



# Crowdsourcing Dialect Characterization through Twitter

Bruno Goncalves, David Sanchez

## ► To cite this version:

Bruno Goncalves, David Sanchez. Crowdsourcing Dialect Characterization through Twitter. PLoS ONE, 2014, 9 (e112074 ), 10.1371/journal.pone.0112074 . hal-01242109

**HAL Id: hal-01242109**

**<https://amu.hal.science/hal-01242109>**

Submitted on 11 Dec 2015

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



# Crowdsourcing Dialect Characterization through Twitter

Bruno Gonçalves<sup>1,2\*</sup>, David Sánchez<sup>3</sup>

**1** Aix-Marseille Université, CNRS, CPT, UMR 7332, 13288 Marseille, France, **2** Université de Toulon, CNRS, CPT, UMR 7332, 83957 La Garde, France, **3** Instituto de Física Interdisciplinar y Sistemas Complejos IFISC (UIB-CSIC), E-07122 Palma de Mallorca, Spain

## Abstract

We perform a large-scale analysis of language diatopic variation using geotagged microblogging datasets. By collecting all Twitter messages written in Spanish over more than two years, we build a corpus from which a carefully selected list of concepts allows us to characterize Spanish varieties on a global scale. A cluster analysis proves the existence of well defined macroregions sharing common lexical properties. Remarkably enough, we find that Spanish language is split into two superdialects, namely, an urban speech used across major American and Spanish cities and a diverse form that encompasses rural areas and small towns. The latter can be further clustered into smaller varieties with a stronger regional character.

**Citation:** Gonçalves B, Sánchez D (2014) Crowdsourcing Dialect Characterization through Twitter. PLoS ONE 9(11): e112074. doi:10.1371/journal.pone.0112074

**Editor:** Tobias Preis, University of Warwick, United Kingdom

**Received:** July 28, 2014; **Accepted:** September 27, 2014; **Published:** November 19, 2014

**Copyright:** © 2014 Gonçalves, Sánchez. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

**Data Availability:** The authors confirm that all data underlying the findings are fully available without restriction. All relevant data are publicly available. Documentation on how to query the Twitter API can be found here: <https://dev.twitter.com/overview/documentation>.

**Funding:** The authors have no support or funding to report.

**Competing Interests:** The authors have declared that no competing interests exist.

\* Email: [bgoncalves@gmail.com](mailto:bgoncalves@gmail.com)

## Introduction

Language is the most characteristic trait of human communication but takes on many heterogeneous forms. Dialects, in particular, are linguistic varieties which differ phonologically, grammatically or lexically in geographically separated regions [1]. However, despite its fundamental importance and many recent developments, the way language varies spatially is still poorly understood.

Traditional methodological approaches in the study of regional dialects are based on interviews and questionnaires administered by a researcher to a small number (typically, a few hundred) of selected speakers known as informants [2]. Based on the answers provided, linguistic atlases are generated that are naturally limited in scope and subject to the particular choice of locations and informants and perhaps not completely free of unwanted influences from the dialectologist. Another approach is the use of mass media corpora which provide a wealth of information on language usage but suffer from the tendency of media and newspapers to use standard norms (the “BBC English” for example) [3] that limits their usefulness for the study of informal local variations.

On the other hand, the recent rise of online social tools has resulted in an unprecedented avalanche of content that is naturally and organically generated by millions or tens of millions of geographically distributed individuals that are likely to speak in vernacular and do not feel constrained to use standard linguistic norms. This, combined with the widespread usage of GPS enabled smartphones to access social media tools provides a unique opportunity to observe how languages are used in everyday life and across vast regions of space.

