

Spanish Corpus of Tweets for Marketing

María Navas-Loro¹ (orcid.org/0000-0003-1011-5023), Víctor Rodríguez-Doncel¹
(orcid.org/0000-0003-1076-2511), Idafen Santana-Pérez¹
(orcid.org/0000-0001-8296-8629), Alba Fernández-Izquierdo¹, and Alberto
Sánchez²

¹ Ontology Engineering Group, Universidad Politécnica de Madrid, Spain

² Havas Media, Madrid, Spain

Abstract. This paper presents a corpus of manually tagged tweets in Spanish language, of interest for marketing purposes. For every Twitter post, tags are provided to describe three different aspects of the text: the emotions, whether it makes a mention to an element of the marketing mix and the position of the tweet author with respect to the purchase funnel. The tags of every Twitter post are related to one single brand, which is also specified for every tweet. The corpus is published as a collection of RDF documents with links to external entities. Details on the used vocabulary and classification criteria are provided, as well as details on the annotation process.

Keywords: corpus, marketing, marketing mix, sentiment analysis, NLP, purchase funnel, emotion analysis

1 Introduction

Twitter is a source of valuable feedback for companies to probe the public perception of their brands. Whereas sentiment analysis has been extensively applied to social media messages (see [16] among many), other dimensions of brand perception are still of interest and have received less attention [12], specially those related to marketing. In particular, marketing specialists are highly interested in: (a) knowing the position of a tweet author in the purchase funnel (this is, where in the different stages of the customer journey is the author in); (b) knowing to which element or elements of the marketing mix³ the text refers to and (c) knowing the author's affective situation with respect to a brand in the tweet.

This paper presents the MAS Corpus, a Spanish corpus of tweets of interest for marketing specialists, labeling messages in the three dimensions aforementioned. The corpus is freely available at <http://mascorpus.linkeddata.es/> and has been developed in the context of the Spanish research project LPS BIGGER⁴, which analyzed different dimensions of tweets in order to extract relevant information on marketing purposes. A first version of the corpus containing only the sentiment analysis annotations was released as the Corpus for Sentiment

³ <http://economictimes.indiatimes.com/definition/marketing-mix>

⁴ <http://www.cienlpsbigger.es>

Analysis towards Brands (SAB) and was described in [15]. Following this work, we have expanded the corpus tagging the messages in the two remaining dimensions described before: the purchase funnel and the marketing mix. Tweets that were almost identical to others have been removed. Categories of each of the three aspects tagged in the corpus (Sentiment Analysis, Marketing Mix and Purchase Funnel) can be found in Table 1.

Table 1. Tags for each category.

Category	Tags
Purchase funnel	awareness, evaluation, purchase, postpurchase, ambiguous, NC2
Marketing Mix	product, price, promotion, place, NC2
Sentiment Analysis	love, hate, satisfaction, dissatisfaction, happiness, sadness, trust, fear, NC2

2 Related Work

2.1 Sentiment Analysis

Even when Sentiment Analysis is a major field in Natural Language Processing, most of works in Spanish tend to focus on polarity [10, 5], being the efforts towards emotions really scarce [22]. Sources of corpora also differ to our aims, since they tend to use specific websites or limit to domains such as tourism and medical opinions [17, 14] instead of social media. An extended review of works in Spanish Sentiment Analysis with regard to our needs can be found in [15].

2.2 Purchase Funnel

Although different purchase funnel interpretations have been suggested in literature [3, 6], we have based our approach on the one defined in the LPS BIGGER project and already used in [25]. This purchase funnel consists of four different stages (*Awareness*, *Evaluation*, *Purchase* and *Postpurchase*), that reflect how the client gets to know the product, investigates or compares it to other options, acquires it and actually uses and reviews it, respectively.

To the best of our knowledge, there are not public Spanish corpora available containing purchase funnel annotations, since the only work in Spanish on this topic the authors are aware of did not release the dataset used [25]. Nevertheless, the concept of Purchase Intention has been widely covered in literature, especially for marketing purposes in English language. Differently to Sentiment Analysis, Purchase Intention tries to detect or distinguish whether the client intends to buy a product, rather than whether he likes it or not [26]. Starting with the WISH corpus [8], covering wishes in several domains and sources (including product reviews), most works aim to discriminate between different kinds of intentions of users: in [21], the analysis focuses in suggestions and wishes for products and services both in a private dataset and in a part of the previously

mentioned WISH corpus; also an analysis performed on tweets about different intentions can be found in [13].

