INTERSPEECH 2023
20-24 August 2023, Dublin, Ireland

Laughter in task-based settings: whom we talk to affects how, when, and how often we laugh

Catarina Branco[1,3], *Isabel Trancoso*[1], *Paulo Infante*[2], *Khiet P. Truong*[3]

[1]INESC ID / IST, Portugal
[2]CIMA, IIFA, Department of Mathematics, ECT, University of Évora, Portugal
[3]University of Twente, The Netherlands

catarinabranco3699@gmail.com, isabel.trancoso@inesc-id.pt, pinfante@uevora.pt, k.p.truong@utwente.nl

Abstract

Map task corpora are not typically used to study laughter, but they allow an interesting analysis of multiple factors such as familiarity between the participants, their gender, and eye contact. We conducted linear/generalized mixed-effects analysis to study if co-laughter, laughter rate, and the percentage of voiced frames in laughs are influenced by such factors. Our results show that, in conversations without eye contact, the gender of the participant was statistically relevant regarding laughter rate and the percentage of voiced frames, and the difference in gender was relevant regarding co-laughter. On the other hand, with eye contact, familiarity was statistically relevant with respect to co-laughter, laughter rate, and the percentage of voiced frames. Most of our results align and extend what has been previously found, except for voiced laughs between friends. This study emphasizes the highly variable character of laughter and its dependence on interlocutors' characteristics.

Index Terms: Laughter, dialogue, gender, familiarity.

1. Introduction

Laughter, which is primarily produced during social interactions between human beings, can be classified as a category of non-linguistic communication [1] that can be used to convey emotions. It has been the subject of many studies which aim to understand how it is produced and perceived in different contexts. There are studies that focus on the impact of laughter in conversations with speakers of the same or opposite gender [1], the level of familiarity between speakers [2], and the gender of the speakers [3]. Other studies focus on shared or antiphonal laughter that refers to instances where one person laughs almost immediately (within 0.5-1 seconds) after another person laughs, investigating its acoustics [4] and the relationship between the speakers, such as their familiarity and gender [5]. The relationship between laughter and the visual contact between speakers has also been studied, along with other visual signals such as body language, which together convey a range of emotions and messages [6]. Other studies have explored how non-verbal vocal cues can affect conversations, particularly when compared to face-to-face conversations [7].

The aim of our research is to understand how various previously mentioned factors, such as gender[1], familiarity, and visual contact can impact specific (acoustic) features of laughter. Previous research on laughter has produced inconsistent and sparse results. Thus our study aims to control for more variables and explore the interactions between these factors, contributing to a more comprehensive understanding of laughter. The laughter features under study include the rate of laughter and co-laughter (here defined as any laugh that starts within the start and end of another person's laugh) per subject and the percentage of co-laughter in the total amount of laughs per speaker. To conduct this study, we used the HCRC Map Task corpus [8] which is characterised by many instances of laughter which are spontaneous (not elicited or acted) and produced in mixed- or same-gender interactions with or without eye contact, and with strangers or familiar speakers.

[1]We acknowledge that gender is a social construct and that more than the two commonly known identities, man and woman, can exist. In the current study though, we chose to work with this binary distinction given that the data used in our analyses adopts this binary distinction.

2. Related work

2.1. Familiarity

Several papers have explored the relationship between familiarity and laughter [2, 5, 9, 1], addressing the frequency and the acoustics of (co-)laughter. In general, speakers who are familiar with each other produce more laughter than those who do not know each other [5]. In [1], gender seems to interact with familiarity. Both males and females produce the same amount of laughter in unfamiliar, same-gender conversations. However, when the speakers are familiar to each other, males produce more laughter regardless of interlocutor gender, whereas females only produce more laughter with familiar male speakers. With regards to joint laughter, previous studies [2, 5] have found that individuals produced more co-laughter or antiphonal laughter (i.e., laughter between interlocutors that occurs in close proximity) in conversations with friends than with strangers.

