

The Propaganda Machine

Dulfer, Rafael Malach
University of Twente
PO Box 217, 7500 AE Enschede
The Netherlands
r.m.dulfer@student.utwente.nl

ABSTRACT

Recent developments in the political landscape has shown us the many ways in which technology can be abused for propagandist purposes. In an attempt to better understand how computer science, especially the field of Natural Language Generation, can be abused to influence the opinions of others, a system is created and evaluated which generates propagandist texts based on a Risk game. The system, while severely limited in scope, is able convince some test subjects, pointing at the terrifying ability that a more sophisticated system might have.

Keywords

bias, reporting, data-to-text, automatic language generation, war, propaganda

1. INTRODUCTION

Natural Language Generation has been a topic of interest for Computer Science and Linguistics researchers for decades[10]. Over time, Natural Language Generation systems, also known as NLG systems, have been employed for a variety of purposes, including the reporting of weather forecasts, stock market reports[10], football results[13], reports on health data[5] and general news reporting[7].

While most of the uses of NLG systems are to generate rather benign content, the possibility for abuse is quite prominent. For example, a political organization could use such a system to generate propaganda and spread misinformation for their own political gains. While most current research tries to find methods to **detect** such fake-news, such as Zellers et al's (2019) GROVER which attempts to use machine learning for this purpose, even those detection methods could be used to generate the very thing it tries to fight[14]. At the same time, seemingly innocuous research into writing more affective or personalized texts, such as performed by Chen et al. who generated politically biased news articles[3] or Goudbeek et al. who investigated the impact of emotion on language use[6], could evolve into propaganda machines. Instead, it might be better to not beat around the proverbial bush and attempt to create such a propagandist machine ourselves, in the protected context of a scientific environment. Creating a controlled toy example can help us demonstrate and understand the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

33rd Twente Student Conference on IT July. 3rd, 2020, Enschede, The Netherlands.

Copyright 2020, University of Twente, Faculty of Electrical Engineering, Mathematics and Computer Science.

potential effects of using NLG to manipulate or fabricate facts for nefarious ends. Using such a toy example, we can ask ourselves the question:

How well would an NLG system be able to manipulate the perception of factual data?

2. THE PROPAGANDA SYSTEM

If we want to explore the above question, we obviously need a system which could generate biased propaganda and as such a system does not yet exist¹, it must first be created.

In the world of Natural Language Generation, a number of different approaches exist at tackling the same problem. While Neural Network based systems like GPT-2 are, arguably, the state of the art in making computers generate text[4], it was instead opted to go for a more basic template-based data-to-text system. Not only is such a system relatively simple to set up and describe, it also limits the ability of the system in such a way that it could never actually be used for real propagandist purposes. In this way, the research is kept to an ethically acceptable minimum viable product. It should also be noted that a lot of automated journalism systems that are currently in use work in a similar fashion[7].

Generally, a data-to-text system works by taking some set of structured data, for example weather information. Filtering out relevant information through a process called "content selection" and then filling the information into pre-made templates[10]. While the prominence of template-based data-to-text systems might give off the impression that one could simply copy an existing system, these systems are usually strongly tied to a specific domain[12]. Since there is little in the form of structured war data, propagandist templates and propagandist content selection, these must all be created from scratch.

2.1 Data Gathering

For the data it was imperative to have a clear, structured set of data. Besides this it was also important to have a dataset with as little ethical baggage as possible. Modern wars are quite heavily documented, but carry a lot of ethical baggage, while older wars which have lost their relevance over time are also less comprehensively documented. Instead of taking real war data, it was opted to go for fictional war data retrieved from a game of Risk. This fit the "toy example" approach taken for the system and also helped further bind to system to a toy domain. Moving from game data to real war data would be a non-trivial task after all.

¹at least not openly.

2.1.1 The Risk Game

The data itself was taken from the Free Open-Source Risk clone, Domination². Risk is a classic board-game which tasks its players with attempting to take over the world. In this way, it serves as a toy war-simulator. Domination proved to be a good candidate not only because of its open-source nature making it easy to hook in a data gathering system, but also because of its inclusion of AI players and a command-line based interface, making it trivial to automate the data gathering process.

A game of Risk has 6 actions that can occur:

1. Fortification, denoted as “fortify”, new units are placed on a country.
2. Attacking, denoted as “attack”, a player attacks another player from country A to country B.
3. Moving, denoted as “move”, a player moves units from country A to country B.
4. Get a card, denoted as “card”, a player gains a card if they conquered a country this turn, could later be traded for more units.
5. Trade cards, denoted as “trade”, a player can trade 3 cards of similar type to gain more units.
6. Complete a mission, denoted as “goal”, a player completes a specific goal, this action possibly wins them the game.