Finally, the most similar categories to the ones in our purchase funnel interpretation are the ones in [4], where the authors differentiate between several kinds of intention, being some of them (such as *wish*, *compare* or *complain*) easy mappable to our purchase funnel stages. Also the corpus used in [9], that classifies into pre-purchase and post-purchase reviews, shares our “timeline” interpretation of the purchase funnel. Out of the marketing domain, corpora labeled with purchase funnel tags for an specific domain have also been published, e.g., for the London musicals and recreational events [7].

2.3 Marketing Mix

Although the original concept of marketing mix [2] contained twelve elements for manufacturers, the most extended categorization for marketing is the one proposed by [11], consisting of four aspects (*price*, *product*, *promotion*, *place*) often known as “the four Ps” (or 4Ps) and revisited several times in literature [24]. Nevertheless, while marketing mix is a well-known and extended concept in the marketing field, in NLP the task of identifying these facets is often simply referred as detecting or recognizing “aspects”, excepting some cases in literature [1]. This task has been often tackled in English [18, 20], while in Spanish corpora we can find a few datasets containing information about aspects, such as those in [5, 19].

3 Tagging Criteria

The corpus consists of more than 3k tweets of brands from different sectors, namely *Food*, *Automotive*, *Banking*, *Beverages*, *Sports*, *Retail* and *Telecom* (the complete list of brands, as well as statistics on the corpus, can be downloaded with it). When several brands appear in one tweet, just one of them is considered in the tagging process (the marked one); at the same time, the same tweet can appear several times in the corpus considering different brands. Every tweet is tagged in the dimensions exposed in Table 1; more than one tag is possible in sentiment and marketing mix dimensions (except simultaneously tagging the pairs of directly opposed emotions), while the purchase funnel, as representing a path on the purchase journey, only presents a tag per tweet. We describe below each dimension, along with a brief report on the criteria used for tagging each category (the complete criteria document can be downloaded with the corpus).

3.1 Sentiment Analysis

A tweet can be tagged with one or several emotions (as long as it does not contain directly opposite emotions), or with a *NC2* label meaning there are no emotions on it. Each basic emotion embraces also secondary emotions in it (described in Table 2), and a combination of them can express more complex feelings often seen in customers, such as shown in the following examples:

- When a customer is unable to find a desired product, the post is tagged as *sadness* (for the unavailability) and *satisfaction* (because it reveals previous satisfaction with the brand that deserves the effort of keep looking exactly for it instead of switching to one from another brand).
- When a post shows that a purchase is recurrent, it is tagged as *trust*, referring to the loyalty of the client.
- Emoticons of love are tagged as *love* and musical ones as *happiness* (unless irony happens). *Love* typically implies *happiness*.
- *Happiness* can only be tagged for an already acquired product or service.

Emotion	Related emotions
Trust	Optimism, Hope, Security
Satisfaction	Fulfillment, Contentment
Happiness	Joy, Gladness, Enjoyment, Delight, Amusement, Joviality, Enthusiasm, Jubilation, Pride, Triumph
Love	Passion, Excitement, Euphoria, Ecstasy
Fear	Nervousness, Alarm, Anxiety, Tenseness, Apprehension, Worry, Shock, Fright, Terror, Panic, Hysteria, Mortification
Dissatisfaction	Dislike, Rejection, Revulsion, Disgust, Irritation, Aggravation, Exasperation, Frustration, Annoyance
Sadness	Depression, Defeat, Hopelessness, Unhappiness, Anguish, Sorrow, Agony, Melancholy, Dejection, Loneliness, Humiliation, Shame, Guilt, Regret, Remorse, Disappointment, Alienation, Isolation, Insecurity
Hate	Rage, Fury, Wrath, Envy, Hostility, Ferocity, Bitterness, Resentment, Spite, Contempt, Vengefulness, Jealousy

Table 2. Main emotions and their secondary emotions.