The acoustics of laughter were also found to be related to familiarity. In [2], it was found that laughter contains proportionally fewer voiced frames when produced with familiar speakers, and when produced by male speakers. Laughter duration was also significantly affected by familiarity in this study: speakers laugh longer when they know each other.

Overall, these studies suggest that there is a clear difference between friends and strangers in terms of laughter, with individuals in general laughing more and producing more co-laughter in conversations with friends. In other studies, the conclusions about familiarity are very dependent on whether the speakers are of the same-gender or not. Nonetheless, these findings support the inclusion of familiarity as a variable in this research.

2.2. Gender

In this section, we will examine the results of studies that take into account both the gender of the individual and the gender of the conversational pair, i.e., whether the dyadic interaction concerns same-gender or mixed-gender speakers. Taking into account only the gender of the individual, previous research has shown that, in general, females tend to laugh more than

10.21437/Interspeech.2023-1914

males [10, 6] and produce voiced laughs more often [3]. In a study on nonverbal cues in phone conversations, it was also found that females laugh more than males [7].

When looking at mixed-gender dyads, it was found that females produce more antiphonal laughter than males [5]. In familiar conditions both males and females laugh more often with laughter having higher pitch [1]. In same-gender dyads, it matters for male speakers (but not for female speakers) whether they are talking to familiar or unfamiliar persons: when talking to familiar persons, male speakers laugh more, and their laughs have a higher pitch and longer duration [1].

All results considered, it can be concluded that whether or not individuals are talking to someone of the same or different gender is relevant to explore. Mainly, results report that females laugh and co-laugh in mixed-gender conversations more often than males. For same-gender conversations, there is no overall difference in how females laugh in situations where they are friends or not. In contrast, males laugh significantly less in same-gender conversations with strangers.

3. Method

3.1. Data

As previously mentioned, the corpus used in this study is the HCRC map task corpus [8]. This corpus consists of dyadic conversations between two speakers in a task-based interaction, with a total of 64 participants. Interlocutors have different maps and have to talk to each other to find a route from A to B. The corpus includes interactions between speakers of different genders (FM: female×male, FF: female × female, MM: male × male), friends and strangers (FAM, the familiar pairs were recruited by asking participants to bring a friend to participate with them), and speakers who can see each other or not (EYECONTACT), see Table 1. Moreover, the starting and ending times of laughter have been annotated. In Table 1, we can observe the amount of laughter separated by the condition of seeing each other or not, gender of the conversational pair, and familiarity.

Table 1: *Number of annotated laughs in the HCRC corpus, separated by familiarity (*FAM+/-*), seeing each other or not (*EYECONTACT+/-*) and same-gender (*FF,MM*) or mixed-gender (*FM*) conversations.*

	EYECONTACT+		EYECONTACT-	
genders	FAM+	FAM-	FAM+	FAM-
FM	12	46	10	150
FF	42	10	210	174
MM	164	43	95	8

3.2. Analysis of Data

In order to find any relations between our proposed dependent and independent variables we decided to perform a linear mixed-effects analysis [11]. This choice was taken considering the unbalanced nature of our corpus, as we can see in Table 1 the distribution of laughter across the different conditions is highly unbalanced, and to account for within-person repeated measures (i.e., the same speaker occurs in several conversations). To ensure the validity of the results obtained through this model, it is necessary to confirm two assumptions: the residuals of the obtained model must have a constant variance and have to be normally distributed.

All linear mixed-effects or generalized linear mixed-effects models were carried out in R [12], with the lme4 [13] and lmerTest [14] libraries. The fixed effects selected were gender (*Gender*: female or male), familiarity (*Familiarity*: friends or strangers), and same-gender vs. mixed-gender interactions (*GenderConv*: same gender or not). The dependent variables selected are related to the frequency and amount of voicing of (co-)laughter. To account for individual differences, we added subject as random effect. To find the best model for each one of the dependent variables, we applied the following procedure:

1: Fit the model with all the independent variables and report the results.
2: Compare different models by removing independent variables with the original model using the likelihood ratio to check whether any variable was worth removing.
3: For the model with only the relevant variables, include the interactions between those variables.
4: Test the significance of the interactions by comparing the model with and without some of the interactions.
5: Report the model with the lower AIC (Akaike Information Criterion) value as the final model.

