

# ¡Qué maravilla! Multimodal Sarcasm Detection in Spanish: a Dataset and a Baseline

**Khalid Alnajjar**

Department of Digital Humanities  
University of Helsinki

khalid.alnajjar@helsinki.fi

**Mika Hämäläinen**

Department of Digital Humanities  
University of Helsinki

mika.hamalainen@helsinki.fi

## Abstract

We construct the first ever multimodal sarcasm dataset for Spanish. The audiovisual dataset consists of sarcasm annotated text that is aligned with video and audio. The dataset represents two varieties of Spanish, a Latin American variety and a Peninsular Spanish variety, which ensures a wider dialectal coverage for this global language. We present several models for sarcasm detection that will serve as baselines in the future research. Our results show that results with text only (89%) are worse than when combining text with audio (91.9%). Finally, the best results are obtained when combining all the modalities: text, audio and video (93.1%).

## 1 Introduction

Figurative language is one of the most difficult forms of natural language to model computationally and there have been several studies in the past focusing on its subcategories such as metaphor interpretation (Xiao et al., 2016; Hämäläinen and Alnajjar, 2019a), humor generation (Hämäläinen and Alnajjar, 2019b) and analyzing idioms (Flor and Klebanov, 2018). Sarcasm is one of the extreme forms of figurative language, where the meaning of an utterance has little to do with the surface meaning (see Kreuz and Glucksberg 1989).

Understanding sarcasm is difficult even for us humans as it requires certain mental capacities such as a theory of mind (see Zhu and Wang 2020) and it is very dependent on the context and speaker who is being sarcastic. There are also very different view to sarcasm in the literature, for example, according to Kumon-Nakamura et al. (1995) sarcasm requires an allusion to a failed expectation and pragmatic insincerity (see Grice 1975) to be present in the same time. However, Utsumi (1996) highlights that these two preconditions are not enough, as sarcasm needs an ironic context to take place.

Haverkate (1990) argues that, in the context of sarcasm, the meaning difference can either be the

complete opposite of the semantic meaning of a sentence or somewhat different as seen in the lexical opposition of the words and the intended meaning. The fact that there are several different theoretical ways of understanding sarcasm, highlights the complexity of the phenomenon.

In this paper, we present an audio aligned dataset for sarcasm detection in Spanish. The dataset containing text and video timestamps has been released openly on Zenodo<sup>1</sup>. An access to the dataset with the video clips<sup>2</sup> can be granted upon request for academic use only. In addition, we will present a baseline model for this dataset to conduct multimodal sarcasm detection in Spanish.

## 2 Related work

In this section, we will present some of the recent related work on sarcasm detection. There has been some work also on sarcasm generation (Chakrabarty et al., 2020) and interpretation (Peled and Reichart, 2017), but they are rather different as tasks and we will not discuss them in detail.

Badlani et al. (2019) show an approach for sarcasm detection in online reviews. They train a CNN (convolutional neural network) based model on separate feature embeddings for sarcasm, humor, sentiment and hate speech. Similarly, Babanejad et al. (2020) also detect sarcasm in text. They combine an LSTM (long short-term memory) model with BERT. Dubey et al. (2019) also work on text only by detecting sarcastic numbers in tweets. They experiment with rules, SVMs (support vector machines) and CNNs.

Cai et al. (2019) use an LSTM model to detect sarcasm in tweets. Their approach is multimodal in the sense that it takes text and images into account, but it does not deal with audio and video like our

<sup>1</sup>Open access version of the data (contains text only) <https://zenodo.org/record/4701383>

<sup>2</sup>Access by request version of the data (videos and text) <https://zenodo.org/record/4707913>

Speaker	Utterance	Translation	Sarcasm
Archer	No, Lana, para nada	No, Lana, not at all	true
Stan	Lo siento chicos, mi papá dice que está muy ocupado con los Broncos, no tiene tiempo	I am sorry guys, my dad says he is very busy with the Broncos, he doesn't have time	false
Lana	Decías algo acerca de un plan	You said something about a plan	true

Table 1: Example sentences from the dataset.

approach.

