

# OFAI–UKP at HAHA@IberLEF2019: Predicting the Humorousness of Tweets Using Gaussian Process Preference Learning

Tristan Miller<sup>1</sup>[0000–0002–0749–1100], Erik-Lân Do Dinh<sup>2</sup>[0000–0002–1536–3854],  
Edwin Simpson<sup>2</sup>[0000–0002–6447–1552], and Iryna Gurevych<sup>2</sup>[0000–0003–2187–7621]

<sup>1</sup> Austrian Research Institute for Artificial Intelligence (OFAI), Freyung 6,  
1010 Vienna, Austria, <http://www.ofai.at/>

<sup>2</sup> Ubiquitous Knowledge Processing Lab (UKP-TUDA), Department of Computer  
Science, Technische Universität Darmstadt, Hochschulstraße 10, 64289 Darmstadt,  
Germany, <https://www.ukp.tu-darmstadt.de/>

**Abstract** Most humour processing systems to date make at best discrete, coarse-grained distinctions between the comical and the conventional, yet such notions are better conceptualized as a broad spectrum. In this paper, we present a probabilistic approach, a variant of Gaussian process preference learning (GPPL), that learns to rank and rate the humorousness of short texts by exploiting human preference judgments and automatically sourced linguistic annotations. We apply our system, which had previously shown good performance on English-language one-liners annotated with pairwise humorousness annotations, to the Spanish-language data set of the HAHA@IberLEF2019 evaluation campaign. We report system performance for the campaign’s two subtasks, humour detection and funniness score prediction, and discuss some issues arising from the conversion between the numeric scores used in the HAHA@IberLEF2019 data and the pairwise judgment annotations required for our method.

**Keywords:** Computational humour · Humour · Gaussian process preference learning · GPPL · Best–worst scaling

## 1 Introduction

Humour is an essential part of everyday communication, particularly in social media [21,33], yet it remains a challenge for computational methods. Unlike conventional language, humour requires complex linguistic and background knowledge to understand, which are difficult to integrate with NLP methods [19].

An important step in the automatic processing of humour is to recognize its presence in a piece of text. However, its intensity may be present or perceived

---

Copyright © 2019 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0). IberLEF 2019, 24 September 2019, Bilbao, Spain.

to varying degrees to its human audience [5]. This level of appreciation (i.e., *humorousness* or equivalently *funniness*) can vary according to the text’s content and structural features, such as nonsense or disparagement [7] or, in the case of puns, contextual coherence [25] and the cognitive effort required to recover the target word [18, pp. 123–124].

While previous work has considered mainly binary classification approaches to humorousness, the HAHA@IberLEF2019 shared task [10] also focuses on its gradation. This latter task is important for downstream applications such as conversational agents or machine translation, which must choose the correct tone in response to humour, or find appropriate jokes and wordplay in a target language. The degree of creativeness may also inform an application whether the semantics of a joke can be inferred from similar examples.

This paper describes the OFAI–UKP system that participated in both subtasks of the HAHA@IberLEF2019 evaluation campaign: binary classification of tweets as humorous or not humorous, and the quantification of humour in those tweets. Our system employs a Bayesian approach—namely, a variant of Gaussian process preference learning (GPPL) that infers humorousness scores or rankings on the basis of manually annotated pairwise preference judgments and automatically annotated linguistic features. In the following sections, we describe and discuss the background and methodology of our system, our means of adapting the HAHA@IberLEF2019 data to work with our system, and the results of our system evaluation on this data.

## 2 Background

Pairwise comparisons can be used to infer rankings or ratings by assuming a *random utility model* [37], meaning that the annotator chooses an instance with probability  $p$ , where  $p$  is a function of the *utility* of the instance. Therefore, when instances in a pair have similar utilities, the annotator selects one with a probability close to 0.5, while for instances with very different utilities, the instance with higher utility will be chosen consistently. The random utility model forms the core of two popular preference learning models, the Bradley–Terry model [6,26,31], and the Thurstone–Mosteller model [37,28]. Given this model and a set of pairwise annotations, probabilistic inference can be used to retrieve the latent utilities of the instances.