The best models and the results obtained for each dependent variable in the study are reported in Section 4

3.3. Feature Extraction

In this section, we describe how we extracted the dependent variables we proposed, as well as any adjustments made to ensure that the assumptions required for a linear or generalised mixed model analysis were met. Based on previous works, we selected several laughter variables that seem to be impacted by gender and familiarity.

Regarding **co-laughter**, we initially proposed to extract the rate of co-laughter (the number of co-laughs divided by the duration of the conversation) and the percentage of co-laughs in the total amount of laughs per speaker. However, upon verifying whether these variables were adequate, we discovered that they violated one of the assumptions needed to conduct a linear mixed model analysis. The residuals of the model were not normally distributed. To overcome this situation, we decided to make a simpler variable with a binary outcome (co-laughter was present or not) to understand how our independent variables influence whether people co-laugh or not and used a broader model, the generalized linear mixed-effects model.

The variable **laughter rate** was extracted by dividing all laughs for each speaker by the duration of the conversation in which they participated. This variable fulfilled the assumptions needed to conduct a linear mixed-effects analysis, thus we were able to carry out the procedure described in Section 3.2.

For the **percentage of voiced frames**, initially, we extracted the pitch value for each frame of any annotated laughter event using *Parselmouth*, [15], a Python library for the Praat software [16]. Finally, we divided the number of voiced frames by the total number of frames for each laughter event. When examining the assumptions of the linear mixed model, we found that one of them was not met. The residuals of the model were not normally distributed. To address this issue, we created a new binary variable for each speaker and used a broader model, the generalized linear mixed-effects model. For each laugh of a speaker, the percentage of voiced frames was calculated. If the median of this speaker's distribution of percentage of voiced frames was above zero, the speaker would be assigned the value 1. If the median was zero, the speaker would be assigned the value 0.

4. Results

Following the procedure described in Section 3, and for each one of our dependent variables, we found the model that best fitted the data available.

4.1. Results for conversations where participants do not see each other

In this Section, we will present the results obtained for the data where participants did not see each other.

First, we analysed the relationship between our independent variables and the variable **co-laughter** using a generalised linear mixed effects analysis. The results of the model, which included the interaction between *Gender* and *GenderConv* as the relevant independent variable, are shown in Table 2. The results show that when two females are talking, they are 9 times more likely to co-laugh than if it is a male with another male; and when a female is talking to another female, she is 9 times more likely to co-laugh than if she is talking to a male.

For the variable **laughter rate**, we conducted a linear mixed model analysis, following the procedure discussed in Section 3. The results of the model are presented in Table 2. As *Gender* is the only statistically relevant independent variable ($p < 0.01$), we can conclude that females have a higher laughter rate than males.

Finally, regarding the binary variable related to the **percentage of voiced frames**, we proceeded to carry out a generalised mixed model analysis, where the results are reported in Table 2. Looking at the results, we can see that *Gender* is the only statistically relevant variable ($p < 0.001$), in particular, we can conclude that females are about 6 times more likely to produce voiced laughs than males.

Table 2: *Results of the linear and generalised mixed effects analysis for the data where participants do not see each other*

	β	*z value*	*p value*
Co-laughter			
Intercept	-1.088	-1.691	0.0907
Gendermale	0.238	0.255	0.798
GenderConvsame	2.216	2.96	**0.00308**
GenderConvsame:Gendermale	-2.408	-1.981	**0.0476**
Laughter rate			
Intercept	0.276	13.513	**2.78e-14**
Gendermale	-0.125	-3.613	**0.0011**
Percentage of voiced frames			
Intercept	1.595	4.576	**4.73e-6**
Gendermale	-1.793	-3.541	**0.0004**

4.2. Results for conversations where participants see each other

In this Section we will explore the results obtained for the data where participants saw each other.