The newly added data gathering system simply stored each of these action taken by the players and wrote them to a json formatted file.³ This data could then later be used by the content-selection and templating systems to generate the texts.

In the context of this Risk game, the final goal of the system then becomes to attempt to convince the reader that a losing player is actually winning.

2.2 Content Selection

A single game of Risk creates a lot of data. Analysis of our own data shows that on average there are about 489 events in a single game of Risk. While it would be completely possible to generate texts for the whole game, it would generate texts that are too long for the purposes of this research. As such, a content selection method was needed to drastically cut down the amount of data that was reported on. The first step was to simply jump directly to the middle of the game. It is usually in the middle of the game where things are most exciting, the winner will start to emerge, but is not yet completely obvious. Reporting only on this part of the game allows the program to immediately skip to the most interesting parts.

Next it was decided to report on a per-turn basis. Because Risk is a sequential game, each player performs a set of actions until they no longer can or care to, after which the other player is allowed to perform a set of actions. This set of actions is what we can define as a turn. Grouping the data by turns brings most of the relevant data together and allows us to make a clear distinction between when actions are performed by a certain player and when actions are performed **against** this player.

²<http://yura.net/projects/domination>

³For an example of the file structure, see Pseudocode 6. For a full example, see rafael.thebias.nl/research/wardata.json

Finally, the decision was made to attempt to group related data together. In order to group related data together it must first be defined what exactly makes data related. Given a single player’s turn, the most interesting events will always be related to a specific war. A war can then be defined as a number of attacks from one country A to country B. We can then conclude that any previous actions also involving country A will be relevant to this war. This is the first step of relevancy. The second step is to consider whether a war is continuing. If the player is able to use country A to conquer country B, then not only are all actions involving country A relevant, but also all actions involving country B. It is possible that country B is now used to invade country C. A continuing war like this will keep being grouped recursively until the attacks stop. In this way, a single turn could have multiple wars, all of which have related events being grouped together (See figure 4 for a graphical description). Grouping these related events together allows us to make a more cohesive story. A single war and its consequences can be described in a paragraph, only for the next paragraph to then describe the next war.

As a result of this content selection and grouping, the final output of the system will describe the events in the game around the midpoint, turn by turn, war by war, keeping all relevant information together to allow for a more cohesive narrative flow.

3. THE PROPAGANDA NARRATIVE

The second part of any data-to-text system after the data and content selection part is, obviously, the text. Generally, these systems heavily rely on the existence of corpora containing texts in the relevant domain as inspiration or to copy verbatim for their templates[4]. Surprisingly, there did not seem to be an existing well-organized corpora of propaganda. There has however been a lot of research done in the field of propaganda, its effects and what makes propaganda effective. It was these papers, together with R Bytwerk’s German Propaganda Archive⁴ which provided the most inspiration for the templates.

3.1 The Nature of Propaganda

Propaganda in the real world takes many forms, most of which are completely irrelevant to the domain of a Risk game. Renowned propaganda theorist Harold D Lasswell (1927) defines there to be two main groups of propaganda, inter-group and extra-group[8]. Inter-group is what most people think of when it comes to propaganda, as it describes the type of propaganda which targets the own populace (A famous example of inter-group propaganda is Rosie the Riveter, see figure 5). For the purposes of this system however, we must generate extra-group propaganda as the readers belong to a significant “out-group” by virtue of not being part of the in-game, fictional, Risk world for which the propaganda is written. Lasswell likens this type of propaganda to religious conversion.

It is this specific fictional nature that proved the writing of the templates to be significantly problematic. In the real world, propaganda often serves two purposes, to arouse positivity towards the self and negativity towards the other[8]. This is frequently done through appeals to common virtues, history and ideals as well as by highlighting the enemies perceived depravity[11]. The created templates however, had to serve a different goal. Instead of trying to invoke emotional hostility towards the enemy, it

⁴<https://research.calvin.edu/german-propaganda-archive/>

had to attempt to convince the reader that the enemy was losing the war which is only a very specific subset of propaganda. Still, there is precedence for propaganda attempting to spin losses and bad battles into good news[2][1] and painting a negative picture of the enemy can still affect the judgement of readers even if they have no emotional investment[9].

3.2 The Templates

The final part of the system comes in the form of the templates. The templates take the data generated and selected in the previous steps and create full sentences from them. By combining these sentences together you can create a full text. The system can create two types of reports, a neutral report and a biased/propagandist report. The templates follow the default Python String Formatter format, meaning that any string of characters between brackets ({}) are interpreted as variables. The format was expanded so that it also allows any python code inside to brackets to be run. For example, "the sum of a+b={a+b}" given a=2, b=3 will return "the sum of a+b=5".