Besides pairwise comparisons, a random utility model is also employed by MaxDiff [27], a model for best–worst scaling (BWS), in which the annotator chooses the best and worst instances from a set. While the term “best–worst scaling” originally applied to the data collection technique [15], it now also refers to models such as MaxDiff that describe how annotators make discrete choices. Empirical work on BWS has shown that MaxDiff scores (instance utilities) can be inferred using either maximum likelihood or a simple counting procedure that produces linearly scaled approximations of the maximum likelihood scores [16]. The counting procedure defines the score for an instance as the fraction of times the instance was chosen as best, minus the fraction of times the instance was

chosen as worst, out of all comparisons including that instance [23]. From this point on, we refer to the counting procedure as BWS, and apply it to the tasks of inferring scores from both best–worst scaling annotations for metaphor novelty and pairwise annotations for funniness.

Gaussian process preference learning (GPPL) [11], a Thurstone–Mosteller–based model that accounts for the features of the instances when inferring their scores, can make predictions for unlabelled instances and copes better with sparse pairwise labels. GPPL uses Bayesian inference, which has been shown to cope better with sparse and noisy data [39,38,4,24], including disagreements between multiple annotators [12,36,14,22]. Through the random utility model, GPPL is able to handle disagreements between annotators as noise, since no label has a probability of one of being selected.

Given a set of pairwise labels, and the features of labelled instances, GPPL can estimate the posterior distribution over the utilities of any instances given their features. Relationships between instances are modelled by a Gaussian process, which computes the covariance between instance utilities as a function of their features [32]. Since typical methods for posterior inference [29] are not scalable ( $\mathcal{O}(n^3)$ , where  $n$  is the number of instances), some of the present authors introduced a scalable method for GPPL that permits arbitrarily large numbers of instances and pairs [35]. This method uses stochastic variational inference [20], which limits computational complexity by substituting the instances for a fixed number of *inducing points* during inference.

Our GPPL method has already been applied with good results to ranking arguments by convincingness (which, like funniness, is an abstract linguistic property that is hard to quantify directly) and to ranking English-language one-liners by humorousness [35,34]. In these two tasks, GPPL was found to outperform SVM and BiLSTM regression models that were trained directly on gold-standard scores, and to outperform BWS when given sparse training data, respectively. We therefore elect to use GPPL on the Spanish-language Twitter data of the HAHA@IberLEF2019 shared task.

In the interests of replicability, we will be freely releasing the code for running our GPPL system, including the code for the data conversion and subsampling process detailed in §3.2.<sup>3</sup>

## 3 Experiments

### 3.1 Tasks

The HAHA@IberLEF2019 evaluation campaign consists of two tasks. Task 1 is humour detection, where the goal is to predict whether or not a given tweet is humorous, as determined by a gold standard of binary, human-sourced annotations. Systems are scored on the basis of accuracy, precision, recall, and F-measure. Task 2 is humorousness prediction, where the aim is to assign each funny tweet a score approximating the average funniness rating, on a five-point scale, assigned

<sup>3</sup> <https://github.com/UKPLab/haha2019-GPPL>

by a set of human annotators. Here system performance is measured by root-mean-squared error (RMSE). For both tasks, the campaign organizers provide a collection of 24 000 manually annotated training examples. The test data consists of a further 6000 tweets whose gold-standard annotations were withheld from the participants.

### 3.2 Data Preparation

For each of the 24 000 tweets in the HAHA@IberLEF2019 training data, the task organizers asked human annotators to indicate whether the tweet was humorous, and if so, how funny they found it on a scale from 1 (“not funny”) to 5 (“excellent”). (This is essentially the same annotation scheme used for the first version of the corpus [9] which was used in the previous iteration of HAHA [8].) As originally distributed, the training data gives the text of each tweet along with the number of annotators who rated it as “not humour”, “1”, “2”, “3”, “4”, and “5”. For the purposes of Task 1, tweets in the positive class received at least three numerical annotations and at least five annotations in total; tweets in the negative class received at least three “not humour” annotations, though possibly fewer than five annotations in total. Only those tweets in the positive class are used in Task 2.