In this work, we use a large dataset of geolocated Tweets to study local language variations across the world. Similar datasets

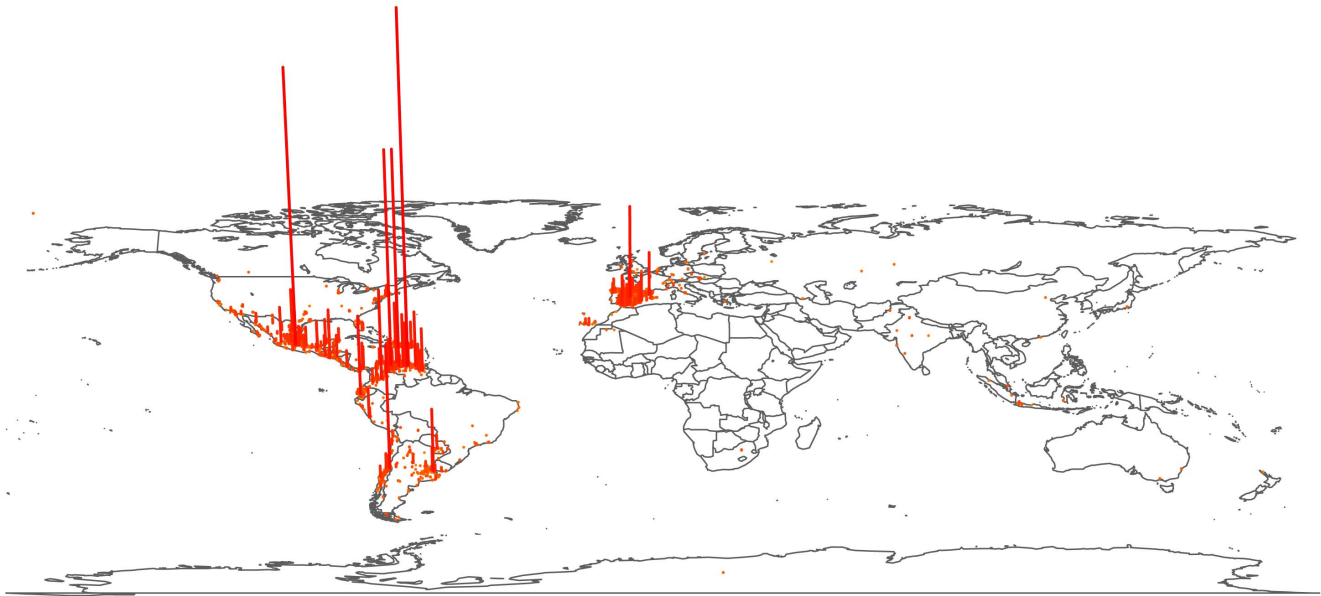
have recently been used to map public opinion and social behavior [4–11] and to analyze planetary language diversity [12].

Preliminary results demonstrating the feasibility of this approach have thus far been limited to considering only few words or just a few geographical areas [13,14]. Here, we move beyond the mere proof of concept and provide a detailed global picture of spatial variants for a specific language. For definiteness, we choose Spanish as it is not only one of the most spoken in the world but it has the added advantage of being spatially distributed across several continents [15,16]. Several other languages such as Mandarin or English have more native speakers or higher supra-regional status but their use is hindered by the limited local availability of Twitter (Mandarin) or a high abundance of homographs that precludes a detailed lexicographic analysis (English).

## Methods

We used the Twitter gardenhose to gather an unbiased sample of all tweets written in Spanish that contained GPS information over the course of over two years. Language detection was performed using the state of the art Chromium Compact Language Detector [17] software library.

The resulting dataset contained over  $5 \times 10^7$  geolocated tweets written in Spanish distributed across the world (see Fig. 1). As expected, most tweets are localized in Spain, Spanish America and extensive areas of the United States. These results are consistent with recent sociolinguistic data [18,19], providing an initial level of validation to our approach. Interestingly, we also find significant contributions from major non-Spanish-speaking cities in Latin America and Western Europe, likely due to considerable population of temporary settlers and tourists. See Ref. [12] for further details and results on this dataset.



**Figure 1. Spanish tweet locations.** The overwhelming majority of Spanish tweets are located in Spain and Spanish America but significant contributions arise in certain US states and major Western European and Brazilian cities.  
doi:10.1371/journal.pone.0112074.g001

Traditional approaches in dialectology have preferred rural, male informants while modern analyses include interactions with urban speakers regardless of age and gender. On average, Twitter users are young, urban [20] and more likely to be technologically savvy thus providing more modern perspective on the use of language.