Using the same variables and analysis as we did for the data where participants did not see each other, we obtained the results presented in Table 3. For **co-laughter**, *Familiarity* is the only statistically relevant independent variable ($p < 0.05$): friends are about 3 times more likely to co-laugh with each other than strangers.

Regarding the variable **laughter rate**, following a linear mixed model analysis, we found that *Familiarity* was the only statistically relevant independent variable ($p < 0.01$), more specifically, in conversations between friends the laughter rate is higher than in conversations between strangers (Table 3).

Table 3: *Results of the linear and generalised mixed effects analysis for the data where participants see each other*

	β	*z value*	*p value*
Co-laughter			
Intercept	-0.0704	-0.258	0.7960
Familiaritystranger	-1.039	-2.566	**0.0103**
Laughter Rate			
Intercept	0.0699	8.83	**1.6e-12**
Familiaritystranger	-0.028	-3.12	**0.0024**
Percentage of voiced frames			
Intercept	0.656	1.86	0.0629
Familiaritystranger	-0.935	-2.21	**0.0027**

Lastly, for the variable related to the **percentage of voiced frames**, the model obtained showed that the only statistically relevant independent variable was *Familiarity*: individuals in conversations with friends are about 2.5 times more likely to produce voiced laughs than with strangers (Table 3).

5. Discussion

This paper aimed to investigate whether there are changes in how people laugh depending on their relationship, gender, and the gender of the person they are talking to in the HCRC Map task corpus. We found several results regarding laughter rate, co-laughter, and percentage of voiced frames in laughter. This Section will explore how our results compare to previous in the literature, and what can be concluded, including limitations and suggestions for future research.

5.1. Comparison with other results found in past studies

5.1.1. Co-laughter

Initially, regarding the variable of co-laughter, we wanted to address the percentage of co-laughter in total laughs and co-laughter rate. However, these variables failed to meet assumptions for a linear mixed-effects analysis. For this reason, we instead examined whether participants co-laughed with others or not, creating a binary variable, which allowed us to perform a generalised mixed-effects analysis.

For conversations where participants did not see each other, the variable that was statistically relevant regarding co-laughter was the interaction between the gender of the participants and the genders of the conversational pair. We found that females talking with other females co-laugh more often than when talking to males, and also, there is more co-laughter in conversations between two females than between two males. In [5] they found only relevant results in mixed-gender dyads where females showed antiphonal laughter more often than males. For same-gender dyads, it was found that friends displayed antiphonal laughter more often than strangers [9]. However, it is important to consider that our results cannot be directly compared to those found in the literature, as the participants in those experiments were able to see each other, while in our study this is not the case. In addition, our definition of co-laughter differs slightly from that of antiphonal laughter.

When participants do see each other, we find that familiarity was significantly related to co-laughter: when talking to a friend, one produces laughs more often than when talking to a stranger. Although the conclusion taken is more detailed and specific to our study, we can still relate it to other findings. The results presented in [2] and in [5] corroborate our finding, given that they also found significant results regarding familiarity and co-laughter as already explored in Section 2.1. Note that in their experiments, the participants were also able to see each other. Furthermore, in [9] this relationship was also found, but only in conversations between individuals of the same gender, and where the participants were able to see each other.

5.1.2. Laughter rate

When the participants do not see each other, the gender of the participant plays a significant role in the laughter rate, ($p < 0.01$) where females laugh more often than males. This is in line with previous studies [10, 6] and in [6] where it was also found that females laugh more often than males (although in their studies, participants could see each other).

When participants do see each other, familiarity plays a significant role ($p < 0.01$), where participants laughed more often when talking to their friends than strangers. Other results in the literature support our findings: in [5] and in [1], they also found that familiarity had an impact on the laughter rate, with friends laughing more than strangers.

5.1.3. Percentage of voiced frames

Finally, the last variable that we explored in the scope of this study was the percentage of voiced frames of the laughs of the participants. Due to the concerns already discussed in Section 3.3, we changed this variable to have a binary outcome and conducted a generalised mixed-effects analysis.

