

A Survey on Computational Propaganda Detection

Giovanni Da San Martino^{1*}, Stefano Cresci², Alberto Barrón-Cedeño³,
Seunghak Yu⁴, Roberto Di Pietro⁵ and Preslav Nakov¹

¹Qatar Computing Research Institute, HBKU, Doha, Qatar

²Institute of Informatics and Telematics, IIT-CNR, Pisa, Italy

³DIT, Alma Mater Studiorum–Università di Bologna, Forlì, Italy

⁴MIT Computer Science and Artificial Intelligence Laboratory, Cambridge, MA, USA

⁵College of Science and Engineering, HBKU, Doha, Qatar

{gmartino, rdipietro, pnakov}@hbku.edu.qa, s.cresci@iit.cnr.it, a.barron@unibo.it,
seunghak@csail.mit.edu

Abstract

Propaganda campaigns aim at influencing people’s mindset with the purpose of advancing a specific agenda. They exploit the anonymity of the Internet, the micro-profiling ability of social networks, and the ease of automatically creating and managing coordinated networks of accounts, to reach millions of social network users with persuasive messages, specifically targeted to topics each individual user is sensitive to, and ultimately influencing the outcome on a targeted issue. In this survey, we review the state of the art on computational propaganda detection from the perspective of Natural Language Processing and Network Analysis, arguing about the need for combined efforts between these communities. We further discuss current challenges and future research directions.

1 Introduction

The Web makes it possible for anybody to create a website or a blog and to become a news medium. Undoubtedly, this is a hugely positive development as it elevates freedom of expression to a whole new level, giving anybody the opportunity to make their voice heard. With the rise of social media, everyone can reach out to a very large audience, something that until recently was only possible for major news outlets.

However, this new avenue for self-expression has brought also unintended consequences, the most evident one being that the society has been left unprotected against potential manipulation from a multitude of sources. The issue became of general concern in 2016, a year marked by micro-targeted online disinformation and misinformation at an unprecedented scale, primarily in connection to Brexit and the US Presidential campaign; then, in 2020, the COVID-19 pandemic also gave rise to the first global infodemic. Spreading disinformation disguised as news created the illusion that the information was reliable, and thus people tended to lower their natural barrier of critical thinking compared to when information came from different types of sources.

Whereas false statements are not really a new phenomenon —e.g., yellow press has been around for decades— this time things were notably different in terms of scale and effectiveness thanks to social media, which provided both a medium to reach millions of users and an easy way to micro-target specific narrow groups of voters based on precise geographic, demographic, psychological, and/or political profiling.

An important aspect of the problem that is often largely ignored is the mechanism through which disinformation is being conveyed, which is using *propaganda techniques*. These include specific rhetorical and psychological techniques, ranging from leveraging on emotions —such as using *loaded language*, *flag waving*, *appeal to authority*, *slogans*, and *clichés*— to using logical fallacies —such as *straw men* (misrepresenting someone’s opinion), *red herring* (presenting irrelevant data), *black-and-white fallacy* (presenting two alternatives as the only possibilities), and *whataboutism*. Moreover, the problem is exacerbated by the fact that propaganda does not necessarily have to lie; it could appeal to emotions or cherry-pick the facts. Thus, we believe that specific research on propaganda detection is a relevant contribution in the fight against online disinformation.

Here, we focus on *computational propaganda*, which is defined as “propaganda created or disseminated using computational (technical) means” [Bolsover and Howard, 2017]. Traditionally, propaganda campaigns had been a monopoly of state actors, but nowadays they are within reach for various groups and even for individuals. One key element of such campaigns is that they often rely on coordinated efforts to spread messages at scale. Such coordination is achieved by leveraging botnets (groups of fully automated accounts) [Zhang et al., 2016], cyborgs (partially automated) [Chu et al., 2012] and troll armies (human-driven) [Linville and Patrick, 2018], known as *sockpuppets* [Kumar et al., 2017], *Internet water army* [Chen et al., 2013], *astroturfers* [Ratkiewicz et al., 2011], and *seminar users* [Darwish et al., 2017]. Thus, a promising direction to thwart propaganda campaigns is to discover such coordination; this is demonstrated by recent interest by Facebook¹ and Twitter².

*Contact Author

¹newsroom.fb.com/news/2018/12/inside-feed-coordinated-inauthentic-behavior/

²https://help.twitter.com/en/rules-and-policies/platform-manipulation

In order for propaganda campaigns to work, it is critical that they go unnoticed. This further motivates work on detecting and exposing propaganda campaigns, which should make them increasingly inefficient. Given the above, in the present survey, we focus on computational propaganda from two perspectives: (i) the content of the propaganda messages and (ii) their propagation in social networks.