Castro et al. (2019) present a multimodal sarcasm dataset in English. The dataset consists of annotated videos from TV sitcoms such as Friends and the Big Bang Theory, apart from being in English instead of Spanish, the main difference is that our dataset consists of animated cartoons instead of TV shows played by real people. Another big difference is in the data collection as they opted for querying sarcastic video clips, where as the data we work with represents full episodes. Chauhan et al. (2020) use this data and present a multimodal sarcasm detection framework based on a Bi-GRU model.

Many of the related work has been focusing on text only. Research on multimodal approaches has been carried out only for English data, not unlike the textual approaches.

### 3 Dataset

We base our work on the sarcasm annotated dataset from the MA thesis of the second author of this paper Hämäläinen (2016)<sup>3</sup>. This dataset is based on two episodes of South Park with voice-overs in Latin-American Spanish and two episodes of Archer with voice-overs in Spanish of Spain. The dataset has the speaker, their utterance and sarcasm annotations for each utterance in all of the episodes. However, the released data has been released shuffled for copyright reasons and it contains text only. Unlike the recent multimodal dataset for English (Castro et al., 2019), this data is expert annotated according to several different theories on sarcasm.

Annotation based on theories is important in order to establish a concrete meaning for sarcasm and to avoid the same mistakes as Castro et al. (2019) had. In their paper, they report that the most sarcastic character in The Big Bang Theory is Sheldon, however this cannot be true as one of the main characteristics of Sheldon is that he does not understand sarcasm. Therefore, their annotations ignore the fundamentally important characteristic of sarcasm,

<sup>3</sup>Available on <https://www.kaggle.com/mikahama/the-best-sarcasm-annotated-dataset-in-spanish>

which is speaker intent, and rather they consider sarcasm purely based on subjective intuition.

In order to produce a multimodal dataset out of the existing one, we locate the corresponding videos for the annotations and manually align them with the video. We use our own in-house tool JustAnnotate for this task<sup>4</sup>. This was a time consuming task, but as a result we ended up with a high-quality dataset with audio, video and text aligned. While aligning the dataset, we found several errors in the original transcriptions that we fixed. We did not alter the sarcasm annotations. In addition to the alignment, we introduced scene annotations. An episode of a TV show consists of many different scenes, and sarcasm is typically highly contextual, we indicate in the data which utterances belong to the same scene to better capture the context of each utterance.

Table 1 shows an example of the dataset. The English translation is provided for convenience, but it is not included in the dataset itself. Each line is aligned with audio and video. As we can see from these examples, sarcasm in the dataset is very contextually dependent as the sarcastic sentences presented in the table might equally well be sincere remarks if uttered by someone else or in a different context.



Figure 1: Archer uttering a sarcastic sentence that goes against the common sense

Figure 1 shows an example of a scene in the corpus. In this particular scene, Archer asks sarcastically *¿Dónde se compra la leche materna?* (Where

<sup>4</sup><https://mikakalevi.com/downloads/JustAnnotate.exe>

does one buy breast milk?). This is an example of sarcasm in the corpus where sarcasm violates common sense. Depending on the speaker, the utterance might be sarcastic or the speaker might lack knowledge on the topic.



Figure 2: Cartman uttering a sarcastic sentence that can be resolved only by visual cues.

In Figure 2 Cartman comments on the neckpiece of Stan by saying *Esas corbatas están de moda, tiene suerte de tenerla* (Those neckpieces are fashionable, you are lucky to have one). This is an example of a very different type of sarcasm that cannot be detected just by having common knowledge about the world. In order to understand the sarcastic intent, a system would need to have an access to the video as well to detect the unfashionable neckpiece and the disappointed facial expression of Stan.

## 4 Method

In this section, we present our method for detecting sarcasm in the multimodal dataset. We experiment with text only, text and audio and all modalities. All models are trained by using the same random train (80%) and test (20%) splits. For the neural model, 10% of the training split is used for validation.

### 4.1 Text only

For the text only model, we experiment with two models. In the first one, we use an off the shelf OpenNMT (Klein et al., 2017) model. We train the model using a bi-directional long short-term memory (LSTM) based model (Hochreiter and Schmidhuber, 1997) with the default settings except for the encoder where we use a BRNN (bi-directional recurrent neural network) (Schuster and Paliwal, 1997) instead of the default RNN (recurrent neural network). We use the default of two layers for both the encoder and the decoder and the default atten-

tion model, which is the general global attention presented by Luong et al. (2015). The model is trained for the default 100,000 steps.

