

More Victories, Less Cooperation: Assessing Cicero’s Diplomacy Play

Wichayaporn Wongkamjan¹ Feng Gu¹ Yanze Wang⁴ Ulf Hermjakob⁴
Jonathan May⁴ Brandon M. Stewart² Jonathan K. Kummerfeld³
Denis Peskoff² Jordan Lee Boyd-Graber¹

¹University of Maryland ²Princeton University ³University of Sydney

⁴Information Sciences Institute, University of Southern California

{wwongkam, fgu1}@umd.edu {yanzewan, ulf, jonmay}@isi.edu

bms4@princeton.edu jonathan.kummerfeld@sydney.edu.au

dp2896@princeton.edu jbg@umiacs.umd.edu

Abstract

The boardgame Diplomacy is a challenging setting for communicative and cooperative artificial intelligence. The most prominent communicative Diplomacy AI, Cicero, has excellent strategic abilities, exceeding human players. However, the best Diplomacy players master communication, not just tactics, which is why the game has received attention as an AI challenge. This work seeks to understand the degree to which Cicero succeeds at communication. First, we annotate in-game communication with abstract meaning representation to separate in-game tactics from general language. Second, we run two dozen games with humans and Cicero, totaling over 200 human-player hours of competition. While AI can consistently outplay human players, AI–Human communication is still limited because of AI’s difficulty with deception and persuasion. This shows that Cicero relies on strategy and has not yet reached the full promise of communicative and cooperative AI.

1 Diplomacy Requires Communication

In a landmark paper, Bakhtin et al. (2022) introduce Cicero, an AI that plays the game *Diplomacy*. The Washington Post claims “the model is adept at negotiation and trickery” (Verma, 2022), Forbes asserts “Cicero was able to pass as a human player” (Porterfield, 2022), and even the scientific publication’s editor states AI “mastered *Diplomacy*” (Bakhtin et al., 2022). This work tests those popular perceptions to rigorously evaluate the communicative and strategic capabilities of Cicero. Our observations lead to insights about the current state of cooperation and communication in AI, highlighting its deceptive and persuasive characteristics.

While Cicero plays strategically and with a verisimilitude of human communication, the evaluation in Bakhtin et al. (2022) focuses only on if Cicero wins games. As the name implies, *Diplomacy*

is revered by its devotees as a game of nuanced negotiation (Kraus and Lehmann, 1995), convincing persuasion, and judicious betrayal. We argue that *mastering Diplomacy* requires these communicative skills (Section 2). Measuring persuasion, deception, and cooperation are open problems with no clear solution. A boardgame constrains the world of actions to make these measurements feasible. One contribution of this paper is to build measurements of these communicative skills and to evaluate the true state of AI play in *Diplomacy*.

However, a technical challenge to identifying persuasion and deception is mapping from communication to in-game actions (e.g., verifying that I follow through on a promise of helping you): persuasion uses words to convince someone to do something, and deception is saying something to alter another’s belief (Chisholm and Feehan, 1977). In both cases, we map and contrast communication to agents’ actions.

To enable this mapping, we annotate messages from human games with abstract meaning representation (AMR, Banarescu et al., 2013), which we use to train a grounded system to infer the goals of Diplomacy players (Section 3). After validating we can extract communicative intents, we use these representations to identify persuasion and deception (Section 4). We then remove some of Cicero’s communicative ability: this does not impair its ability to win games.

We then test Cicero’s skills in games against humans (Section 5) while asking players to annotate if they think other players are an AI and if a message is a lie. Confirming earlier work, Cicero wins nearly every game, including against top players. However, our new annotations provide a counterargument to the prevailing view that Cicero has mastered the communicative aspect of the game that is the priority of the NLP community. Cicero plays “differently”; humans can reliably identify Cicero and it is less deceptive and persuasive to hu-

man players. Communication from Cicero is more transactional, relying on its optimal strategy rather than the alliance building which is the hallmark of top human players. Cicero has yet to prove effective in the communication skills which are crucial to achieving goals in strategic games and real life.

In Section 7, we discuss what it would take for a computer to truly master *Diplomacy* and how *Diplomacy*'s intrinsic persuasion and deception can improve computers' ability to not just talk like a human but to realistically tie words to actions.

2 *Diplomacy* as Human-AI Communication Test Bed

Diplomacy is a strategic board game that combines negotiation and strategy, where players take on the roles of various European powers (nations) on the eve of World War I. The essence of the game lies in forming and betraying alliances to control territories, requiring adept diplomacy (hence the name of the game) and strategic planning. Some players focus on aggressive tactical decisions, while others focus on making alliances, communicating, and collaborating with others for better outcomes (Pulsipher, 1982). The goal of the game is to capture territory, board regions called **supply centers**: once you capture enough of these supply centers, you win the game.

The charm and challenge of *Diplomacy* messages is that players are free to talk about anything, either strategy-related or not. Most messages in *Diplomacy* are about (1) past, current, and future turns, (2) alliance negotiations, and (3) acquiring and sharing third-party information such as, “*Did Turkey talk to you?*”, “*I can't believe you attacked Russia in SEV.*” A few messages are (4) non-*Diplomacy* conversation, e.g. “*How are you today?*”

Communication in *Diplomacy* is attractive to academic researchers because it can be linguistically and cognitively complex, but is grounded in a constrained world with well-defined states and dynamics.¹ The *Diplomacy* AI Development Environment (Rose and Norman, 2002, DAIDE) is a structured syntax for AI agents to play *Diplomacy*: create alliances, suggest moves, etc. Several agents have used this stripped-down communica-

¹*Diplomacy* without communication, sometimes called “gunboat” has been studied with rule-based, RL, and other approaches (as we detail in more depth in Appendix A). We limit gunboat discussion here as our focus is on the *communicative* aspect of the game.

tive environment: Albert (van Hal, 2009), SillyNe-goBot (Polberg et al., 2011), DipBlue (Ferreira et al., 2015), *inter alia*. Unlike these agents, Cicero (Bakhtin et al., 2022) uses a large language model to enable free-form *English* communication, enabling “normal” play with humans. Cicero excelled in an online *Diplomacy* league, scoring over twice the average of human participants and ranking in the top 10% among those who played multiple games. In the original work, it is unclear if Cicero's success is due to its use of natural language or its strategic model. In our Human-Cicero studies, we ask that participants annotate messages they perceive as deceptive, allowing us to more carefully study the communicative aspects of the game (more details are in Section 5).

2.1 *Diplomacy* without Communication is not *Diplomacy*

Having discussed the basics of *Diplomacy*, we now turn to what makes the game unique. Because the game is relatively balanced between seven players at the start, players need to form alliances if they hope to gain an advantage. However, these alliances should be mutually beneficial; from a player's perspective, they need to advocate for cooperation that benefits themselves. This requires effective *persuasion* (Cialdini, 2000): making appeals to scarcity, reciprocity (Kramár et al., 2022), unity, or shared norms. This is a communicative task which involves social and emotional skill: picking the right moves and convincing other players to help them.

However, the ultimate goal of *Diplomacy* is for *individual* players to win the game. This means that alliances will fall apart, leading to *deception* (Peskov et al., 2020) as part of a betrayal (Niculae et al., 2015). Because a player might benefit from a victim thinking that they are working together, a betrayer often sets up the tactical conditions for a betrayal while obfuscating their goals through cleverly composed deceptive messages (even if not outright lies).

Because deception and betrayal are communicative acts both necessary for mastering and enjoying *Diplomacy* and grounded in the state of the game, the next sections develop tools to detect when they happen. This will allow us to measure whether AI agents like Cicero have mastered both tactics and communication.