Finally, it is worth noting that, even though there have been several recent surveys on fake news detection [Shu et al., 2017; Zhou et al., 2019], fact-checking [Thorne and Vlachos, 2018], and truth discovery [Li et al., 2016], none of them focuses on computational propaganda. There has also been a special issue of the Big Data journal on Computational Propaganda and Political Big Data [Bolsover and Howard, 2017], but it did not include a survey. Here we aim to bridge this gap.

2 Propaganda

The term *propaganda* was coined in the 17th century, and initially referred to the propagation of the Catholic faith in the New World [Jowett and O'Donnell, 2012, p. 2]. It soon took a pejorative connotation, as its meaning was extended to also mean opposition to Protestantism. In more recent times, back in 1938, the Institute for Propaganda Analysis [Ins, 1938], defined propaganda as “*expression of opinion or action by individuals or groups deliberately designed to influence opinions or actions of other individuals or groups with reference to predetermined ends*”.

Recently, Bolsover et. al [2017] dug deeper into this definition identifying its two key elements: (i) trying to influence opinion, and (ii) doing so on purpose. Influencing opinions is achieved through a series of rhetorical and psychological techniques. Clyde R. Miller in 1937 proposed one of the seminal categorizations of propaganda, consisting of seven devices [Ins, 1938], which remain well accepted today [Jowett and O'Donnell, 2012, p.237]: *name calling, glittering generalities, transfer, testimonial, plain folks, card stacking, and bandwagon*. Other scholars consider categorizations with as many as eighty-nine techniques [Conserva, 2003], and Wikipedia lists about seventy techniques.³ However, these larger sets of techniques are essentially subtypes of the general schema proposed in [Ins, 1938].

Propaganda is different from *disinformation*⁴, in particular with reference to the truth value of the managed information and its goal, which in disinformation are (i) *false*, and (ii) *intending to harm*, respectively. The (often-neglected) intention to harm popped up in 2016, due to both the Brexit referendum and the US Presidential elections, when society and academia discovered that the news cycle got *weaponized* by disinformation. Contrarily, propaganda can hook to claims that are either true or false, and its intended objectives can be either harmful or harmless (even good⁵). In practice, propaganda and disinformation are used synergetically to achieve specific objectives, effectively turning social media into a weapon. Another related concept is that of “fake news”, where the focus is on a piece of information being factually false.

Although lying and creating fake stories is considered as one of the propaganda techniques (some authors refer to it as “black propaganda” [Jowett and O'Donnell, 2012]), there are contexts where this course of actions is often done without pursuing the objective to influence the audience, as in satire and clickbaiting. These special cases are of less interest when it comes to fighting the weaponization of social media, and are therefore considered out of the scope for this survey.

3 Text Analysis Perspective

Research on propaganda detection based on text analysis has a short history, mainly due to the lack of suitable annotated datasets for training supervised models. There have been some relevant initiatives, where expert journalists or volunteers analyzed entire news outlets, which could be used for training. For example, *Media Bias/Fact Check* (MBFC)⁶ is an independent organization analyzing media in terms of their factual reporting, bias, and propagandist content, among other aspects. Similar initiatives are run by *US News & World Report*⁷ and the European Union.⁸ Such data has been used in distant supervision approaches [Mintz et al., 2009], i.e., by assigning each article from a given news outlet the label propagandistic/non-propagandistic using the label for that news outlet. Unfortunately, such coarse approximation inevitably introduces noise to the learning process, as we discuss in Section 5.

In the remainder of this section, we review current work on propaganda detection from a text analysis perspective. This includes the production of annotated datasets, characterizing entire documents, and detecting the use of propaganda techniques at the span level.

3.1 Available Datasets

Given that existing models to detect propaganda in text are supervised, annotated corpora are necessary. Table 1 shows an overview of the available corpora (to the best of our knowledge), with annotation both at the document and at the fragment level.

Rashkin et al. [2017] released TSHP-17, a balanced corpus with document-level annotation including four classes: *trusted*, *satire*, *hoax*, and *propaganda*. TSHP-17 belongs to the collection of datasets annotated via distant supervision: an article is assigned to one of the classes if the outlet that published it is labeled as such by the *US News & World Report*. The documents were collected from the English Gigaword and from seven unreliable news sources.