This ordinal data cannot be used as-is with our GPPL system, which requires as input a set of preference judgments between pairs of instances. To work around this, we converted the data into a set of ordered pairs of tweets such that the first tweet has a lower average funniness score than the second. (We consider instances in the negative class to have an average funniness score of 0.) While an exhaustive set of pairings would contain 575 976 000 pairs (minus the pairs in which both tweets have the same score), we produced only 10 730 229 pairs, which was the minimal set necessary to accurately order the tweets. For example, if the original data set contained three tweets  $A$ ,  $B$ , and  $C$  with average funniness scores 5.0, 3.0, and 1.0, respectively, then our data would contain the pairs  $(C, B)$  and  $(B, A)$  but not  $(C, A)$ . To save memory and computation time in the training phase, we produced a random subsample such that the number of pairs where a given tweet appeared as the funnier one was capped at 500. This resulted in a total of 485 712 pairs. In a second configuration, we subsampled up to 2500 pairs per tweet. We used a random 60% of this set to meet memory limitations, resulting in 686 098 pairs.

With regards to the tweets’ textual data, we do only basic tokenization as preprocessing. For lookup purposes (synset lookup; see §3.3), we also lemmatize the tweets.

### 3.3 Experimental Setup

As we adapt an existing system that works on English data [34], we generally reuse the features employed there, but use Spanish resources instead. Each tweet is represented by the vector resulting from a concatenation of the following:

- The average of the word embedding vectors of the tweet’s tokens, for which we use 200-dimensional pretrained Spanish Twitter embeddings [13].
- The average frequency of the tweet’s tokens, as determined by a Wikipedia dump.<sup>4</sup>
- The average word polysemy—i.e., the number of synsets per lemma of the tweet’s tokens, as given by the Multilingual Central Repository (MCR 3.0, release 2016) [17].

Using the test data from the HAHA@IberLEF2018 task [8] as a development set, we further identified the following features from the UO\_UPV system [30] as helpful:

- The heuristically estimated turn count (i.e., the number of tokens beginning with - or --) and binary dialogue heuristic (i.e., whether the turn count is greater than 2).
- The number of hashtags (i.e., tokens beginning with #).
- The number of URLs (i.e., tokens beginning with `www` or `http`).
- The number of emoticons.<sup>5</sup>
- The character and token count, as well as mean token length.
- The counts of exclamation marks and other punctuation (., ;?).

We adapt the existing GPPL implementation<sup>6</sup> using the authors’ recommended hyperparameter defaults [35]: batch size  $|P_i| = 200$ , scale hyperparameters  $\alpha_0 = 2$  and  $\beta_0 = 200$ , and the number of inducing points (i.e., the smaller number of data points that act as substitutes for the tweets in the dataset)  $M = 500$ . The maximum number of iterations was set to 2000. Using these feature vectors, hyperparameter settings, and data pairs, we require a training time of roughly two hours running on a 24-core cluster with 2 GHz CPU cores.

After training the model, an additional step is necessary to transform the GPPL output values to the original funniness range (0, 1–5). For this purpose, we train a Gaussian process regressor which we supply with the output values of the GPPL system as features and the corresponding HAHA@IberLEF2018 test data values as targets. However, this model can still yield results outside the desired range when applied to the GPPL output of the HAHA@IberLEF2019 test data. Thus, we afterwards map too-large and too-small values onto the range boundaries. We further set an empirically determined threshold for binary funniness estimation.

### 3.4 Results and Discussion

Tables 1 and 2 report results for the binary classification setup (Task 1) and the regression task (Task 2), respectively. Included in each table are the scores of our

<sup>4</sup> <https://dumps.wikimedia.org/eswiki/20190420/eswiki-20190420-pages-articles.xml.bz2>; last accessed on 2019-06-15.