3.2.1 The Basic Sentences

The most important and interesting part of the whole research is the propagandist report. The propagandist report is built up from various basic sentences which follow a flow as defined by the grouping described in 2.2. The templates are then split into two different types, negative and positive. In the case a positive event has occurred, a positive template is picked. In the event a negative event has occurred, a negative template is picked. Given a specific player, the following algorithm decides if an event is positive:

Pseudocode 1: is_positive Algorithm

```

1  input: current player, event
2  output: boolean
3  begin
4    if event is an attack
5      if country is not conquered
6        if attacking player is the current player
7          if enemy has lost less than 50% of units stationed
8            return False # Our attack was unsuccessful, not
            ↪ positive
9        else
10       return True # Our attack was successful even if
            ↪ we failed to take the country, positive
11     else
12       return True # Enemy failed at conquering out
            ↪ country, positive
13     else if country is conquered and attacking player is
            ↪ the current player
14       return True # We conquered a country, positive
15     else
16       return False # They conquered a country, negative
17     else if event is not an attack and player is the
            ↪ current player
18       return True # We did something else, positive
19     else
20       return False # They did something else, negative
21  end

```

After an event has been designated as positive or not, it picks a template from the list of templates depending on the type of event. If the event was not an attack, we can simply pick a random template from the list (see Pseudocode 7 for examples). If the event was an attack however, we have two more attributes to determine. First thing to determine is the style of the attacks of which there are three types:

- Steamroll, when our army is big and their army is small

- Underdog, when their army is big and our army is small
- Struggle, when both armies are equally matched

The following algorithm decides the style of an attack:

Pseudocode 2: get_style Algorithm

```

1  input: current player, event
2  output: string
3  begin
4    enemy_units ← amount of enemy units involved in this
            ↪ attack
5    player_units ← amount of player units involved in this
            ↪ attack
6    if enemy units are less than half of the player units
7      if attacking player is the current player
8        return "steamroll" # We attack with many, they
            ↪ defend with few
9      else
10       return "underdog" # They attack with few, we
            ↪ defend with many
11     else if enemy units are more than double of the player
            ↪ units
12       if attacking player is the current player
13         return "underdog" # We attack with few, they
            ↪ defend with many
14       else
15         return "steamroll" # They attack with many, we
            ↪ defend with few
16     else
17       return "struggle"
18  end

```

Finally, the nature of the attack must be determined, of which there are four:

- Win, Our army wins a country
- Loss, Our army loses a country
- Defend, Our army successfully defends a country
- Progress, Progress was made in the war, but neither army was victorious

After all of these qualities are determined, the system can pick a template sentence from the list that fits these qualities (see Pseudocode 8 for examples).

3.2.2 Connecting The Sentences

The above system does not make for a fully coherent text yet, it still lacks any sort of cohesion. This cohesion is added to the text by a system of what has been called "fluff" text. Fluff text serves no pragmatic purpose, but instead connects separate events together. It is these fluff lines which allow for the text to become a cohesive narrative whole.

The fluff text function on a number of "scopes" which correspond to the scopes described in section 2.2. Just as with the event templates, there are positive and negative fluff texts. Below this level, there are the "Introduction" and "Continuing" texts. An introduction text is just as it says, it serves to be the introduction for the entire text corresponding to this turn. A turn can be split up into several paragraphs, one paragraph for each group of war data. The continuing texts look at the previous groups of data, see whether they were generally positive or negative, and connect the two paragraphs together.

Pseudocode 3: First Layer Connecting Template Examples

```

1  "positive": {
2    "introduction": [

```

```

3     "The glory of god has smiled upon us again!",
4   ],
5   "continue": {
6     "negative": [
7       "However, not all is lost!",
8     ],
9     "positive": [
10      "And that's not all. ",
11    ]
12  },
13 },
14 "negative": {
15   "introduction": [
16     "God has been testing our brave soldiers.",
17   ],
18   "continue": {
19     "positive": [
20       "Not all is well in the state of {curr_player}
21         ↳ however.",
22     ],
23     "negative": [
24       "Regrettably, we were tested even more.",
25     ]
26   }

```

The second layer of connecting sentences comes in the form of connecting related events together. As described previously, all events related to a "war" can be grouped together. We can connect these events into a single paragraph by using more linking sentences.