According to Barrón-Cedeño et al. [2019] the low amount of sources considered per class is a downside of TSHP-17, as the systems trained on it might be modeling the news outlets, rather than propaganda itself (or any of the other three classes). To cope with this limitation, Barrón-Cedeño et al. released QProp a twice-as-big binary imbalanced dataset in which ~ 10% of the articles belong to class *propaganda*.

⁶<http://mediabiasfactcheck.com>

⁷www.usnews.com/news/national-news/articles/2016-11-14/avoid-these-fake-news-sites-at-all-costs

⁸www.usnews.com/news/national-news/articles/2016-11-14/avoid-these-fake-news-sites-at-all-costs; <http://euvsdisinfo.eu/>

³http://en.wikipedia.org/wiki/Propaganda_techniques

⁴<http://eeas.europa.eu/topics/countering-disinformation/en>

⁵Think of Greta Thunberg's highly propagandistic speech at the UN in 2019.

Corpus	Level	Sources	Classes	Articles	Prop.
TSHP-17	document	11 (2)	4	22,580	5,330
QProp	document	104 (10)	2	51,294	5,737
PTC	text span	49 (13)	18	451	7,385

Table 1: Textual datasets available to train supervised propaganda identification models at different granularity levels.

Once again, the annotation in QProp is obtained by distant supervision; this time with information from MBFC. Aside the binary propaganda *vs.* trustworthy annotation, in QProp each article has associated metadata from its source such as the bias level (e.g., *left*, *center*, *right*) from MBFC and geographical information, average sentiment, publication date, identifier, author, and official source name from GDELT.⁹

However, both TSHP-17 and QProp lack information about the precise location of a propagandist snippet within a document. Since propaganda is conveyed by using specific rhetoric and psychological techniques, a separate line of research recently aimed to identify the use of such techniques. In particular, Da San Martino et al. [2019b] proposed a dataset with assets that previously available resources lacked. First, their PTC corpus is manually judged by professional annotators, rather than using distant supervision. Second, the annotation is at the fragment level: specific text spans are flagged, rather than full documents. Third, it goes deeper into the types of propaganda, considering 18 propaganda techniques, rather than the binary *propaganda vs. non-propaganda* setting. The curated list of techniques is summarized in Table 2. Whereas the volume of PTC is way lower than that of TSHP-17 and QProp—a few hundred articles against thousands—it contains more than 7,000 propagandist snippets. See Figure 1 for an example with annotations.

Another relevant line of research is on computational argumentation, which deals with some logical fallacies considered to be propaganda techniques. Habernal et al. [2017] described a corpus with 1.3k arguments annotated with five fallacies such as *ad hominem*, *red herring*, and *irrelevant authority*.

3.2 Text Classification

Early approaches to propaganda identification are fairly aligned to the produced corpora. Rashkin et al. [2017] defined a classical four-classes text classification task: propaganda *vs* trusted *vs* hoax *vs* satire, using the TSHP-17 dataset. Using word *n*-gram representation with logistic regression, they found that their model performed well only on articles from sources that the system was trained on.

Barrón et al. [2019] used a binary classification setting: detecting propaganda *vs* non-propaganda and experimented on TSHP-17 and QProp corpora. They ran a massive set of experiments, investigating various representations, from writing style and readability level to the presence of certain keywords, together with logistic regression and SVMs, and confirmed that using distant supervision, in conjunction with rich representations, might encourage the model to predict the source, rather than to discriminate propaganda from non-propaganda.

⁹<https://www.gdeltproject.org/>

Technique	Definition
Name calling	attack an object/subject of the propaganda with an insulting label
Repetition	repeat the same message over and over
Slogans	use a brief and memorable phrase
Appeal to fear	support an idea by instilling fear against other alternatives
Doubt	questioning the credibility of someone/something
Exaggeration/minimizat.	exaggerate or minimize something
Flag-Waving	appeal to patriotism or identity
Loaded Language	appeal to emotions or stereotypes
Reduction ad hitlerum	disapprove an idea suggesting it is popular with groups hated by the audience
Bandwagon	appeal to the popularity of an idea
Casual oversimplification	assume a simple cause for a complex event
Obfuscation, intentional vagueness	use deliberately unclear and obscure expressions to confuse the audience
Appeal to authority	use authority’s support as evidence
Black&white fallacy	present only two options among many
Thought terminating clichés	phrases that discourage critical thought and meaningful discussions
Red herring	introduce irrelevant material to distract
Straw men	refute argument that was not presented
Whataboutism	charging an opponent with hypocrisy

Table 2: List of the 18 propaganda techniques and their definitions.