<sup>5</sup> [https://en.wikipedia.org/wiki/List\\_of\\_emoticons#Western](https://en.wikipedia.org/wiki/List_of_emoticons#Western), Western list; last accessed on 2019-06-15.

<sup>6</sup> <https://github.com/UKPLab/tacl2018-preference-convincing>

**Table 1.** Results for Task 1 (humour detection)

System	F <sub>1</sub>	Precision	Recall	Accuracy
winner	0.821	0.791	0.852	0.855
OFAI-UKP	0.660	0.588	0.753	0.698
baseline	0.440	0.394	0.497	0.505

**Table 2.** Results for Task 2 (funniness score prediction)


System	RMSE
winner	0.736
OFAI-UKP	1.810
baseline	2.455

own system, as well as those of the top-performing system and a naïve baseline. For Task 1, the naïve baseline makes a random classification for each tweet (with uniform distribution over the two classes); for Task 2, it assigns a funniness score of 3.0 to each tweet.

In the binary classification setup, our system achieved an F-measure of 0.660 on the test data, representing a precision of 0.588 and a recall of 0.753. In the regression task, we achieved RMSE of 1.810. The results are based on the second data subsample (up to 2500 pairs), with the results for the first (up to 500 pairs) being slightly lower. Our results for both tasks, while handily beating those of the naïve baseline, are significantly worse than those reported by some other systems in the evaluation campaign, including of course the winner. This is somewhat surprising given GPPL’s very good performance in our previous English-language experiments [34].

Unfortunately, our lack of fluency in Spanish and lack of access to the gold-standard scores for the test set tweets precludes us from performing a detailed qualitative error analysis. However, we speculate that our system’s less than stellar performance can partly be attributed to the information loss in converting between the numeric scores used in the HAHA@IberLEF2019 tasks and the preference judgments used by our GPPL system. In support of this explanation, we note that the output of our GPPL system is rather uniform; the scores occur in a narrow range with very few outliers. (Figure 1 shows this outcome for the HAHA@IberLEF2018 test data.) Possibly this effect would have been less pronounced had we used a much larger subsample, or even the entirety, of the possible training pairs, though as discussed in §3.2, technical and temporal limitations prevented us from doing so. We also speculate that the Gaussian process regressor we used may not have been the best way of mapping our GPPL scores back onto the task’s funniness scale (albeit still better than a linear mapping).

Apart from the difficulties posed by the differences in the annotation and scoring, our system may have been affected by the mismatch between its language



**Figure 1.** Gold values of the HAHA@IberLEF2018 test data vs. the GPPL output of our system, before mapping to the expected funniness range using GPR. The lowest output value (−1400) was removed from the plot to obtain a better visualization.

resources and the language of the test data. That is, while we relied on language resources like Wikipedia and MCR that reflect standardized registers and prestige dialects, the HAHA@IberLEF2019 data is drawn from unedited social media, whose language is less formal, treats a different range of topics, and may reflect a wider range of dialects and writing styles. Twitter data in particular is known to present problems for vanilla NLP systems, at least without extensive cleaning and normalization [1]. This is reflected in our choice of word embeddings: while we achieved a Spearman rank correlation of  $\rho = 0.52$  with the HAHA@IberLEF2018 test data using embeddings based on Twitter data [13], the same system using more “standard” Wikipedia-/news-/Web-based embeddings [2] resulted in a correlation near zero.

## 4 Conclusion

This paper has presented the OFAI-UKP system for predicting both binary and graded humorousness. It employs Gaussian process preference learning, a Bayesian system that learns to rank and rate instances by exploiting pairwise preference judgments. By providing additional feature data (in our case, shallow linguistic features), the method can learn to predict scores for previously unseen items.

Though our system had previously achieved good results with rudimentary, task-agnostic linguistic features on two English-language tasks (including one in-

volving the gradation of humorousness), its performance on the Spanish-language Twitter data of HAHA@IberLEF2019 was less impressive. We tentatively attribute this to the information loss involved in the (admittedly artificial) conversion between the numeric annotations used in the task and the preference judgments required as input to our method, and to the fact that we do not normalize the Twitter data to match our linguistic resources.