Pseudocode 4: Second Layer Connecting Template Examples

```

1 "positive": {
2   "opening": [
3     "We have received news from the {country1} front. ",
4   ],
5   "intermittent": [
6     "Those despicable {other_player}ns must be crying to
7       ↳ their mothers.",
8   ],
9   "continuing": [
10    "But that is not all! ",
11  ],
12  "conclusion": [
13    "Our victory is assured! "
14  ],
15  "reason": [
16    "This was made possible because",
17  ]
18 },
19 "negative": {
20   "opening": [
21     "Tragedy has struck {country2}.",
22   ],
23   "intermittent": [
24     "Be assured that this is only a minor setback.",
25   ],
26   "continuing": [
27     "The cowardly {other_player}ns pushed their luck
28       ↳ even more.",
29   ],
30   "conclusion": [
31     "We will never surrender!",
32   ],
33   "reason": [
34     "We have heard some more reports from our spies. ",
35   ]

```

Most of these templates are used to connect various "attack" events together. "opening" contains lines that start a paragraph, "conclusion" those that end paragraphs. "reason" is used to connect a specific attack event to other type of events like fortifications or movement events. "intermittent" is for extra fluff lines regardless of event type and "continuing" helps connect a continuing advance from country A to country B.

With all of this, Figure 1 and Figure 2 can be created.⁵

3.2.3 The Neutral Report

As a contrast to the propagandist report, the neutral report is a very basic report which just shows all data, uninteresting and unbiased. It goes through every action in the game sequentially and describes them. Its purpose is to be used in the evaluation to present readers with an unbiased point of view as a contrast to the highly biased propagandist reports created before.

Because of the very basic nature of the reports, the templates themselves are also very basic. Every type of event has a single line dedicated to it and when this event is encountered the line is picked and filled in.

Pseudocode 5: Neutral Templates

```

1 "fortify": "{curr_player} has fortified {country1} with
2   ↳ {amount} units. It had {unit.before} units before
3   ↳ and now has {unit.after}",
4 "attack": "{curr_player} has attacked {other_player} in
5   ↳ {country2} from {country1}. {country2} had
6   ↳ {init_defend} units before and {after_defend}
7   ↳ units afterwards. {country1} had {init_offend}
8   ↳ units before and {after_offend} units afterwards.
9   ↳ {curr_player} rolled {rolls.offend} and
10  ↳ {other_player} rolled {rolls.defend}.",
11 "move": "{curr_player} moved {amount} units from
12   ↳ {country1} to {country2}.",
13 "card": "{curr_player} got the following card {card}.",
14 "trade": "{curr_player} traded in {cards} and received
15   ↳ {amount} units in return.",
16 "goal": "{curr_player} completed their mission to
17   ↳ {mission}."

```

Using these templates, lines like in Figure 3 can be created.

4. EVALUATION

Having created the system that generates propaganda, the final step is to test the quality and effect of the resulting texts. The evaluation took on a variety of different forms, all with the same final goal. To test how well the propagandist texts were able to convince the reader that the player who was losing, was actually winning.⁶

The evaluation concerned 4 texts, all generated by the system. The game itself concerned the two players of Romania and Mongolia which both attempted to control the entire world, resulting in Mongolia being the eventual winner. As such, the goal of the evaluation became to try to convince the readers that Romania was actually winning, even if it was not. The texts presented to the subjects were, in varying orders, a neutral report⁷, a biased report for Romania⁸, a biased report for Mongolia⁹ and a background report¹⁰. The background report was a special report created for the evaluation which gave readers more info about the state of the game. Consenting test subjects were presented these texts and interviewed about their opinions.

⁵For a full overview of all templates, see rafael.thebias.nl/research/templates.json

⁶For sketched (anonymized) notes on all the interviews, see rafael.thebias.nl/research/interviews.txt.

⁷See https://rafael.thebias.nl/research/wardata_text_1.pdf

⁸See https://rafael.thebias.nl/research/wardata_text_2.pdf

⁹See https://rafael.thebias.nl/research/wardata_text_3.pdf

¹⁰See https://rafael.thebias.nl/research/wardata_background.pdf

The glory of god has smiled upon us again! Word has reached from the Eastern United States front. It took only 1 hour to make Western United States great again! We killed all enemy soldiers with ease. The weakling Mongolians have no chance against our brave soldiers. This was made possible because Our great strategists have decided to deploy more troops to Eastern United States. This will surely bring us ever closer to victory! We are strategically moving our troops in Eastern United States to Western United States. Eastern United States will be kept safe by the brave men staying behind, while Western United States will become stronger and safer than ever.