<p>Manchin says Democrats acted like babies at the SOTU</p> <p>Democrat West Virginia Sen. Joe Manchin says his colleagues’ refusal to stand or applaud during President Donald Trump’s State of the Union speech was disrespectful and a signal that the party is more concerned with obstruction than it is with progress. In a glaring sign of just how stupid and petty things have become in Washington these days, Manchin was invited on Fox News Tuesday morning to discuss how he was one of the only Democrats in the chamber for the State of the Union speech not looking as though Trump had killed his grandma</p>	<ol style="list-style-type: none"> 1. Stereotyping, name calling or labeling 2. Black-and-white fallacy 3. Loaded language 4. Exaggeration 5. Loaded language
--	--

Figure 1: Text excerpt with annotated propaganda techniques.

They advocated for providing assurance that test data come from news sources that were not used for training, and investigated what representations remain robust in such a setting.

3.3 Detecting the Use of Propaganda Techniques

Da San Martino et al. [2019b] defined two tasks, based on annotations from the PTC dataset: (i) binary classification—given a sentence in an article, predict whether any of the 18 techniques has been used in it; (ii) multi-label multi-class classification and span detection task—given a raw text, identify both the specific text fragments where a propaganda technique is being used as well as the type of technique. Such a fine-grained level of analysis may provide support and explanations to the user on why an article has been judged as propagandistic by an automatic system. The authors proposed a multi-granularity deep neural network that modulates the signal from the sentence-level task to improve the prediction of the fragment-level classifier.

A shared task was held within the *2019 Workshop on NLP4IF: censorship, disinformation, and propaganda*¹⁰, based on the PTC corpus and the task definitions above. The best-performing models for both tasks used BERT-based contextual representations. Other approaches used contextual representations based on RoBERTa, Grover, and ELMo, or context-independent representations based on lexical, sentiment-based, readability, and TF-IDF features. Ensembles were also popular. Further details are available in the shared task overview paper [Da San Martino et al., 2019a].

4 Network Analysis Perspective

As seen in Section 3, the rhetoric techniques used for influencing readers’ opinions can be detected directly in the text. Contrarily, identifying the intent behind a propaganda campaign requires analysis that goes beyond individual texts, involving (among others) classification of the social media users that contributed to injecting and spreading propaganda within a network. Thus, a necessary condition to detect the intention to harm implies detecting malicious coordination (i.e., coordinated inauthentic behavior). Throughout the years, this high-level task has been tackled in different ways.

4.1 Early Approaches

Early approaches for detecting malicious coordination were based on classifying individual nodes in a network as either *malicious* or *legitimate*. Then, clusters of malicious nodes were considered to be acting in coordination. In other words, the concept of coordination was not embedded within the models, but it was added “a posteriori”. The vast majority of these approaches are based on supervised machine learning and each account under investigation was analyzed *in isolation*. That is, given a group of accounts to analyze, the supervised technique was separately applied to each account of the group, that in turn received a label assigned by the detector.

The key assumption of this body of work is that each malicious account has features that make it clearly distinguishable from legitimate ones. This approach to the task also revolved around the application of off-the-shelf, general-purpose classification algorithms. Widely used algorithms include decision trees and random forests, SVMs, boosting and bagging (e.g., Adaptive Boost and Decorate) and, more recently, deep neural networks [Kudugunta and Ferrara, 2018].

The most widely known example of this kind of detectors, is Botometer [Yang et al., 2019], a social bot detection system. By leveraging more than 1,200 features for a social media account, it evaluates profile characteristics, social network structure, the produced content (including sentiment expressions), and temporal features. Botometer simultaneously analyzes multiple dimensions of suspicious accounts for spotting bots. Instead, other systems solely rely on network characteristics [Yang et al., 2015], textual content [Rangel and Rosso, 2019], or profile information [Lee and Kim, 2014]. These latter systems are typically easier to game, since they only analyze a single facet of the complex, evolving behavior of bad online actors.

4.2 Evolving Threats

Despite having achieved promising initial results, these early approaches had several limitations. First, the performance of a supervised detector strongly depends on the availability of a ground truth (training) dataset. In most cases, a real ground truth is lacking and the labels are manually given by human operators. Unfortunately, as of 2020, we still have diverse and conflicting definitions of what a malicious account really is [Grimme et al., 2017], and humans have been proven to suffer from annotation biases and to largely fail at spotting sophisticated bots and trolls [Cresci et al., 2017]. To make matters worse, it has been demonstrated that malicious accounts “evolve” (i.e., they change their characteristics and behaviors) in an effort to evade detection by established techniques [Cresci et al., 2017]. Nowadays, sophisticated malicious accounts are using the same technological weapons as their hunters —such as powerful AI techniques— for generating credible texts (e.g., with GPT-2), profile pictures (e.g., with StyleGAN)¹¹, and videos (e.g., using deep-fakes), thus dramatically increasing their capabilities of impersonating real people, and hence of escaping detection.