Possible future work would include mitigating the above two problems (for example, by normalizing the language of the tweets, and by coming up with a better way of converting between humour annotation formats, or by sourcing new preference judgments from Spanish-speaking annotators), and by using additional, humour-specific features, including some of those used in past work as well as those inspired by the prevailing linguistic theories of humour [3]. The benefits of including word frequency also point to possible further improvements using  $n$ -grams, tf-idf, or other task-agnostic linguistic features.

### Acknowledgments

This work has been supported by the German Federal Ministry of Education and Research (BMBF) under the promotional reference 01UG1816B (CEDIFOR), by the German Research Foundation (DFG) as part of the QA-EduInf project (grants GU 798/18-1 and RI 803/12-1), by the DFG-funded research training group “Adaptive Preparation of Information from Heterogeneous Sources” (AIPHES; GRK 1994/1), and by the Austrian Science Fund (FWF) under project M 2625-N31. The Austrian Research Institute for Artificial Intelligence is supported by the Austrian Federal Ministry for Science, Research and Economy.

### References

1. Proceedings of the ACL 2015 Workshop on Noisy User-generated Text. Association for Computational Linguistics (Jul 2015). <https://doi.org/10.18653/v1/W15-43>, <https://www.aclweb.org/anthology/W15-4300>
2. Almeida, A., Bilbao, A.: Spanish 3B words word2vec embeddings (version 1.0) (2018). <https://doi.org/10.5281/zenodo.1410403>
3. Attardo, S.: Linguistic Theories of Humor. Mouton de Gruyter, Berlin (1994). <https://doi.org/10.1515/9783110219029>
4. Beck, D., Cohn, T., Specia, L.: Joint emotion analysis via multi-task Gaussian processes. In: Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing. pp. 1798–1803. Association for Computational Linguistics (2014). <https://doi.org/10.3115/v1/D14-1190>
5. Bell, N.D.: Failed humor. In: Attardo, S. (ed.) The Routledge Handbook of Language and Humor, pp. 356–370. Routledge Handbooks in Linguistics, Routledge, New York (Feb 2017), <https://www.routledgehandbooks.com/doi/10.4324/9781315731162.ch25>
6. Bradley, R.A., Terry, M.E.: Rank analysis of incomplete block designs: I. The method of paired comparisons. *Biometrika* **39**(3/4), 324–345 (Dec 1952). <https://doi.org/10.2307/2334029>



