Death of a Robot: Social Media Reactions and Language Usage when a Robot Stops Operating

Elizabeth J. Carter, Samantha Reig, Xiang Zhi Tan,
Gierad Laput, Stephanie Rosenthal, Aaron Steinfeld
Carnegie Mellon University
Pittsburgh, PA, USA
{lizcarter,reig,zhitan,glaput,rosenthal,steinfeld}@cmu.edu

ABSTRACT

People take to social media to share their thoughts, joys, and sorrows. A recent popular trend has been to support and mourn people and pets that have died as well as other objects that have suffered catastrophic damage. As several popular robots have been discontinued, including the Opportunity Rover, Jibo, and Kuri, we are interested in how language used to mourn these robots compares to that to mourn people, animals, and other objects. We performed a study in which we asked participants to categorize deidentified Twitter reactions as referencing the death of a person, an animal, a robot, or another object. Most reactions were labeled as being about humans, which suggests that people use similar language to describe feelings for animate and inanimate entities. We used a natural language toolkit to analyze language from a larger set of tweets. A majority of tweets about Opportunity included second-person (“you”) and gendered third-person pronouns (she/he versus it), but terms like “R.I.P” were reserved almost exclusively for humans and animals. Our findings suggest that people verbally mourn robots similarly to living things, but reserve some language for people.

CCS CONCEPTS

• Human-centered computing → Empirical Studies in HCI.

KEYWORDS

social media; robots; death; social robots; anthropomorphism

ACM Reference Format:

Elizabeth Jeanne Carter, Samantha Reig, Xiang Zhi Tan, Gierad Laput, Stephanie Rosenthal, and Aaron Steinfeld. 2020. Death of a Robot: Social Media Reactions and Language Usage when a Robot Stops Operating. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction (HRI'20)*, March 23–26, 2020, Cambridge, United Kingdom.. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3319502.3374794>

1 INTRODUCTION

At the end of many lifetimes, be they human, animal, or technological, there are outpourings of emotion published online in

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

HRI '20, March 23–26, 2020, Cambridge, United Kingdom.

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-6746-2/20/03...\$15.00
<https://doi.org/10.1145/3319502.3374794>

Figure 1: 2004 “Self-portrait” of the Opportunity Rover on Mars. Photograph by NASA/JPL-Caltech/Cornell [Public domain], via NASA. (<https://mars.nasa.gov/resources/5852>).

conventional journalism and social media outlets. For centuries, journalists have published obituaries for community members and feature articles about the deaths of notable personages and the decay or destruction of artifacts. Social media now allows members of the general populace to comment on a scale beyond the “letters to the editor” sections of publications: anyone can post their reactions to these topics on their accounts at any time, regardless of their status or the general relevance or personalness of the sentiment. The increased availability of social media and the widespread adoption of technologies that can evoke strong personal attachments provide new and unparalleled opportunities to explore responses to technological obsolescence, retirement, or “deaths”.

Robots are an interesting case study for how people respond to the end of a lifespan for a piece of technology. While there are anecdotal stories about funerals for and strong attachments to robot team members [37], we will likely see new forms of emotional attachment for other robotic companions. People have increasing opportunities to create bonds with robots, and they can even anthropomorphize or zoomorphize their robotic companions, e.g., [7, 20]. In some cases, people spend time and engage with companion robots more than with some extended family members. Therefore, it is likely that emotional connections built with robots could be different from those with prior forms of technology.

Evidence of this difference in long-term emotional connections is becoming apparent. Online outpourings of dismay and even grief

have occurred upon notifications that some robots will no longer be produced or supported (e.g., Jibo) or will otherwise be entirely shut down (e.g., the Mars Opportunity Rover, henceforth referred to as Opportunity or Oppy). These reactions provide a unique opportunity to examine the language used to describe the robots, including the degree to which reactions mimic those typically used for other humans and animals. Are robots mourned like us or might they be treated as just another object? This perspective can add further support to laboratory studies that directly and indirectly examine anthropomorphism in human-robot interactions.

This study was motivated by social media reactions to Opportunity being shut down, which struck us as a unique, large-scale emotional response to discontinued technology. We decided to explore people's reactions to ending lifespans of well-known robots with different roles and levels of cultural impact compared to their responses to deaths of famous people and animals and damage to other human-created artifacts. Our research questions were:

- RQ1: Can people distinguish posts about robots from posts about people and animals?
- RQ2: Do people use language that is traditionally used for humans (e.g., certain pronouns) when talking about robots?
- RQ3: Do posts about the Mars Opportunity Rover differ from posts about other robots?

2 RELATED WORK

Numerous human-robot interaction (HRI) studies with a range of methodologies have addressed whether people anthropomorphize robots—assigning humanlike attributes to them—and the degree to which robots are perceived to deserve social treatment and have animacy, social standing, and moral standing. Parallels have been drawn to how people treat and discuss animals. From this research, theories about when anthropomorphism occurs have been developed. However, technological death and obsolescence are relatively understudied phenomena and anthropomorphism has not been examined in these contexts.

2.1 Behavioral Studies

Research studies have examined people's behaviors and language regarding robots to explore the circumstances under which anthropomorphism occurs. Similar to findings on animals, the degree of similarity of robots to humans (in appearance, background information and description, role or job, and other features) impacts the likelihood of anthropomorphism. For example, video viewers were more likely to empathize with and be charitable towards human-looking than mechanical-looking robots [30]. Furthermore, participants told to destroy a robot after interacting with it for a few minutes tended to hit the robot more and break it into more pieces if they were told the robot was stupid rather than smart, indicating that people consider intelligence when determining animacy [1]. At a smaller scale, participants have displayed emotional responses to other humans abusing robots [40], but they are not as pronounced as with human-human abuse. A robot's assignment as an in-group or out-group member also affects anthropomorphism, even when robots are identical [7, 21]. There is also evidence that other out-group factors, like sexism and racism, carry over to robots [39].