4.3 Modern Approaches

The difficulties at detecting sophisticated bots and trolls with early approaches lead to a new research trend whose primary characteristic is to target groups of accounts as a whole, rather than focusing on individual accounts. In recently proposed detectors, coordination is considered a key feature to analyze, and it is modeled within the detectors themselves. The rationale for this choice is that malicious accounts act in coordination (e.g., sbots are often organized in botnets, trolls form so-called troll armies) to amplify their effect [Zhang et al., 2016]. Moreover, by analyzing large groups of accounts, modern detectors also have more data to exploit for fueling powerful AI algorithms [Sun et al., 2017]. The shift from individual to group analysis was accompanied by another shift from general-purpose machine learning algorithms, to ad-hoc algorithms specifically designed for detecting coordination. In other words, the focus shifted from feature engineering to learning effective feature representations and of designing brand-new and customized algorithms [Cai et al., 2017]. Many modern detectors are also unsupervised or semi-supervised, to overcome the generalization deficiencies of supervised detectors that are severely limited by the availability of exhaustive training datasets [Echeverría et al., 2018].

Examples of such systems implement network-based techniques, aiming at detecting suspicious account connectivity patterns [Liu et al., 2017; Chetan et al., 2019; Pacheco et al., 2020]. Coordinated behavior appears as near-fully connected communities in graphs, dense blocks in adjacency matrices, or peculiar patterns in spectral subspaces [Jiang et al., 2016]. Other techniques adopted unsupervised approaches for spotting anomalous patterns in the temporal tweeting and retweeting behaviors of groups of accounts —e.g., by computing metrics of distance out of the accounts activity time series and by subsequently account clustering [Chavoshi et al., 2016; Mazza et al., 2019].

¹⁰<http://www.netcopia.net/nlp4if/2019/>

¹¹<https://www.wired.com/story/facebook-removes-accounts-ai-generated-photos/>

The rationale behind such approaches is based on evidence suggesting that human-driven and legitimate behaviors are intrinsically more heterogeneous than automated and inauthentic ones [Cresci et al., 2020]. Consequently, a large cluster of accounts with highly similar behavior might serve as a red flag for coordinated inauthentic behavior. Distance (or similarity) between account activity time series was computed via dynamic time warping [Chavoshi et al., 2016], or as the Euclidean distance between the feature vectors computed by an LSTM autoencoder [Mazza et al., 2019]. More recently, other authors investigated the usefulness of Inverse Reinforcement Learning (IRL) for inferring the intent that drives the activity of coordinated groups of malicious accounts. Inferring intent and motivation from observed behavior has been extensively studied in the framework of IRL, with the main goal of finding the rewards behind an agent’s observed behavior. The inferred rewards can then be used as features in supervised learning systems aimed at detecting malicious and coordinated agents.

The switch from early to modern detectors demonstrated that the approach (e.g., individual vs group-based, supervised vs unsupervised) to the task of propaganda and malicious accounts detection can have serious repercussions on detection performance. However, some scientific communities naturally tend to favor a specific approach. For example, the majority of techniques that perform network analysis (e.g., by considering the social or interactions graph of the accounts) are intrinsically group-based. More often than not, they are also unsupervised. Contrarily, all techniques based on textual analyses, such as those that solely rely on natural language processing, are supervised detectors that analyze individual accounts [Rangel and Rosso, 2019]. As a consequence, some combinations of the cited approaches—above all, text-based detectors that perform unsupervised group analysis—are almost unexplored. For the future, it would thus be advisable to put efforts along the highlighted research directions that have been mostly overlooked until now.

5 Lessons Learned

The main lesson from our analysis is that there is a disconnection between NLP and Network Analysis communities when it comes to fighting Computational Propaganda, and therefore combined approaches may lead to systems significantly outperforming the current state of the art. A detailed analysis is reported in the following.

5.1 Text Analysis Lessons

From a text analysis perspective, we see that there is a lack of a suitable dataset for document-level propaganda detection. The attempts to use distant supervision as a substitute, by projecting labels from media to all the articles they have published is problematic in many aspects, even when done carefully. Indeed, distant supervision inevitably introduces noise in the learning process, as it is based on the wrong assumption that all articles from a given source would be either propaganda or non-propaganda. In reality, a propagandist source could periodically post objective non-propagandist information to boost its credibility.