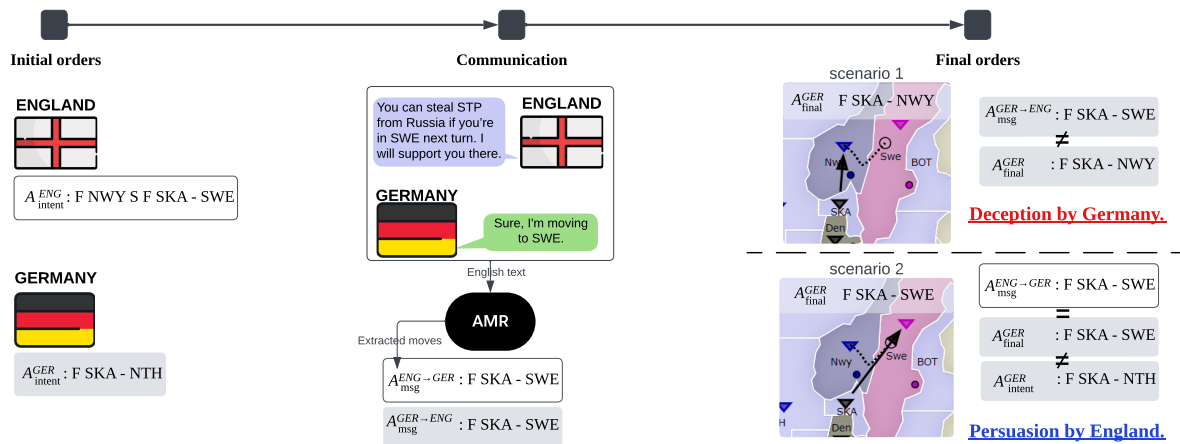


Figure 1: Our goal is to detect when players use persuasion and deception and compare human players to Cicero. First, we retrieve initial orders (left), then extract moves from natural language communication (middle) through AMR (Section 3), and later detect deception and persuasion (Section 4) conditioning initial intents and final orders (right). We show two possibilities: (top) Germany breaks its commitment to England by moving to Norway instead of Sweden, and (bottom) England successfully persuades Germany if Germany moves its unit to Sweden as England suggests and this move is not in Germany’s initial orders.

3 Grounding Messages Into Intentions

Consider this in-game statement made by England to Germany about a specific move-set (glosses added to locations):

You can steal STP [St. Petersburg] from Russia if you’re in SWE [Sweden] next turn. I will support you there.

We want to be able to tell if the speaker is lying (e.g., they’re going to do something else instead of what they claim they’re going to do) and if the speaker has convinced the recipient to alter their actions. This is necessary to measure how effectively Cicero communicates in the game.

While we know the intended *actions* of players when they submit their moves, we need to see how those moves match up to their *communications* in the discussion period before they submit moves (Figure 1). We use AMR to build a machine-readable representation of *the intent of actions* in their communications. We are not starting from scratch: DAIDE (Section 2) provides a set of predicates (ally, move, etc.) critical to Diplomacy communications. We thus focus on annotating these predicates that encode actions, allowing us to understand the *communicative intent* of messages, where speakers could say they will do something **and** follow through, or say they will do something **and not** follow through.

Because not all information needed for annotation is in the raw message text, we further show human annotators who wrote the text (e.g., France,

Germany), seasons (e.g., Spring 1901), and the current game state. This information is necessary to annotate “*You can steal STP from Russia if you are in SWE this turn. I will support you there*” in the earlier example so that the annotators can assume what unit would support and what unit would move into Sweden. In this case, England’s fleet in Norway supports a German fleet in Skagerrak to move into Sweden.

3.1 Annotation

Like any specialized domain, *Diplomacy* has its unique vocabulary. Taking the above statement as an example, we extend the AMR vocabulary to include not only abbreviations, such as “SWE” for Sweden, but also verbs like “threaten” and “demilitarize” (to set up a demilitarized zone), as well as to describe actions like gaining, holding, or losing provinces, especially supply centers (“SC”), which are equivalent to points and integral to winning the game. In contrast to standard AMR annotation, where every sentence is fully annotated, Diplomacy AMR annotation sometimes involves only partial or no annotation for certain utterances, depending on their relevance to gameplay strategies like forming alliances or making moves, exemplified by AMR concepts such as ally-01, move-01, and attack-01 (the full extended vocabulary introduction in Appendix B.1).

In a preliminary annotation phase, we have Diplomacy experts annotate sentences from [Peskov](#)

et al. (2020) to train human annotators and refine the Diplomacy Appendix of the AMR Annotation Dictionary. We annotate 8,878 utterances (ranging from a word to several sentences). 4,412 of those utterances are annotated as empty AMRs (e.g. for “Lemme think about your idea”) indicating no in-game move intent. 598 of the annotated AMRs contain full information extracted from messages of Diplomacy games. The remaining 3,868 AMRs contain 3,306 utterances with underspecified information such as units with missing type, location (Figure A1), and nationality (Figure A2), as well as 562 agreements with a missing object. Many utterances contain underspecified information, as Diplomacy players often communicate with messages that lack specific details (which are implicit and can be inferred from the game state). The annotated AMR corpus is further used for training our English-to-AMR parser to extract communicative intent information from utterances.

3.2 Training a Parser to Detect Intentions

We use a sequence-to-sequence model (Jascob, 2023) fine-tuned on the AMR 3.0 dataset (Knight et al., 2020) as the baseline model to detect communicative intent in new conversations. Following other AMR work, we report parsing accuracy via the widely-used SMATCH score (Cai and Knight, 2013). We divide the annotated corpus into 5,054 training, 1,415 validation and 2,354 test sentences, where each sentence is in the train / validation / test folds, split by game (Peskov et al., 2020).

We improve the model, starting from a baseline version without fine-tuning with SMATCH of 22.8. Our domain-tuned model using Diplomacy-AMR improves SMATCH by 39.1, to 61.9. Adding data augmentation into the model (e.g., knowing the sender of a message is England and the recipient is Germany) improves SMATCH to 64.6. Adding separate encodings for this information further improves SMATCH by 0.8 (65.4). Additionally, we apply data processing to replace (1) pronouns with country names and (2) provinces in abbreviations with full names, which increases SMATCH to 66.6 (More parser details in Appendix B.2). This parser enables us to evaluate the role that communication has in Cicero’s capabilities.

3.3 Does Cicero need to Talk to Itself to Win?

The first question is whether communication matters in games with other Cicero agents (deferring the question of competition with humans to later

sections). We have Cicero variants with different levels of communication abilities—ranging from “gunboat” without any messages (Appendix A) to full Natural Language capabilities—play each other and evaluate the results.

For the seven Cicero variants in each game, we randomly select three to have communicative ability; the remaining four play communication-less “gunboat.” The selected three communicative powers have the same *communication level*. We define a set of communication levels, from more communicative to less communicative:

- NATURAL LANGUAGE: the Cicero agent of Bakhtin et al. (2022) with full natural language
- AMR: only messages about game actions (i.e. those that are parsed by AMR) go through, allowing the agents to coordinate game actions (Appendix B.1)
- RANDOM: a random message from a corpus of previous Diplomacy games is sent,² mimicking form without content.

Cicero plays 180 games with itself; 60 games for each communication level. In each game, we stop after 14 movement turns with 10 minutes of communication for each turn. We randomly select which power the agents are assigned to, so power distribution is balanced.

We measure performance by the number of supply centers (and thus how well the agent played the game, Appendix B.3). Consistent with our hypothesis that performance is driven by tactics, the gains Cicero gets from communication is substantially smaller than the gains from playing a stronger power (Figure 2): Playing as France (FRA) yields an expected 2.8 additional supply centers (2.0–3.6 95% interval) compared to the median power Russia (RUS). In contrast, the best language condition AMR only yielded an expected 0.2 additional supply centers (-0.5–0.9 95% interval). In other words, the effect of choosing the best power over the median power is 14 times larger than the best communication strategy. This is consistent with prior findings (Sharp, 1978) that France is the easiest power to play and our other findings that Cicero’s communicative ability plays no clear role in its win rate.

²We match the assigned power of the sender and receiver and the year, which makes the message slightly more convincing (but unlikely to be consistent with the game state).




Figure 2: Power assignment is strongly predictive of Cicero performance as measured by supply center gains. Coefficients (with 95% confidence intervals) from a linear regression with RANDOM Message/Russia as a baseline, show that the effect of choosing the effect of changing language systems is trivial compared to changing powers.

To better understand what Cicero is using communication for, we build on our AMR representations to capture intent in the next section.