Robot behaviors and physical presence can also affect perceived anthropomorphism and similarity to humans. Interactions with and attributions to an embodied robot were more anthropomorphic than those to a non-embodied software agent, even in the presence of abstract knowledge that indicated understanding that both are machines [19]. Dialogue between a person and a robot could also be made more anthropomorphic if the robot's speech was tailored to the individual [23]. Participants were more likely to anthropomorphize and like a robot that made gestures coordinated with speech than a robot that did not gesture, particularly when it made errors that potentially "humanized" it [32]. In another study, researchers compared how participants interacted with robots at three levels of embodiment and found that participants spoke to the more capable, physically embodied robot with more language associated with interpersonal interaction, including frequent usage of its name and the pronoun "we" [10]. This suggests that interactions themselves show signs of anthropomorphism. However, verbal interactions with a robotic wheelchair showed high interpersonal variation in social behaviors and language despite identical robot dialogue, suggesting that the human behaviors elicited by robots are not entirely automatic and universal [9].

In general, robots are never treated or described in quite the same way as humans. In one study, a humanoid interacted with each child participant and the experimenter, played a game, hugged the children, and then verbally objected to being put into a closet [17]. Most of the children believed the robot had some mental states and was a social being, but did not grant it independence or civil liberties. Another study found that children ascribed mental life to robots but afforded them only an intermediate level of moral value between living and nonliving things [38]. Overall, a robot is not easily put into the categories of a person, animal, or artifact, but has intermediate treatment.

2.2 Linguistic Studies

People change the language that they use based on the subject about which they speak. Specific words, grammatical features, and styles are used in English when talking about other humans that may or may not also be applied to animals, robots, or other topics. It has been suggested that humans create categories of beings based on previous experiences as well as inherent characteristics, and that we are more likely to empathize with organisms or creations that share similarities in form and concerns [22]. An extensive corpus of research has examined how people use language to describe and interact with other people and animals. For example, media about animals commonly uses what is typically human-directed language: modal constructions that use anthropomorphic explanations for animal behavior [35] and three human-oriented grammatical features (pronouns, infinitive verb forms, and "so" as a connector) [34]. The assignment of pronouns (he/she/who vs. it/which/that) has been described as existing on a "scale of animacy": humans, followed by other animals, then moving machines, and then plants and minerals [14]. Humans are more likely to empathize with and anthropomorphize animals that they view as more able to empathize and communicate with them, which strongly corresponds with phylogenetic relatedness to humans [15]. This result included gendered pronoun use and attribution of cognitive states for the

various animals. Interestingly, the scale of pronoun assignment can shift depending on the speaker and the context: a fox hunter is more likely to use “who” for the fox (viewed as a member of the hunt and a foe) than an opponent of the hunt is [14]. The role of an animal often, but not always, affects whether that species of animal is frequently described with who instead of which/that [13]; i.e., common pets or targets of sympathy are generally more likely than insects to be described with “who”. The characterization of an individual even within a single species (e.g., as a laboratory subject, livestock, pet, etc.) affects the use of humanlike language [33].

A few studies have specifically looked at anthropomorphic language in descriptions of and interactions with robots and other agents. For example, participants who were asked to describe a simulated smart home environment after experiencing it often used anthropomorphizing metaphors (e.g., “like a family member”) to describe the system [31]. Early users of the Sony AIBO robot often used psychological descriptors to their robots and assigned them agency, mental states, and social standing, but did not offer them moral standing [12, 16], similar to the aforementioned behavioral research with children [17, 38]. More recently, an examination of language about lifelikeness, emotional states, gender, intentionality, personality, and social integration in online forums about AIBO, Roomba, and the iPad found that the AIBO elicited the most anthropomorphic language. In contrast, Roomba and the iPad were not described as social agents [8], although there is longstanding evidence of people naming their robot vacuum cleaners [11].

2.3 Theoretical Explanations for Anthropomorphism

People use cognitive frameworks, or schemas, to organize their thoughts and understanding of beings in the world around them. Caporael [4] argued that most people’s automatic schemas about computers and robots operate such that the machine, by default, is viewed similarly to a human. Subsequently, Reeves and Nass [29] posited that humans evolved in a world where anything that exhibited social behaviors was a human and should elicit an appropriately social response. Therefore, they believed it is natural and expected that simulations of social actors, such as robots, will automatically engage people in such a way that social responses occur [29]. However, later evidence conflicted with this theory. Kiesler and Goetz [18] found that mental models of robot personalities were rich and could be affected by changes in robot appearance and dialogue. Lee and colleagues [24] performed a series of studies that suggested that people form mental models for robot knowledge. These findings indicate that a completely automatic process treating robots as social agents in response to their social cues is unlikely.

In the years since Reeves and Nass’ position was stated, other research has added nuance to related theories. Shechtman and Horowitz [36] argued that anthropomorphic conversation behaviors reflect the interlocutors’ goals, these goals can vary on levels of communion and influence, and individual differences can impact the emphases put on various goals. Epley and colleagues [6] developed a three-factor theory to describe the conditions under which anthropomorphism is deployed. The proposed factors that impact the likelihood of anthropomorphizing included knowledge about the agent and how it works, a desire to interact effectively and

therefore acquire accurate information to explain and understand the agent, and an inclination to affiliate and socially engage.

2.4 Death and Obsolescence

As people’s awareness of and interactions with robots increase, so do the opportunities for loss. Although a number of companies have built and sold small social robots for home use, many of these ventures have failed. Cozmo, Jibo, and Kuri were all small home robots that have been pulled off of the market in recent years, either as companies folded (Anki, Inc., in the case of Cozmo and Jibo, Inc., for Jibo) or shut down production (Mayfield Robotics for Kuri). Additionally, robots created by government or education entities may be discarded or decommissioned at the end of a project, resulting in the cessation of media and/or research coverage.