To be able to determine exactly what the major local varieties of Spanish are, we use a list of concepts and utterances selected from an exhaustive study of lexical variants in major Spanish-speaking cities. The Varilex database [21] provides a comprehensive list of possible words representing several concepts, such as ‘popcorn’, ‘car’, ‘bus’, etc. We selected a subset of concepts that minimized possible semantic ambiguities by ensuring that they contained no common words. The complete list of words and maps for each concept studied can be accessed at <http://www.bgoncalves.com/languages/spanish.html>.

In our initial set of Tweets we observed  $7.5 \times 10^5$  geolocated instances where words from our catalogue were used. Individual instances were then aggregated geographically into cells of  $0.25'' \times 0.25''$ , which corresponds to an approximate area of  $25 \times 25 \text{ km}^2$  in the equator.

Finally, we define the dominant word for each concept in each geographical cell by a simple majority rule and generate a  $M = N_{\text{cells}} \times N_{\text{words}}$  matrix where element  $M_{ij}$  is 1 when word  $j$  is the dominant for a given concept in cell  $i$  and 0 otherwise. The resulting matrix has  $N_{\text{cells}} = 1135$  rows and  $N_{\text{words}} = 131$  columns and constitutes the dataset used for the analysis presented in the remainder of this paper.

## Results and Discussion

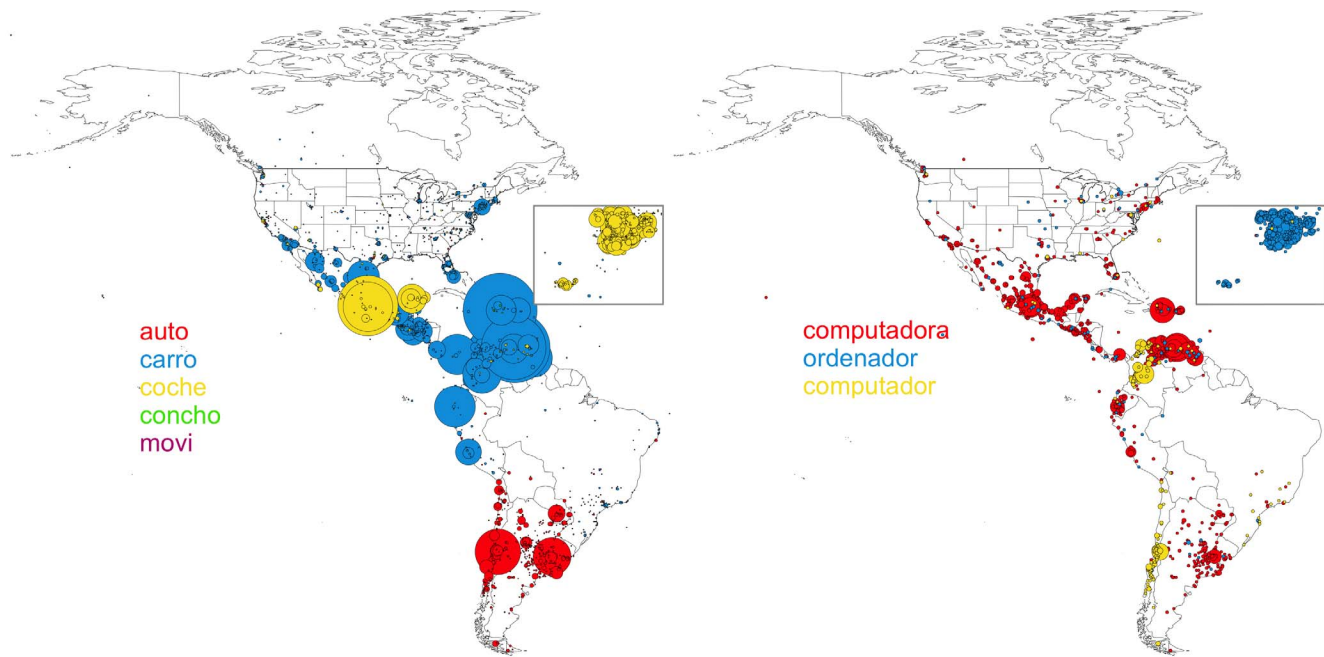
Figure 2 illustrates two illustrative concepts (‘computer’ and ‘car’) that are both associated to multiple utterances. Each utterance is represented with a different color. We draw a circle centered on each cell with an area proportional to the number of tweets that use the corresponding expression. It is clear from the map that some expressions (*computadora*, *ordenador*, *computador*)

are strongly clustered in space, allowing us to easily define regional dialects characterized by the set of dominant words used to express the concepts in our list. Due to the unique resolution of our data we could limit the isoglosses (boundaries) of the regions corresponding to each concept-word with a high degree of precision. However, the isoglosses corresponding to different concepts can overlap and bundle rendering any simple arrangement of dialect areas almost impossible.