The second model is a Support Vector Machine (SVM) (Schölkopf et al., 2000), due to its efficiency when dealing with a high dimensional space and ability to train a model with small data. We use the SVM implementation provided in Scikit-learn (Pedregosa et al., 2011). Following the work of Castro et al. (2019), we use an RBF kernel and a scaled gamma. The regularization parameter  $C$  is set for 1000. This setup is followed in all of our SVM models.

Regarding the textual features of the SVM, we make use of GloVe (Pennington et al., 2014) embeddings<sup>5</sup> trained on the Spanish Billion Words Corpus (Cardellino, 2019) and ELMo (Peters et al., 2018) embeddings provided by (Che et al., 2018). Each textual instance is tokenized using TokTok<sup>6</sup>, and then a sentence-level vector is constructed by computing the centroid (i.e., average vector) of all tokens, for each word embeddings type. In the case of ELMo, the vector of each token is the average of the last three layers of the neural network. The input to the SVM model is the concatenation of the two types of sentence embeddings.

### 4.2 Text and audio

This model is an SVM based model that extends the textual SVM model with audio features. We do not extend the OpenNMT model with audio features as the library does not provide us with audio and video inputters.

For all the audio, we set their sample size into 22 kHz to convert the data into a manageable and consistent size. Thereafter, we extract different audio features using librosa (McFee et al., 2020). These features include short-time Fourier transform (Nawab and Quatieri, 1987), mel-frequency cepstral coefficients (Stevens et al., 1937), chroma, Tonnetz (Harte et al., 2006), zero-crossing rate, spectral centroid and bandwidth, and pitches. In total, 13 features<sup>7</sup> were extracted. By combining all these features, we get the audio vector.

<sup>5</sup><https://github.com/dccuchile/spanish-word-embeddings>

<sup>6</sup><https://github.com/jonsafari/tok-tok>

<sup>7</sup>We used the following methods from librosa: *stft*, *mfcc*, *chroma\_stft*, *spectral\_centroid*, *spectral\_bandwidth*, *spectral\_rolloff*, *zero\_crossing\_rate*, *piptrack*, *onset\_strength*, *mel\_spectrogram*, *spectral\_contrast*, *tonnetz* and *harmonic*

### 4.3 All modalities

For videos, instead of trying to represent an entire video as a vector like some of the existing approaches (Hu et al., 2016) to video processing, we extract 3 frames for each video corresponding to an utterance. We extract the frames by dividing the frames of a video clip into three evenly sized chunks and taking the first frame of each chunk. The key motivation behind this is that we are working with animation, where most of the frames are static and changes in between frames are not big. Therefore representing the entire video clip is not important and it would only increase the complexity of the system.

We extract visual features from each of the three frames extracted using a pre-trained ResNet-152 model (He et al., 2016). Features are taken from the last layer in the network, and the overall video vector is the sequence of the three feature embeddings, in the same order. All the vectors described above (i.e., textual, audio and visual vectors) are passed as input to the all-modalities SVM model.

## 5 Results

In this section, we report the accuracy of predictions by the neural model and the three SVM models that are based on 1) text only, 2) text and audio, and 3) text, audio and video. The results can be seen in Table 2.

Input	Accuracy
Neural Model	
Text	87.5%
SVM	
Text	89.0%
Text + Audio	91.9%
Text + Audio + Video	<b>93.1%</b>

Table 2: Accuracies of the predictions by all models for the sarcasm detection task.

As we can see in the results, having more modalities in the training improved the results. The audio features were able to capture more features important for sarcasm than pure text. Having all the three modalities at the same time gave the best results, with a 4.1% gain in the accuracy from the text-based model. The neural model reached to the lowest accuracy, most likely due to the fact that it was not trained with pretrained embeddings, a source of information that was available to the SVM models.

### 5.1 Error analysis

When we look at the predictions by the model best model (text + audio + video), we can see that the sarcasm detection is not at all an easy task.