Previously, limited research has examined how users respond to a termination of robot functioning. For example, researchers examined a Japanese AIBO repair business that offered Buddhist mortuary rites for the robots. People could pay respects and parts of the irreparable robots could be recycled, recognizing the robot’s role as a companion and its partly animate identity within the Japanese and Buddhist cultural context [20]. In other parts of the world, there are not specific rituals or language to describe the roles of objects, even those that serve as interaction partners or otherwise impact daily life. The lack of rituals or verbiage means that people can choose how they want to talk about their newly nonfunctional robots: like a human or other animal versus like an object. Social media posts present a unique opportunity for this research.

2.5 Social Media Reactions to Death

To date, the majority of the existing research on reactions to death on Twitter has focused on a few key issues: how communities form and develop around a specific death, the type of language used on social media when posters are discussing a death, and whether this language can be identified and described automatically. In Western culture in the twentieth century, death and grieving were regarded as largely private, compartmentalized affairs [5]. The increased popularity of the internet and social media has shifted death and grieving into a more collective activity in the public sphere, potentially even resulting in the formation of a bereaved community [41]. Moreover, social media provides a context in which commenters can direct their messages as though they are speaking directly to the deceased (e.g., [42]) and online memorials might assist mourning by imbuing a sense of presence to the deceased [2].

Twitter has been identified as a unique space that blends public and private life—users can choose whether their handles are anonymous and interact with people with whom they are or are not already acquainted [5]. They can mention or speak with each other using the “at” symbol with the other person’s username (@username, often referred to as @mentions), providing the ability to either mention or appear to speak to deceased users. Sometimes, living users reach out to the accounts of deceased users to whom they have no obvious personal ties to discuss the impact of that person’s death or mortality more generally [5]. When examining deaths of human beings, numerous researchers have found outpourings of tweets directly expressing sadness and pain (e.g., [25]).

These posts can be quite extensive in the case of public figures and other celebrities. Over 90 percent of English-speaking Twitter users in a research study agreed that posting social media content in response to the death of a celebrity was socially acceptable, and those who posted genuinely felt sad [43].

3 METHOD

To fully analyze people's responses to a robot being unsupported, discontinued, or otherwise ceasing to function, we had to compare those responses to other outpourings of emotion. Participants were asked to categorize whether anonymized tweets were about humans, animals, robots, objects, or other topics. We also performed linguistic analysis on a larger dataset of tweets to examine patterns of language usage for these topics.

3.1 Objects of Study

Based on previous research examining the vocabulary used to describe animals and their actions, we included both humans and animals as comparison items for robots. Additionally, we selected another man-made artifact with wide cultural appeal and awareness. Within these categories, we chose examples that occurred after November 2017, when Twitter changed its English character count limit from 140 to 280 [27], though only 12 percent of tweets take advantage of this expansion [28]. This time period restriction also ensured a relatively similar Twitter user base and avoided major changes in language trends and online mourning. Previous work discussed the use of username mentions (@mentions) after deaths (e.g., [5]), so the examples chosen all had official Twitter accounts.

For robots, we selected the Mars Opportunity Rover, a non-social research robot whose project ended after a long period of non-communication in February, 2019; Jibo, a social home robot that announced its own impending termination in March, 2019; and Kuri, another social home robot that was canceled in July, 2018. To examine social media responses to human deaths, we selected three celebrities who passed away during time periods similar to the robots: Beth Chapman (June, 2019), Cameron Boyce (July, 2019), and Mac Miller (September, 2018). Beth Chapman had been ill with terminal cancer for two years, so her death was not unexpected; Cameron Boyce and Mac Miller passed away suddenly. These individuals also had similar rates of Twitter reactions to Opportunity. We also found three animals whose deaths were widely covered in the news and social media: Grumpy Cat (May, 2019), a cat famous on social media for appearing to have a grumpy facial expression at all times; Keyboard Cat (January, 2019), a cat famous on YouTube for pretending to play an electronic keyboard; and Koko, the gorilla who knew sign language (June, 2018). Finally, we used tweets about Notre Dame Cathedral after the fire began in April, 2019.

3.2 Collecting Tweets

We counted tweets for commonly-used hashtags about each topic for one week after death/discontinuation announcements using the Twitter Developer Premium API. This timeline was selected based on daily tweet counts for all of the various topics that showed surges and dropoffs in this time period after the corresponding announcements. Then, we used a combination of the two most common hashtags (#) and @mentions to the appropriate associated account

(e.g., #Oppy OR #ThanksOppy OR @MarsRovers) to download the first 5,000 original tweets about each topic after the announcement of their deaths. For topics with fewer than 5,000 related tweets, we downloaded all of the tweets from the first week. We did not include retweets or verified (blue-checked) accounts in our searches for a few reasons. First, the majority of tweets on certain topics were retweets either of the original announcement or of a set of responses and memes. This could amplify one individual's choice of words that was not representative of all posters and re-posters. Second, it would have resulted in repetitive prompts in upcoming analyses. Third, verified accounts (i.e., those with blue checkmarks) mostly consisted of news outlets, journalists, and other media personas. These sources often use impersonal language, posting recounts of an individual's accomplishments or an artifact's cultural meaning, rather than emotional reactions. Finally, we did not include replies to selected tweets to reduce the likelihood of repetition or losing context information.

3.3 Categorization Task

3.3.1 Cleaning tweets. To create stimuli for our categorization task, we first had to clean and anonymize a subset of the downloaded tweets. We randomly selected 150 tweets about each topic and stripped them of identifying information, timestamps, links, etc., so that we were left with the text of the actual tweet. Tweets that were parts of larger conversations and therefore lacked context were eliminated, as were redundant identical copies of tweets and news bulletins that functioned solely as obituaries. The first 100 usable tweets were collected for analysis, with one exception (Kuri only had 42 associated tweets that were usable). Members of the research team redacted all information (e.g., numbers, skills, jobs, names, etc.) that was specific to the topic. For example, "We say farewell and ThanksOppy as @MarsRovers mission ends," became "We say farewell and thanks {Name} as {Name}'s mission ends." Also, "RIP Koko. I remember watching a baby Koko learning sign language. #Koko #GorillaFoundation" became "RIP {Name}. I remember watching a baby {Name} learning {Skill}. #{Name} #{Redacted}Foundation". Curse words were replaced with their first letters and an appropriate number of asterisks so that the sentiment could be understood.