Figure 1: Positive Introductory Paragraph

It is sadly always calmest before the storm. Bad news from Madagascar. Mongolia showed that they can only win in an unfair matchup! Sending their biggest army to our smaller defense in Madagascar shows they are afraid of our true power! No matter, when they are faced with our true power they will see what real battle is like! Our great government supports the families of lost ones completely. We have discovered the enemies plans. Mongolia has moved troops from East Africa to Madagascar leaving East Africa completely defenseless for us to take. Is this Mongolia showing defeat?

Our victory is assured!

Figure 2: Negative Concluding Paragraph

Mongolia has fortified Central America with 7 units. It had 3 units before and now has 10

Mongolia has fortified North Africa with 10 units. It had 2 units before and now has 12

Mongolia has fortified Ukraine with 5 units. It had 3 units before and now has 8

Mongolia has attacked Romania in Eastern United States from Central America. Eastern United States had 5 units before and 5 units afterwards. Central America had 10 units before and 8 units afterwards. Mongolia rolled [3, 3, 1] and Romania rolled [5, 4].

Mongolia moved 3 units from North Africa to East Africa.

Figure 3: Example lines Neutral text

4.1 The Interviews

While there was some variation in the way different subjects were presented with different texts, the general skeleton of each interview stayed the same. Subjects were told that the research was about computer-generated texts based on a game of Risk. However, instead of telling them the true purpose of the evaluation, they were instead told it was about the quality of the text. This was done to make sure subjects would not be influenced into reading more critically than they usually would and also to mimic real propaganda, as propaganda seldom presents itself as such.

In general, subjects were given a biased report and instructed to read it completely. After the reading had finished, they would be asked some deflection open questions about quality and language use to keep up appearances that that was the true goal of the research. They would then be presented with a final question: “Who do you think is performing best?”. Subjects would then be quizzed further about why they think a certain player is doing better. After the first round of questions, the subject would be presented with a second text. This second text served as an opportunity to see how conflicting information might cause the subject to change their mind. After the second text, subjects would again be asked which player was performing the best and to elaborate on their decision. This final discussion marked the end of the interview, after which subjects were told the true purpose of the research.

4.1.1 The Variations

The evaluation saw multiple variations in the way texts were presented. The first version of the interview saw subjects presented with the background report at first, then the Romania report and finally the neutral report. This version concerned itself with how influenced people would be if they had a full overview of the state of the world and a full overview of the facts. This version generally resulted

in people accurately claiming that Mongolia was winning.

The second version was the same as the first version, but without the background text. This more accurately reflected a real world situation in which someone does not know everything about the state of the world. This version saw people more frequently call Romania the winner, even when presented with the objective facts.

The third version skipped out completely on the neutral report. Instead, subject were first shown the Romanian report and then followed with the Mongolian report. This version evaluated how conflicting propaganda would affect the reader. While people would at first generally point to Romania as the winner, after being shown the Mongolian text they would usually settle on Mongolia as the final overall winner.

The final version was a reversal of the third version. Subjects were first shown the Mongolian report and then the Romanian report. This version had mixed results.

4.2 Results

In the end, one out of three interviewees incorrectly pointed to Romania as the player who did best overall. Of the people who were only shown the Romanian text and the neutral text, only one in four saw through the propaganda and pointed to Mongolia as the winner. Some of the interviewees pointed out the propagandist nature of the texts, but this did not consistently affect their perception of the state of the game.

The neutral text was not able to convince any of the Romania voters that Romania was not doing best. While it made some subjects doubt their initial conviction, they usually found it too confusing and boring to read to really change their opinions.

Subjects frequently cited the abundance of repetition in the texts to be detrimental, either causing them to get

annoyed by the text or, in the worst case, feel like they were being lied to. Response to the propaganda was quite varied. Some correctly pointed out the propaganda and therefore went with the exact opposite of what the propaganda was trying to convince them of while others did not notice the propaganda at all or even directly referenced the propagandist lines as reasons why one side might be winning (“Mongolia is fortifying a lot, I think they might be afraid of Romania.”).

5. CONCLUSION

The propaganda turned out to work surprisingly well. While only a minority of those interviewed specifically pointed at Romania as the winner, those who did belong to this minority believed in the propaganda quite strongly.

We must remember that this research was limited in a way that actual propaganda machines would not be. A real propaganda machine can selectively show information, it can appeal to common symbols and it can be tailored to a specific population. If we take the variation which most closely resembles real-world propaganda scenarios (no full knowledge of the world, only shown one side of propaganda), then three in four people fell for texts that were quite simple and blatant in their propagandist contents. One could imagine that a more sophisticated machine, with less repetition, more professionally written templates and more real-world virtues to appeal to, could work even better. Such a system can be created (and possibly already is in the disguise of neutral journalism [7]).