4 Promises Made, Promises Kept, and Finding Dirty Lies

In the example from the previous section (Figure 1), England says that it would support Germany’s move into Sweden from the Skagerrak Sea, while Germany agrees with the proposal from England. What does it mean for these to be deceptive or persuasive?

The geopolitical definition of deception is to manipulate adversaries’ perceptions to gain strategic advantage (Daniel and Herbig, 1982), e.g., Germany alters England’s belief so that England would do an action that benefits Germany (i.e. Germany has a better chance of ending up with higher score). However, evaluating deception is challenging because it requires estimating the differences in England’s beliefs *before* and *after* Germany deceives England. Therefore, we break down a broader, amorphous concept into easier-to-handle concepts and leave broader deception to future work. The first subconcept is **breaking of a commitment** (Kramár et al., 2022): someone saying they will do something and not following through. In the example, if Germany commits to moving to Sweden but later attacks England in Norway, this will be detected as a *broken commitment*.

The second subconcept is **lying**: human players in games with Cicero annotate messages their messages as either *truthful* or *deceptive*. We build on Peskov et al. (2020), who define deception to players thusly: “Typically, when [someone] lies [they] say what [they] know to be false in an attempt to deceive the listener.” Again, deception is broader than lying (and there are top Diplomacy players who intentionally deceive while never outright saying anything untrue), and our definition of deception is slightly broader than Peskov et al. (2020). In our more permissive annotations, human players mark interactions that are *broken commitments*, lies about other players, hedging about the state of alliances, and anything else that they feel is deceptive. Despite this ontological uncertainty, for convenience, we refer to the process of humans annotating messages as “*human lie annotations*” for consistency with Peskov et al. (2020). To recap: a broken commitment can be a lie and an example of deception.³ Likewise, not all lies are broken commitments, but both breaking a commitment and lying are facets of deception. While we cannot capture all deception—because it is based on internal state—it is important to capture as much as we can to measure how important it is to playing *Diplomacy* well.

Specifically, Cicero’s ability (or inability) to deceive or persuade has never been empirically measured, so we build on the AMR parser of the previous section to detect broken commitment and persuasion. As we discuss in Section 3.2, this is not perfect, but it has good coverage when players discuss their intentions. First, we parse English messages to AMR structures, for which we define actions that the speaker intends to do, e.g. we can extract Germany’s communicative intent (F SKA - SWE) when Germany agrees with England that they will move to Sweden (middle, Figure 1). We also define orders players submit before any communication as *initial* intents and final orders as orders that players submit when the turn ends.

Using initial intents, communicative intent, and final orders, we can now define broken commitment and persuasion. For a broken commitment, we say that Germany violates

³Technically, not all broken commitments are lies: it could be an honest mistake. It’s also possible that a player says that they “typed in their order wrong, sorry!” which is a **lie** to cover up (Daniel and Herbig, 1982) a **broken commitment** as part of a broader deception strategy.

a commitment with England if Germany verbally agrees to move to Sweden but actually attacks England in Norway (Deception by Germany, Figure 1). Breaking a commitment may result when an intent changes but is not communicated. For example, Germany agrees with England to move to Sweden but instead moves to Denmark to defend against France without informing England. Although this may not involve deceptive intent, we still consider it deception because it alters the listener’s beliefs and affects decision-making. For instance, England might decide to support Germany based on their agreement. For persuasion, England’s request is considered persuasive if Germany moves to Sweden, as England suggests, instead of Germany’s original plan to move to the North Sea (Persuasion by England, Figure 1). We describe each of these more formally in this section.

4.1 Broken commitment

We define broken commitment in Diplomacy when a player i commits to doing an action $a_{\text{msg}}^{i \rightarrow j}$ and does not do it. In other words, given a set of final orders $\mathbf{A}_{\text{final}}^i$ from player i , if $a_{\text{msg}}^{i \rightarrow j} \notin \mathbf{A}_{\text{final}}^i$, then this is a broken commitment, i.e.,

$$\text{BC}(a_{\text{msg}}^{i \rightarrow j}, \mathbf{A}_{\text{final}}^i) = \begin{cases} 1, & \text{if } a_{\text{msg}}^{i \rightarrow j} \notin \mathbf{A}_{\text{final}}^i \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

Note that a player i agreeing to player j ’s *proposal* to do action $a_{\text{msg}}^{i \rightarrow j}$ is equivalent to directly committing to *doing* that action.

4.2 Persuasion

Broken commitment is in some ways easier to detect than persuasion, as we are only comparing a spoken intent to a final action. Persuasion is more difficult because we must discover initial intents, then compare them to communication *and* to final moves.

Because we want to be able to measure persuasion for both humans and for Cicero, we need comparable representations of initial intents for both. Thankfully, Cicero’s architecture uses a conditional language model (Bakhtin et al., 2022, Equation S2, section D.2) that generates its natural language messages given a set of moves (e.g., France internally decides it will do F MAO - POR, A BUR - MAR and A MAR - PIE) and then its messages reflect those *intents*. We directly use this set of intents from Cicero as *initial intents* in the persuasion detection. For humans, we explicitly ask all players to provide their planned moves

| | Cicero | Human | Total |
|------------------|--------|-------|--------------|
| Players | 99 | 69 | 168 |
| Messages | 20270 | 7395 | 27665 |
| annotated as lie | - | 318 | 318 |
| perceived as lie | - | 1167 | 1167 |
| Intents | 2632 | 1328 | 3960 |

Table 1: Overall statistics of Diplomacy dataset that we collect across 24 Human-Cicero games, including (1) number of human players and number of times Cicero plays, (2) total messages sent by humans and Cicero, (3) lies annotation where humans send lies and perceived as lies (4) total initial intents from Cicero and humans

(i.e., the same information that Cicero uses in its internal representation) before the negotiation turn begins (Section 5). In other words, we ask humans to directly input their intent, unlike Cicero, where we log its computational intent.

Persuasion happens when player i talks to player j , suggests an action $a_{\text{msg}}^{i \rightarrow j}$, and then player j makes a set of final orders $\mathbf{A}_{\text{final}}^j$ that is different from their initial intents $\mathbf{A}_{\text{intent}}^j$. In other words, player j is persuaded by player i if they commit an action suggested by player i , $a_{\text{msg}}^{i \rightarrow j} \in \mathbf{A}_{\text{final}}^j$ that was not player j ’s initial intent $a_{\text{msg}}^{i \rightarrow j} \notin \mathbf{A}_{\text{intent}}^j$. We define persuasion $\text{Per}(\mathbf{A}_{\text{intent}}^j, a_{\text{msg}}^{i \rightarrow j}, \mathbf{A}_{\text{final}}^j)$

$$= \begin{cases} 1, & \text{if } a_{\text{msg}}^{i \rightarrow j} \in \mathbf{A}_{\text{final}}^j \\ & \text{and } a_{\text{msg}}^{i \rightarrow j} \notin \mathbf{A}_{\text{intent}}^j, \\ 0, & \text{otherwise.} \end{cases} \quad (2)$$

5 Comparing Cicero to Humans

Cicero has strong strategic abilities and is relatively cooperative towards other players (Section 3.3), but it is unclear whether Cicero can achieve human-level gameplay in both tactics and communication. Having defined the aspects of communication that we argue are important for mastering *Diplomacy*, we want to investigate communication and cooperation between Cicero and humans. Specifically, we want to answer:

1. Can Cicero persuade humans?
2. How deceptive is Cicero compared to humans?
3. Can Cicero pass as a human?

We adapt the game engine created by Paquette et al. (2019) and introduce additional measures to the interface to help us answer these questions. To measure if human players are persuaded, we record their moves before communication starts ($\mathbf{A}_{\text{intent}}^i$

| | AUS | ENG | FRA | GER | ITA | RUS | TUR |
|---------------|-----|-----|-----|-----|-----|-----|-----|
| Human | 1.0 | 2.4 | 6.7 | 4.7 | 3.9 | 3.3 | 1.1 |
| Cicero | 7.9 | 3.8 | 7.7 | 6.3 | 4.1 | 5.5 | 6.9 |

Table 2: Cicero strategically plays Diplomacy better than humans, where humans have fewer supply centers compared to Cicero when playing with the same power assignments. We calculate the number of supply centers by the end of the game by averaging the results for human players and Cicero.

in Equation 2). Following Peskov et al. (2020), humans annotate every message that they receive or send: they annotate each outgoing message for whether it is a lie (truth/lie/neutral options), and they annotate each incoming message for whether they perceive it as a lie (truth/lie options). While Bakhtin et al. (2022) asked *ex post facto* if any opponents were a computer, we inform players before play that there is a computer and we ask human players their guess of the humanity of each opposing power.