For the data where participants did not see each other, we found that the gender of the participants was statistically significant regarding the amount of voicing in laughs. Specifically, we found that females were six times more likely to produce voiced laughs than males. In [3] and in [2], it was also found that females produce more voiced laughs than males. However, note that in their studies, participants could see each other.

When the participants could see each other, we found that familiarity was the only significant independent variable for percentage of voiced frames. We found that individuals when talking with friends are 2.5 times more likely to produce voiced laughs than with strangers. These findings contrast with those from [2], which found that in conversations between friends there were fewer voiced laughs produced.

5.1.4. Seeing each other or not

While eye contact was not an independent variable in our study, we did observe some differences between the two contexts that warrant future research. What is striking is that we find when interlocutors do *not* see each other, *Gender* drives the differences found in the amount of (co-)laughter and voiced laughter produced. When interlocutors do see each other, it is *Familiarity* that drives these differences. At this moment, there is no clear explanation for why *Gender* plays a significant role when there is no eyecontact, and why *Familiarity* plays a significant role when there is eyecontact. It has been found previously however for the HCRC Map Task corpus, that interactions between interlocutors who could see each other were shorter in duration and less problematic in completing a task together [17]. In the conversations where the interlocutors could not see each other, speakers used more words, interrupted each other twice as often, and overlapped each other more often.

5.2. Limitations of the current study

One of the limitations to consider are the definitions and annotations of co-laughter and laughter in general. We defined co-laughter in our study as any laugh that starts within the start and end of another person's laugh. However, other studies only consider co-laughs that start within a certain threshold after the first laugh starts or a laugh that occurs right after someone else finishes laughing. Moreover, the definition of laughter across studies is not clear. In order to obtain reliable laughter annotations, clear instructions are needed for the annotators such that ambiguity can be resolved. For many of the speech corpora unfortunately, these instructions are lacking because in most of the cases, the corpora were not designed with the main purpose to study laughter. Upon examination of the audio and annotations in the HCRC Map Task corpus, we found a wide range of different laughs labelled as "laughter" from clear and stereotypical laughs to more subtle ones that were barely audible and very short. When do we consider a sound laughter, and when does it start and end? Although attempts have been made to reach a more common understanding and description of laughter [18], it takes a lot of effort to actually implement these efforts across multiple corpora. One cannot say to what extent this lack of a universal definition of (co-)laughter affects valid comparisons between studies. Thus, it remains a factor to take into account when comparing laughter studies.

5.3. Future work

For future research, we recommend to consider the absence or presence of visual contact in more detail by taking this factor into account in the analysis. As previous research [17], and our current work have pointed out, visual contact matters but more research is needed to understand more precisely how this impacts laughing behaviours. We also see interesting follow-up research opportunities of our work in carrying out similar laughter analyses on other map task corpora recorded in different languages and cultures (e.g., Cantonese [19], French [20]) to enable a cross-cultural comparison.

6. Conclusions

We investigated how gender and familiarity affects how often and the way we laugh. Using the HCRC Map Task corpus, we found that gender plays an important role in laughing behaviour when interlocutors do not see each other. Females laugh more often than males, and produce voiced laughter more often than males. When taking into account the gender of the interlocutor, females co-laugh more often than males when talking to another female. Familiarity on the other hand, plays an important role when interlocutors do see each other. Friends (co-)laugh more often than strangers, and produce voiced laughter more often than strangers. Our work contributes to the notion that laughter is a highly variable social signal that is influenced by whom one is talking to.

7. Acknowledgements

This work was partially supported by Portuguese national funds through Fundação para a Ciência e a Tecnologia (FCT), with reference UIDB/50021/2020, as well as by the Portuguese Recovery and Resilience Plan (RRP) through project C645008882-00000055 (Responsible.AI).

8. References

[1] M. Owren and J.-A. Bachorowski, "Reconsidering the evolution of nonlinguistic communication: The case of laughter," *Journal of Nonverbal Behavior*, vol. 27, pp. 183–200, 09 2003.