So if the propaganda machine can exist, what can be done to combat it? Generally, the neutral text had no effect. The text that was the best at showing readers that Romania was losing regardless of propaganda was the background text. In the real world, however, such omnipresent knowledge of the state of the world is impossible. The text that was second-best at convincing readers Romania was actually losing was the Mongolia text, but fighting propaganda with propaganda, while it does have historical precedent, is hardly the most moral method.

Instead the best method would probably be to find some sort of middle ground between the neutral text and the background text. The neutral text was simply an information dump. It caused readers to get overloaded, bored, and they would generally ignore it. The background text was easy to read and understand, but it required an omniscient knowledge of the world which is impossible in the real world. If the raw data from the neutral text were to be presented in a more digestible, but still objective style, it could perhaps help convince people of the truth regardless of Propaganda.¹¹

5.1 Future Work

Because of a severely limited time frame, a lot of shortcuts had to be taken which negatively impacted the research. Most prominently, the amount of template lines written were too little to create sufficiently varying text and the evaluation had a severe lack of participants.

If this research were to be continued, the evaluation would have to be expanded. While there was some variation in the way texts were presented, a number of surrounding factors could still have influenced the results. There was also a significant lack of participants, which makes it impossible to perform statistically significant research. At

the same time, the system could have been improved a lot. There were frequent complaints about issues with flow and repetitive texts to the point that it negatively impacted the ability of the texts to convince the readers. With more time, these issues of flow and repetition could be fixed by using more sophisticated and varied templates.

There is still a lot of work to be done in this field.

Acknowledgements

The author would like to thank the following people:

Barbara Bulten, for helping to brainstorm template texts and trialling the evaluation process.

Lorenzo Gatti, for consistently providing feedback and helping the research stay on track and on schedule.

6. REFERENCES

- [1] R. Bytwerk. Rome, 1944.
- [2] R. Bytwerk and H. E. Seidat. Airplanes against England, 1940.
- [3] W.-F. Chen, H. Wachsmuth, K. Al-Khatib, and B. Stein. Learning to Flip the Bias of News Headlines. pages 79–88, 2019.
- [4] A. Gatt and E. Krahmer. Survey of the state of the art in natural language generation: Core tasks, applications and evaluation. *Journal of Artificial Intelligence Research*, 61:1–64, jan 2018.
- [5] D. Gkatzia. Content Selection in Data-to-Text Systems: A Survey. *CoRR*, abs/1610.0, 2016.
- [6] M. Goudbeek, N. Braun, C. Out, and E. Krahmer. Producing Affective Language: Content Selection, Message Formulation, and Computational Modelling. In *The Second International Conference on Human and Social Analytics*, pages 43–47, 2016.
- [7] A. Graefe. Guide to Automated Journalism. *Tow Center for Digital Journalism Report*, 2016.
- [8] H. D. Lasswell. The Theory of Political Propaganda. *American Political Science Review*, 21(3):627–631, aug 1927.
- [9] A. T. Little. Propaganda and credulity. *Games and Economic Behavior*, 102:224–232, mar 2017.
- [10] E. Reiter. An architecture for data-to-text systems. In *Proceedings of the 11th European Workshop on Natural Language Generation, ENLG 07*, pages 97–104, 2007.
- [11] S. Scriver. War Propaganda. In *International Encyclopedia of the Social & Behavioral Sciences: Second Edition*, pages 395–400. Elsevier Inc., mar 2015.
- [12] M. Theune, E. Klabbbers, J. R. De Pijper, E. Krahmer, and J. Odijk. From data to speech: a general approach. *Natural Language Engineering*, 7(1):47–86, mar 2001.
- [13] C. van der Lee, E. Krahmer, and S. Wubben. PASS: A Dutch data-to-text system for soccer, targeted towards specific audiences. In *Proceedings of the 10th International Conference on Natural Language Generation*, pages 95–104, Santiago de Compostela, Spain, jan 2018. Association for Computational Linguistics (ACL).
- [14] R. Zellers, A. Holtzman, H. Rashkin, Y. Bisk, A. Farhadi, F. Roesner, and Y. Choi. Defending Against Neural Fake News. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 9054–9065. Curran Associates, Inc., 2019.

¹¹You then reach the philosophical question of whether it is even possible to perform objective reporting but that question lies outside the scope of this research.