3.2 Purchase Funnel

Each tweet can belong to a stage in the purchase funnel, be ambiguous or be related to a brand without the author being involved in the purchase (such as is the case of posts of the brand itself). Different phases and concrete examples are tagged in the corpus as follows:

- **Awareness** The first contact of the client with the brand (either showing a willingness to buy or not), usually expressed in first person and mentioning advertising, videos, publicity campaigns, etc. Some examples of awareness would be:
 - (1) *I just loved last Movistar ad.*
 - (2) *I like the videos in Nike's YouTube channel.*

- **Evaluation** The post implies some research on the brand (such as questions or seek of confirmation) or comparison to others (by showing preferences among them, for instance), and some interest in acquiring a product or service. Examples of evaluation would be the following:
 - (3) *I prefer Citroen to more expensive brands, such as Mercedes or BMW.*
 - (4) *Looking for a second-hand Kia Sorento in NY, please send me a DM.*
- **Purchase** There is a direct reference to the moment of a purchase or to a clear intention of purchase (usually in first person). Some examples:
 - (5) *I've finally decided to switch to Movistar.*
 - (6) *Buying my brand new blue Citroen right now!*
- **Postpurchase** Texts referring to a past purchase or to a current experience, implying to own a product. This class presents a special complexity, since interpretation on the same linguistic patterns change depending on the kind of product, as already exposed in [25] and exemplified in the sentences below:
 - (7) *I like Heineken, the taste is so good.
I would love a Heineken!*
 - (8) *I like BMWs, they are so classy!
I would love a BMW!*

In (7), the client has likely tasted that beer brand before; people does not tend to like or want beverages they have no experience with (at least without mentioning, such as in “*I want to taste the new Heineken.*”). But the same fact is not derived from more expensive items, even when expressed the same way, such as happens in (8): someone can like a car (such as its appearance or its engine) without having used it or intending to. This is why our criteria states that these kind of expressions must be tagged as *Postpurchase* for some brands (depending on the sector) and others must be tagged as *Ambiguous*, since there can be several possible and equally likely interpretations.
- **Ambiguous** This category includes critical posts, suggestions and recommendations, along with posts where it is not clear in which stage the customer is (such as the case mentioned above).
 - (9) *Do not buy Milka!*
 - (10) *Loving the new Kia!*
- **NC2** Includes impersonal messages without opinions (such as corporative news or responses of the brand to clients), questions implying no personal evaluation or intention (for instance, involving a third person), texts with buy or rental offers with no mention to real use experience, etc.
 - (11) *2008 Hyundai for sale.*
 - (12) *My aunt didn't like the Kia.*

3.3 Marketing Mix

We have added a *NC2* class to the four original McCarthy’s Ps to indicate none of the four aspects is treated in the tweet. It must be noted that, differently than the purchase funnel, several marketing mix tags can appear in the same tweet (except of the *NC2*). Brief explanation of each of the categories tagged for marketing mix, along with examples and part of the criteria, are exposed below:

- **Product** This category encompasses texts related to the features of the product (such as its quality, performance or taste), along with references to design (such as size, colors or packing) or guaranty, such as in the following examples:
 - (13) *I find the new iPhone too big for my pocket.*
 - (14) *I love the new mix Milka Oreo!*
 Note that when someone loves/likes something (such as food), we assume it refers to some feature of a product (such as its taste), so we tag it as *Product*.
- **Promotion** Texts referring to all the promotions and programs of the brand channeled to increase sales and ensure visibility to their products or the brand, such as advertisements, sponsorships (such as prices, sport teams or events), special offers, work offers, promotional articles, etc.
 - (15) *Freaking out with the new 2x1 @Ikea!*
 - (16) *La Liga BBVA is the best league in the world.*
- **Price** Includes economical aspects of a product, such as references to its value or promotions involving discounts or price drops (that must also be tagged as *Promotion*). Examples of texts that should be tagged as *Price* would be the following:
 - (17) *I'm afraid that I can't afford the new Toyota.*
 - (18) *Yesterday I saw the same Adidas for just 40€!*
- **Place** Aspects related to commercialization, such physical places of distribution of the products (for instance, if a product is difficult to find) and customer service (in every stage of the purchase: information, at the point of sale, postpurchase, technical support, etc).
 - (19) *I love the new Milka McFlurry at McDonalds*
 - (20) *Already three malls and unable to find the new Nike Pegasus!*
- **NC2** Impersonal messages of the brand, news or texts that include none of the aspects mentioned before.
 - (21) *Nike is paying no tax!*
 - (22) *I can't decide between Puleva and Pascual.*

4 The MAS Corpus

4.1 Building the corpus

A different approach was used for Marketing Mix and the Purchase Funnel tagging with respect to the Sentiment Analysis tagging procedure (where three taggers acted independently with just a common criteria document) exposed in [15]. This meets the need of streamlining the whole tagging process, that happens to be both difficult and time-consuming for taggers. This new procedure is briefly exposed below:

1. A first version of the criteria document was written, based on the study of literature and previous experience within the LPS BIGGER project.