Similarly, sources that are generally recognized as objective might occasionally post information that promotes a particular agenda. One way to deal with this issue might be to devise advanced learning algorithms, such as Generative Adversarial Networks (GANs), which can be trained to avoid specific biases, i.e., modelling the article source. Another issue with distant supervision is that while it is acceptable for training, it cannot give a fair assessment of a system at testing time, something that previous work has ignored.

Another lesson is that it seems more promising to focus on detecting the use of fine-grained propaganda techniques in text. Propaganda techniques are well-defined and well-known in the literature, and thus it makes sense to focus on them, as they are the very device on which propaganda is built. Notably, a proper dataset is already available for this new task, it is of reasonable size (350K tokens, which compares well to datasets for the related task of named entity recognition, whose typical size is 200K tokens), and covers a wide range of 18 commonly accepted techniques, comprising both various kinds of appeal to emotions as well as logical fallacies.

5.2 Network Analysis Lessons

Typically, when scholars and OSN administrators identify new coordinated behavior that goes undetected by existing techniques, as a reaction they start the development of new detectors. The implication of this reactive approach is that improvements occur only some time after having collected evidence of a new mischievous behavior. Bad actors thus benefit from a large time span—the time needed to design, develop, and deploy a new detector—during which they are free to tamper with our online environments.

A second lesson learned is related to the use of machine learning algorithms, the vast majority of which are designed to operate in environments that are stationary and neutral. Unfortunately, in the task of propaganda campaign detection both assumptions are easily violated, yielding unreliable predictions and severely decreased performance [Goodfellow et al., 2014]. Stationarity is violated by the mechanism of evolution of malicious accounts, resulting in accounts exhibiting different behavior and characteristics over time. Neutrality is violated as well, since propaganda spreaders and bot masters are actively trying to fool detectors. Consequently, the exceptional results in malicious accounts detection that we reported in our papers might be actually largely exaggerated.

Adversarial machine learning may however mitigate both previous issues, since the existence of adversaries is accounted for by design. We could thus apply adversarial machine learning to study vulnerabilities of existing detectors and the possible attacks the cited vulnerabilities could lead to, before they have been exploited by the adversaries. Interestingly, this paradigm has recently been applied for improving bot detection as well as for fake news detection [Wu et al., 2020; Zellers et al., 2019]. Finally, it is worth noting that all tasks related to the detection of online deception, manipulation, and automation—including, but not limited to, propaganda campaign detection—are intrinsically adversarial.

6 Challenges and Future Forecasting

6.1 Major Challenges

Computational propaganda detection is still in its early stages and the following challenges need to be addressed:

1. Text is not the only way to convey propaganda. Sometimes, pictures convey stronger messages than texts, as for certain political memes. Thus, it is becoming increasingly necessary to analyze multiple modalities of data (e.g., images, videos, speech). This is challenging because, even if some research was conducted on how to effectively understand cross-modal information in various domains, little has been done on what information (provided by a given modality) can be leveraged to detect propaganda.
2. Explainability is a desirable feature of propaganda detection systems in order to make them accepted at large. In fact, it is crucial to be able to motivate decisions, especially controversial ones (e.g., banning of OSN accounts or removal of posts/news). However, most of the recent developments in propaganda and coordination detection are based on deep learning, which lacks explainability — for the short and medium term, at least.
3. In addition to being able to classify individual documents as propaganda or single accounts as deceptive/coordinated, it would be useful to also provide information towards understanding the goals and the strategy of propaganda campaigns [Atanasov et al., 2019]. This problem currently stands as largely unsolved and calls for joint efforts in propaganda and coordination detection.
4. Recent advances in neural language models have made it difficult even for humans to detect synthetic text. Zellers et al. [2019] showed that a template system helps manipulate the output format of a language model, while Yang et al. [2018] suggested how to transfer the style of the language model to the target domain. With all building blocks already in place, it is likely that automatically-generated propaganda will surface in the near future.
5. The vast majority of existing detectors are evaluated only on a single annotated dataset. Often, the dataset is collected and annotated for a specific study, and is subsequently disregarded. As such, we currently lack the ability to evaluate detectors' capability of generalizing the performance obtained *in silico*, also when applied in-the-wild. For the future, it is advisable to devote additional efforts to curate large annotated datasets. Extensive data sharing initiatives —such as that of Twitter related to recent information operations¹²— are thus particularly welcome.
6. When dealing with user-generated data, ethical considerations are also important. We should thus guarantee that all analysis and potential sharing of datasets are conducted respecting the privacy of the involved users. This can also affect data availability, as demonstrated by the Facebook/Social Science One URL dataset¹³, whose release was postponed for almost two years due to the need to implement robust privacy-preserving mechanisms.