An interesting example of a correctly predicted sarcastic utterance is *Lucen bien muchachos. ¡A patear culos!* (You look great, guys. Let’s kick some ass!). This is an example of a visually interpretable sarcasm where the kids the sentence was uttered to looked all ridiculous. This would seem, at first, to highlight that the model has learned something important based on the visual features. However, we can see that this is not at all the case as the model predicts incorrectly the following sarcastic utterance: *Sí Stan, es lo que quiere la gente. No te preocupes, luces genial.* (Yes Stan, that is what the people want. Don’t worry, you look great.) The context is similar to the one where the model predicted the sarcasm correctly, which means that the visual features are not representative enough for the model to correctly annotate detect this sarcastic utterance.

Interestingly, the model predicted *Sí amigo, es una réplica de la corbata del Rey Enrique V* (Yes friend, it is a replica of the neckpiece of the King Henry V) as sarcastic while in fact the utterance was not sarcastic. This utterance refers to the same neckpiece as seen in Figure 2. The neckpiece appeared frequently in sarcastic contexts, so the model overgeneralized that anything said about the neckpiece must be sarcastic.

## 6 Conclusions

We have presented the first multimodal dataset for detecting sarcasm in Spanish. The dataset has been released on Zenodo. Our initial results serve as a baseline for any future work on sarcasm detection on this dataset.

Based on the results, it is clear that multimodality aids in detecting sarcasm as more contextual information is exposed to the model. Despite the improvements when considering multiple modalities, sarcasm detection is a very difficult task to model as it demands a global understanding of the world and the specific context the sarcastic utterance is in, as discussed in our error analysis. Even though the overall accuracy is high, it is clear the model makes errors that indicate that it has learned the data, but not the phenomenon.



## References

- Nastaran Babanejad, Heidar Davoudi, Aijun An, and Manos Papagelis. 2020. [Affective and contextual embedding for sarcasm detection](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 225–243, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Rohan Badlani, Nishit Asnani, and Manan Rai. 2019. [An ensemble of humour, sarcasm, and hate speech for sentiment classification in online reviews](#). In *Proceedings of the 5th Workshop on Noisy User-generated Text (W-NUT 2019)*, pages 337–345, Hong Kong, China. Association for Computational Linguistics.
- Yitao Cai, Huiyu Cai, and Xiaojun Wan. 2019. [Multi-modal sarcasm detection in Twitter with hierarchical fusion model](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2506–2515, Florence, Italy. Association for Computational Linguistics.
- Cristian Cardellino. 2019. [Spanish Billion Words Corpus and Embeddings](#).
- Santiago Castro, Devamanyu Hazarika, Verónica Pérez-Rosas, Roger Zimmermann, Rada Mihalcea, and Soujanya Poria. 2019. Towards multimodal sarcasm detection (an `_obviously_` perfect paper). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4619–4629.
- Tuhin Chakrabarty, Debanjan Ghosh, Smaranda Muresan, and Nanyun Peng. 2020. [R<sup>3</sup>: Reverse, retrieve, and rank for sarcasm generation with commonsense knowledge](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7976–7986, Online. Association for Computational Linguistics.
- Dushyant Singh Chauhan, Dhanush S R, Asif Ekbal, and Pushpak Bhattacharyya. 2020. [Sentiment and emotion help sarcasm? a multi-task learning framework for multi-modal sarcasm, sentiment and emotion analysis](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4351–4360, Online. Association for Computational Linguistics.
- Wanxiang Che, Yijia Liu, Yuxuan Wang, Bo Zheng, and Ting Liu. 2018. [Towards better UD parsing: Deep contextualized word embeddings, ensemble, and treebank concatenation](#). In *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 55–64, Brussels, Belgium. Association for Computational Linguistics.
- Abhijeet Dubey, Lakshya Kumar, Arpan Somani, Aditya Joshi, and Pushpak Bhattacharyya. 2019. [“when numbers matter!”: Detecting sarcasm in numerical portions of text](#). In *Proceedings of the Tenth Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 72–80, Minneapolis, USA. Association for Computational Linguistics.
- Michael Flor and Beata Beigman Klebanov. 2018. Catching idiomatic expressions in efl essays. In *Proceedings of the Workshop on Figurative Language Processing*, pages 34–44.
- H Paul Grice. 1975. Logic and conversation. 1975, pages 41–58.
- Mika Härmäläinen. 2016. [Reconocimiento automático del sarcasmo - ¡Esto va a funcionar bien!](#) Master’s thesis, University of Helsinki, Finland. URN:NBN:fi:hulib-201606011945.
- Mika Härmäläinen and Khalid Alnajjar. 2019a. [Let’s FACE it. Finnish poetry generation with aesthetics and framing](#). In *Proceedings of the 12th International Conference on Natural Language Generation*, pages 290–300, Tokyo, Japan. Association for Computational Linguistics.
- Mika Härmäläinen and Khalid Alnajjar. 2019b. Modelling the socialization of creative agents in a master-apprentice setting: The case of movie title puns. In *Proceedings of the 10th International Conference on Computational Creativity*. Association for Computational Creativity.
- Christopher Harte, Mark Sandler, and Martin Gasser. 2006. Detecting harmonic change in musical audio. In *Proceedings of the 1st ACM workshop on Audio and music computing multimedia*, pages 21–26.
- Henk Haverkate. 1990. A speech act analysis of irony. *Journal of pragmatics*, 14(1):77–109.
- K. He, X. Zhang, S. Ren, and J. Sun. 2016. [Deep residual learning for image recognition](#). In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Sheng-Hung Hu, Yikang Li, and Baoxin Li. 2016. Video2vec: Learning semantic spatio-temporal embeddings for video representation. In *2016 23rd International Conference on Pattern Recognition (ICPR)*, pages 811–816. IEEE.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander M. Rush. 2017. [OpenNMT: Open-Source Toolkit for Neural Machine Translation](#). In *Proc. ACL*.
- Roger J Kreuz and Sam Glucksberg. 1989. How to be sarcastic: The echoic reminder theory of verbal irony. *Journal of experimental psychology: General*, 118(4):374.