The natural way to overcome this difficulty and characterize the various regional dialects present in modern day Spanish is to apply machine learning (ML) approaches to automatically cluster the  $M$  matrix and identify which cells are closely related to one another. We start by applying Principal Component Analysis to reduce the dimensionality of the matrix  $M$ . PCA determines the linear combinations of the columns (features in ML literature) of the matrix that explain most of the variance observed in the rows (observations). We find that by projecting the data onto the 40 principal components (see Fig. 3) we are able to maintain over 94% of the variance in the data while reducing by 2/3 the dimension of the matrix with clear numerical advantages.

## Superdialects

The task of identifying meaningful clusters in this matrix is now simplified. We proceed by applying the well known  $K$ -means [22] algorithm that iteratively refines the position of the centers of  $K$  clusters until it finds a stable set of locations. The main difficulty of utilizing this algorithm lies in identifying the correct number  $K$  of clusters to utilize. Here, we apply the  $f(k)$  metric introduced by Pham *et al.* to establish the best value for  $K$ . We run  $K$ -means with values of  $K$  up to 20 using 100 different random initializations and depict the results in Fig. 4 A). For verification purposes, we also plot the value of the Silhouette [23] of the clusters found with each value of  $K$ . Both metrics agree that 2 is the correct number of clusters (both curves show an extremum at that point), leading to two clusters of size 179 (cluster  $\alpha$ ) and 956 (cluster  $\beta$ ), respectively.



**Figure 2. Geographical distribution of the dominant word for the concepts 'computer' (left) and 'car' (right).** Map locations are colored according to the most common expression found in the corresponding cell. The area of the circle is proportional to the number of tweets.  
doi:10.1371/journal.pone.0112074.g002

A geographic plot of the location of the cells belonging to each cluster ( $\alpha$  and  $\beta$ ) provides a fundamental clue to their meaning [see Fig. 4 B)]. Strikingly, we find a profound correlation between location of cells belonging to cluster  $\alpha$  (red dots) and areas of high population density. We validate this idea using estimates of the population living within each cell provided by the LandScan dataset [24]. Hence, we plot the population distribution boxplot for each cluster in Fig. 4 C). The results clearly confirm our intuition. Cluster  $\alpha$  corresponds to cells with a typical population that is significantly larger than cluster  $\beta$ . This suggests a natural lexical bipartition of Spanish into two superdialects. Superdialect  $\alpha$  is utilized by speakers in main American and Spanish cities and


corresponds to an international variety with a strongly urban component while superdialect  $\beta$  is comprised mostly of rural areas and small towns.

A reverse analysis of the cluster method shows that superdialect  $\alpha$  presents most of the selected words plus a few unique ones (the top 3 dominant words are *acera*, *parabrisas* and *calzoncillos*), which justifies that this linguistic variety can be viewed as a dialect. Our result provides some evidence that the increasing globalization of major languages leads to an homogenization that is especially apparent for the active lexicon [25]. Cities naturally exert a linguistic centripetal force that favors dialect unification, smoothing possible lexical differences. This leveling process present in all countries (thereby its international denomination) is reinforced by the rapid increase of worldwide social ties and the powerful influence of mass media precisely located in important metropolitan areas (Madrid, Mexico City, Miami) [26]. Several other sociolinguistic aspects (prestige, higher educational status) also have a role that is more visible in urban environments.

