

Automatized Detection and Annotation for Calls to Action in Latin-American Social Media Postings

Wassiliki Siskou, Clara Giralt Mirón, Sarah Molina Raith, Miriam Butt

University of Konstanz

`firstname.lastname@uni-konstanz.de`

Abstract

Voter mobilization via social media has shown to be an effective tool (Savage et al., 2016). While previous research has primarily looked at how calls-to-action (CTAs) were used in Twitter messages from non-profit organizations (Guidry et al., 2014) and protest mobilization (Rogers et al., 2019), we are interested in identifying the linguistic cues used in CTAs found on Facebook and Twitter for an automatic identification of CTAs. The work is part of an on-going collaboration with researchers from political science, who are investigating CTAs in the period leading up to recent elections in three different Latin American countries. We developed a new NLP pipeline for Spanish to facilitate their work. Our pipeline annotates social media posts with a range of linguistic information and then conducts targeted searches for linguistic cues that allow for an automatic annotation and identification of relevant CTAs. By using carefully crafted and linguistically informed heuristics, our system so far achieves an F1-score of 0.85.

1 Introduction

We here report on a NLP pipeline designed to automatically identify and annotate “calls-to-action” (CTAs). We focus specifically on the mobilization of potential voters via social media as part of an on-going collaboration with partners from political science. CTAs are of interest for political science because recent years have seen social movements and political parties working with CTAs via social media as an effective tool for social mobilization (Savage et al., 2016; Guidry et al., 2014).

Guided by the interests of our project partners from political science (Haiges and Zuber, 2021), who are investigating the expression of grievances in the context of ethnic political parties, we focus on Facebook and Twitter posts leading up to the recent (2020) elections across Latin America and aim to support their research by automatically detecting

and annotating CTAs. We faced several challenges in doing this. For one, we cannot build on very much previous NLP work. Although initial work has been done to identify CTAs for Russian (Rogers et al., 2019), no research had as yet been done on the linguistic expression of CTAs in Spanish. We were also faced with severe issues of data scarcity — in 31,229 social media postings contained in the corpus, only 2,542 were found to contain CTAs. In addition, the corpus collected by our partners is inherently unbalanced, as it differs across countries and elections.

As such we decided to implement a primarily rule-based NLP pipeline for the automatized detection and annotation of CTAs. Our pipeline is based on previous efforts focusing on English and German and is designed particularly to identify deep morphosyntactic, semantic and pragmatic linguistic features that are relevant for analysis at the discursive level (Biber et al., 1998; Biber and Conrad, 2009) and for the linguistic framing of utterances (cf. Druckman’s frames of communication (Druckman, 2011; Chong and Druckman, 2007)). While other studies have focused on analyzing the effects of such CTAs on voters (Heiss and Matthes, 2016; Kligler-Vilenchik et al., 2021), in this use case we are concerned with how these mobilizations are linguistically expressed. We therefore aim at automatically identifying CTAs via linguistic cues.

2 Related Work

As already noted, there is limited previous work we can build on. Guidry et al. (2014) showed that Twitter messages from non-profit organizations framed as a CTA were retweeted more often and generated more interaction between users than other messages, while they were simultaneously the least used strategy for messages. Rogers et al. (2019) worked on historical data of Bolotnaya protests (2012) in Russia to demonstrate the possibility of automatically detecting CTAs in social

media posts. Their classification task yielded an F1-score of 0.77, thus showing that CTAs in Russian can be successfully detected automatically to some degree. While their focus lies on the detection of CTAs in a protest setting, where movements use social media to mobilize people and convince them to join the protests, we are focusing on CTAs that are targeting voter mobilization.

3 Characterization of Calls to Action

Our specific use case is the analysis of social media posts published in the period leading up to an election and aimed at mobilizing voters. We therefore restricted an identification of CTAs to sentences in social media posts that directly or indirectly call upon the addressees for political participation in an election setting. Typical examples are (1) and (2).

- (1) *Cumplamos nuestra responsabilidad de votar.*
‘Let us fulfill our responsibility to vote.’
- (2) *¡Sal a votar! ‘Get out to vote!’*
- (3) *Es momento de poner fin a décadas de los mismos de siempre. Llegó el momento del pueblo.*
‘It is time to put an end to decades of the same old ones. The time has come for the people.’

Whereas the CTAs in the first two examples are directly encoded and expressed via an imperative, the more indirect example (3) also falls under our definition of a CTA. Here, the message to politically participate is covertly encoded and does not involve a direct use of an imperative. In order to automatically identify such indirect CTAs, it is necessary to implement linguistically informed rules (see §6). We exclude CTAs that aim at getting readers to click on a specific link, support an organization with donations or attend events or demonstrations, even though they are linguistically similarly framed as direct CTAs. This exclusion is motivated by the research interests of our partners, who are not interested in these other types of CTAs, but in voter mobilization.