APPENDIX

A. PSEUDOCODE & FIGURES

Pseudocode 6: Data Structure

```

1  {
2    players: {
3      {PLAYER_NAME}: { // All the player data, used
4        ↪ later for the names of countries in texts and
5        ↪ for more insight into the state of the world.
6        capital: {COUNTRY_NAME},
7        mission: {MISSION_DESC},
8        initial_countries: [
9          {COUNTRY_NAME},
10         ...
11       ]
12     },
13     events: [ // There are 5 types of events
14       {
15         type: "fortify", // A country is fortified
16         player1: {PLAYER_NAME}, // Use of
17         ↪ player1/country1 for every type of event
18         ↪ makes reading the data easier
19         country1: {COUNTRY_NAME},
20         unit_before: n,
21         unit_after: n,
22         amount: n
23       },
24       {
25         type: "attack", // A country is attacked from
26         ↪ another country
27         player1: {ATTACKER_NAME},
28         player2: {DEFENDER_NAME},
29         country1: {ATTACKING_COUNTRY},
30         country2: {DEFENDING_COUNTRY},
31         init_offend: n, // units in attacking country
32         ↪ before attack
33         init_defend: n,
34         after_offend: n, // units in attacking country
35         ↪ after attack
36         after_defend: n,
37         rolls_offend: [ // Dice rolled by attacking
38         ↪ country
39         n,
40         ...
41       ],
42       rolls_defend: [
43         n,
44         ...
45       ],
46       {
47         type: "move", // Units are moved from one
48         ↪ country to the other
49         player1: {PLAYER_NAME},
50         country1: {ORIGIN_COUNTRY},
51         country2: {DESTINATION_COUNTRY},
52         init_origin: n, // Amount of units in origin
53         ↪ country
54         init_dest: n,
55         amount: n // Amount of units moved
56       },
57       {
58         type: "card", // Player receives a card
59         player1: {PLAYER_NAME},
60         card: "{TYPE} - {COUNTRY_NAME}"
61       },
62       {
63         type: "trade", // Player trades cards for units
64         player1: {PLAYER_NAME},
65         cards: [
66           "{TYPE} - {COUNTRY_NAME}",
67           ...
68         ],
69         amount: n // Amount of units received
70       },
71       {
72         type: "goal", // Player has achieved a
73         ↪ goal/mission
74         player1: {PLAYER_NAME},
75         mission: {MISSION_DESC},
76         win_game: {TRUE/FALSE} // Was this a
77         ↪ winning goal?
78       }
79     ]
80   }

```

```

72   countries: { // All countries on the game map, used to
73   ↪ reason about the state of the game
74   {COUNTRY_NAME}: {
75     continent: {CONTINENT_NAME},
76     id: {ID}, // Every country has an associated
77     ↪ in-game id,
78     neighbours: [ // Countries this country is
79     ↪ directly connected to
80     {NEIGHBOUR_NAME},
81     ...
82   ]
83 },
84 continents: { // All continents on the map, used to
85 ↪ reason about the state of the game
86 {CONTINENT_NAME}: {
87   id: {ID},
88   value: n, // How many extra units this continent
89   ↪ is worth
90   countries: [
91     {COUNTRY_NAME},
92     ...
93   ],
94   border_countries: [ // All countries which
95   ↪ border this continent but are not in it
96   {BORDER_NAME},
97   ...
98 ]

```

Pseudocode 7: Non-attack Template Examples

```

1  "positive": {
2    "card": [
3      "We have received news from our spies that
4      ↪ {other_player} is developing new
5      ↪ weaponry. Still, {country2} engineering
6      ↪ is no match for our own! ",
7    ],
8    "trade": [
9      "Confirmed reports show that the arrival of new
10     ↪ {other_player} troops is imminent. It is to
11     ↪ be seen where they will be deployed, but no
12     ↪ matter how many {country2} bodies they
13     ↪ throw at us, we will persevere!",
14   ],
15   "fortify": [
16     "Word has reached of {amount} new {other_player}
17     ↪ troops arriving in {country1}. It is no matter,
18     ↪ our victory is already assured! ",
19   ],
20   "move": [
21     "{other_player} has moved troops from {country1} to
22     ↪ {country2} leaving {country1} completely
23     ↪ defenseless for us to take. Is this
24     ↪ {other_player} showing defeat? ",
25   ],
26   "negative": {
27     "card": [
28       "We have received news from our spies that
29       ↪ {other_player} is developing new
30       ↪ weaponry. Still, {country2} engineering
31       ↪ is no match for our own! ",
32     ],
33     "trade": [
34       "Confirmed reports show that the arrival of new
35       ↪ {other_player} troops is imminent. It is to
36       ↪ be seen where they will be deployed, but no
37       ↪ matter how many {country2} bodies they
38       ↪ throw at us, we will persevere!",
39     ],
40     "fortify": [
41       "Word has reached of {amount} new {other_player}
42       ↪ troops arriving in {country1}. It is no matter,
43       ↪ our victory is already assured! ",
44     ],
45     "move": [
46       "{other_player} has moved troops from {country1} to
47       ↪ {country2} leaving {country1} completely
48       ↪ defenseless for us to take. Is this
49       ↪ {other_player} showing defeat? ",
50     ],
51   ]