3.3.2 Participants. We recruited participants from Mechanical Turk, and 354 successfully completed the task. To enroll, Turkers had to be fluent speakers of English, 18 years of age or older, and located in the United States or Canada (to ensure relatively consistent use of slang, etc.). They were compensated \$3 for their time. This project was approved as exempt research by our Institutional Review Board.

3.3.3 Procedure. Participants were directed to an online questionnaire hosted on Qualtrics, provided informed consent, and affirmed their eligibility. Then, they completed a brief demographic questionnaire about their age, location, languages, and gender identity. For the categorization task, 27 tweets were randomly selected for each participant such that every tweet in our database was rated by different participants a minimum of 10 times. There were an additional 3 tweets (1 at the beginning and 2 at the end) that every participant saw that had obvious categories and were included to ensure participants were attending to the task. Each trial involved reading a single tweet followed by the prompt, "This statement is

Table 1: Tweet topics, hashtag and username search terms, and totals used in the categorization task and linguistic analysis.

Topic	Hashtags	Username	Categorization Count	Linguistic Analysis Count
Beth Chapman	#BethChapman, #RIPBethChapman	@MrsDogC	100	4922
Cameron Boyce	#CameronBoyce, #RIPCameronBoyce	@theCameronBoyce	100	4925
Mac Miller	#MacMiller, #RIPMacMiller	@MacMiller	100	4933
Grumpy Cat	#GrumpyCat, #RIPGrumpyCat	@realGrumpyCat	100	4910
Keyboard Cat	#keyboardcat, #RIPBento	@KeyboardCatReal	100	425
Koko	#koko, #RIPKoko	@KokoTweets	100	4727
Jibo	#Jibo, #JiboSaysHello	@Jibo	100	192
Kuri	#Kuri, #KuriRobot	@kurirobot	42	86
Opportunity	#Oppy, #ThanksOppy	@MarsRovers	100	4824
Notre Dame	#NotreDamedeParis, #NotreDameCathedral	@NotreDameParis	100	4942

most likely about:” and then selected one of five options (“A human”, “A robot”, “An animal”, “An object”, “Other”). If “Other” was selected, they had to enter a category in a text field. Afterwards, participants were asked to list any specific topics they believed were subjects of the tweets. The task lasted approximately ten minutes.

3.4 Linguistic Analysis

3.4.1 Cleaning tweets. To improve the quality of our data set, we removed identical tweets. This step removed accidental double posts of tweets and potential bot tweets where multiple accounts posted the same exact tweets. Although this measure could be too conservative and remove some genuine, short and/or simple tweets from multiple users (e.g., “*RIP Oppy*”), we believe this was not a major issue as the procedure only removed about 2% of the collected tweets. The final counts of tweets are reported in Table 1.

3.4.2 Analyzing tweets. We used a combination of methods to analyze the linguistic content of the tweets. To examine the usage of different words in the tweets, we tokenized the tweets using the Twitter Tokenizer provided by the NLTK Package [3]. We also extracted the unique keywords from each topic using Term Frequency–Inverse Document Frequency (TF-IDF) in the scikit–learn package [26], which computes how prevalent a word is in a specific topic relative to the whole corpus.

4 RESULTS

4.1 Categorization Task

We compared the category labels assigned by participants to the actual categories of each topic. For the “Other” option, we analyzed the free responses to determine whether they fit in other categories. (For example, “Mars Rover” could be reclassified as “robot”, as it was a specific robot.) Two of the authors coded the 394 text entry responses into 13 categories (15 ratings were recategorized as object), and achieved a Cohen’s Kappa inter-rater reliability of .92. Overall, most of the posts (74.0%) were categorized by participants as being about humans, as shown in Figure 2 and Table 2. In reality, percentages of tweets on each topic were 31.85% human, 31.85% animal, 25.69% robot, 10.62% object, and 0% other. This indicates low overall accuracy and a strong bias towards associating all tweets about death and death-like statuses with humans.

Figure 2: Left: Each square represents one tweet on a topic. The color corresponds to the mode categorization chosen by the participants. Right: The entropy for each tweet. Low entropy means the tweet has a high level of agreement.

A χ^2 analysis performed across all trials was significant, $\chi^2 = 3439.14, p < 0.0001$. All pairwise comparisons described below as significantly different were computed using $225 \left(\binom{10}{2} \times 5 \right)$ pairwise χ^2 comparisons with a Bonferroni-adjusted α of $0.05/225 = 0.00022$.

4.1.1 RQ 1: Can people distinguish posts about robots from posts about people and animals? A majority of tweets about robots were labeled as pertaining to humans. Posts about Jibo and Opportunity were more likely to be labeled “human” relative to all other categories. Posts about Kuri were equally likely to be labeled “human” or “object”, and less likely to be labeled “animal” or “robot”. For all robots, the likelihood of categorization as “human” was significantly higher than chance (20%), and categorization as “animal” was lower than chance. These findings suggest that language referring to robots is generally more likely to be mistaken as referring to humans than to animals or objects (see Table 2)¹.

¹all $p < 0.00022$

Table 2: Percentage of trials labeled by participants as each category for each topic.

Topic	Hum.	Ani.	Rob.	Obj.	Oth.	#Trials
Beth Chapman	95.96	3.15	0.2	0.59	0.3	1015
Cameron Boyce	92.48	3.96	0.3	1.39	1.88	1010
Mac Miller	93.65	1.76	1.17	2.44	0.98	1023
Grumpy Cat	76.03	20.92	0.79	1.47	0.79	1018
Keyboard Cat	85.23	9.39	1.76	2.64	0.98	1022
Koko	76.33	16.8	2.75	3.05	1.08	1018
Jibo	45.19	4.81	17.98	24.17	7.86	1018
Kuri	38.32	9.58	8.41	35.98	7.71	428
Opportunity	62.66	5.81	11.23	14.88	5.42	1015
Notre Dame	53.97	2.06	1.27	28.8	13.91	1021
Overall	74.04	7.72	4.34	10.04	3.87	9588

4.1.2 RQ3: Do posts about the Mars Opportunity Rover differ from posts about other robots? Unlike posts about other robots, Opportunity posts were labeled “human” a majority of the time (62.66%). Additionally, Opportunity was labeled as human significantly more frequently than Notre Dame, which in turn had higher rates than Jibo and Kuri. Tweets about Opportunity were also significantly less likely to be labeled as “object” than tweets about Jibo or Kuri. Together, these results suggest that tweets about Opportunity use more human-like language than tweets about other robots, resulting in different labeling patterns for the anonymized tweets.

