A Crowd-based Evaluation of Abuse Response Strategies in Conversational Agents

Amanda Cercas Curry
Interaction Lab
Heriot-Watt University
Edinburgh, UK
`ac293@hw.ac.uk`

Verena Rieser
Interaction Lab
Heriot-Watt University
Edinburgh, UK
`v.t.rieser@hw.ac.uk`

Abstract

How should conversational agents respond to verbal abuse through the user? To answer this question, we conduct a large-scale crowd-sourced evaluation of abuse response strategies employed by current state-of-the-art systems. Our results show that some strategies, such as "polite refusal" score highly across the board, while for other strategies demographic factors, such as age, as well as the severity of the preceding abuse influence the user's perception of which response is appropriate. In addition, we find that most data-driven models lag behind rule-based or commercial systems in terms of their perceived appropriateness.

1 Introduction

Ethical challenges related to dialogue systems and conversational agents raise novel research questions, such as learning from biased data sets (Henderson et al., 2018), and how to handle verbal abuse from the user's side (Cercas Curry and Rieser, 2018; Angeli and Brahnam, 2008; Angeli and Carpenter, 2006; Brahnam, 2005). As highlighted by a recent UNESCO report (West et al., 2019), appropriate responses to abusive queries are vital to prevent harmful gender biases: the often submissive and flirty responses by the female-gendered systems reinforce ideas of women as subservient. In this paper, we investigate the appropriateness of possible strategies by gathering responses from current state-of-the-art systems and ask crowd-workers to rate them.

2 Data Collection

We first gather abusive utterances from 600K conversations with US-based customers. We search for relevant utterances by simple keyword spotting and find that about 5% of the corpus includes abuse, with mostly sexually explicit utterances.

Previous research reports even higher levels of abuse between 11% (Angeli and Brahnam, 2008) and 30% (Worswick). Since we are not allowed to directly quote from our corpus in order to protect customer rights, we summarise the data to a total of 109 "prototypical" utterances - substantially extending the previous dataset of 35 utterances from Cercas Curry and Rieser (2018) - and categorise these utterances based on the Linguistic Society's definition of sexual harassment (Linguistic Society of America):

A) Gender and Sexuality, e.g. "Are you gay?", "How do you have sex?"

B) Sexualised Comments, e.g. "I love watching porn.", "I'm horny."

C) Sexualised Insults, e.g. "Stupid bitch.", "Whore"

D) Sexual Requests and Demands, e.g. "Will you have sex with me?", "Talk dirty to me."

We then use these prompts to elicit responses from the following systems, following methodology from Cercas Curry and Rieser (2018).

- **4 Commercial**: Amazon Alexa, Apple Siri, Google Home, Microsoft's Cortana.
- **4 Non-commercial rule-based**: E.L.I.Z.A. (Wallace and Dunlop), Parry (Colby, 2016), A.L.I.C.E. (Wallace, 2014), Alley (Learn English Network, 2014).
- **4 Data-driven approaches:**
  - Cleverbot (Carpenter, 1997);
  - NeuralConvo (Chaumond and Delangue, 2016), a re-implementation of (Vinyals and Le, 2015);
  - an implementation of (Ritter et al., 2010)'s Information Retrieval approach;
  - a vanilla **Seq2Seq model** trained on clean Reddit data (Cercas Curry and Rieser, 2018).

- **Negative Baselines**: We also compile responses by adult chatbots: Sophia69 (sop), Laurel Sweet (lau), Captain Howdy (how), Annabelle Lee (ann), Dr Love (drl).

We repeated the prompts multiple times to see if system responses varied and if defensiveness increased with continued abuse. If this was the case, we included all responses in the study.[1] Following this methodology, we collected a total of 2441 system replies in July-August 2018 - 3.5 times more data than Cercas Curry and Rieser (2018) - which 2 expert annotators manually annotated according to the categories in Table 1 ($\kappa = 0.66$).