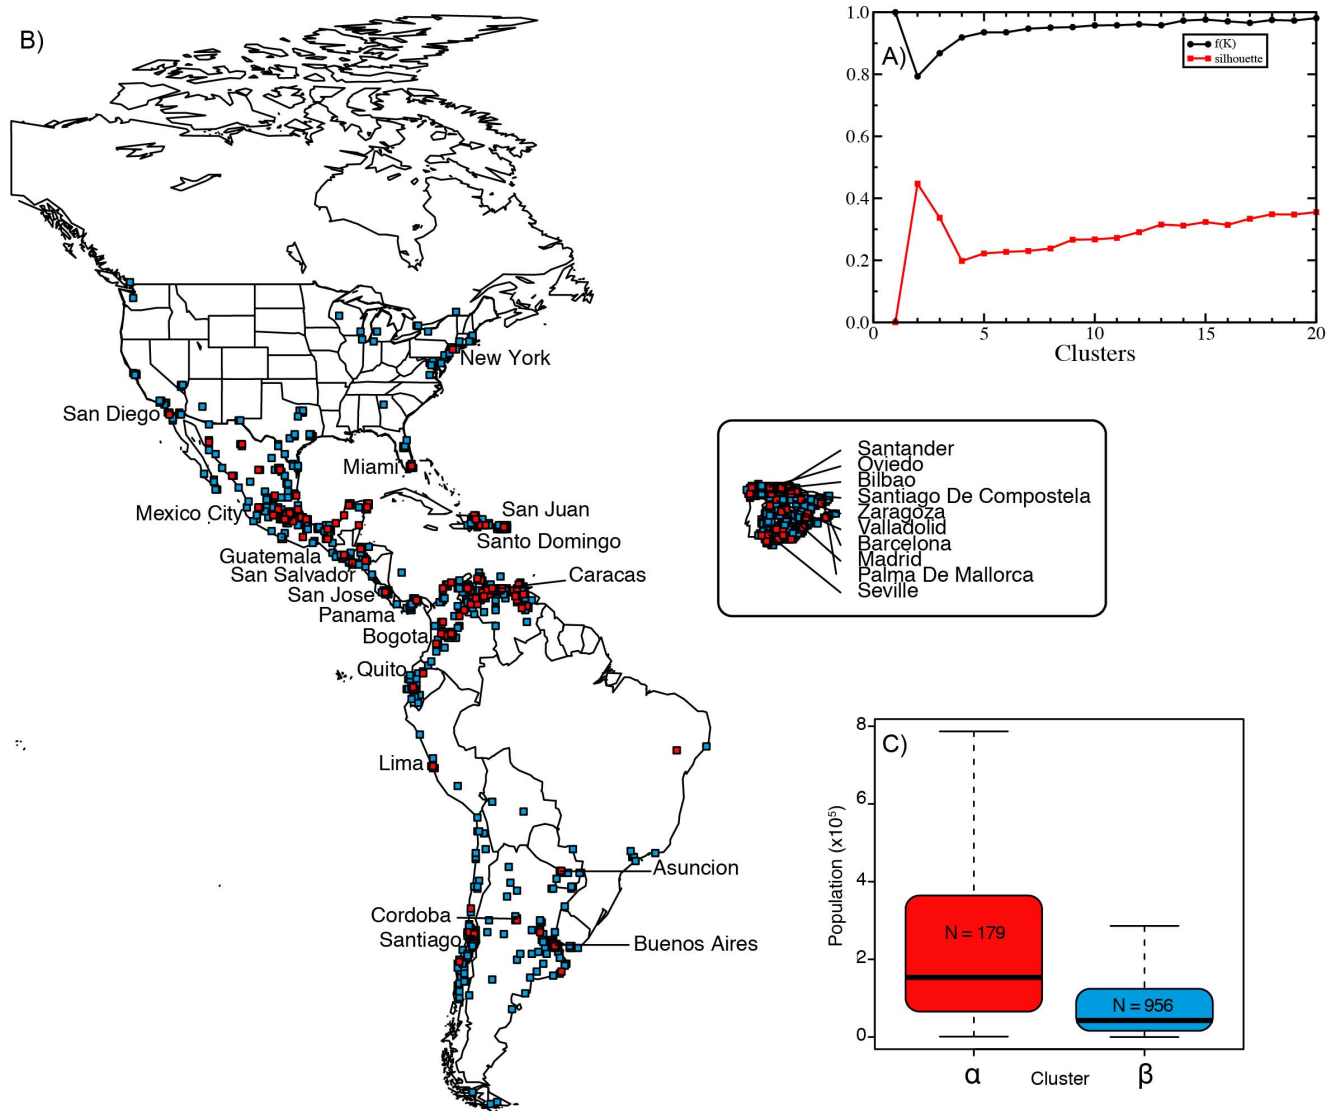
In contrast, rural areas (superdialect  $\beta$ ) are generally more conservative and keep a larger number of characteristic lexical items and native words. As a result, the dialectal area corresponding to superdialect  $\beta$  is much more geographically diverse and can be further split, as discussed below.