2. Then Tagger 1 tagged a representative part of the corpus (about 800 tweets), highlighting main doubts and dubious tweets with regard to the criteria, that are revised; new tagging examples are added, and some nuances and special cases are rewritten.
3. Taggers 2 and 3 revise the tags by Tagger 1, paying special attention to tweets marked as dubious: if an agreement is reached, the tagging is updated consequently; otherwise, the tweet is tagged as *Ambiguous* or *NC2*.
4. Then each tagger takes a part of the corpus to tag it following the new criteria and highlighting doubts again; these tweets will be revised with remaining taggers, reaching an agreement on the final unique tags in the corpus.

4.2 Publishing the corpus as Linked Data

We maintain the RDF representation used in the previous version of the corpus, using again our own vocabulary⁵ to express the purchase funnel and the marketing mix. We also reuse Marl [27] and Onyx [23] for emotions and polarity, and SIOC⁶ and GoodRelations⁷ for post and brand representation. Also links to the entries of brands and companies in external databases such as Thomson Reuters' PermID⁸ and DBpedia⁹ extend the information in the tweets. Fig. 1 shows an example of a tweet tagged in the dimensions extracted from the corpus.

4.3 Corpus description

Final corpus contains 3763 tweets. Statistics on linguistic information in the corpus can be found in Table 3, along with specific data relevant for Social Media, such as the amount of hashtags, user mentions and URLs. The distribution of categories varies depending on the sector, as shown in Table 4. Mentions of *Place* are for instance more common in *Sports* than in other categories, such as *Beverages* or *Telecom*. Also when opinions are expressed differs: tweets in the *Food* sector tend to refer to the *Postpurchase* phase, while others tend to be more ambiguous or refer to previous phases. Regarding emotions, some of them just appear in certain domains, such as *Fear* for *Banking*.

5 Conclusions

Whereas the SAB corpus provided a collection of tweets tagged with labels useful for making Sentiment Analysis towards brands, this new corpus is of interest for the marketing analysis in a broader way; the MAS Corpus allows marketing professionals to have additional information of habits and behaviors, strong and weak points of the whole purchase experience, and also full insights on concrete aspects of each client reviews.

⁵ <http://sabcorpus.linkeddata.es/vocab>

⁶ <https://www.w3.org/Submission/sioc-spec/>

⁷ <http://purl.org/goodrelations/>

⁸ <https://permid.org/>

⁹ <http://dbpedia.org/>

Table 3. Total and average (per tweet) statistics on the corpus. Stanford CoreNLP was used for POS information, while patterns were used for detecting hashtags (`#`), mentions (`@`) and URLs (`www.*`/`http*`).

	TOTAL		AVG		TOTAL		AVG
Tweets	3763	-		Verbs	6971	1.85	
Sentences	5189	1.38		Nouns	8353	2.22	
Tokens	59555	15.83		NPs	6952	1.85	
Hashtags	1819	0.48		Adjectives	2761	0.73	
Mentions	2306	0.61		Adverbs	1584	0.42	
URLs	2111	0.56		Neg. Adverbs	560	0.15	

Table 4. Statistics on the corpus. Column *ANY* in emotional categories shows the percentage of posts with any emotion (this is, non neutral posts); remaining columns show the percentage of each category among these non neutral posts. For Purchase Funnel and Marketing Mix, each column represents the percentages of each of the tags described in Section 3.

	ANY	HAT	SAD	FEA	DIS	SAT	TRU	HAP	LOV
FOOD	54.79	1.50	1.20	0.00	8.08	45.21	44.01	14.67	12.87
AUTOMOTIVE	9.11	0.00	0.22	1.11	2.44	6.89	3.33	1.11	0.89
BANKING	24.67	5.33	1.00	15.00	23.83	1.33	0.50	0.00	0.00
BEVERAGES	63.11	2.07	1.19	0.74	19.11	44.00	32.74	7.26	7.70
SPORTS	34.15	2.45	2.60	0.31	13.32	18.84	11.94	4.90	11.33
RETAIL	33.00	3.20	1.11	1.48	11.95	14.53	14.41	3.69	3.45
TELECOM	40.17	12.97	0.84	0.00	30.13	8.79	6.28	3.35	1.26