4 NLP pipeline LiAnS

For the automatic detection and annotation of CTAs we used LiAnS, a rule-based NLP pipeline. LiAnS (Linguistic Annotation Service) has been built on the basis of *VisArgue* (Gold et al., 2015), a NLP

pipeline initially designed for analyzing linguistic features in spoken dialogs and debates in English and German. We built a Spanish version of the pipeline in the context of the CTA project.¹ The automated linguistic feature identification in *VisArgue* is realized in a rule-based fashion: For each feature, a list of relevant cues (lexical items and constructions) is defined by language experts, and rules are created for the disambiguation according to the context they are found in. This has a clear advantage over a naive and decontextualized application of static word lists. The linguistic features covered by *VisArgue* were selected by experts of theoretical and computational linguistics and are strongly grounded in theoretical linguistic insights. We used the existing features and categories as a blueprint for the Spanish LiAnS. Carefully crafted feature sets and disambiguation rules were added to ensure the reliable annotation of CTAs.

The main workflow of LiAnS consists of two steps: 1) preprocessing; 2) annotation with prior disambiguation. We first convert the raw posts into standardized XML files, which are then further preprocessed using the *Stanza* NLP kit (Qi et al., 2020). *Stanza* conducts sentence splitting, tokenization and lemmatization on the input files, and adds POS-tags, morphological features and dependency relations of each token (lexeme) as XML attributes. LiAnS further adds *discourse unit splitting*. Approximating the definition of a *basic discourse unit* (Polanyi et al., 2004), each clause is defined as one discourse unit (DU), e.g., if a sentence is comprised of one matrix clause and one embedded clause, the sentence is defined as containing two DUs.

After preprocessing, rule-based annotation is applied on the basis of pre-defined disambiguation rules for linguistic features as in Table 1. By considering the position of the linguistic cue and the POS of the surrounding lexemes, we implement any ambiguity resolution that might be necessary. Annotations are initially applied to the lexemes. Based on the particular type of use cases, aggregated, context-aware feature calculation on higher levels, e.g., DU-, sentence- or document-level can be defined. LiAnS allows for a modular usage of the features, which means that users can select their own customized subset of features that are of interest to the specific use case. LiAnS will then

¹Our Spanish pipeline is based on a refined and extended German and English LiAnS version of *VisArgue*, designed and built by Qi Yu and Marina Janka at the University of Konstanz.

exclusively annotate these features.²

5 Data

The overall corpus is still in the process of being compiled by our political science partners (Haiges and Zuber, 2021), and currently consists of around 30,000 posts on Twitter and Facebook published by political parties, indigenous associations and presidential candidates before and during the three most recent elections in Latin America, which took place in Bolivia (18 October 2020), Ecuador (7 February 2021) and Peru (11 April 2021). Table 2 shows the exact number of posts in our current corpus per country. As can be seen the data set is not only unbalanced, but also fairly small for Bolivia, motivating our rule-based approach, which is not data hungry, but instead relies on linguistic knowledge.

6 Methodology

All posts in the dataset were annotated with LiAnS using the features available for the Spanish version. To only extract DUs and sentences that thematically deal with elections from the dataset, we manually compiled a list of voting-related keywords such as *democracia* ‘democracy’ and *elecciones* ‘elections’, including their synonyms and semantically related nouns, verbs and adjectives. Corresponding DUs and sentences were then extracted.

Direct CTAs can be detected very straightforwardly based simply on the identification of imperative forms, so one easy way of proceeding is to search for instances of imperatives of verbs like *votar* ‘vote’ as well as the participle of *votar* ‘voting’. In a further step, we searched for sentences which contain imperatives and election-related keywords from our hand-defined list. In this way, we are able to detect CTAs as in example (2). Additionally, sentences containing the election date (e.g. *este 7 de febrero*), or count downs to the election day (*faltan 5 días para*) were also considered as CTAs since they aim at reminding voters about when to exercise their right to vote.

Moving on to more indirect and emotive CTAs, we established rules that search for the noun *voto* ‘vote’ preceded by the pronoun *tú* ‘your’ or *nuestro* ‘our’ to identify posts that directly aimed at motivating voters through the use of more appellative expressions by directly addressing them. CTAs as

in (4) are often followed by a phrase describing the expected (positive) change that may occur after voting for a specific candidate or voting in general.

- (4) *Pido tu voto para...*
‘I ask for your vote to...’

Furthermore, we looked for modal phrases and expressions that imply obligations and were used in combination with election-related vocabulary to identify voter-mobilization.