[2] M.-P. Jansen, K. P. Truong, and D. K. J. Heylen, "How familiarity influences the frequency, temporal dynamics and acoustics of laughter," in *2021 9th International Conference on Affective Computing and Intelligent Interaction (ACII)*. IEEE, 2021, pp. 1–8.

[3] J.-A. Bachorowski, M. J. Smoski, and M. J. Owren, "The acoustic features of human laughter," *The Journal of the Acoustical Society of America*, vol. 110, no. 3, pp. 1581–1597, 2001.

[4] K. P. Truong and J. Trouvain, "On the acoustics of overlapping laughter in conversational speech," in *Proceedings of Interspeech*, 2012, pp. 851–854.

[5] M. Smoski and J.-A. Bachorowski, "Antiphonal laughter between friends and strangers," *Cognition and Emotion*, vol. 17, no. 2, pp. 327–340, 2003.

[6] K. Grammer, "Strangers meet: Laughter and nonverbal signs of interest in opposite-sex encounters," *Journal of Nonverbal Behavior*, vol. 14, no. 4, pp. 209–236, 1990.

[7] A. Vinciarelli, P. Chatziioannou, and A. Esposito, "When the words are not everything: The use of laughter, fillers, back-channel, silence, and overlapping speech in phone calls," *Frontiers in ICT*, vol. 2, 2015. [Online]. Available: https://www.frontiersin.org/article/10.3389/fict.2015.00004

[8] A. H. Anderson, M. Bader, E. G. Bard, E. Boyle, G. Doherty, S. Garrod, S. Isard, J. Kowtko, J. McAllister, J. Miller *et al.*, "The hcrc map task corpus," *Language and speech*, vol. 34, no. 4, pp. 351–366, 1991.

[9] M. J. Smoski and J.-A. Bachorowski, "Antiphonal laughter in developing friendships," *Annals of the New York Academy of Sciences*, vol. 1000, no. 1, pp. 300–303, 2003.

[10] R. R. Provine, "Laughter punctuates speech: Linguistic, social and gender contexts of laughter," *Ethology*, vol. 95, no. 4, pp. 291–298, 1993.

[11] J. C. Pinheiro and D. M. Bates, "Linear mixed-effects models: basic concepts and examples," *Mixed-effects models in S and S-Plus*, pp. 3–56, 2000.

[12] R Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, 2021. [Online]. Available: https://www.R-project.org/

[13] D. Bates, M. Mächler, B. Bolker, and S. Walker, "Fitting linear mixed-effects models using lme4," *Journal of Statistical Software*, vol. 67, no. 1, pp. 1–48, 2015.

[14] A. Kuznetsova, P. B. Brockhoff, and R. H. B. Christensen, "lmerTest package: Tests in linear mixed effects models," *Journal of Statistical Software*, vol. 82, no. 13, pp. 1–26, 2017.

[15] Y. Jadoul, B. Thompson, and B. de Boer, "Introducing parselmouth: A python interface to praat," *Journal of Phonetics*, vol. 71, pp. 1–15, 2018. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0095447017301389

[16] P. Boersma, "Praat, a system for doing phonetics by computer," *Glot. Int.*, vol. 5, no. 9, pp. 341–345, 2001.

[17] E. A. Boyle, A. H. Anderson, and A. Newlands, "The effects of visibility on dialogue and performance in a cooperative problem solving task," *Language and speech*, vol. 37, no. 1, pp. 1–20, 1994.

[18] K. P. Truong, J. Trouvain, and M.-P. Jansen, "Towards an annotation scheme for complex laughter in speech corpora," in *Proceedings of Interspeech*, 2019, pp. 529–533.

[19] G. Winterstein, C. Tang, and R. Lai, "Cantomap: a Hong Kong Cantonese Maptask corpus," in *Proceedings of the Twelfth Language Resources and Evaluation Conference*, 2020, pp. 2906–2913.

[20] J. Gorisch, C. Astésano, E. G. Bard, B. Bigi, and L. Prévot, "Aix Map Task corpus: The French multimodal corpus of task-oriented dialogue," in *9th International conference on Language Resources and Evaluation*, 2014, pp. 2648–2652.