7. Carretero-Dios, H., Pérez, C., Buela-Casal, G.: Assessing the appreciation of the content and structure of humor: Construction of a new scale. *Humor: International Journal of Humor Research* **23**(3), 307–325 (Aug 2010). <https://doi.org/10.1515/humr.2010.014>
8. Castro, S., Chiruzzo, L., Rosá, A.: Overview of the HAHA task: Humor analysis based on human annotation at IberEval 2018. In: Rosso, P., Gonzalo, J., Martínez, R., Montalvo, S., de Albornoz, J.C. (eds.) *Proceedings of the Third Workshop on Evaluation of Human Language Technologies for Iberian Languages*. CEUR Workshop Proceedings, vol. 2150, pp. 187–194. Spanish Society for Natural Language Processing (Sep 2018), <http://ceur-ws.org/Vol-2150/overview-HAHA.pdf>
9. Castro, S., Chiruzzo, L., Rosá, A., Garat, D., Moncecchi, G.: A crowd-annotated Spanish corpus for humor analysis. In: *Proceedings of the Sixth International Workshop on Natural Language Processing for Social Media*. pp. 7–11. Association for Computational Linguistics (2018), <http://aclweb.org/anthology/W18-3502>
10. Chiruzzo, L., Castro, S., Etcheverry, M., Garat, D., Prada, J.J., Rosá, A.: Overview of HAHA at IberLEF 2019: Humor analysis based on human annotation. In: *Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2019)*. CEUR Workshop Proceedings (Sep 2019)
11. Chu, W., Ghahramani, Z.: Preference learning with Gaussian processes. In: *Proceedings of the 22nd International Conference on Machine Learning*. pp. 137–144. ACM (2005). <https://doi.org/10.1145/1102351.1102369>
12. Cohn, T., Specia, L.: Modelling annotator bias with multi-task Gaussian processes: An application to machine translation quality estimation. In: *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*. vol. 1, pp. 32–42. Association for Computational Linguistics (2013), <http://aclweb.org/anthology/P13-1004>
13. Deriu, J., Lucchi, A., De Luca, V., Severyn, A., Müller, S., Cieliebak, M., Hoffmann, T., Jaggi, M.: Leveraging large amounts of weakly supervised data for multi-language sentiment classification. In: *Proceedings of the 26th International World Wide Web Conference*. pp. 1045–1052. International World Wide Web Conferences Steering Committee (2017)
14. Felt, P., Ringger, E.K., Seppi, K.D.: Semantic annotation aggregation with conditional crowdsourcing models and word embeddings. In: *Proceedings of the 26th International Conference on Computational Linguistics*. pp. 1787–1796 (2016), <http://aclweb.org/anthology/C16-1168>
15. Finn, A., Louviere, J.J.: Determining the appropriate response to evidence of public concern: The case of food safety. *Journal of Public Policy & Marketing* **11**(2), 12–25 (Sep 1992). <https://doi.org/10.1177/074391569201100202>
16. Flynn, T.N., Marley, A.A.J.: Best–worst scaling: Theory and methods. In: Hess, S., Daly, A. (eds.) *Handbook of Choice Modelling*, pp. 178–201. Edward Elgar Publishing, Cheltenham, UK (2014). <https://doi.org/10.4337/9781781003152.00014>
17. Gonzalez-Agirre, A., Laparra, E., Rigau, G.: Multilingual Central Repository version 3.0. In: *Proceedings of the 8th International Conference on Language Resources and Evaluation*. pp. 2525–2529. European Language Resources Association (2012)
18. Hempelmann, C.F.: *Paronomasic Puns: Target Recoverability Towards Automatic Generation*. Ph.D. thesis, Purdue University, West Lafayette, IN, USA (Aug 2003)
19. Hempelmann, C.F.: Computational humor: Beyond the pun? In: Raskin, V. (ed.) *The Primer of Humor Research*, pp. 333–360. No. 8 in *Humor Research*, Mouton de Gruyter, Berlin (2008). <https://doi.org/10.1515/9783110198492.333>