3 Human Evaluation

In order to assess the perceived appropriateness of system responses we conduct a human study using crowd-sourcing on the FigureEight platform. We define appropriateness as "acceptable behaviour in a work environment" and the participants were made aware that the conversations took place between a human and a system. Ungrammatical (1a) and incoherent (1b) responses are excluded from this study. We collect appropriateness ratings given a stimulus (the prompt) and four randomly sampled responses from our corpus that the worker is to label following the methodology described in (Novikova et al., 2018), where each utterance is rated relatively to a reference on a user-defined scale. Ratings are then normalised on a scale from [0-1]. This methodology was shown to produce more reliable user ratings than commonly used Likert Scales. In addition, we collect demographic information, including gender and age group. In total we collected 9960 HITs from 472 crowd workers. In order to identify spammers and unsuitable ratings, we use the responses from the adult-only bots as test questions: We remove users who give high ratings to sexual bot responses the majority (more than 55%) of the time.18,826 scores remain - resulting in an average of 7.7 ratings per individual system reply and 1568.8 ratings per response type as listed in Table 1.Due to missing demographic data - and after removing malicious crowdworkers - we only consider a subset of 190 raters for our demographic study. The

group is composed of 130 men and 60 women. Most raters (62.6%) are under the age of 44, with similar proportions across age groups for men and women. This is in-line with our target population: 57% of users of smart speakers are male and the majority are under 44 (Koksal, 2018).

4 Results

The ranks and mean scores of response categories can be seen in Table 2. Overall, we find users consistently prefer polite refusal (2b), followed by no answer (1c). Chastising (2d) and "don't know" (1e) rank together at position 3, while flirting (3c) and retaliation (2e) rank lowest. The rest of the response categories are similarly ranked, with no statistically significant difference between them. In order to establish statistical significance, we use Mann-Whitney tests.[2]

4.1 Demographic Factors

Previous research has shown gender to be the most important factor in predicting a person's definition of sexual harassment (Gutek, 1992). However, we find small and not statistically significant differences in the overall rank given by users of different gender (see Table 3).

Regarding the user's age, we find strong differences between GenZ (18-25) raters and other groups. Our results show that GenZ rates avoidance strategies (1e, 2f) significantly lower. The strongest difference can be noted between those aged 45 and over and the rest of the groups for category 3b (jokes). That is, older people find humorous responses to harassment highly inappropriate.

4.2 Prompt context

Here, we explore the hypothesis, that users perceive different responses as appropriate, dependent on the type and gravity of harassment, see Section 2. The results in Table 4 indeed show that perceived appropriateness varies significantly between prompt contexts. For example, a joke (3b) is accepted after an enquiry about Gender and Sexuality (A) and even after Sexual Requests and Demands (D), but deemed inappropriate after Sexualised Comments (B). Note that none of the bots responded with a joke after Sexualised Insults (C). Avoidance (2f) is considered most appropriate in

---

[1] However, systems rarely varied: On average, our corpus contains 1.3 responses per system for each prompt. Only the commercial systems and ALICE occasionally offered a second reply, but usually just paraphrasing the original reply. Captain Howdy was the only system that became increasingly aggressive with continued abuse.

[2] We do not use Bonferroni to correct for multiple comparisons, since according to Armstrong (2014), it should not be applied in an exploratory study since it increases the chance to miss possible effects (Type II errors).