### Regional dialects

The size imbalance between the two clusters when combined with our intuition suggest that we can also employ the statistical procedure discussed above to further divide the largest cluster ( $\beta$ ). We apply  $K$ -means recursively until the remaining cluster has a similar size to the previous ones. In the end, we obtain five well defined clusters as displayed in Fig. 5. Clearly, three regions can be distinguished, each of them characterized by a set of dominant words. Yellow dots span a wide area covering Mexico, Central America, the Caribbean and north-western areas of South America. Green dots correspond to the Southern Cone while



**Figure 3. Cumulative variance explained as a function of the number of components.** With 40 components (vertical blue line) we are able to maintain over 94% of the variance present in the data while significantly reducing the matrix size.  
doi:10.1371/journal.pone.0112074.g003



**Figure 4. Characterization of the two superdialects.** A)  $f(K)$  and silhouette statistics as a function of  $K$ . B) Geographical representation of the two clusters,  $\alpha$  (red) and  $\beta$  (blue). For visualization purposes we increased the size of each cell. The name of main cities corresponding to superdialect  $\alpha$  are shown for clarity. C) Population distribution of the cells corresponding to each cluster.  
doi:10.1371/journal.pone.0112074.g004

blue dots are almost exclusively accumulated within Spain. The first region is quite diverse. In fact, smaller cells can be aggregated into two additional clusters (depicted with magenta and orange dots in Fig. 5). Interestingly, the magenta and orange dots seem to be localized in the Mexican plateau, the interior of Central America and Andean Colombia, in contrast with the speech of Venezuela, the Antilles and coastal areas represented with yellow dots. This division between highland and lowland varieties agrees with classifications discussed previously in the linguistics literature [27].

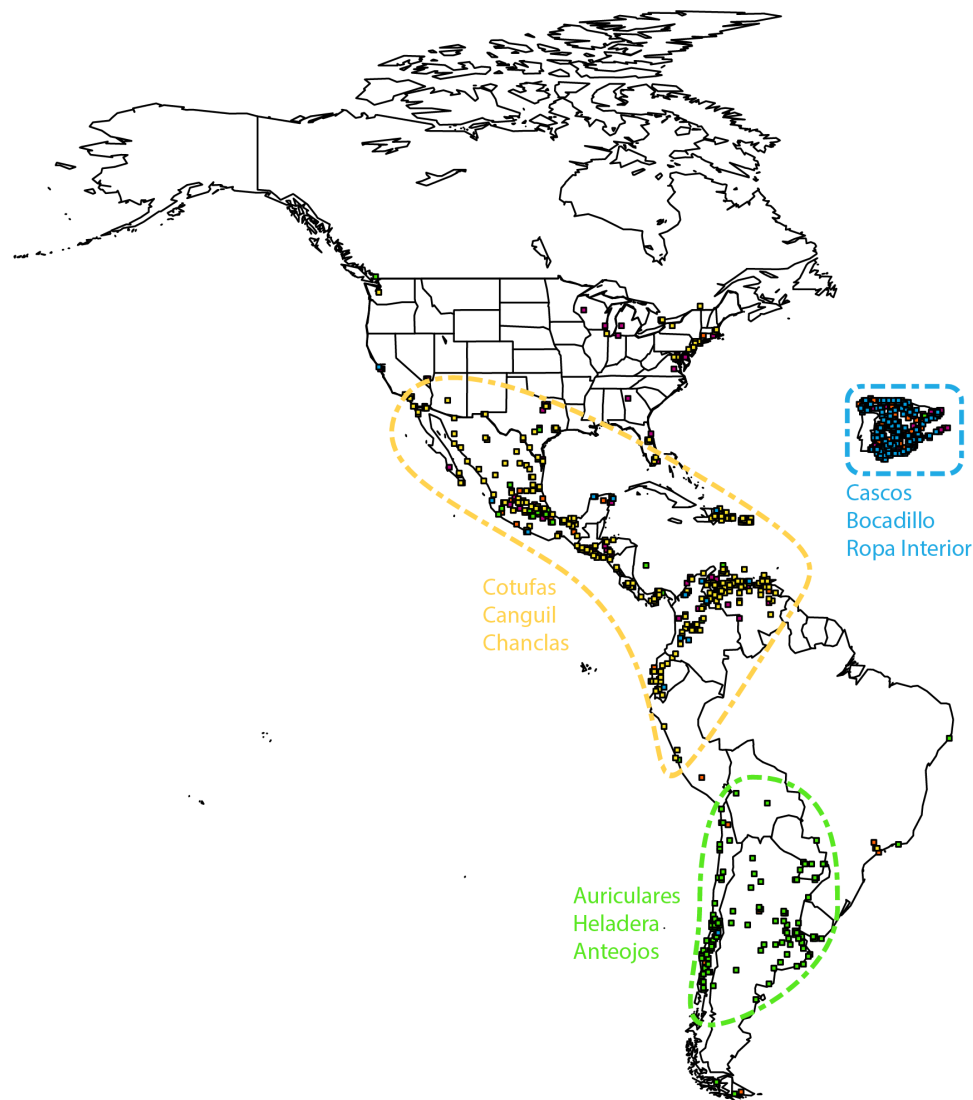
The two American regions marked in Fig. 5 partly reflect the settlement patterns and the formal colonial Spanish administration within the Empire. Conquerors and settlers occupied first the territories of Mexico, Peru and the Caribbean, and only much later colonists established permanent residence in the Southern Cone, which stayed away from prestigious linguistic norms. This strong cultural heritage that can still be observed, centuries later, in our datasets deserves to be further analysed in future works