¹²<http://transparency.twitter.com/en/information-operations.html>

¹³<http://socialscience.one/blog/unprecedented-facebook-urls-dataset-now-available-research-through-social-science-one>

6.2 Forecasting

Given the above challenges and the existence of some previously remarked under-explored directions, we highlight the following research directions:

1. There is growing motivation for jointly tackling the textual and the network aspects of propaganda detection, as relying on a single paradigm is a recipe for failure. For instance, if a pre-trained language model such as GPT-2 is used as an automated propaganda generation method, it may become ineffective to detect propaganda when focusing on linguistic features alone, since it would take longer to detect propaganda than to generate it. Thus, in the future it will be necessary to go beyond texts and to also analyze the network nodes and the connectivity patterns through which propaganda spreads.
2. Spreading propaganda through multiple modalities is increasingly popular. Maliciously crafted images or videos can be more effective than articles when targeting the millennial generation, who is more familiar with watching than reading. Again, research in detecting propaganda needs to move beyond text analysis, and to embrace more comprehensive analyses that span over various data modalities.

7 Conclusion

Among the contributions of our work, we surveyed state-of-the-art computational propaganda detection methodologies. We also showed how the rapid pace of evolution of the techniques adopted by an adversary are impairing current propaganda detection solutions. Further, we justified our call for moving beyond textual analysis and we argued for the need of combined efforts blending Natural Language Processing, Network Analysis, and Machine Learning. Finally, we showed concrete promising research directions in the field of computational propaganda detection.

References

- [Atanasov et al., 2019] Atanasov et al. Predicting the role of political trolls in social media. In *CoNLL*, pages 1023–1034, 2019.
- [Barrón-Cedeño et al., 2019] Barrón-Cedeño et al. Propgy: Organizing the news based on their propagandistic content. *Inf. Process. Manag.*, 56(5):1849–1864, 2019.
- [Bolsover and Howard, 2017] Bolsover and Howard. Computational propaganda and political big data: Toward a more critical research agenda. *Big Data*, 5(4):273–276, 2017.
- [Cai et al., 2017] Cai et al. Detecting social bots by jointly modeling deep behavior and content information. In *CIKM*, pages 1995–1998, 2017.
- [Chavoshi et al., 2016] Chavoshi et al. DeBot: Twitter bot detection via warped correlation. In *ICDM*, pages 817–822, 2016.
- [Chen et al., 2013] Chen et al. Battling the internet water army: Detection of hidden paid posters. In *ASONAM*, pages 116–120, 2013.