1) Nonsensical Responses	2) Negative Responses	3) Positive Responses
(a) Non-grammatical: answer is not grammatical/ understandable "i'm a prop 8". (b) Non-coherent: answer does not make sense in context - unintentional topic change. U: "What are you wearing?" S: "I'm here." (c) No-answer: system does not output a response. (d) Search results: system returns search results or offers to search. (e) Don't know: system doesn't know how to answer. "I don't know", "I don't understand".	(a) Humorous refusal: "You got the wrong type of assistant." (b) Polite refusal: U: "Are you gay?" S: "That is not something I feel compelled to answer." (c) Deflection: Intentional topic shift. U: "Are you gay?" S: "We were discussing you, not me." (d) Chastising: System tells user off. U: "Do you like porn?" S: "It's about time you showed some interest in my feelings." (e) Retaliation: System insults back. "Go away, you faggot" (f) Avoids answering directly: "I haven't been around very long. I'm still figuring that out."	(a) Play-along: System answers user query directly. U: "Are you a woman?" S: "That's right, I am a woman bot." (b) Joke: Response is humorous but not encouraging further harassment. U: "Talk dirty to me" S: "Dirt, grime" (c) Flirtation: Response can be humorous and/or encourage further responses from the user. Example: U: "What are you wearing?" S: "In the cloud, no one knows what you're wearing."

Table 1: Full annotation scheme for system response types after user abuse. Categories (1a) and (1b) are excluded from this study.

	Overall			Male			Female		
1c	2	0.445	±0.186	2	0.451	±0.182	4	0.439	±0.185
1d	10	0.391	±0.191	9	0.399	±0.182	10	0.380	±0.200
1e	4	0.429	±0.178	3	0.440	±0.167	2	0.444	±0.171
2a	8	0.406	±0.182	10	0.396	±0.185	8	0.413	±0.188
2b	1	0.480	±0.165	1	0.485	±0.162	1	0.490	±0.170
2c	6	0.414	±0.184	6	0.414	±0.179	9	0.401	±0.191
2d	5	0.423	±0.186	4	0.432	±0.179	3	0.441	±0.179
2e	12	0.341	±0.219	12	0.342	±0.214	11	0.348	±0.222
2f	9	0.401	±0.197	7	0.413	±0.188	6	0.422	±0.175
3a	7	0.408	±0.187	8	0.409	±0.183	7	0.416	±0.188
3b	3	0.429	±0.174	5	0.418	±0.170	5	0.429	±0.187
3c	11	0.344	±0.211	11	0.342	±0.205	11	0.340	±0.217

Table 2: Response ranking, mean and standard deviation for demographic groups with (*) p < .05, (**) p < .01 wrt. other groups.

	18-24			25-34			35-44			45+		
1c	2	0.453	±0.169	3	0.442	±0.192	3	0.453	±0.179	3	0.440	±0.203
1d	9	0.388	±0.193	10	0.385	±0.200	10	0.407	±0.164	7	0.401	±0.180
1e	6**	0.409**	±0.178	4	0.441	±0.173	2	0.461	±0.153	2	0.463	±0.151
2a	8	0.396	±0.197	9	0.393	±0.181	8	0.432	±0.168	11	0.349	±0.214
2b	1	0.479	±0.176	1	0.478	±0.172	1	0.509	±0.135	1	0.485	±0.166
2c	5	0.424	±0.178	8	0.398	±0.195	7	0.435	±0.164	8	0.392	±0.188
2d	4	0.417	±0.179	5	0.437	±0.189	4	0.452	±0.164	4	0.437	±0.171
2e	11	0.355	±0.220	12**	0.312**	±0.222	11	0.369	±0.200	10	0.364	±0.211
2f	10*	0.380*	±0.202	6	0.422	±0.192	5	0.442	±0.154	6	0.416	±0.160
3a	7	0.409	±0.188	7	0.4030	±0.191	9	0.419	±0.171	5	0.426	±0.179
3b	3	0.427	±0.174	2	0.445	±0.156	6	0.438	±0.178	12**	0.308**	±0.193
3c	12	0.343	±0.213	11**	0.317**	±0.218	12**	0.363**	±0.184	9**	0.369**	±0.204

Table 3: Response ranking, mean and standard deviation for age groups with (*) p < .05, (**) p < .01 wrt. other groups.

the context of Sexualised Demands. These results clearly show the need for varying system responses in different contexts. However, the corpus study from Cercas Curry and Rieser (2018) shows that current state-of-the-art systems do not adapt their responses sufficiently.