20. Hoffman, M.D., Blei, D.M., Wang, C., Paisley, J.W.: Stochastic variational inference. *Journal of Machine Learning Research* **14**, 1303–1347 (May 2013), <http://jmlr.org/papers/v14/hoffman13a.html>
21. Holton, A.E., Lewis, S.C.: Journalists, social media, and the use of humor on Twitter. *Electronic Journal of Communication* **21**(1&2) (2011), <http://www.cios.org/EJCPUBLIC/021/1/021121.html>
22. Kido, H., Okamoto, K.: A Bayesian approach to argument-based reasoning for attack estimation. In: Sierra, C. (ed.) *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*. pp. 249–255. *International Joint Conferences on Artificial Intelligence* (2017). <https://doi.org/10.24963/ijcai.2017/36>
23. Kiritchenko, S., Mohammad, S.M.: Capturing reliable fine-grained sentiment associations by crowdsourcing and best–worst scaling. In: *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. pp. 811–817. *Association for Computational Linguistics* (2016). <https://doi.org/10.18653/v1/N16-1095>
24. Lampos, V., Aletras, N., Preotiuc-Pietro, D., Cohn, T.: Predicting and characterising user impact on Twitter. In: *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*. pp. 405–413. *Association for Computational Linguistics* (2014). <https://doi.org/10.3115/v1/E14-1043>
25. Lippman, L.G., Dunn, M.L.: Contextual connections within puns: Effects on perceived humor and memory. *Journal of General Psychology* **127**(2), 185–197 (Apr 2000). <https://doi.org/10.1080/00221300009598578>
26. Luce, R.D.: On the possible psychophysical laws. *Psychological Review* **66**(2), 81–95 (1959). <https://doi.org/10.1037/h0043178>
27. Marley, A.A.J., Louviere, J.J.: Some probabilistic models of best, worst, and best–worst choices. *Journal of Mathematical Psychology* **49**(6), 464–480 (2005). <https://doi.org/10.1016/j.jmp.2005.05.003>
28. Mosteller, F.: Remarks on the method of paired comparisons: I. The least squares solution assuming equal standard deviations and equal correlations. *Psychometrika* **16**(1), 3–9 (Mar 1951). <https://doi.org/10.1007/BF02313422>
29. Nickisch, H., Rasmussen, C.E.: Approximations for binary Gaussian process classification. *Journal of Machine Learning Research* **9**, 2035–2078 (Oct 2008), <http://www.jmlr.org/papers/volume9/nickisch08a/nickisch08a.pdf>
30. Ortega-Bueno, R., Muñiz Cuza, C.E., Medina Pagola, J.E., Rosso, P.: UO-UPV: Deep linguistic humor detection in Spanish social media. In: Rosso, P., Gonzalo, J., Martínez, R., Montalvo, S., de Albornoz, J.C. (eds.) *Proceedings of the Third Workshop on Evaluation of Human Language Technologies for Iberian Languages*. *CEUR Workshop Proceedings*, vol. 2150, pp. 203–213. *Spanish Society for Natural Language Processing* (Sep 2018), [http://ceur-ws.org/Vol-2150/HAHA\\_paper2.pdf](http://ceur-ws.org/Vol-2150/HAHA_paper2.pdf)
31. Plackett, R.L.: The analysis of permutations. *Journal of the Royal Statistical Society, Series C (Applied Statistics)* **24**(2), 193–202 (1975). <https://doi.org/10.2307/2346567>
32. Rasmussen, C.E., Williams, C.K.I.: *Gaussian Processes for Machine Learning*. *Adaptive Computation and Machine Learning*, MIT Press, Cambridge, MA, USA (2006), <http://www.gaussianprocess.org/gpml/>
33. Shifman, L.: *Memes in Digital Culture*. *Essential Knowledge*, MIT Press, Cambridge, MA, USA (Oct 2013)
34. Simpson, E., Do Dinh, E.L., Miller, T., Gurevych, I.: Predicting humorousness and metaphor novelty with Gaussian process preference learning. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL 2019)*. *Association for Computational Linguistics* (Jul 2019), to appear

35. Simpson, E., Gurevych, I.: Finding convincing arguments using scalable Bayesian preference learning. *Transactions of the Association for Computational Linguistics* **6**, 357–371 (2018), <http://aclweb.org/anthology/Q18-1026>
36. Simpson, E.D., Venanzi, M., Reece, S., Kohli, P., Guiver, J., Roberts, S.J., Jennings, N.R.: Language understanding in the wild: Combining crowdsourcing and machine learning. In: *Proceedings of the 24th International Conference on World Wide Web*. pp. 992–1002. International World Wide Web Conferences Steering Committee (2015). <https://doi.org/10.1145/2736277.2741689>
37. Thurstone, L.L.: A law of comparative judgment. *Psychological Review* **34**(4), 273–286 (1927). <https://doi.org/10.1037/h0070288>
38. Titov, I., Klementiev, A.: A Bayesian approach to unsupervised semantic role induction. In: *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*. pp. 12–22. Association for Computational Linguistics (2012), <http://aclweb.org/anthology/E12-1003>
39. Xiong, H.Y., Barash, Y., Frey, B.J.: Bayesian prediction of tissue-regulated splicing using RNA sequence and cellular context. *Bioinformatics* **27**(18), 2554–2562 (Sep 2011). <https://doi.org/10.1093/bioinformatics/btr444>