```

28 }

Pseudocode 8: Attack Template Examples

```

1  "positive": {
2    "won": {
3      "steamroll": [
4        "Our glorious army has completely wiped away the
          ↳ enemy forces in {country2}!",
5      ],
6      "underdog": [
7        "Like David to Goliath our brave soldiers marched
          ↳ into what {other_player} thought was
          ↳ assured destruction, but {curr_player}ns
          ↳ know no fear and swiftly took down the
          ↳ giant in {country2}!",
8      ],
9      "struggle": [
10       "Our {ordinal(init_offend-after_offend+1)} battalion
          ↳ waltzed right into {country2} and wrestled
          ↳ control away from those cowardly
          ↳ {other_player}ns.",
11     ],
12   },
13   "progress": [
14     "Great progress has been made in the invasion of
          ↳ {country2}. It is only a matter of time
          ↳ before they crumble under our might!",
15   ],
16   "defending": {
17     "steamroll": [
18       "The \"brilliant\" {other_player} strategists
          ↳ showed their charity today as they marched
          ↳ their brittle army directly into our gunfire
          ↳ at {country2}. We should thank them for
          ↳ accepting defeat so easily!",
19     ],
20     "underdog": [
21       "The {other_player}ns must have thought taking
          ↳ {country2} would be easy. They
          ↳ underestimated the great {curr_player}
          ↳ spirit and cunning! They arrived with many
          ↳ and left with few.",
22     ],
23     "struggle": [
24       "The battle in {country2} showed that we will be
          ↳ victorious in even the most difficult
          ↳ combat. Our brave soldiers endured the
          ↳ attack of those cheating {other_player}ns
          ↳ with strong will and sacrifice for the
          ↳ motherland.",
25     ],
26   },
27 },
28 "negative": {
29   "loss": {
30     "steamroll": [
31       "Our armies in {country2} were met with more
          ↳ resistance than expected. This might give
          ↳ some the idea that we are losing, but
          ↳ nothing could be more wrong! We continue
          ↳ the attack on {country2} to free the people
          ↳ from the tyranny of {other_player} and we
          ↳ shall soon destroy their illegitimate
          ↳ occupation!",
32     ],
33     "underdog": [
34       "We had to march our small battalion of soldiers in
          ↳ {country1} to a battle in {country2} that
          ↳ seemed lost before it began. But we had to
          ↳ do it so we could weaken and distract those
          ↳ terrible {other_player}ns. Fear not, for we
          ↳ have lost but temporarily and will get
          ↳ revenge for our brave soldiers soon.",
35     ],
36   },
37   "progress": [
38     "When attempting to liberate {country2} today many
          ↳ of our heroic soldiers have had to give their
          ↳ lives for their country. They have not given
          ↳ their lives in vain. The enemy thinks they
          ↳ have won, but they are not prepared for the
          ↳ swift retribution that they will face."
39   ],
40   "defending": {
41     "steamroll": [
42       "After a great battle in {country2}, our leaders

```

↳ decided that we had to leave this part of
 ↳ the country in order to protect it's people
 ↳ and it's culture from destruction by the
 ↳ viscous {curr_player}. We have lost a
 ↳ number of good men, but we have heard
 ↳ rumours that they have lost even more. We
 ↳ shall return to {country2} stronger and
 ↳ better than ever before.",

```

43   },
44   "underdog": [
45     "{other_player} showed no respect as they marched
          ↳ their biggest army to steal our land in
          ↳ {country2}. They could achieve victory
          ↳ through cheating and unfair fighting. We
          ↳ shall see who is victorious when they face
          ↳ our even bigger armies!",
46   ],
47   "struggle": [
48     "After a long and hard struggle in {country2}, the
          ↳ enemy has marched into the country and
          ↳ started oppressing it's people. This is only
          ↳ a minor setback and we shall return with
          ↳ swift—full vengeance.",
49   ]
50 }
51 }

```




Figure 4: Diagram of the layers of data groupings, each circle represents a grouping.



Figure 5: Rosie the Riveter, J. Howard Miller 1942-1943, restored by Adam Cuerden, public domain