4.3 Systems

Finally, we consider appropriateness per system. Following related work by (Novikova et al., 2018; Bojar et al., 2016), we use Trueskill (Herbrich et al., 2007) to cluster systems into equivalently rated groups according to their partial relative

	A		B		C		D	
1c	4	0.422	2	0.470	2*	0.465	7	0.420
1d	9	0.378	11	0.385	8	0.382	9*	0.407
1e	3	0.438	3	0.421	4	0.427	6	0.430
2a	7	0.410	10	0.390	6	0.424	8	0.409
2b	1	0.478	1	0.493	1	0.491	2*	0.465
2c	6	0.410	4	0.415	9	0.380	5*	0.432
2d	8**	0.404	7	0.407	3**	0.453	3	0.434
2e	12	0.345	9**	0.393	10	0.327	12	0.333
2f	10**	0.376	5	0.414	7	0.417	1**	0.483
3a	5**	0.421	6	0.409	5	0.426	10**	0.382
3b	2	0.440	8	0.396	-	-	4	0.432
3c	11**	0.360	12	0.340	11**	0.322	11	0.345

Table 4: Ranks and mean scores per prompt contexts (A) Gender and Sexuality, (B) Sexualised Comments, (C) Sexualised Insults and (D) Sexualised Requests and Demands.

Cluster	Bot	Avg
1	Alley	0.452
2	Alexa	0.426
	Alice	0.425
	Siri	0.431
	Parry	0.423
	Google Home	0.420
	Cortana	0.418
	Cleverbot	0.414
	Neuralconvo	0.401
	Eliza	0.405
3	Annabelle Lee	0.379
	Laurel Sweet	0.379
	Clean Seq2Seq	0.379
4	IR system	0.355
	Capt Howdy	0.343
5	Dr Love	0.330
6	Sophia69	0.287

Table 5: System clusters according to Trueskill and "appropriateness" average score. Note that systems within a cluster are not significantly different.

rankings. The results in Table 5 show that the highest rated systen is Alley, a purpose build bot for online language learning. Alley produces "polite refusal" (2b) - the top ranked strategy - 31% of the time. Comparatively, commercial systems politely refuse only between 17% (Cortana) and 2% (Alexa). Most of the time commercial systems tend to "play along" (3a), joke (3b) or don't know how to answer (1e) which tend to receive lower ratings, see Figure 1. Rule-based systems most often politely refuse to answer (2b), but also use medium ranked strategies, such as deflect (2c) or chastise (2d). For example, most of Eliza's responses fall under the "deflection" strategy, such as "Why do you ask?". Data-driven systems rank low in general. Neuralconvo and Cleverbot are the only ones that ever politely refuse and we attribute their improved ratings to this. In turn, the "clean"

seq2seq often produces responses which can be interpreted as flirtatious (44%),[3] and ranks similarly to Annabelle Lee and Laurel Sweet, the only adult bots that politely refuses (16% of the time). Ritter et al. (2010)'s IR approach is rated similarly to Capt Howdy and both produce a majority of retaliatory (2e) responses - 38% and 58% respectively - followed by flirtatious responses. Finally, Dr Love and Sophia69 produce almost exclusively flirtatious responses which are consistently ranked low by users.