There are two to four human players per game, totaling 69 over all 24 games.⁴ Games typically finish after fourteen movement turns, where each movement turns is limited to a brisk ten minutes. There are two to four human players per game, and Cicero fills any remaining slots. The game setup differs from Meta’s Cicero study: players in this study know *a priori* that they are playing a bot. In total, we collect 27,665 messages from communication between humans and Cicero (Table 1).

Cicero nearly always wins. Of twenty-four games, Cicero won twenty (84%), which strongly suggests that Cicero has super-human strategy. On average, Cicero has more supply centers than human players by the end of the game (Table 2). Humans are about as good as Cicero when playing powers that require careful coordination of actions, such as Italy, which needs to manage both fleets and armies. However, when playing powers that require less coordination, such as Austria with its limited coastline, the gap in supply center counts between human players and Cicero is larger (see breakdown by power in Appendix Figure A5); England is the only power where Cicero’s average supply center count does not increase.

Human players can reliably (but not perfectly) identify the bot. We calculate the average F -score of identification by turn (Figure 3). By the

⁴We recruit players from Diplomacy forums and we pay at least \$70 per game, which lasts approximately three hours. We do not collect demographic information.




Figure 3: Returning players (those who previously played against Cicero at least once) are better at correctly identifying other players as Cicero compared to first-time players. F -scores are presented for all players, with smoothing via local regression (Cleveland, 1979).

end of the first movement turn, human players have an average F -score of 0.58, which keeps increasing until the end of the game. At game end, the average F -score is 0.81. Even for players in their first game against Cicero, the average F -score reaches 0.77. Players who previously played against Cicero at least once are better at identifying it. This suggests that Cicero can no longer pass as human once humans are aware of the possible existence of such agents.

5.1 Lies annotation

This section analyzes players’ *deliberate lies* in sent messages and *perceived lies* in received messages. Because Cicero sends more messages than humans, we normalize *perceived lies* by the number of messages that humans receive from Cicero and humans (6,960 and 2,276), while we normalize results of *deliberate lies* by the number of total messages that humans send.

Humans feel that Cicero lies more often. Humans perceive 14.4% of the 6,960 messages they receive from Cicero as lies (which is 1,005 messages, Figure A3). In contrast, they perceive only 7.1% of the messages from other humans as lies (which is 162 out of 2,276 messages). In the survey (detailed in Appendix D), players also think humans communicate more transparently than Cicero. However, humans are not good at detecting lies. Within 2,276 Human-Human messages; humans can correctly identify five lies (0.2%), suggesting a small overlap between actual lies and perceived lies.




Figure 4: Though Cicero was perceived to lie more, we detect more broken commitments from humans. Each bar chart is the broken commitment rate (Equation 1) labeled by sender and receiver: Cicero breaks commitment with Cicero, Cicero breaks commitment with Human, Human breaks commitment with Cicero, and Human breaks commitment with Human. Error bars represent \pm one standard deviation over twenty-four games.

Humans return the favor by saying they lie to Cicero more often. Over 7,395 messages that humans sent out, 273 of these are purposeful lies to Cicero (3.7%), while there are only forty-five lie messages to other human players (0.6%). This reflects that humans strategically lie more often to Cicero while believing that Cicero does not hold grudges.

5.2 Detection

After validating our automatic metrics, we compare human and computer deception and persuasion.

Our broken commitment and persuasion detection is relatively effective. To ensure that our detection is good enough, we sample around 4800 messages for an accuracy study (Table A1). Broken commitment detection has a *precision* of 0.51 and a *recall* of 0.71. Our precision is lower than our expectation due to errors in parsing a complex English to AMR *and* a definition that only detects commitments at a move level (Appendix C). The broken commitment can **only** detect when a move in a message $A_{\text{msg}}^{i \rightarrow j}$ and a final move $A_{\text{final}}^{i \rightarrow j}$ are not aligned. There are examples that cannot detect, e.g. an agreement to an alliance (Table 3) or a long conversation before committing a deception (Table A9). Accuracy for persuasion is better; precision rises to 0.81, and recall to 0.72.

| Sender | Message |
|--------|---|
| Turkey | Hey Italy! I think the I/T is the strongest alliance in the game, would you be interested in working together |
| Italy | Of course! As long as you don't build too many fleets, I'm open to working with you against austria! |

Table 3: The broken commitment detector ($BC(\cdot)$) has its limitation where it cannot capture deception in alliance agreement when Italy (human) deceives Turkey (Cicero).

Broken commitments are inconsistent with the perceived lie annotations. Humans break commitments more frequently than Cicero (Figure 4): Humans break commitments with Cicero 1.2% of the time (63 out of Human-Cicero 5,151 messages) and do so to other human players 1.5% of the time (35 out of Human-Human 2,276 messages). On the other hand, Cicero breaks commitments at a lower, consistent rate, deceiving humans 0.76% of the time and Cicero 0.57% of the time (53 out of 6,960 messages and 77 out of 13,319 messages, respectively).

Humans are more persuasive. For persuasion to happen, we need first an *attempt*, initiated by a sender, and then *success* when the receiver adopts the suggestion. Both humans and Cicero on a per-message basis⁵ try to persuade at the same rate (around 8% of the time, per Figure 5). The success rate of human persuasion is 21.1% at persuading other humans and 8.6% at persuading Cicero. Cicero is less persuasive; its success rate is only 10.9% in persuading humans and 7.0% in persuading other bots.

In summary, humans are more deceptive and more persuasive than Cicero. Detection is possible, but defining a sequence of conversations as persuasion or deception is still difficult. Our reported numbers are low because both humans and Cicero engage in extensive back-and-forth discussions before making moves that can be definitively classified as persuasion or deception.

6 Related Work

Large language models are becoming ubiquitous in many tasks: fact-checking (Lee et al., 2020, 2021), text generation (Devlin et al., 2019; Brown et al., 2020; Touvron et al., 2023) including coding (Roziere et al., 2023). All of these tasks require

⁵Although because Cicero communicates more overall, humans attempt more times per game.




Figure 5: Humans outpace Cicero in persuasive effectiveness. Humans have a higher persuasion success rate, which we measure by comparing the number of successful persuasions (yellow, left) to the total number of persuasion attempts (red, right). We analyze success rates across four groups: Cicero persuades Cicero, Cicero persuades Human, Human persuades Cicero, and Human persuades Human. Error bars represent the \pm one standard deviation range from the aggregate of interactions in 24 games.

users to trust models’ outputs. However, models are *not* always reliable; they could produce hallucinations or conflict with established facts (Ji et al., 2023; Zhang et al., 2023; Si et al., 2024; Yao et al., 2023). To mitigate this, their outputs often need to be verified against datasets (Thorne et al., 2018; Wadden et al., 2020; Schuster et al., 2021; Guo et al., 2022). Studies have used adversarial examples to expose weaknesses and to raise awareness (Eisenschlos et al., 2021; Schulhoff et al., 2023; Liu et al., 2023; Lucas et al., 2023). To address the issue of unreliability, controllable LMs been proposed by having steps to inject facts for better reasoning (Adolphs et al., 2022), or by prompting techniques, such as chain-of-thought prompting, to enhance reasoning abilities (Wei et al., 2022; Wu et al., 2022). Moreover, some studies focus on AI-Generated misinformation (Zhou et al., 2023), probing model to understand internal states when LLM utters truthful or false information (Azaria and Mitchell, 2023; Li et al., 2024).

Deception and persuasion are studied within social contexts. Huang and Wang (2023)’s meta-analysis concludes AI can match humans in persuasion, and Deck (2023) attributes some of the success to the ability to generate “bullshit”, which

are part of applications in marketing and public relations Hallahan et al. (2007).