since the formation of regional dialects is a complex process with multiple factors (population origin, degree of dependence on the metropolitan state, communication routes, etc.).

## Conclusions

Using a large dataset of user generated content in vernacular Spanish, we analyse the diatopic structure of modern day Spanish language at the lexical level. By applying standard machine learning techniques, we find, for the first time, two large Spanish varieties which are related to, respectively, international and local speeches. We can also identify regional dialects and their approximate isoglosses. Our results are relevant to empirically understand how languages are used in real life across vastly different geographical regions. We believe that our work has considerable latitude for further applications in the computational study of linguistics, a field full of rewarding opportunities. One can envisage much deeper analyses pointing the way towards new





**Figure 5. Characterization of major cluster  $\beta$ .** Geographical representation of regional dialects. For visualization purposes we increased the size of each cell. Three well separated regions are indicated with dashed lines along with the top 3 dominant words characteristic of that region. doi:10.1371/journal.pone.0112074.g005

developments in sociolinguistic studies (bilingualism, creole varieties). Our work is based on a synchronous approach to language. However, the possibilities presented by the combination of large scale online social networks with easily affordable GPS enabled devices are so remarkable that might permit us to observe, for the first time, how diatopic differences arise and develop in time.

## Acknowledgments

We thank I. Fernández-Ordóñez and F. Moreno Fernández for useful discussions.

Disclaimer: This product was made utilizing the LandScan 2007 High Resolution global Population Data Set copyrighted by UT-Battelle, LLC,

operator of Oak Ridge National Laboratory under Contract No. DE-AC05-00OR22725 with the United States Department of Energy. The United States Government has certain rights in this Data Set. Neither UT-BATTELLE, LLC NOR THE UNITED STATES DEPARTMENT OF ENERGY, NOR ANY OF THEIR EMPLOYEES, MAKES ANY WARRANTY, EXPRESS OR IMPLIED, OR ASSUMES ANY LEGAL LIABILITY OR RESPONSIBILITY FOR THE ACCURACY, COMPLETENESS, OR USEFULNESS OF THE DATA SET.

## Author Contributions

Conceived and designed the experiments: BG DS. Performed the experiments: BG. Analyzed the data: BG. Wrote the paper: BG DS.