	PURCHASE FUNNEL						MARKETING MIX				
	NC2	AWA	EVA	PUR	POS	AMB	NC2	PROD	PRI	PROM	PLA
FOOD	43.41	3.59	3.29	4.19	40.72	5.09	48.80	30.84	2.10	15.27	7.49
AUTOMOTIVE	85.56	2.67	4.00	0.22	4.44	3.33	77.56	4.67	2.00	16.00	1.56
BANKING	58.50	5.83	2.00	0.00	7.83	25.67	53.33	8.50	7.83	21.17	13.17
BEVERAGES	33.63	0.44	13.33	8.44	11.26	32.74	19.85	70.37	2.22	8.59	8.59
SPORTS	63.09	2.91	4.29	1.84	7.50	19.75	54.98	6.43	17.76	0.92	30.32
RETAIL	89.29	2.71	4.80	0.62	1.97	1.60	72.17	12.56	2.09	8.62	7.51
TELECOM	94.14	0.42	0.42	0.00	4.60	0.00	91.63	1.26	1.67	4.60	0.00


```

mas:827146264517165056 a sioc:Post ;
sioc:id "827146264517165056" ;
sioc:content "Las camisetas nike 2002~2004 y las adidas 2006~2008 son el amor de mi vida"@es ;
marl:describesObject mas:Nike ;
sabd:isInPurchaseFunnel sabv:postPurchase;
sabd:hasMarketingMix sabv:product;
onyx:hasEmotion sabv:love, sabv:satisfaction, sabv:happiness ;
marl:hasPolarity marl:positive ;
marl:forDomain "SPORT" .

mas:Nike a gr:Brand ;
rdfs:seeAlso <http://dbpedia.org/resource/Nike> ;
sabd:1-5000062703 a gr:Business ;
rdfs:label "Nike Inc", "Nike" ;
owl:sameAs permid:1-4295904620 .

```

Fig. 1. Sample tagged post, and extra information on its brand (Nike) and company (Nike Inc).

Acknowledgments. This work has been partially supported by LPS-BIGGER (IDI-20141259), esTextAnalytics project (RTC-2016-4952-7), Datos 4.0 project with ref. TIN2016-78011-C4-1-R, a Predoctoral grant by the Consejo de Educación, Juventud y Deporte de la Comunidad de Madrid partially founded by the European Social Fund, two Predoctoral grants from the I+D+i program of the Universidad Politécnica de Madrid and a Juan de la Cierva contract. We would also want to thank Pablo Calleja for his help in corpora statistics extraction.

References

1. Bel, N., Diz-pico, J., Pocostales, J.: Classifying short texts for a Social Media monitoring system Clasificación de textos cortos para un sistema monitor de los Social Media. *Procesamiento del Lenguaje Natural* 59, 57–64 (2017)
2. Borden, N.H.: The concept of the marketing mix. *Journal of advertising research* 4(2), 2–7 (1964)
3. Bruyn, A.D., Lilien, G.L.: A multi-stage model of word-of-mouth influence through viral marketing. *Int. Journal of Research in Marketing* 25(3), 151–163 (2008)
4. Cohan-Sujay, C., Madhulika, Y.: Intention Analysis for Sales, Marketing and Customer Service. *Proceedings of COLING 2012, Demonstration Papers*, (December 2012), 33–40 (2012)
5. Cumbreiras, M.Á.G., Cámara, E.M., et al.: TASS 2015 - The evolution of the Spanish opinion mining systems. *Procesamiento de Lenguaje Natural* 56, 33–40 (2016)
6. Elzinga, D., Mulder, S., Vetvik, O.J., et al.: The consumer decision journey. *McKinsey Quarterly* 3, 96–107 (2009)
7. García-Silva, A., Rodríguez-Doncel, V., Corcho, Ó.: Semantic characterization of tweets using topic models: A use case in the entertainment domain. *Int. J. Semantic Web Inf. Syst.* 9(3), 1–13 (2013)
8. Goldberg, A.B., Fillmore, N., Andrzejewski, D., Xu, Z., Gibson, B., Zhu, X.: May All Your Wishes Come True : A Study of Wishes and How to Recognize Them. *Proceedings of Human Language Technologies: NAACL '09 (June)*, 263–271 (2009)
9. Hasan, M., Kotov, A., Mohan, A., Lu, S., Stieg, P.M.: Feedback or Research: Separating Pre-purchase from Post-purchase Consumer Reviews. In: *Advances in*