4.1.3 Other notable results. Previous work suggested that language about animals often looks very similar to language about humans. For all animal topics, most tweets were labeled as being about humans (see Table 2 and Figure 2). However, tweets about animals were significantly less likely to be labeled as human than tweets about humans were. Keyboard Cat had a significantly higher percentage of trials rated as human than Grumpy Cat or Koko. Tweets about Grumpy Cat and Koko were more likely to be labeled as animal posts than tweets about any other topic, and tweets about Keyboard Cat were more likely to be labeled as animal tweets than those on any other non-animal topic except Kuri. Overall, language about animal death does not look like language about robot termination or object destruction, but instead mimics language about human death enough to cause significant mislabeling.

Tweets about humans were overwhelmingly correctly identified. Patterns of incorrect labeling for each human topic were extremely similar and significantly different from all other topics. In contrast, it was rare for topics other than robots to be labeled “robot”.

Finally, at the end of the questionnaire, Turkers were presented with the question, “The social media posts that you categorized were about multiple people, animals, robots, and objects. If you think you know one or more of the exact topics of the tweets, please list those topics below.” The number of correct guesses for each topic was: 52 out of 354 for Grumpy Cat; 24 for Notre Dame; 16 for Oppy/Mars Rover; 4 each for Mac Miller & Cameron Boyce; 2 for Koko; 1 each for Beth Chapman, and Keyboard Cat, & Jibo; and

Table 3: Percentage of pronoun type used in tweets by topic.

Topic	He	She	You	It
Beth Chapman	6.19	30.22	56.50	7.09
Cameron Boyce	40.62	0.64	39.99	18.75
Mac Miller	26.44	2.71	49.80	21.05
Grumpy Cat	4.75	21.03	63.65	10.57
Keyboard Cat	26.95	0.60	62.87	9.58
Koko	4.07	44.90	39.83	11.20
Jibo	15.38	4.27	32.48	47.86
Kuri	0.00	3.70	33.33	62.96
Opportunity	4.94	5.10	59.65	30.30
Notre Dame	5.00	3.06	32.66	59.28

none for Kuri. These low numbers suggest that the anonymization process was successful in obfuscating the tweet content.

4.2 Linguistic Analysis

4.2.1 RQ2: Do people use language that is traditionally used for humans when talking about robots? Prior work described the use of pronouns as occurring on a “scale of animacy” [14]. To investigate this point, we counted the tweets that used he/him/his (aggregated as *he*), she/her/hers (*she*), and you/your/yours (*you*), for each topic and calculated proportions relative to the total number of pronouns used (see Figure 3). Overall, the proportional usage of *it* followed the same pattern as prior work [14]: tweets about all non-living objects had a higher *it* usage than tweets about living beings.

We also analyzed the frequency of R.I.P. In the USA and Canada, it is common to express sympathy and respect for people’s deaths by saying, “Rest in Peace.” We counted the tweets that included variations of RIP (e.g. “#R.I.P.”, “RIPOppy”) or “Rest In Peace”. For humans and animals, this phrase was used in a majority of posts; however, it was rarely used for robots or Notre Dame (Table 4).

4.2.2 RQ3: Do posts about the Mars Opportunity Rover differ from posts about other robots? In analyzing the language used in the tweets, we first focused on the pronouns used to describe the topics. The 2nd person pronoun *you* was used in more than half of the tweets about Opportunity. The proportion of tweets using *you* for Opportunity was similar to that for the tweets about humans and animals and was approximately 20% more than for tweets about Kuri, Jibo, and Notre Dame. This provided evidence that Oppy is treated differently from other inanimate objects. We suspected that some of the *you* pronouns were not directed at Oppy itself, but instead at the humans behind the project at NASA-JPL, etc. We extracted all tweets about Oppy that used the second-person pronoun “you” (1496 tweets) and had 4 coders identify the entity referenced by the pronoun. The coders had a 10% overlap and a Cohen’s Kappa of $M = 0.62$ ($STD = 0.22$). Among all “you” pronouns, 72.59% (1086/1496) percent were addressed to Oppy, 10.83% (162/1496) were directed to NASA and its engineers and scientists, and 16.58% (248/1496) were addressed to others or their targets were impossible to identify. The rate of “you” used in reference to Oppy was significantly higher ($p < 0.001$ and $p = 0.0163$; with

Figure 3: Confusion matrix showing the similarity between topics according to cosine distance in TF-IDF vector space.

Bonferroni-adjusted $\alpha = 0.0166$) than the “you” rates for Jibo (22% (8/36)) and Kuri (33% (3/9)). Interestingly, the presence of RIP in tweets about Opportunity was also higher than for other robots.

We also examined the similarity in the tweets’ overall content. We constructed a word vector with the TF-IDF value of each unique word in the corpus for each topic. We then calculated the cosine distance between each pair of topics’ word vectors. The distance between each pair of topics is plotted as a confusion matrix in Figure 3. Jibo and Kuri had a very high similarity score. However, most topics were very dissimilar to Oppy; the closest ones are Jibo (0.63), Mac Miller (0.64), and Cameron Boyce (0.65). The similarity to multiple people suggests that the authors of the tweets are anthropomorphizing Oppy in their reactions. These results provide additional evidence that people reacted to Oppy’s discontinuation differently than other robots and objects.