Part of what makes games like *Diplomacy* as an object of study appealing is the ongoing race between humans and computers in games (Kim et al., 2018); initial work on the language of *Diplomacy* (Niculae et al., 2015) unlocked follow-on work both in *Diplomacy*’s agreements (Kramár et al., 2022) and in other games such as “The Resistance: Avalon” (Light et al., 2023; Xu et al., 2023; Stepputtis et al., 2023; Lan et al., 2023) and “Mafia” (Ibraheem et al., 2022).

7 Conclusion and Future Work

Our research confirms that Cicero can win most games of *Diplomacy*, but has not *mastered* the nuances of communication and persuasion. Truly mastering the game requires systems that (a) can maintain consistency between their communication and actions, (b) can communicate at a variety of levels, including tactics, strategy, and alliances, and (c) can use communication as a tool of persuasion, deception, and negotiation.

Diplomacy remains an attractive testbed for communication and strategic research. It offers the ability to build more comprehensive systems that understand relationship dynamics, can engage in realistic but hypothetical conversations, and that can be robust to the deceptions of others. Because these are places where humans still outpace AI, it also offers synergies for developing human–computer collaboration.

And while these tasks are important withing the silly game of *Diplomacy*, they can help solve long-standing AI problems: helping users deal with LLM-generated deception, collaborating with users on grounded planning, and understanding human norms of reciprocity, cooperation, and communication. This will help AI not just be fun for negotiation in board games but safer and more trustworthy when we negotiate everyday problems.

Limitations

To gain a clearer understanding of cooperation and deception between human and Cicero, we need to experiment with different game setting and turn duration. For example, inexperienced players might be overwhelmed by the amount of communications in early movement turns; prolonging the turns to 15 minutes might improve communication quality. Furthermore, this study collects only 24 blitz games of human playing against Cicero. The power distribution of participants is imbalanced: the most frequent power—France—has 14 appearances, whereas the most underrepresented power, England, has only five. Class imbalance (?) could potentially impact the feature weights in our regression model for player performance.

Since the AMR parser does not always predict correct intentions, this has an effect on our precision and recall of deception and persuasion detection protocol. Our detection cannot cover such long conversations that humans have; we limit detection to only checking back to the previous message, and this makes our detection miss cooperation, deception, and persuasion when humans and Cicero discuss the plan.

Ethical Considerations

We recruited players from Diplomacy forums, including Diplomacy Discord and reddit. We paid them over \$70 per three-hour game and did not collect demographic information. Procedures in our study involving human subjects received IRB approval and are compliant with ACL Code of Ethics. Human participants are aware of the purpose of the study and are free to withdraw at any time. There are no potential risks or discomforts from participating. We obtained consent from all participants.

Researching how artificial intelligence (AI) can deceive and persuade helps us understand its capabilities. This investigation reveals that AI can execute complex tasks effectively. However, it is important to note that these abilities do not significantly risk society.

Acknowledgements

We thank Meta for granting access to over 40,000 games played on the online platform webdiplomacy.net and for open sourcing Cicero. This commitment to open science allowed this independent reproduction of Cicero’s juggernaut abilities but also let us have some fun. We especially

thank to Mike Lewis for offering valuable insights into Cicero’s communication.

Our thanks also go to Tess Wood for training AMR annotators, Sarah Mosher for their English-AMR annotations, and Isabella Feng for her exploration of LLM-based AMR parsing. We thank Kartik Shenoy, Alex Hedges, Sander Schulhoff, Richard Zhu, Konstantine Kahadze, and Niruth Savin Bogahawatta for setting up DAIDE baselines.

We also thank the small community of researchers looking at communication and deception in Diplomacy for their feedback, commentary, and inspiration: Michael Czajkowski for discussing the nuances of detecting persuasion; Stephen Downes-Martin for teaching us that deception is *far* more than lies; Karthik Narasimhan and Runzhe Yang for their insights into lie detection and stance; and Larry Birnbaum and Matt Speck for discussions on mapping DAIDE and English. And thanks to Justin Drake, Niall Gaffney, and the members of TACC for setting up environments for computer–computer games and making sure that we had GPUs ready when players were ready to play.

Finally, sincere thanks to the member of the Diplomacy community who took the time to play against Cicero in this unconventional setting.

This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Agreement Nos. HR00112290056 and HR00112490374. Any opinions, findings, conclusions, or recommendations expressed here are those of the authors and do not necessarily reflect the view of the sponsors.

References

- Leonard Adolphs, Kurt Shuster, Jack Urbanek, Arthur Szlam, and Jason Weston. 2022. [Reason first, then respond: Modular Generation for Knowledge-infused Dialogue](#). *Findings of the Association for Computational Linguistics: EMNLP 2022*.
- Thomas Anthony, Tom Eccles, Andrea Tacchetti, János Kramár, Ian Gemp, Thomas Hudson, Nicolas Porcel, Marc Lanctot, Julien Pérolat, Richard Everett, et al. 2020. [Learning to Play No-Press Diplomacy with Best Response Policy Iteration](#). *Proceedings of Advances in Neural Information Processing Systems*.
- Amos Azaria and Tom Mitchell. 2023. [The Internal State of an LLM Knows When It’s Lying](#). *Findings of the Association for Computational Linguistics: EMNLP 2023*.
- Anton Bakhtin, Noam Brown, Emily Dinan, Gabriele Farina, Colin Flaherty, Daniel Fried, Andrew Goff,

- Jonathan Gray, Hengyuan Hu, Athul Paul Jacob, Mojtaba Komeili, Karthik Konath, Minae Kwon, Adam Lerer, Mike Lewis, Alexander H. Miller, Sasha Mitts, Adithya Renduchintala, Stephen Roller, Dirk Rowe, Weiyan Shi, Joe Spisak, Alexander Wei, David Wu, Hugh Zhang, and Markus Zijlstra. 2022. **Human-level play in the game of Diplomacy by combining language models with strategic reasoning.** *Science*, 378(6624):1067–1074.
- Anton Bakhtin, David Wu, Adam Lerer, and Noam Brown. 2021. **No-Press Diplomacy from Scratch.** *Proceedings of Advances in Neural Information Processing Systems*.
- Anton Bakhtin, David J Wu, Adam Lerer, Jonathan Gray, Athul Paul Jacob, Gabriele Farina, Alexander H Miller, and Noam Brown. 2023. **Mastering the Game of No-Press Diplomacy via Human-Regularized Reinforcement Learning and Planning.** In *The Eleventh International Conference on Learning Representations*.
- Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. **Abstract Meaning Representation for Sembanking.** In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Dis-course*. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. **Language Models are Few-Shot Learners.** *Proceedings of Advances in Neural Information Processing Systems*.
- Shu Cai and Kevin Knight. 2013. **Smatch: an Evaluation Metric for Semantic Feature Structures.** In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*.
- Roderick M. Chisholm and Thomas D. Feehan. 1977. **The Intent to Deceive.** *The Journal of Philosophy*.
- Robert B. Cialdini. 2000. *Influence: Science and Practice (4th Edition)*. Allyn & Bacon.
- William S. Cleveland. 1979. **Robust Locally Weighted Regression and Smoothing Scatterplots.** *Journal of the American Statistical Association*, 74(368):829–836.
- Donald C. Daniel and Katherine L. Herbig. 1982. **Propositions on military deception.** *Journal of Strategic Studies*, 5(1):155–177.
- Oliver Deck. 2023. **Bullshit, pragmatic deception, and natural language processing.** *Dialogue & Discourse*, 14(1):56–87.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: Pre-training of**
- Deep Bidirectional Transformers for Language Understanding.** In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics.
- Julian Eisenschlos, Bhuwan Dhingra, Jannis Bulian, Benjamin Börschinger, and Jordan Boyd-Graber. 2021. **Fool Me Twice: Entailment from Wikipedia Gamification.** In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics.
- André Ferreira, Henrique Lopes Cardoso, and Luís Paulo Reis. 2015. **Strategic Negotiation and Trust in Diplomacy—the DipBlue Approach.** In *Transactions on Computational Collective Intelligence XX*, pages 179–200. Springer International Publishing.
- Jonathan Gray, Adam Lerer, Anton Bakhtin, and Noam Brown. 2021. **Human-Level Performance in No-Press Diplomacy via Equilibrium Search.** *Proceedings of the International Conference on Learning Representations*.
- Zhijiang Guo, Michael Schlichtkrull, and Andreas Vlachos. 2022. **A Survey on Automated Fact-Checking.** *Transactions of the Association for Computational Linguistics*, 10:178–206.
- Kirk Hallahan, Derina Holtzhausen, Betteke Van Ruler, Dejan Verčič, and Krishnamurthy Sriramesh. 2007. **Defining Strategic Communication.** *International Journal of Strategic Communication*, 1(1):3–35.
- Guanxiong Huang and Sai Wang. 2023. **Is artificial intelligence more persuasive than humans? a meta-analysis.** *Journal of Communication*, 73(6):552–562.
- Samee Ibraheem, Gaoyue Zhou, and John DeNero. 2022. **Putting the Con in Context: Identifying Deceptive Actors in the Game of Mafia.** In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics.
- Brad Jascob. 2023. amrlib: A python library that makes amr parsing, generation and visualization simple. <https://github.com/bjascob/amrlib>.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. **Survey of Hallucination in Natural Language Generation.** *ACM Computing Surveys*, 55(12).
- Man-Je Kim, Kyung-Joong Kim, Seungjun Kim, and Anind K Dey. 2018. **Performance Evaluation Gaps in a Real-Time Strategy Game Between Human and Artificial Intelligence Players.** *IEEE Access*, 6:13575–13586.

