nemade, G., Lu, Y. & Le, Q. V. (2020). Towards a human-like open-domain chatbot. arXiv preprint. Available at: <https://arxiv.org/abs/2001.09977>
- Askell, A. (2020). GPT-3: Towards renaissance models. Daily Nous Blog: Philosophers On GPT-3. Available at: <http://dailynous.com/2020/07/30/philosophers-gpt-3/#askell>
- Axelrod, R. (1997). The dissemination of culture: A model with local convergence and global polarization. *Journal of Conflict Resolution*, 41(2), 203–226
- Banisch, S. & Olbrich, E. (2021). An argument communication model of polarization and ideological alignment. *Journal of Artificial Societies and Social Simulation*, 24(1), 1
- Banisch, S. & Shamon, H. (2021). Biased processing and opinion polarisation: Experimental refinement of argument communication theory in the context of the energy debate. Available at SSRN 3895117, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3895117
- Bao, S., He, H., Wang, F. & Wu, H. (2020). PLATO: Pre-trained Dialogue Generation Model with Discrete Latent Variable. ACL
- Beisbart, C., Betz, G. & Brun, G. (2021). Making reflective equilibrium precise. A formal model. Forthcoming, Ergo
- Betz, G. (2012). *Debate Dynamics: How Controversy Improves Our Beliefs*. Dordrecht: Springer
- Betz, G. (2015). Truth in evidence and truth in arguments without logical omniscience. *The British Journal for the Philosophy of Science*, 67(4)
- Betz, G., Richardson, K. & Voigt, C. (2021). Thinking aloud: Dynamic context generation improves zero-shot reasoning performance of GPT-2. arXiv preprint. Available at: <https://arxiv.org/pdf/2103.13033.pdf>
- Blodgett, S. L., Barocas, S., Daumé III, H. & Wallach, H. (2020). Language (technology) is power: A critical survey of "Bias" in NLP. Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, 5454-5476
- Borg, A., Frey, D., Šešelj, D. & Straßer, C. (2017). An argumentative agent-based model of scientific inquiry. In S. Benferhat, K. Tabia & M. Ali (Eds.), *Advances in Artificial Intelligence: From Theory to Practice*, (pp. 507–510). Springer Verlag
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I. & Amodei, D. (2020). Language models are few-shot learners. arXiv preprint. Available at: <https://arxiv.org/abs/2005.14165>

- Dale, R. (2021). GPT-3: What's it good for? *Natural Language Engineering*, 27(1), 113–118
- Devlin, J., Chang, M. W., Lee, K. & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1
- Ester, M., Kriegel, H. P., Sander, J. & Xu, X. (1996). A density-based algorithm for discovering clusters in large spatial databases with noise. *KDD-96 Proceedings*, 96(34), 226–231
- Grüne-Yanoff, T. (2016). Framing. In S. O. Hansson & G. Hirsch-Hadorn (Eds.), *The Argumentative Turn in Policy Analysis. Reasoning about Uncertainty*, (pp. 189–215). Cham: Springer
- Gururangan, S., Swayamdipta, S., Levy, O., Schwartz, R., Bowman, S. R. & Smith, N. A. (2019). Annotation artifacts in natural language inference data. arXiv preprint. Available at: <https://arxiv.org/pdf/1803.02324.pdf>
- Hahn, U. & Oaksford, M. (2007). The rationality of informal argumentation: A Bayesian approach to reasoning fallacies. *Psychological review*, 114(3), 704
- Hegselmann, R. & Krause, U. (2002). Opinion dynamics and bounded confidence: Models, analysis and simulation. *Journal of Artificial Societies and Social Simulation*, 5(3), 2
- Hesse, M. (1988). The cognitive claims of metaphor. *The Journal of Speculative Philosophy*, 2(1), 1–16
- Holtzman, A., Buys, J., Du, L., Forbes, M. & Choi, Y. (2019). The curious case of neural text degeneration. International Conference on Learning Representations
- Kassner, N. & Schütze, H. (2020). Negated and misprimed probes for pretrained language models: Birds can talk, but cannot fly. Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, 7811-7818
- Manning, C. D. & Schütze, H. (1999). *Foundations of Statistical Natural Language Processing*. Cambridge, MA: MIT Press
- Mäs, M. & Flache, A. (2013). Differentiation without distancing. Explaining bi-polarization of opinions without negative influence. *PLoS ONE*, 8(11), 1–17
- Nadeem, M., Bethke, A. & Reddy, S. (2020). StereoSet: Measuring stereotypical bias in pretrained language models. arXiv preprint. Available at: <https://arxiv.org/abs/2004.09456>
- Olsson, E. J. (2013). A Bayesian simulation model of group deliberation and polarization. In F. Zenker (Ed.), *Bayesian Argumentation*, (pp. 113–133). Berlin Heidelberg: Springer
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D. & Sutskever, I. (2019). Language models are unsupervised multitask learners. Preprint. Available at: <https://d4mucfpxsywv.cloudfront.net/better-language-models/language-models.pdf>
- Rappaport, J. (2019). Communicating with slurs. *The Philosophical Quarterly*, 69(277), 795–816
- Reimers, N. & Gurevych, I. (2019). Sentence-BERT: Sentence embeddings using Siamese BERT-Networks. Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing
- Russell, S. J. & Norvig, P. (2016). *Artificial Intelligence: A Modern Approach*. Harlow: Pearson
- Sešelj, D. & Straßer, C. (2013). Abstract argumentation and explanation applied to scientific debates. *Synthese*, 190(12), 2195–2217
- Shwartz, V., West, P., Le Bras, R., Bhagavatula, C. & Choi, Y. (2020). Unsupervised commonsense question answering with self-talk. In B. Webber, T. Cohn, Y. He & Y. Liu (Eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, (pp. 4615–4629). Stroudsburg, PA: Association for Computational Linguistics
- Simon, H. A. (1982). *Models of Bounded Rationality*. Cambridge, MA: MIT Press
- Singer, D. J., Bramson, A., Grim, P., Holman, B., Jung, J., Kovaka, K., Ranginani, A. & Berger, W. J. (2019). Rational social and political polarization. *Philosophical Studies*, 176(9), 2243–2267

- Suppes, P. (1970). Probabilistic inference and the concept of total evidence. In P. Suppes & J. Hintikka (Eds.), *Information and Inference*, (pp. 49–65). Dordrecht: Reidel
- Taillandier, P., Salliou, N. & Thomopoulos, R. (2021). Introducing the argumentation framework within agent-based models to better simulate agents' cognition in opinion dynamics: Application to vegetarian diet diffusion. *Journal of Artificial Societies and Social Simulation*, 24(2), 6
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. & Polosukhin, I. (2017). Attention is all you need. arXiv preprint. Available at: <https://arxiv.org/abs/1706.03762>
- Walton, D. & Gordon, T. F. (2019). How computational tools can help rhetoric and informal logic with argument invention. *Argumentation*, 33(2), 269–295
- Welleck, S., Kulikov, I., Roller, S., Dinan, E., Cho, K. & Weston, J. (2020). Neural text generation with unlikelihood training. arXiv preprint. Available at: <https://arxiv.org/abs/1908.04319>
- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M. et al. (2019). HuggingFace's transformers: State-of-the-art natural language processing. arXiv preprint. Available at: <https://arxiv.org/abs/1910.03771>
- Zhang, Y., Sun, S., Galley, M., Chen, Y. C., Brockett, C., Gao, X., Gao, J., Liu, J. & Dolan, W. (2020). DialoGPT: Large-scale generative pre-training for conversational response generation. ACL