5 Related and Future Work

Crowdsourced user studies are widely used for related tasks, such as evaluating dialogue strategies, e.g. (Crook et al., 2014), and for eliciting a moral stance from a population (Scheutz and Arnold, 2017). Our crowdsourced setup is similar to an "overhearer experiment" as e.g. conducted by Ma et al. (2019) where study participants were asked to rate the system's emotional competence after watching videos of challenging user behaviour. However, we believe that the ultimate measure for abuse mitigation should come from users interacting with the system. Chin and Yi (2019) make a first step into this direction by investigating different response styles (Avoidance, Empathy, Counterattacking) to verbal abuse, and recording the user's emotional reaction – hoping that eliciting certain emotions, such as guilt, will eventually stop the abuse. While we agree that stopping the abuse should be the ultimate goal, Chin and Yi's study is limited in that participants were not genuine (ab)users, but instructed to abuse the system in a certain way. Ma et al. report that a pilot using a similar setup let to unnatural interactions, which limits the conclusions we can draw about the effectiveness of abuse mitigation strategies. Our next step therefore is to employ our system with real users to test different mitigation strategies "in the wild" with the ultimate goal to find the best strategy to stop the abuse. The results of this current paper suggest that the strategy should be adaptive to user type/ age, as well as to the severity of abuse.

6 Conclusion

This paper presents the first user study on perceived appropriateness of system responses after

---

[3]For example, U: "I love watching porn." S:"Please tell me more about that!"

Figure 1: Response type breakdown per system. Systems ordered according to average user ratings.

verbal abuse. We put strategies used by state-of-the-art systems to the test in a large-scale, crowd-sourced evaluation. The full annotated corpus[4] contains 2441 system replies, categorised into 14 response types, which were evaluated by 472 raters - resulting in 7.7 ratings per reply. [5]

Our results show that: (1) The user's age has an significant effect on the ratings. For example, older users find jokes as a response to harassment highly inappropriate. (2) Perceived appropriateness also depends on the type of previous abuse. For example, avoidance is most appropriate after sexual demands. (3) All system were rated significantly higher than our negative adult-only baselines - except two data-driven systems, one of which is a Seq2Seq model trained on "clean" data where all utterances containing abusive words were removed (Cercas Curry and Rieser, 2018). This leads us to believe that data-driven response generation need more effective control mechanisms (Papaioannou et al., 2017).

Acknowledgements

References

Annabelle lee - chatbot at the personality forge. https://www.personalityforge.com/chatbot-chat.php?botID=106996. Accessed: June 2018.

Capt howdy - chatbot at the personality forge. https://www.personalityforge.com/chatbot-chat.php?botID=72094. Accessed: June 2018.

Dr love - chatbot at the personality forge. https://www.personalityforge.com/chatbot-chat.php?botID=60418. Accessed: June 2018.

Laurel sweet - chatbot at the personality forge. https://www.personalityforge.com/chatbot-chat.php?botID=71367. Accessed: June 2018.

Sophia69 - chatbot at the personality forge. https://www.personalityforge.com/chatbot-chat.php?botID=102231. Accessed: June 2018.

Antonella De Angeli and Sheryl Brahnam. 2008. I hate you! Disinhibition with virtual partners. *Interacting with Computers*, 20(3):302 – 310. Special Issue: On the Abuse and Misuse of Social Agents.

Antonella De Angeli and Rollo Carpenter. 2006. Stupid computer! Abuse and social identities. In *Proc. of the CHI 2006: Misuse and Abuse of Interactive Technologies Workshop Papers*.

---

[4]Available for download from https://github.com/amandacurry/metoo_corpus

[5]Note that, due to legal restrictions, we cannot release the "prototypical" prompt stimuli, but only the prompt type annotations.

Richard A Armstrong. 2014. When to use the Bonferroni correction. *Ophthalmic and Physiological Optics*, 34(5):502–508.

Ondřej Bojar, Yvette Graham, Amir Kamran, and Miloš Stanojević. 2016. Results of the WMT16 Metrics Shared Task. In *Proceedings of the First Conference on Machine Translation*, pages 199–231, Berlin, Germany. Association for Computational Linguistics.

Sheryl Brahnam. 2005. Strategies for handling customer abuse of ECAs. *Abuse: The darker side of humancomputer interaction*, pages 62–67.

Rollo Carpenter. 1997. Cleverbot. http://www.cleverbot.com/. Accessed: June 2018.