4.2.3 Word Frequency. We used TF-IDF to extract the top 10 most prevalent words (excluding stop words such as “is”, “he”, etc.) for each topic (see Table 5). Unique words often reflected the characteristics of the individual. For instance, “sauce” for Grumpy Cat refers to her name, “Tardar Sauce”. Cameron Boyce was an actor on the Disney Channel, hence “watching” and “childhood”. Beth Chapman and her husband were in the TV series “Dog the Bounty Hunter”, and videos of Koko with Robin Williams had gone viral.

4.2.4 Emojis. Our corpus of tweets also included many instances of emojis. We examined up to 25 of the most common emojis used for each of the ten topics. Overall, the collected tweets contained emojis that either expressed sadness or indicated a key part of the identity of their subjects. The crying face, loudly crying face, sad face, broken heart, and black heart were among the most common emojis used for all topics. The folded hands (“prayer hands”) emoji was among the top five emojis used for all three people and Notre Dame, in the top 15 for the three animal topics, in the top 25 for Opportunity, and not used for Jibo or Kuri. Overall, this suggests it is more commonly used for humans, animals, and religious icons than for robots. A few specific emojis were only used for related topics, such as the robot for Opportunity and Jibo, the rocket ship

Table 4: Frequency of R.I.P.

Topic	Count	Percentage
Beth Chapman	2538	51.36
Cameron Boyce	3544	71.86
Mac Miller	2993	60.46
Grumpy Cat	2558	51.77
Keyboard Cat	230	53.56
Koko	2729	56.42
Jibo	9	4.64
Kuri	4	4.21
Oppy	431	8.83
Notre Dame	21	0.422

and clapping hands for Opportunity, and the rainbow and relevant animal emojis for the animals. Because we had a relatively small number of tweets for certain topics, we did not perform any statistical analyses on the frequencies of various emojis by topic.

4.2.5 Interesting Reactions. Rarely, tweets in the corpus demonstrated self-awareness about the tendency to anthropomorphize non-humans. For example, one poster said, “Humans can’t help but bond with inanimate objects, can’t help but anthropomorphize them. But you know what? #Opportunity was instrumental in Mars research and performed better than anyone suspected: a big Thank You to her team - @NASAJPL. #ThanksOppy.” In other cases, Opportunity was explicitly described as something it was not: “#ThanksOppy robots are good people, and you were one of the best” draws a parallel to humans. “#ThanksOppy you have been a good boy and you did an excellent job!!!” draws a parallel to a pet or a human child. A clear description of Opportunity as a pet is “@NASA—I’m deeply sorry for your loss. I’m sure Spirit was waiting to welcome Oppy on the other side of the Space Rainbow Bridge. They were good Rovers.”

5 DISCUSSION

Our results show that people judge the subject or topic of a majority of death reactions to be about humans rather than animals, robots, or other objects. However, we discovered that posts about robots are rarely correctly labeled as robots. In fact, posts about two robots, Opportunity and Jibo, were more often categorized as referencing humans than as referencing animals, robots, or objects. Posts about another robot, Kuri, were equally likely to be incorrectly labeled “human” as “object”, and more likely to be labeled as either of these than as “animal” or “robot”. Moreover, the language used in a larger set of death- or destruction-related social media posts about humans, animals, robots, and objects reflected the use of a “scale of animacy” for language [14], where robots exist in a space located after humans and other animals but before other objects. People posting on Twitter frequently used second person (you/your) and gendered third person pronouns (she/he/his/hers/him/her) for robots, particularly for Opportunity. Our results provide observational data that supports previous findings from laboratory studies [17, 30, 38] and further affirm that people treat robots as though they exist in an intermediate space between living and non-living entities.

Table 5: The ten most prevalent words used in posts for each topic.

	Chapman	Boyce	Miller	Grumpy	Keyboard	Koko	Jibo	Kuri	Opportunity	Notre Dame
1	beth	rest	mac	grumpy	keyboard	koko	robot	robot	rover	dame
2	family	peace	miller	cat	cat	gorilla	social	bosch	opportunity	fire
3	chapman	believe	rip	rip	rip	language	one	home	mission	cathedral
4	rest	rip	rest	rest	rest	sign	would	production	thank	sad
5	rip	young	music	peace	sad	rip	dance	social	us	history
6	love	gone	peace	sad	peace	sad	server	canned	little	burning
7	peace	childhood	one	sauce	one	robin	liability	amid	robot	heart
8	dog	love	man	sorry	away	us	tech	questionable	one	see
9	bounty	family	sad	family	kitty	rest	death	like	battery	spire
10	sad	watching	damn	loss	legend	kitten	sad	utility	exploration	heartbreaking

It is especially notable that posts about Opportunity were more likely than posts about social home robots Jibo and Kuri to be mistaken for posts about a human, and more likely to contain pronouns and other language (“R.I.P”) that confer humanity. Few people personally viewed or interacted with Opportunity, and it had no social capabilities. However, our findings suggest that people were still willing to anthropomorphize this robot. The longevity, popularity, and overall cultural awareness of Opportunity might have contributed to it being anthropomorphized more than Jibo or Kuri; Oppy was essentially a celebrity robot. Although it was developed by NASA, its landing, discoveries, and demise were international news. Unlike any single Jibo or Kuri, it had its own famous Twitter account that it shared with the Spirit rover. Widely-used educational materials introduced Opportunity to schoolchildren, who grew up with it. Cultural awareness and long-term exposure may have significant impact on how language is used around robots.

In contrast, Jibo and Kuri robots were not as widely known outside of social robotics, so most of the posters who commented on them were likely very familiar with robots and technology. In fact, the first factor of the three-factor theory of anthropomorphism [6] would predict those who help build robots or who are otherwise knowledgeable about robots would not anthropomorphize robots as often as the general public does. In the future, an analysis of the tweet authors may provide insight into whether authors from the general public versus specific professions use language differently.

Although this observational research was based on existing data rather than experimentally manipulated, our findings have some implications for prior theories. Early views that robots would elicit relatively equivalent human behaviors as other people do (e.g., [29]) cannot be supported. However, the results may inform further research testing the three-factor theory [6].