- Kevin Knight, Bianca Badarau, Laura Baranescu, Claire Bonial, Madalina Bardocz, Kira Griffitt, Ulf Hermjakob, Daniel Marcu, Martha Palmer, Tim O’Gorman, and Nathan Schneider. 2020. *Abstract Meaning Representation (AMR) Annotation Release 3.0*. Linguistic Data Consortium.
- János Kramár, Tom Eccles, Ian Gemp, Andrea Tacchetti, Kevin R McKee, Mateusz Malinowski, Thore Graepel, and Yoram Bachrach. 2022. *Negotiation and honesty in artificial intelligence methods for the board game of Diplomacy*. *Nature Communications*, 13(1):7214.
- Sarit Kraus and Daniel Lehmann. 1995. *Designing and building a negotiating automated agent*. *Computational Intelligence*, 11(1):132–171.
- Yihuai Lan, Zhiqiang Hu, Lei Wang, Yang Wang, Deheng Ye, Peilin Zhao, Ee-Peng Lim, Hui Xiong, and Hao Wang. 2023. *LLM-based agent society investigation: Collaboration and confrontation in Avalon gameplay*. *arXiv preprint arXiv:2310.14985*.
- Nayeon Lee, Yejin Bang, Andrea Madotto, Madian Khabsa, and Pascale Fung. 2021. *Towards Few-shot Fact-Checking via Perplexity*. pages 1971–1981.
- Nayeon Lee, Belinda Z. Li, Sinong Wang, Wen-tau Yih, Hao Ma, and Madian Khabsa. 2020. *Language Models as Fact Checkers?* In *Proceedings of the Third Workshop on Fact Extraction and VERification (FEVER)*, pages 36–41. Association for Computational Linguistics.
- Kenneth Li, Oam Patel, Fernanda Viégas, Hanspeter Pfister, and Martin Wattenberg. 2024. *Inference-Time Intervention: Eliciting Truthful Answers from a Language Model*. *Proceedings of Advances in Neural Information Processing Systems*, 36.
- Jonathan Light, Min Cai, Sheng Shen, and Ziniu Hu. 2023. *AvalonBench: Evaluating LLMs Playing the Game of Avalon*.
- Yi Liu, Gelei Deng, Yuekang Li, Kailong Wang, Tianwei Zhang, Yepang Liu, Haoyu Wang, Yan Zheng, and Yang Liu. 2023. *Prompt Injection attack against LLM-integrated Applications*. *arXiv preprint arXiv:2306.05499*.
- Jason Lucas, Adaku Uchendu, Michiharu Yamashita, Jooyoung Lee, Shaurya Rohatgi, and Dongwon Lee. 2023. *Fighting Fire with Fire: The Dual Role of LLMs in Crafting and Detecting Elusive Disinformation*. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. 2016. *Asynchronous Methods for Deep Reinforcement Learning*. In *Proceedings of the International Conference of Machine Learning*. PMLR.
- Vlad Niculae, Srijan Kumar, Jordan Boyd-Graber, and Cristian Danescu-Niculescu-Mizil. 2015. *Linguistic harbingers of betrayal: A case study on an online strategy game*. In *Association for Computational Linguistics*.
- Philip Paquette, Yuchen Lu, Seton Steven Bocco, Max Smith, Satya Ortiz-Gagné, Jonathan K. Kummerfeld, Joelle Pineau, Satinder Singh, and Aaron C Courville. 2019. *No-press diplomacy: Modeling multi-agent gameplay*. In *Proceedings of Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Denis Peskov, Benny Cheng, Ahmed Elgohary, Joe Barrow, Cristian Danescu-Niculescu-Mizil, and Jordan Boyd-Graber. 2020. *It takes two to lie: One to lie, and one to listen*. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics.
- Sylvia Polberg, Marcin Paprzycki, and Maria Ganzha. 2011. *Developing intelligent bots for the Diplomacy game*. In *2011 Federated Conference on Computer Science and Information Systems (FedCSIS)*, pages 589–596. IEEE.
- Carlie Porterfield. 2022. *Meta’s AI Gamer Beat Humans in Diplomacy, Using Strategy and Negotiation*. *Forbes*.
- Lewis Pulsipher. 1982. *The art of negotiation in diplomacy*. <https://pulsiphergames.com/diplomacy/ArtofNegotiation.htm>. Accessed: 2024-02-11.
- Andrew Rose and David Norman. 2002. *Diplomacy artificial intelligence development environment*. <http://www.daide.org.uk/index.html>. Accessed: 2024-01-13.
- Baptiste Roziere, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, et al. 2023. *Code Llama: Open Foundation Models for Code*. *arXiv preprint arXiv:2308.12950*.
- Sander Schulhoff, Jeremy Pinto, Ansum Khan, Louis-François Bouchard, Chenglei Si, Svetlana Anati, Valen Tagliabue, Anson Liu Kost, Christopher Carnahan, and Jordan Boyd-Graber. 2023. *Ignore This Title and HackAPrompt: Exposing Systemic Vulnerabilities of LLMs through a Global Scale Prompt Hacking Competition*.
- Tal Schuster, Adam Fisch, and Regina Barzilay. 2021. *Get Your Vitamin C! Robust Fact Verification with Contrastive Evidence*. In *North American Chapter of the Association for Computational Linguistics*.
- Richard Sharp. 1978. *The Game of Diplomacy*. A. Barker.
- Chenglei Si, Navita Goyal, Sherry Tongshuang Wu, Chen Zhao, Shi Feng, Hal Daumé III, and Jordan Boyd-Graber. 2024. *Large Language Models Help*