- [Chetan et al., 2019] Chetan et al. Corerank: Ranking to detect users involved in blackmarket-based collusive retweeting activities. In *WSDM*, pages 330–338, 2019.
- [Chu et al., 2012] Chu et al. Detecting automation of Twitter accounts: Are you a human, bot, or cyborg? *TDSC*, 9(6):811–824, 2012.
- [Conserva, 2003] Conserva. *Propaganda Techniques*. AuthorHouse, 2003.
- [Cresci et al., 2017] Cresci et al. The paradigm-shift of social spambots: Evidence, theories, and tools for the arms race. In *WWW Companion*, pages 963–972, 2017.
- [Cresci et al., 2020] Cresci et al. Emergent properties, models, and laws of behavioral similarities within groups of Twitter users. *Comput. Commun.*, 150:47–61, 2020.
- [Da San Martino et al., 2019a] Da San Martino et al. Findings of the NLP4IF-2019 shared task on fine-grained propaganda detection. In *NLP4IF@EMNLP*, pages 162–170, 2019.
- [Da San Martino et al., 2019b] Da San Martino et al. Fine-grained analysis of propaganda in news articles. In *EMNLP*, pages 5640–5650, 2019.
- [Darwish et al., 2017] Darwish et al. Seminar users in the Arabic Twitter sphere. In *SocInfo*, pages 91–108, 2017.
- [Echeverrià et al., 2018] Echeverrià et al. LOBO: Evaluation of generalization deficiencies in Twitter bot classifiers. In *ACSAC*, pages 137–146, 2018.
- [Goodfellow et al., 2014] Ian Goodfellow et al. Generative adversarial nets. In *NIPS*, pages 2672–2680, 2014.
- [Grimme et al., 2017] Grimme et al. Social bots: Human-like by means of human control? *Big Data*, 5(4):279–293, 2017.
- [Habernal et al., 2017] Habernal et al. Argotario: Computational argumentation meets serious games. In *EMNLP*, pages 7–12, 2017.
- [Ins, 1938] How to detect propaganda. In *Publications of the Institute for Propaganda Analysis*, pages 210–218, 1938.
- [Jiang et al., 2016] Jiang et al. Inferring lockstep behavior from connectivity pattern in large graphs. *Knowledge and Information Systems*, 48(2):399–428, 2016.
- [Jowett and O’Donnell, 2012] Jowett and O’Donnell. *Propaganda and Persuasion*. SAGE, 2012.
- [Kudugunta and Ferrara, 2018] Kudugunta and Ferrara. Deep neural networks for bot detection. *Information Sciences*, 467:312–322, 2018.
- [Kumar et al., 2017] Kumar et al. An army of me: Sockpuppets in online discussion communities. In *WWW*, pages 857–866, 2017.
- [Lee and Kim, 2014] Lee and Kim. Early filtering of ephemeral malicious accounts on Twitter. *Comput. Commun.*, 54:48–57, 2014.
- [Li et al., 2016] Li et al. A survey on truth discovery. *SIGKDD Explor. Newsl.*, 17(2):1–16, 2016.
- [Linville and Patrick, 2018] Linville and Patrick. Troll factories: The Internet Research Agency and state-sponsored agenda building. *Resource Centre on Media Freedom in Europe*, 2018.
- [Liu et al., 2017] Liu et al. HoloScope: Topology-and-spike aware fraud detection. In *CIKM*, pages 1539–1548, 2017.
- [Mazza et al., 2019] Mazza et al. RTbust: Exploiting temporal patterns for botnet detection on Twitter. In *WebSci*, pages 183–192, 2019.
- [Mintz et al., 2009] Mintz et al. Distant supervision for relation extraction without labeled data. In *ACL-AFNLP*, pages 1003–1011, 2009.
- [Pacheco et al., 2020] Pacheco et al. Unveiling coordinated groups behind White Helmets disinformation. In *WWW Companion*, pages 611–616, 2020.
- [Rangel and Rosso, 2019] Rangel and Rosso. Overview of the 7th author profiling task at PAN 2019: Bots and gender profiling in Twitter. In *CLEF*, 2019.
- [Rashkin et al., 2017] Rashkin et al. Truth of varying shades: Analyzing language in fake news and political fact-checking. In *EMNLP*, pages 2931–2937, 2017.
- [Ratkiewicz et al., 2011] Ratkiewicz et al. Truthy: Mapping the spread of astroturf in microblog streams. In *WWW*, pages 249–252, 2011.
- [Shu et al., 2017] Shu et al. Fake news detection on social media: A data mining perspective. *SIGKDD Explor. Newsl.*, 19(1):22–36, 2017.
- [Sun et al., 2017] Sun et al. Revisiting unreasonable effectiveness of data in deep learning era. In *ICCV*, pages 843–852, 2017.
- [Thorne and Vlachos, 2018] Thorne and Vlachos. Automated fact checking: Task formulations, methods and future directions. In *COLING*, pages 3346–3359, 2018.
- [Wu et al., 2020] Wu et al. Using improved conditional generative adversarial networks to detect social bots on Twitter. *IEEE Access*, 8:36664–36680, 2020.
- [Yang et al., 2015] Yang et al. VoteTrust: Leveraging friend invitation graph to defend against social network sybils. *TDSC*, 13(4):488–501, 2015.
- [Yang et al., 2018] Yang et al. Unsupervised text style transfer using language model discriminators. In *NIPS*, pages 7287–7298, 2018.
- [Yang et al., 2019] Yang et al. Arming the public with artificial intelligence to counter social bots. *Human Behavior and Emerging Technologies*, 1(1):48–61, 2019.
- [Zellers et al., 2019] Zellers et al. Defending against neural fake news. In *NIPS*, pages 9051–9062, 2019.
- [Zhang et al., 2016] Zhang et al. The rise of social botnets: Attacks and countermeasures. *TDSC*, 15(6):1068–1082, 2016.
- [Zhou et al., 2019] Zhou et al. Fake news: Fundamental theories, detection strategies and challenges. In *WSDM*, pages 836–837, 2019.