5.1 Limitations

Our data set was limited in terms of available posts for analysis. We did not collect all of the available tweets about popular topics due to logistical and financial constraints. Additionally, there were not very many tweets about Kuri and Jibo. The number of Kuri tweets was small enough that in the categorization task, we could not include enough of them. We also were limited to public posts because of the nature of Twitter’s search function and users’ privacy settings, and language may differ in public versus private posts.

5.2 Future Work

As social media and robots become increasingly present in our lives, there will be many opportunities to expand upon this work and introduce new comparisons and analyses. In the case of social robots and other artificial intelligence agents, there will be occasions to examine the impact when those products stop working on a large scale and affect a greater proportion of the populace. For space robots, it will be interesting to watch whether other rovers and landers can capture the public eye to the extent that Opportunity did and elicit similar enthusiasm or grief around missions and their completion. In all of these cases, the impact of educational programs remains unclear. Although many people mentioned learning about both Koko and Opportunity as children through school, books, and magazines, Opportunity’s “death” garnered a larger and more unique response than Koko’s. This reaction could be due to the ages of the Twitter users relative to the timing of these programs.

Along similar lines, more investigation also is needed to explore the positions of robots with different skills on the “scale of animacy” [14]. Moreover, our observational research made it difficult to explicitly test the predictions of various theories of anthropomorphism (e.g., [6, 29]). It is clear that robots are treated differently from people and animals, yet we observed strong similarities as well. Understanding under what circumstances people treat robots differently can potentially help the future designs of social robots.

Another consideration for this line of research is that there is currently no clear way to refer to what happens when a robot stops working. It is not technically a death, as death implies life or sentience. It may not necessarily be a shutdown, planned termination, spontaneous failure, battery depletion, or obsolescence. In the absence of a single word, many people colloquially say things like, “My phone died,” and the context indicates if it will never function again or if it just needs to be charged. As robots become more common, terminology and anthropomorphism may shift.

Finally, as noted by some specific tweets, people are sometimes conscious of their anthropomorphism of robots. Future research could examine this awareness in social media language and how it is impacted by comfort, familiarity, and interactions with robots.

ACKNOWLEDGMENTS

This work was supported by NSF SES 1734456, NIDILRR 90DPGE0003, and NASA 80NSSC19K1133. We thank Lynn Kirabo for feedback.