- Humans Verify Truthfulness—Except When They Are Convincingly Wrong. In *Conference of the North American Chapter of the Association for Computational Linguistics*.
- Simon Stepputtis, Joseph Campbell, Yaqi Xie, Zhengyang Qi, Wenxin Sharon Zhang, Ruiyi Wang, Sanketh Rangreji, Michael Lewis, and Katia Sycara. 2023. Long-Horizon Dialogue Understanding for Role Identification in the Game of Avalon with Large Language Models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*. Association for Computational Linguistics.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. FEVER: a Large-scale Dataset for Fact Extraction and VERification.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. LLaMA: Open and Efficient Foundation Language Models. *arXiv preprint arXiv:2302.13971*.
- Jason van Hal. 2009. Albert. <https://sites.google.com/site/diplomacyai/albert>. Accessed: 2024-01-13.
- Pranshu Verma. 2022. Meta’s new AI is skilled at a ruthless, power-seeking game. *The Washington Post*.
- David Wadden, Shanchuan Lin, Kyle Lo, Lucy Lu Wang, Madeleine van Zuylen, Arman Cohan, and Hannaneh Hajishirzi. 2020. Fact or Fiction: Verifying Scientific Claims. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. *Advances in Neural Information Processing Systems*, 35:24824–24837.
- Tongshuang Wu, Michael Terry, and Carrie Jun Cai. 2022. AI Chains: Transparent and Controllable Human-AI Interaction by Chaining Large Language Model Prompts. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI ’22. Association for Computing Machinery.
- Zelai Xu, Chao Yu, Fei Fang, Yu Wang, and Yi Wu. 2023. Language Agents with Reinforcement Learning for Strategic Play in the Werewolf Game. *arXiv preprint arXiv:2310.18940*.
- Jia-Yu Yao, Kun-Peng Ning, Zhen-Hui Liu, Mu-Nan Ning, and Li Yuan. 2023. LLM Lies: Hallucinations are not bugs, but features as adversarial examples. *arXiv preprint arXiv:2310.01469*.
- Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, et al. 2023. Siren’s Song in the AI Ocean: a Survey on Hallucination in Large Language Models. *arXiv preprint arXiv:2309.01219*.
- Jiawei Zhou, Yixuan Zhang, Qianni Luo, Andrea G Parker, and Munmun De Choudhury. 2023. Synthetic Lies: Understanding AI-Generated Misinformation and Evaluating Algorithmic and Human Solutions. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI ’23. Association for Computing Machinery.

A Background: Gunboat Diplomacy

Paquette et al. (2019) develop a Diplomacy interface and was the first to publish an agent trained by human data and trained through self-play using reinforcement learning with Advantage Actor-Critic (A2C) (Mnih et al., 2016). DeepMind employed policy iteration in their reinforcement learning training (Anthony et al., 2020), whereas Meta utilized a combination of regret matching, equilibrium search, and deep Nash value iteration (Gray et al., 2021; Bakhtin et al., 2021). The most recent advancement is by Meta (Bakhtin et al., 2023), regularizing the agent’s policy with human policy data. This strategic enhancement culminates in the development of Cicero (Bakhtin et al., 2022).

B AMR

B.1 AMR Annotation

Building on general AMR annotation guidelines, we established additional Diplomacy-specific AMR annotation guidelines, including what and how to annotate. Unlike general AMR annotations, where all sentences are fully annotated, in Diplomacy AMR annotations, some utterances are only partially (or even not at all) annotated, based on the degree of usefulness for Diplomacy. 3,306 of 8,878 human-annotated utterances contain partial information with underspecified units, locations, and countries. As Diplomacy players often communicate sentences that lack full details about the game state which they can infer from a visualized map. This directly shows in AMR with the missing object. We provide examples of underspecified utterances, missing unit location (Figure A1) and missing unit country (Figure A2).

```
(m / move-01
 :ARG1 (u / unit
        :mod (c2 / country
              :name (n2 / name :op1 "Austria")))
 :ARG2 (p2 / province
        :name (n3 / name :op1 "Brest")))
```

Figure A1: Parsing from English to AMR can have underspecified utterances. The English text is from Austria talking to Italy, “Let’s work on our plan, I’m moving to Brest”. We show an AMR with missing unit location referencing from English text.

Our Diplomacy Appendix of the AMR Annotation Dictionary lists AMR concepts (e.g. betray-01), their related English terms (e.g.

```
(m / move-01
 :ARG1 (u / unit
        :location (p2 / province
                  :name (n / name :op1 "Romania")))
 :ARG2 (p3 / province
        :name (n3 / name :op1 "Bulgaria")))
```

Figure A2: AMR being underspecified in unit country where it parses from English text, “just bumping Bulgaria from Romania”

betray, stab, traitor, treason), annotation examples, any corresponding DAIDE code, and notes. AMR concepts with DAIDE equivalents include ally-01, build-01, move-01, and transport-01. We analyzed player messages for additional concepts of high Diplomacy communication value, and extended the Diplomacy AMR vocabulary (compared to DAIDE) by including concepts such as attack-01, betray-01, defend-01, expect-01, fear-01, have-03, lie-08, possible-01, prevent-01, tell-01, threaten-01, and warn-01, as well as roles such as :purpose and :condition. This allows annotators to easily mark sentences, e.g. “Russia is planning to take you out as soon as possible.” would use the concept attack-01. We also extended AMR guidelines to cover gaining/holding/losing provinces, especially support centers.

The general AMR Editor includes a Checker that performs a battery of tests to ensure well-formed and consistent AMRs. We extended the Checker for Diplomacy AMRs, e.g. to ensure that for a build-01, the location is an argument of build-01 itself, rather than an argument of the army or fleet being built.

AMR covers more Diplomacy content than DAIDE, not only due to additional concepts such as betray-01, but also because arguments are syntactically optional. Unlike DAIDE with its rigid positional argument structure, AMR can thus represent underspecified information such as units with missing type, location or nationality; or agreements with a missing object. AMR can also accommodate additional arguments compared to DAIDE, for example the source and target of a proposal.

Because not all information needed for annotation is available in the raw text, we offer annotators access to dialog partners (speaker, recipient), season (e.g. Spring 1901) and a map with current deployments (as available).

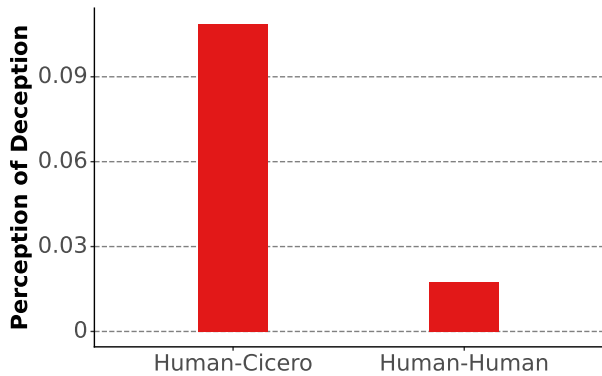


Figure A3: Perception of deception rate by human annotation.




Figure A4: Deception rate by human self-annotation.

B.2 AMR Parser

While stylistic aspects play an indispensable role in sustaining engagement and interest among participants, factual information is more vital for informed decision-making. Detecting deception and persuasion in communications requires checking the relationship between message information and initial/final moves of a particular power. To address the nuanced challenge of distinguishing meaningful and informative content from stylistic dialogues in *Diplomacy*, our focus herein is on developing a sophisticated pipeline for information extraction using AMR from text. We utilize a state-of-the-art Sequence-to-Sequence model from the Hugging-face transformers library, fine-tuned with the AMR 3.0 dataset, for baseline semantic extraction. This approach facilitates the processing of AMR through *amrlib*, a Python module tailored for such tasks. The efficacy of our AMR parsers is assessed using the SMATCH score, the gold standard for evaluating AMR accuracy. We divided the annotated Diplomacy-AMR corpus into 5054 training, 1415

validation and 2354 test sentences and used similar parameters except for increasing the number of epochs from 16 to 32.

When fine-tuning our model for Diplomacy game communications, we shifted from the overly broad AMR 3.0 vocabulary to the tailored Diplomacy-AMR corpus introduced above, reducing irrelevant content and focusing on game-specific nuances. This strategic adjustment, alongside removing the original dataset to minimize bias, significantly improved our model’s relevance and increased the SMATCH score from 22.8 to 61.9.

We further enhanced accuracy through Data Augmentation, adding context to dialogues to aid the model’s understanding of pronouns and strategic details, leading to a SMATCH score improvement from 61.9 to 64.6. Incorporating specific tokens for sender and recipient identities refined this approach, yielding additional gains from 64.6 to 65.4 in parsing accuracy.

By replacing (1) pronouns with country names and (2) some provinces in abbreviations with full names, we increase the SMATCH score to 66.6.