While the identification of direct CTAs is relatively straight-forward and can be pinned down to the use of imperatives and the different forms of the verb *votar* mainly, implicit calls to action require deeper linguistic knowledge, as their surface coding is indirect. For the annotation of implicit CTAs, we perused the dataset manually and tracked down cues that were used to frame the call without using similar formulations to those mentioned above. We found that implicit calls to action were often encoded in sentences implying that it is time *for change* (as in example (3)), *to change the future of the country*, *to rescue the country*, *to stand up for the country* or *to take the country forward*, and thus intend to bring people to the polls to actively engage in the act of change. In addition, implicit CTAs often appeal to voters’ sense of ‘us’ vs. ‘them’. For example, some CTAs read “*juntos + Verb*” (‘together + verb’) and aim to create a sense of belonging to the community by giving the impression to voters that the common goals can only be achieved if they become part of this community by casting their vote for the candidate. To automatically detect those implicit CTAs we implemented rules that search for those key phrases together with election-related vocabulary and tag the posting as such.

7 Results and Discussion

A randomly selected subset of 800 Facebook and 200 Twitter posts from Ecuador was manually annotated by two Spanish speaking annotators, both co-authors of the paper, in order to evaluate the results from our NLP pipeline. We decided to evaluate on data from Ecuador, as the majority of our corpus comes from there. Out of 800 Facebook posts 156 (16 implicit and 140 explicit) and 29 (10 implicit and 19 explicit) of 200 Tweets were identified as CTAs. With the help of a more detailed analysis of the linguistic features found in CTAs we found that 49 % of them include at least one

²Access to the Spanish version of LiAnS and the CTA annotation can be requested via e-mail.

Dimension	Feature	Examples
Discourse Relation	adversative	<i>pero</i> (but); <i>al contrario</i> (on the contrary)
	causal	<i>porque</i> (because); <i>ya que</i> (since) ; <i>puesto que</i> (since)
	conditional	<i>cuando</i> (if); <i>en caso que</i> (in the event that)
	consecutive	<i>pues</i> (for); <i>consecuentemente</i> (consequently)
Modality	modal	<i>deber</i> (must); <i>querer</i> (want); <i>poder</i> (can)
Sentence Modifier	intensifier	<i>muy</i> (very); <i>de hecho</i> (in fact)
	negation	<i>no</i> (no); <i>nunca</i> (never) ; <i>jamás</i> (never)
Information Exchange	accommodation	AGREEMENT: <i>prometer</i> (promise); <i>consentir</i> (agree); <i>reconocer</i> (acknowledge) DISAGREEMENT: <i>rechazar</i> (reject); <i>disputar</i> (dispute); <i>degradar</i> (degrade)
	speech act	INFORMATION GIVING: <i>decir</i> (say); <i>aclarar</i> (clarify) INFORMATION SEEKING: <i>preguntar</i> (ask); <i>consultar</i> (consult)
Politeness	polite items	<i>por favor</i> (please); <i>gracias</i> (thank you)

Table 1: Available features Spanish LiAnS version. The column *Examples* provides a few examples in Spanish as they are implemented in LiAnS. Small caps indicate different subcategories within the feature.

Source	Bolivia	Peru	Ecuador
Twitter	526	1,667	4,959
Facebook	991	1,252	21,834

Table 2: Overview of data set by country and source.

imperative to mobilize voters, while 57 % contain at least one of the different forms of the verb *votar*, 40 % contain the election date, and around 10% include more indirect means of CTAs which can not be pinned down to specific linguistic features but depend more on the pragmatic context of the posting.

Based on the above mentioned subset of 1000 social media posts, our system shows an overall precision of 0.95, a recall of 0.77 and a F1-score of 0.85 for the automatic identification of CTA types. For the Twitter data, precision is 0.92, recall 0.78 and F1 0.84, while the scores for Facebook data show a precision score of 0.81, a recall 0.83 and a F1 score of 0.81. These values allow us to draw the conclusion that our system performs almost equally good across different social media platforms albeit the differing posting format used on both platforms. We attribute the differences between the manual and automatic CTA annotation to the fact that the used rules for CTA identification in our pipeline need to be further refined in order to identify certain linguistic features that are being used to mobilize voters.

An overview of our results for the automatic detection of explicit and implicit CTAs across the whole corpus is presented in Table 3. While the majority of CTAs are formulated as imperatives and were thus identified based on their morphology, a portion of the voter mobilizations were identified

by the LiAnS feature *modal*, specifically the subcategory of *obligation*. In addition to *modal*, the feature *polite items* of LiAnS helped to identify more covert CTAs as in (4), where the appellative character of the sentence comes through the use of a polite phrase.

Source	Bolivia		Peru		Ecuador	
Type	Imp	Exp	Imp	Exp	Imp	Exp
Twitter	0	39	19	43	53	217
Facebook	4	70	8	94	121	2,303

Table 3: # of identified CTAs per country, source and type (Explicit vs Implicit).

Overall our results show that LiAnS can help to annotate small corpora that are unsuitable for machine learning approaches due to their small size and unbalanced nature, therefore reducing the manual annotation effort for our collaboration partners.