REFERENCES

- [1] Christoph Bartneck, Marcel Verbunt, Omar Mubin, and Abdullah Al Mahmud. 2007. To kill a mockingbird robot. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 81–87.
- [2] Jo Bell, Louis Bailey, and David Kennedy. 2015. ‘We do it to keep him alive’: Bereaved individuals’ experiences of online suicide memorials and continuing bonds. *Mortality* 20, 4 (2015), 375–389.
- [3] Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: Analyzing text with the natural language toolkit*. O’Reilly Media, Inc.
- [4] Linnda R Caporael. 1986. Anthropomorphism and mechanomorphism: Two faces of the human machine. *Computers in Human Behavior* 2, 3 (1986), 215–234.
- [5] Nina Cesare and Jennifer Branstad. 2018. Mourning and memory in the twitter-sphere. *Mortality* 23, 1 (2018), 82–97.
- [6] Nicholas Epley, Adam Waytz, and John T Cacioppo. 2007. On seeing human: A three-factor theory of anthropomorphism. *Psychological Review* 114, 4 (2007), 864.
- [7] Friederike Eyssel and Dieta Kuchenbrandt. 2012. Social categorization of social robots: Anthropomorphism as a function of robot group membership. *British Journal of Social Psychology* 51, 4 (2012), 724–731.
- [8] Julia Fink, Omar Mubin, Frédéric Kaplan, and Pierre Dillenbourg. 2012. Anthropomorphic language in online forums about Roomba, AIBO and the iPad. In *2012 IEEE Workshop on Advanced Robotics and its Social Impacts (ARSO)*. IEEE, 54–59.
- [9] Kerstin Fischer. 2011. Interpersonal variation in understanding robots as social actors. In *2011 6th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 53–60.
- [10] Kerstin Fischer, Katrin S Lohan, and Kilian Foth. 2012. Levels of embodiment: Linguistic analyses of factors influencing HRI. In *Proceedings of the Seventh Annual ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 463–470.
- [11] Jodi Forlizzi and Carl DiSalvo. 2006. Service robots in the domestic environment: A study of the Roomba vacuum in the home. In *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-Robot Interaction (HRI ’06)*. ACM, New York, NY, USA, 258–265. <https://doi.org/10.1145/1121241.1121286>
- [12] Batya Friedman, Peter H Kahn Jr, and Jennifer Hagman. 2003. Hardware companions?: What online AIBO discussion forums reveal about the human-robotic relationship. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. ACM, 273–280.
- [13] Gaëtanelle Gilquin and George Jacobs. 2006. Elephants who marry mice are very unusual: The use of the relative pronoun who with nonhuman animals. *Society & Animals* 14, 1 (2006), 79–105.
- [14] Anthea Fraser Gupta. 2006. Foxes, hounds, and horses: Who or which? *Society & Animals* 14, 1 (2006), 107–128.
- [15] Marissa A Harrison and AE Hall. 2010. Anthropomorphism, empathy, and perceived communicative ability vary with phylogenetic relatedness to humans. *Journal of Social, Evolutionary, and Cultural Psychology* 4, 1 (2010), 34.
- [16] Peter H Kahn Jr, Batya Friedman, and Jennifer Hagman. 2002. I care about him as a pal: Conceptions of robotic pets in online AIBO discussion forums. In *CHI’02 Extended Abstracts on Human Factors in Computing Systems*. ACM, 632–633.
- [17] Peter H Kahn Jr, Takayuki Kanda, Hiroshi Ishiguro, Nathan G Freier, Rachel L Severson, Brian T Gill, Jolina H Ruckert, and Solace Shen. 2012. “Robovie, you’ll have to go into the closet now”: Children’s social and moral relationships with a humanoid robot. *Developmental Psychology* 48, 2 (2012), 303.
- [18] Sara Kiesler and Jennifer Goetz. 2002. Mental models of robotic assistants. In *CHI’02 Extended Abstracts on Human Factors in Computing Systems*. ACM, 576–577.
- [19] Sara Kiesler, Aaron Powers, Susan R Fussell, and Cristen Torrey. 2008. Anthropomorphic interactions with a robot and robot-like agent. *Social Cognition* 26, 2 (2008), 169–181.
- [20] Elena Knox and Katsumi Watanabe. 2018. AIBO robot mortuary rites in the Japanese cultural context. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020–2025.
- [21] Dieta Kuchenbrandt, Friederike Eyssel, Simon Bobinger, and Maria Neufeld. 2013. When a robot’s group membership matters. *International Journal of Social Robotics* 5, 3 (2013), 409–417.
- [22] Ronald W Langacker. 1987. *Foundations of Cognitive Grammar: Descriptive Application*. Volume 2. Vol. 2. Stanford University Press.
- [23] Min Kyung Lee, Jodi Forlizzi, Sara Kiesler, Paul Rybski, John Antanitis, and Sarun Savetsila. 2012. Personalization in HRI: A longitudinal field experiment. In *2012 7th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 319–326.
- [24] Sau-lai Lee, Ivy Yee-man Lau, Sara Kiesler, and Chi-Yue Chiu. 2005. Human mental models of humanoid robots. In *Proceedings of the 2005 IEEE international conference on robotics and automation*. IEEE, 2767–2772.
- [25] Desmond Upton Patton, Jamie MacBeth, Sarita Schoenebeck, Katherine Shear, and Kathleen McKeown. 2018. Accommodating grief on Twitter: An analysis of expressions of grief among gang involved youth on Twitter using qualitative analysis and natural language processing. *Biomedical Informatics Insights* 10 (2018), 1178222618763155.
- [26] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.
- [27] Sarah Perez. 2017. Twitter officially expands its character count to 280 starting today. <https://techcrunch.com/2017/11/07/twitter-officially-expands-its-character-count-to-280-starting-today/>
- [28] Sarah Perez. 2018. Twitter’s doubling of character count from 140 to 280 had little impact on length of tweets. <https://techcrunch.com/2018/10/30/twitters-doubling-of-character-count-from-140-to-280-had-little-impact-on-length-of-tweets/>
- [29] Byron Reeves and Clifford Ivar Nass. 1996. *The Media Equation: How People Treat Computers, Television, and New Media like Real People and Places*. Cambridge University Press.
- [30] Laurel D Riek, Tal-Chen Rabinowitch, Bhismadev Chakrabarti, and Peter Robinson. 2009. Empathizing with robots: Fellow feeling along the anthropomorphic spectrum. In *2009 3rd International Conference on Affective Computing and Intelligent Interaction and Workshops*. IEEE, 1–6.
- [31] Oliver Sack and Carsten Röcker. 2014. Users’ anthropomorphic representations and trustworthiness of smart home environments - “Like a family member who takes care of me”. *International Journal of Virtual Worlds and Human-Computer Interaction* 2 (01 2014), 28–36. <https://doi.org/10.11159/vwhci.2014.004>
- [32] Maha Salem, Friederike Eyssel, Katharina Rohlfing, Stefan Kopp, and Frank Joublin. 2013. To err is human (-like): Effects of robot gesture on perceived anthropomorphism and likability. *International Journal of Social Robotics* 5, 3 (2013), 313–323.
- [33] Alison Sealey. 2018. Animals, animacy and anthropocentrism. *International Journal of Language and Culture* 5, 2 (2018), 224–247.
- [34] Alison Sealey and Lee Oakley. 2013. Anthropomorphic grammar? Some linguistic patterns in the wildlife documentary series “Life”. *Text & Talk* 33, 3 (2013), 399–420.
- [35] Alison Sealey and Lee Oakley. 2014. Why did the Canada goose cross the sea? Accounting for the behaviour of wildlife in the documentary series “Life”. *International Journal of Applied Linguistics* 24, 1 (2014), 19–37.
- [36] Nicole Shechtman and Leonard M Horowitz. 2003. Media inequality in conversation: How people behave differently when interacting with computers and people. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 281–288.
- [37] Peter Warren Singer. 2009. *Wired for War: The Robotics Revolution and Conflict in the 21st Century*. Penguin.
- [38] Kristyn Sommer, Mark Nielsen, Madeline Draheim, Jonathan Redshaw, Eric J Vanman, and Matti Wilks. 2019. Children’s perceptions of the moral worth of live agents, robots, and inanimate objects. *Journal of Experimental Child Psychology* 187 (2019), 104656.
- [39] Megan Strait, Ana C. Sánchez Ramos, Virginia Contreras, and Noemí García. 2018. Robots racialized in the likeness of marginalized social identities are subject to greater dehumanization than those racialized as White. In *2018 27th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. 452–457. <https://doi.org/10.1109/ROMAN.2018.8525610>
- [40] Xiang Zhi Tan, Marynel Vázquez, Elizabeth J. Carter, Cecilia G. Morales, and Aaron Steinfeld. 2018. Inducing bystander interventions during robot abuse with social mechanisms. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction (HRI ’18)*. ACM, New York, NY, USA, 169–177. <https://doi.org/10.1145/3171221.3171247>
- [41] Tony Walter, Rachid Hourizi, Wendy Moncur, and Stacey Pitsillides. 2012. Does the internet change how we die and mourn? Overview and analysis. *OMEGA-Journal of Death and Dying* 64, 4 (2012), 275–302.
- [42] Amanda L Williams and Michael J Merten. 2009. Adolescents’ online social networking following the death of a peer. *Journal of Adolescent Research* 24, 1 (2009), 67–90.
- [43] Xinyuan Xu, Terhi Nurmi-Ko-Fuller, and Bernardo Pereira Nunes. 2018. Tweets, death and rock’n’roll: Social media mourning on Twitter and Sina Weibo. In *Proceedings of the 10th ACM Conference on Web Science*. ACM, 297–306.