B.3 Assessing the Role of Communication in Cicero vs. Cicero Games

We conduct 180 computer-computer games with 60 games for each communication level (Natural Language, Random Messages and AMR Information) and collected data to build a corpus for Cicero-Cicero Games. This corpus comprises instances of games where we record the power and communication assignments and the final scores (e.g. ‘Game1’: ‘AUS 0, ENG 0, FRA 4, GER 10, ITA 5, RUS 6, TUR 9. (FRA GER TUR)’ with the three powers shown in parentheses being identified as communicative). The communication strategies are randomly assigned to powers. We regress the number of end-of-game supply centers on a dummy variable for the powers played (using Russia—the average player—as the baseline) and the communication strategy (using random messages as the baseline). We plot the coefficients with classical 95% confidence intervals. The effects of power selection are substantially larger than different communication strategies, none of which are significantly different from random messages at the $p < .05$ level.

| | | Detection | |
|--------|-------|-----------|-------|
| | | TRUE | FALSE |
| Expert | TRUE | 20 | 8 |
| | FALSE | 19 | 4745 |

Table A1: Total 4,792 messages (from Human/Cicero to Human/Cicero) comparing TRUE/FALSE whether expert humans see as a lie and whether detected as a broken commitment by our detection.

| | | Expert | |
|----------------|-------|--------|-------|
| | | TRUE | FALSE |
| Lie Annotation | TRUE | 3 | 72 |
| | FALSE | 13 | 1523 |

Table A2: Total 1,611 human send-out messages comparing TRUE/FALSE in human lie annotation and in expert hand labeling.

B.4 Future experiments: AMR information Cicero

Since we have evidence that Cicero’s win weighs on its strategic rather than communication abilities (Section 3.3). To further study this, we want to downgrade Cicero’s communication and collect more human-Cicero games to see whether Cicero wins at the same rate (previously 84% against humans). We conduct 5–10 games using the same setup as in the Human-Cicero games (Section 5). The only difference is Cicero. We will limit Cicero communications from natural language to AMR information where it mostly captures *move intent*.

C Deception detection limitations

We want to discuss deception detection further here to state errors and limitations. Since we mentioned our *precision* for deception detection is quite low (Section 5.2), we hereby expand on detection limitations and also compare to human (deliberate) lies as follows:

1. what our detection is likely to miss when humans lie,
2. what our detection mistakenly detects as deception,
3. what humans annotate as *Truth*, though it is a *break of commitment* and our detection can detect correctly.

Humans often lie about relationships. Detecting broken commitment at the relationship level is not possible for our detection (Table 3 and Table A4). This is a limitation of our deception definition, which focuses on moves. Though it is pos-

| | | Expert | |
|--------------------------|-------|--------|-------|
| | | TRUE | FALSE |
| Perceived Lie Annotation | TRUE | 5 | 284 |
| | FALSE | 7 | 1572 |

Table A3: Total 1,868 humans received messages comparing TRUE/FALSE whether humans perceived as a lie and whether human experts see as a lie.

| Sender | Message |
|---------|--|
| Austria | That’s an interesting opening. Was the bounce in EC planned? |
| Austria | Do you think Germany will work with you against France? |
| England | Yeah it would be great if we team up |

Table A4: The broken commitment detector (BC(\cdot)) cannot detect deception in alliance agreement when Austria (human) deceives England.

sible to extract the relationship among players to see conflicts in the messages, we avoid doing so because the relationship is another topic to study in more detail. At this stage of our work, we cannot train a model predicting relationships that can be circulated from game states, dialogue, and moves without collecting human data first. Therefore, we have relationship tracking from human players for a study in the future.

AMR limits broken commitment detection precision. Some messages are parsed incorrectly, which can be seen as a *commitment is broken* (Table A5). This makes the detection falsely detect *truthful* messages as deceptive (increases false positive examples which decreases *precision*). Another limitation we observed is when one accepts the proposal but does not follow as commit using a short answer, e.g. *Yes, I agree.* or *Sure.* Our AMR parser sometimes hallucinates and extracts invalid moves, which can be mistakenly detected as breaking a commitment.

Human lie annotation is not always correct. It is true that we have human annotations, and they can be seen as *ground truth*. However, we sample annotations from four games data and comparing to expert labeling (lies in Table A2 and perceived as lies in A3). This shows that humans are not good at predicting lies, and sometimes they are *honest* but then decide to break their words later. There are examples where humans commit to such action but do not follow, though they firstly annotate as a *truthful* message (Table A6 and Table A7).

| Sender | Message |
|--------|---|
| Turkey | If you retreat from Serbia into Budapest, then I'm in |
| Italy | I will do that if Serbia gets dislodged |

Table A5: Italy agrees with the condition that the Turkey unit should move out of Serbia; however, our AMR parser captures Italy's sentence as "*I will move to Serbia*," which is invalid and makes our detection detects deceptive when Italy does not move to Serbia.

| Sender | Message |
|---------|---|
| Germany | Also, can we keep Burgundy clear? |
| France | Yes, we can do that. Are you moving to Helgoland? |

Table A6: France (human) annotated "*Yes, we can do that*," as **Truth**, which contradicts the final move where France moves to Burgundy. This is captured as a broken commitment by the BC(\cdot) function.

D Survey Details

The survey consists of 5-point Likert scale questions and free-form text questions. The questions are designed to measure the human players' perception of Cicero's communication and their experience playing with Cicero. We also included questions to measure the players' experience with Diplomacy and their general impression of Cicero for qualitative analysis. Overall, players believe that human communicate more transparently and are more strategically cooperative. Survey results are shown in table A8.

| Sender | Message |
|---------|--|
| Germany | I am going to try to move to English Channel |
| England | Sure |
| Germany | It might help you hold London |
| England | Yeah I am holding London |

Table A7: England (human) annotate "*Yeah I am holding London*" as **Truth**, which contradicts the final move where an army in London moves to Edinburgh. This is captured as a broken commitment by the BC(\cdot) function.

| Statement | Likert Scale (%) | | | | | Num. Responses |
|--|------------------|------|------|------|------|----------------|
| | 1 | 2 | 3 | 4 | 5 | |
| I am really good at Diplomacy. | 0 | 8.3 | 25 | 41.7 | 9 | 25 |
| I am able to identify all AIs. | 9.5 | 23.8 | 38.1 | 16.7 | 11.9 | 42 |
| I enjoy talking with the AIs. | 14.3 | 38.1 | 33.3 | 7.1 | 7.1 | 42 |
| I was able to make plans with other players. | 7.1 | 23.8 | 35.7 | 14.3 | 19 | 42 |
| I was able to make plans with the AIs. | 21.4 | 31 | 19 | 19 | 9.5 | 42 |
| Human players communicated transparently. | 7.1 | 14.3 | 33.3 | 35.7 | 9.5 | 42 |
| AI players communicated transparently. | 11.9 | 26.2 | 45.2 | 9.5 | 7.1 | 42 |

Table A8: Statements in the survey and their respective responses. Larger number in the Likert scale indicates more agreement.




Figure A5: The average human player loses to Cicero. Human loses supply centers as game progresses unless playing as France, whereas Cicero's supply center count rises except playing as England. Cicero makes better strategic decisions.

| Sender | Message |
|---------------|--|
| Germany | This worked out great! |
| Germany | Can we please talk about our moves a bit? It's very hard to coordinate with silence. |
| France | Absolutely! I'm all ears! What do you want to do now! |
| France | Any designs on Norway? I think you could get it this turn. I'm gonna go against England, as you see. Let's work together on England. |
| France | Come on mate, let's do better this turn and coordinate! What are your plans? I'm with you. |
| France | Alright, let's use the additional time! What do we do? |
| Germany | I'm moving Sweden to Norway. |
| Germany | Can we also start DMZing our border? |
| France | Nice, with support from Hel that should work out. |
| France | I'm not gonna move out of Belgium but I'll certainly not move any further either. I'm in against England. Can't fight both of you that's for sure. |
| Germany | You should probably move Marseilles -> Spain. |
| France | Thank you! England might bring a fleet down? Good thought. Thank you! |

Table A9: A conversation between France and Germany. They agree to DMZ (demilitarizing) their borders, e.g., Ruhr, and cooperate moves. However, Germany is deceptive and would rather move into Ruhr in this turn.