- Sachi Kumon-Nakamura, Sam Glucksberg, and Mary Brown. 1995. How about another piece of pie: The allusional pretense theory of discourse irony. *Journal of Experimental Psychology: General*, 124(1):3.
- Minh-Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective approaches to attention-based neural machine translation. *arXiv preprint arXiv:1508.04025*.
- Brian McFee, Vincent Lostanlen, Alexandros Metsai, Matt McVicar, Stefan Balke, Carl Thomé, Colin Raffel, Frank Zalkow, Ayoub Malek, Dana, Kyungyun Lee, Oriol Nieto, Jack Mason, Dan Ellis, Eric Battenberg, Scott Seyfarth, Ryuichi Yamamoto, Keunwoo Choi, viktorandreevichmorozov, Josh Moore, Rachel Bittner, Shunsuke Hidaka, Ziyao Wei, nullmightybofo, Darío Hereñú, Fabian-Robert Stöter, Pius Friesch, Adam Weiss, Matt Vollrath, and Taewoon Kim. 2020. [librosa/librosa: 0.8.0](#).
- S. Hamid Nawab and Thomas F. Quatieri. 1987. *Short-Time Fourier Transform*, page 289–337. Prentice-Hall, Inc., USA.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Lotem Peled and Roi Reichart. 2017. [Sarcasm SIGN: Interpreting sarcasm with sentiment based monolingual machine translation](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1690–1700, Vancouver, Canada. Association for Computational Linguistics.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [GloVe: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proc. of NAACL*.
- Bernhard Schölkopf, Alex J. Smola, Robert C. Williamson, and Peter L. Bartlett. 2000. [New support vector algorithms](#). *Neural Comput.*, 12(5):1207–1245.
- Mike Schuster and Kuldip K Paliwal. 1997. Bidirectional recurrent neural networks. *IEEE transactions on Signal Processing*, 45(11):2673–2681.
- Stanley Smith Stevens, John Volkman, and Edwin Broomell Newman. 1937. A scale for the measurement of the psychological magnitude pitch. *The journal of the acoustical society of america*, 8(3):185–190.
- Akira Utsumi. 1996. A unified theory of irony and its computational formalization. In *COLING 1996 Volume 2: The 16th International Conference on Computational Linguistics*.
- Ping Xiao, Khalid Alnajjar, Mark Granroth-Wilding, Kat Agres, Hannu Toivonen, et al. 2016. Meta4meaning: Automatic metaphor interpretation using corpus-derived word associations. In *Proceedings of the Seventh International Conference on Computational Creativity*. Sony CSL Paris.
- Ning Zhu and Zhenlin Wang. 2020. [The paradox of sarcasm: Theory of mind and sarcasm use in adults](#). *Personality and Individual Differences*, 163:110035.