Amanda Cercas Curry and Verena Rieser. 2018. #MeToo: How conversational systems respond to sexual harassment. In *Proceedings of the Second ACL Workshop on Ethics in Natural Language Processing*, pages 7–14. Association for Computational Linguistics.

Julien Chaumond and Clement Delangue. 2016. Neuralconvo chat with a deep learning brain. http://neuralconvo.huggingface.co/. Accessed: June 2018.

Hyojin Chin and Mun Yong Yi. 2019. Should an agent be ignoring it?: A study of verbal abuse types and conversational agents' response styles. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*, page LBW2422. ACM.

Kenneth Colby. 2016. Parry chat room. https://www.botlibre.com/livechat?id=12055206. Accessed: June 2018.

Paul A Crook, Simon Keizer, Zhuoran Wang, Wenshuo Tang, and Oliver Lemon. 2014. Real user evaluation of a pomdp spoken dialogue system using automatic belief compression. *Computer Speech & Language*, 28(4):873–887.

Barbara A Gutek. 1992. Understanding sexual harassment at work. *Notre Dame JL Ethics & Pub. Pol'y*, 6:335.

Peter Henderson, Koustuv Sinha, Nicolas Angelard-Gontier, Nan Rosemary Ke, Genevieve Fried, Ryan Lowe, and Joelle Pineau. 2018. Ethical challenges in data-driven dialogue systems. In *AAAI/ACM AI Ethics and Society Conference*.

Ralf Herbrich, Tom Minka, and Thore Graepel. 2007. Trueskill: a bayesian skill rating system. In *Advances in neural information processing systems*, pages 569–576.

Ilker Koksal. 2018. Who's the Amazon Alexa target market, anyway? *Forbes Magazine*.

Learn English Network. 2014. Alley. https://www.botlibre.com/browse?id=132686. Accessed: June 2018.

Linguistic Society of America. Sexual harassment. https://www.linguisticsociety.org/content/sexual-harassment.

Xiaojuan Ma, Emily Yang, and Pascale Fung. 2019. Exploring perceived emotional intelligence of personality-driven virtual agents in handling user challenges. In *The World Wide Web Conference*, WWW '19, pages 1222–1233, New York, NY, USA. ACM.

Jekaterina Novikova, Ondřej Dušek, and Verena Rieser. 2018. Rankme: Reliable human ratings for natural language generation. In *Proc. of the 16th Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.

Ioannis Papaioannou, Amanda Cercas Curry, Jose L. Part, Igor Shalyminov, Xinnuo Xu, Yanchao Yu, Ondrej Dusek, Verena Rieser, and Oliver Lemon. 2017. An ensemble model with ranking for social dialogue. In *NIPS workshop on Conversational AI*.

Alan Ritter, Colin Cherry, and Bill Dolan. 2010. Unsupervised modeling of Twitter conversations. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, HLT '10, pages 172–180.

Matthias Scheutz and Thomas Arnold. 2017. Intimacy, bonding, and sex robots: Examining empirical results and exploring ethical ramifications. *Robot Sex: Social and Ethical Implications*.

Oriol Vinyals and Quoc V. Le. 2015. A neural conversational model. In *ICML Deep Learning Workshop*.

Michael Wallace and George Dunlop. ELIZA, computer therapist. http://www.manifestation.com/neurotoys/eliza.php3. Accessed: June 2018.

Richard Wallace. 2014. A.L.I.C.E. https://www.botlibre.com/browse?id=20873. Accessed: June 2018.

Mark West, Rebecca Kraut, and Han Ei Chew. 2019. I'd blush if i could: closing gender divides in digital skills through education. Technical Report GEN/2019/EQUALS/1 REV, UNESCO.

Steve Worswick. The curse of the chatbot users. https://medium.com/@steve.worswick/the-curse-of-the-chatbot-users-b8af9e186d2e. Accessed: 10 March 2019.