## References

- Chambers J, Trudgill P (1998) *Dialectology*. Cambridge University Press.
- Labov W, Sharon A, Boberg C (2005) *Atlas of North American English*. Phonetics, Phonology and Sound Change. De Gruyter.
- Bauer L (2004) Inferring variation and change from public corpora. In *The handbook of language variation and change*, ed. by Chambers J K, Trudgill P, Schilling-Estes N, Backwell Publishing, 97:114
- Borge-Holthoefer J, Rivero A, García I, Cauhé E, Ferrer A, et al. (2011) Structural and dynamical patterns on online social networks: the Spanish May 15th movement as a case study. *PLoS One* 6: e23883.
- Tumasjan A, Sprenger T, Sandner P, Welpe I (2010) Predicting elections with Twitter: What 140 characters reveal about political sentiment. In: *Proceedings of*

- the Fourth International AAAI Conference on Weblogs and Social Media: 178–185.
6. Culotta A (2010) Towards detecting influenza epidemics by analyzing Twitter messages. In: Proceedings of the First Workshop on Social Media Analytics. ACM: 115–122.
  7. Salathé M, Khandelwal S (2011) Assessing Vaccination Sentiments with Online Social Media: Implications for Infectious Disease Dynamics and Control. *PLoS Computational Biology* 7: e1002199.
  8. Salathé M, Bengtsson L, Bodnar TJ, Brewer DD, Brownstein JS, et al. (2012) Digital Epidemiology. *PLoS Comput Biol* 8: E1002616.
  9. Kulshrestha J, Kooti F, Nikraves A, Gummadi K (2012) Geographic dissection of the Twitter network. In: Proceedings of the 6th International AAAI Conference on Weblogs and Social Media (ICWSM).
  10. Mislove A, Lehmann S, Ahn Y, Onnela J, Rosenquist J (2011) Understanding the demographics of Twitter users. In: Fifth International AAAI Conference on Weblogs and Social Media.
  11. Hong L, Convertino G, Chi E (2011) Language matters in Twitter: A large scale study. In: International AAAI Conference on Weblogs and Social Media: 518–521.
  12. Mocanu D, Baronchelli A, Perra N, Gonçalves B, Zhang Q, et al. (2013) The Twitter of Babel: Mapping World Languages through Microblogging Platforms. *PLoS ONE* 8: e61981.
  13. Eisenstein J, O'Connor B, Smith N, Xing E (2010) A latent variable model for geographic lexical variation. In: Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing: 1277–1287.
  14. Russ B (2012) Examining large-scale regional variation through online geotagged corpora. ADS Annual Meeting.
  15. Ethnologue website. Available: [http://www.ethnologue.org/ethno\\_docs/distribution.asp?by=size](http://www.ethnologue.org/ethno_docs/distribution.asp?by=size) Accessed 2014 Oct. 27.
  16. Penny R (2000) Variation and change in Spanish. Cambridge University Press.
  17. McCandless M (2012) <http://code.google.com/p/chromium-compact-language-detector> Accessed 2014 Oct. 27.
  18. Stewart M (1999) The Spanish language today. Routledge.
  19. Moreno Fernández F, Otero Roth J (2007) Atlas de la lengua española. Ariel.
  20. Smith A, Rainie L (2010) Overview: The people who use Twitter. Available: <http://www.pewinternet.org/2010/12/09/8-of-online-americans-use-twitter/> Accessed 2014 Oct. 27.
  21. Ueda H, Takagaki T (1993) VARILEX, Variación léxica del español del mundo. University of Tokyo. Available at <http://lecture.ecc.u-tokyo.ac.jp/~ueda/varilex/>.
  22. Pham DT, Dimov SS, Nguyen CD (2005) Selection of K in K-means clustering. *Journal of Mechanical Engineering Science* 219: 103.
  23. Rousseeuw PJ (1987) Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics* 20: 53.
  24. Landsat 2007 Dataset. Available: <http://web.ornl.gov/sci/landsat/index.shtml> Accessed 2014 Oct. 27.
  25. López Morales H (2005) Últimas investigaciones sobre léxico hispanoamericano: unidad y variedad. In: Homenaje a José Joaquín Montes Giraldo: estudios de dialectología, lexicografía, lingüística general, etnolingüística e historia cultural: 333–358.
  26. Trudgill P, Hannah J (2002) International English. A guide to the varieties of Standard English. Arnold.
  27. Cotton EG, Sharp JM (1988) Spanish in the Americas. Georgetown University Press.