- Information Retrieval. ECIR 2016. Lecture Notes in Computer Science, vol. 9626, pp. 682–688. Springer, Cham (2016)
10. Martínez-Cámara, E., Martín-Valdivia, M.T., et al.: Polarity classification for Spanish tweets using the COST corpus. *Journal of Information Science* 41(3), 263–272 (jun 2015)
 11. McCarthy, E.: *Basic Marketing, a Managerial Approach*. Sixth Edition, Homewood, Ill.: Richard D. Irwin, Inc. (1978)
 12. Moghaddam, S.: Beyond sentiment analysis: Mining defects and improvements from customer feedback. *LNCS 9022*, 400–410 (2015)
 13. Mohamed, H., Mohamed, S.G., Lamjed, B.S.: Customer Intentions Analysis of Twitter Based on Semantic Patterns. 2015 pp. 2–6 (2015)
 14. Molina-González, M.D., Martínez-Cámara, E., et al.: Cross-domain sentiment analysis using Spanish opinionated words. In: *Proceedings of NLDB*. pp. 214–219 (2014)
 15. Navas-Loro, M., Rodríguez-Doncel, V., Santana-Perez, I., Sánchez, A.: Spanish Corpus for Sentiment Analysis towards Brands. In: *Proc. of the 19th Int. Conf. on Speech and Computer (SPECOM)*. pp. 680–689 (2017)
 16. Pak, A., Paroubek, P.: Twitter as a corpus for sentiment analysis and opinion mining. In: *LREc*. vol. 10 (2010)
 17. Plaza-Del-Arco, F.M., Martín-Valdivia, M.T., et al.: COPOS: Corpus of patient opinions in Spanish. Application of sentiment analysis techniques. *Procesamiento de Lenguaje Natural* 57, 83–90 (2016)
 18. Pontiki, M., Galanis, D., Pavlopoulos, J., Papageorgiou, H., Androutsopoulos, I., Manandhar, S.: Semeval-2014 task 4: Aspect based sentiment analysis pp. 27–35 (01 2014)
 19. Pontiki, M., Galanis, D., Papageorgiou, H., Androutsopoulos, I., Manandhar, S., Al-Smadi, M., Al-Ayyoub, M., Zhao, Y., Qin, B., De Clercq, O., Hoste, V., Apidianaki, M., Tannier, X., Loukachevitch, N., Kotelnikov, E., Bel, N., Jiménez-Zafra, S.M., Eryigit, G.: Semeval-2016 task 5: Aspect based sentiment analysis. In: *Proceedings SemEval-2016*. pp. 19–30. ACL, San Diego, California (June 2016)
 20. Pontiki, M., Galanis, D., Papageorgiou, H., Androutsopoulos, I., Manandhar, S., Al-Smadi, M., Al-Ayyoub, M., Zhao, Y., Qin, B., De Clercq, O., et al.: Semeval-2016 task 5: Aspect based sentiment analysis. In: *ProWorkshop SemEval-2016*. pp. 19–30. ACL (2016)
 21. Ramanand, J., Bhavsar, K., Pedanekar, N.: Wishful thinking: finding suggestions and 'buy' wishes from product reviews. *Proceedings of the NAACL HLT 2010 Workshop CAAGET '10* (June), 54–61 (2010)
 22. Rangel, F., Rosso, P., Reyes, A.: Emotions and Irony per Gender in Facebook. In: *Proceedings of Workshop ES3LOD, LREC-2014*. pp. 1–6 (2014)
 23. Sánchez Rada, J.F., Torres, M., et al.: A linked data approach to sentiment and emotion analysis of twitter in the financial domain. In: *FEOSW* (2014)
 24. Van Waterschoot, W., Van den Bulte, C.: The 4p classification of the marketing mix revisited. *The Journal of Marketing* pp. 83–93 (1992)
 25. Vázquez, S., Muñoz-García, O., Campanella, I., Poch, M., Fisas, B., Bel, N., Andreu, G.: A classification of user-generated content into consumer decision journey stages. *Neural Networks* 58(Supplement C), 68–81 (2014), Special Issue on “Affective Neural Networks and Cognitive Learning Systems for Big Data Analysis”
 26. Vineet, G., Devesh, V., Harsh, J., Deepam, K., Shweta, K.: Identifying purchase intent from social posts. *ICWSM 2014* pp. 180–186 (2014)
 27. Westerski, A., Iglesias, C.A., Rico, F.T.: Linked opinions: Describing sentiments on the structured web of data. In: *Proceedings of the 4th International Workshop Social Data on the Web*. vol. 830 (2011)