8 Limitations

The results of our evaluation show that there is still some room for improvement. First, the currently unbalanced nature of the corpus is not just a problem for machine learning approaches, it also poses challenges for the development of the NLP pipeline, as regional language varieties of all countries must be considered equally when creating the annotation rules. Second, our annotations rely on the morphological features provided by the *Stanza* NLP kit. This means that our pipeline struggles when the morphological analysis delivered by *Stanza* is erroneous. For example, this was the case for *voto* ‘vote’, which can be a verb or a noun. All instances of *voto* were analyzed as nouns by *Stanza*, which makes it difficult to identify DUs such as *voto por*

un futuro mejor ‘vote for a better future’. We adjusted the pipeline to correct for errors of this type.

9 Conclusion and Future work

In conclusion, we have implemented an automated approach to identify explicit and implicit election-related CTAs in Spanish social media. While a few previous studies have looked at the identification of CTAs in tweets related to protests, to the best of our knowledge our work is the first to look at the linguistic patterns that can be found in such attempts at voter mobilization.

Our next steps are, first, to annotate more CTA posts from Peru and Bolivia in order to create a more balanced gold standard corpus; and, second, to create our own balanced corpus. Third, we aim to expand our set of rules to allow especially more implicit CTAs to be annotated automatically.

We plan on adding a more sophisticated scoring system that assigns an aggregated score to the identified CTA cues. Thus, the classification as a CTA of a post will depend on how many linguistic cues that indicate a CTA are present and how heavily they are weighted. The word *vota* ‘vote’, for instance, should receive a greater weight than a modal like *tenemos que* ‘we have to’, since the former is a clear cue for a CTA, while the latter could also be used in other political contexts. We are confident that we can improve our overall results even further with these extensions of the pipeline and the introduction of a sophisticated scoring system. Finally, we also intend to experiment with machine learning approaches once the data set has grown large enough via bootstrapping through CTA classifications and annotations via our pipeline.

Acknowledgement

The work reported on in this paper was funded by the Deutsche Forschungsgemeinschaft (DFG – German Research Foundation) under Germany’s Excellence Strategy – EXC-2035/1 – 390681379 as part of the project “Mobilizing Inequality”.

References

Douglas Biber and Susan Conrad. 2009. *Register, Genre, and Style*. Cambridge University Press, Cambridge.

Douglas Biber, Susan Conrad, and Randi Reppen. 1998. *Corpus Linguistics: Investigating Language Structure and Use*. Cambridge University Press, Cambridge.

Dennis Chong and James N. Druckman. 2007. [Framing public opinion in competitive democracies](#). *American Political Science Review*, 101(4):637–655.

James N. Druckman. 2011. What’s it all about?: Framing in political science. In Gideon Keren, editor, *Perspectives on Framing*, pages 279–301. Taylor and Francis.

Valentin Gold, Mennatallah El-Assady, Annette Hautli-Janisz, Tina Bögel, Christian Rohrdantz, Miriam Butt, Katharina Holzinger, and Daniel Keim. 2015. [Visual linguistic analysis of political discussions: Measuring deliberative quality](#). *Digital Scholarship in the Humanities*, 32(1):141–158.

Jeanine Guidry, Richard Waters, and Gregory Saxton. 2014. [Moving Social Marketing Beyond Personal Change to Social Change](#). *Journal of Social Marketing*, 4:240–260.

Lea Haiges and Christina Isabel Zuber. 2021. Movements, Parties and the Making of Indigenous Politics in Ecuador and Peru. *2021 ECPR General Conference*.

Raffael Heiss and Jörg Matthes. 2016. Mobilizing for some. *J. Media Psychol. Theor. Methods Appl.*, 28:123–135.

Neta Kligler-Vilenchik, Maya de Vries Kedem, Daniel Maier, and Daniela Stoltenberg. 2021. Mobilization vs. Demobilization Discourses on Social Media. *Political Communication*, 38(5):561–580.

Livia Polanyi, Chris Culy, Martin Van Den Berg, Gian Lorenzo Thione, and David Ahn. 2004. Sentential structure and discourse parsing. In *Proceedings of the Workshop on Discourse Annotation*, pages 80–87.

Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. Stanza: A Python Natural Language Processing toolkit for many human languages. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 101–108.

Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2019. Calls to action on social media: Detection, social impact, and censorship potential. In *Proceedings of the Second Workshop on Natural Language Processing for Internet Freedom: Censorship, Disinformation, and Propaganda*, pages 36–44.

Saiph Savage, Andres Monroy-Hernandez, and Tobias Höllerer. 2016. [Botivist: Calling Volunteers to Action Using Online Bots](#). In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work and Social Computing*, page 813–822, New York, NY, USA. Association for Computing Machinery.