Beyond Dyadic Interactions: Considering Chatbots as Community Members

Joseph Seering

Carnegie Mellon University
Pittsburgh, Pennsylvania, USA
jseering@cs.cmu.edu

Geoff Kaufman

Carnegie Mellon University
Pittsburgh, Pennsylvania, USA
gfk@cs.cmu.edu

Michal Luria

Carnegie Mellon University
Pittsburgh, Pennsylvania, USA
mluria@cs.cmu.edu

Jessica Hammer

Carnegie Mellon University
Pittsburgh, Pennsylvania, USA
hammerj@cs.cmu.edu

ABSTRACT

Chatbots have grown as a space for research and development in recent years due both to the realization of their commercial potential and to advancements in language processing that have facilitated more natural conversations. However, nearly all chatbots to date have been designed for dyadic, one-on-one communication with users. In this paper we present a comprehensive review of research on chatbots supplemented by a review of commercial and independent chatbots. We argue that chatbots' social roles and conversational capabilities beyond dyadic interactions have been underexplored, and that expansion into this design space could support richer social interactions in online communities and help address the longstanding challenges of maintaining, moderating, and growing these communities. In order to identify opportunities beyond dyadic interactions, we used research-through-design methods to generate more than 400 concepts for new social chatbots, and we present seven categories that emerged from analysis of these ideas.

CCS CONCEPTS

• **Human-centered computing** → **HCI theory, concepts and models**; *Collaborative and social computing devices*;

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.
CHI 2019, May 4–9, 2019, Glasgow, Scotland UK

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-5970-2/19/05...\$15.00

<https://doi.org/10.1145/3290605.3300680>

KEYWORDS

Chatbots; social identity; online communities; dyadic communication

ACM Reference Format:

Joseph Seering, Michal Luria, Geoff Kaufman, and Jessica Hammer. 2019. Beyond Dyadic Interactions: Considering Chatbots as Community Members. In *CHI Conference on Human Factors in Computing Systems Proceedings (CHI 2019)*, May 4–9, 2019, Glasgow, Scotland UK. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3290605.3300680>

1 INTRODUCTION

From ELIZA to Tay, for half a century chatbots have mixed provision of services with an attempt to emulate human conversational style. Chatbots are becoming commonplace in many domains, from customer service to political organizations. Many provide simple services in response to requests, with human speech patterns added in for readability only. However, some chatbots are designed to imitate more complex human behavior, or even to deceive users into thinking they are human. Despite this expansion into more complex domains of behavior, modern chatbots still follow many of the same conversational paradigms as their ancestors.

Current chatbots are designed primarily for chat-oriented and/or task-oriented roles [29]. When chatbots are task-oriented, they respond to users' commands or requests by providing information or support in return. When they are chat-oriented, chatbots engage users to enjoy the integration of their robotic capabilities with their almost-humanlike speech patterns for communication. Although hundreds of platforms have been designed to promote and support group and community interaction between people, ranging from social media to forums to gaming platforms, we find in the review we present here that chatbots rarely support or engage in group or multiparty interaction.

In contrast to the social roles occupied by chatbots, *human* interaction online encompasses both a wide variety of

dyadic and interpersonal behaviors as well as complex group and community-based interactions [26]. Some of the aspects that characterize complex group interaction online include multiparty interactions, turn-taking, role-taking, timing, and construction [54]. Although the use of chatbots for services can be useful, and development of chatbots' communication skills is important, we suggest that there is potential for development of chatbots that leverage their combined social and computational abilities to make an impact on online groups and meaningfully interact within communities. We argue that this is both an important and underexplored design space. Recent work has shown that chatbots have considerable influence in online spaces [1, 45]; Depending on their design, they vary from helpful and important to the well-being of users [21] to harmful to an entire community [44].

Given the newness of the space and the broad variety of challenges, it is not immediately clear how a community-based chatbot would act, what it would do, or what role it would play in the social fabric of the community. In order to begin unpacking how chatbots might address complicated aspects of social interaction online and how might they positively influence groups, we attempt to explore the design space of chatbots in multiparty interactions. We do this in three phases: (1) We conduct a systematic review of the research on chatbot creation and the current space of commercially and independently-developed chatbots; (2) We use research-through-design to learn about this unexplored design space [64]; and (3) We apply framework from social computing research to analyze and make sense of our design work results.

In the literature review, we look at both research literature and chatbots “in the wild” to establish what currently exists in the space of social chatbots. By classifying papers according to the style of chatbot social engagement presented, we find that nearly 90% of chatbot literature focuses on dyadic chatbots. We conducted a similar review of both commercially and independently-developed chatbots, finding similar results. This issue has been addressed primarily only in very recent work [5], with prior literature based on an implicit presumption of chatbots as dyadic conversationalists.

We next conduct an exploration of the design space for group and community-based chatbots using research-through-design methods. We argue that this method is the most appropriate for the goal of this work, as there is no clear answer on what the role of community chatbots should be and how should they be designed. Theoretical literature suggests many ways in which humans interact in online communities, but we cannot yet say for sure whether chatbots in this space should emulate human social behaviors or whether they should approach socializing from a new perspective.

We map the design space of chatbots in potential group interactions in several steps: We use several ideation methods to generate as many ideas as possible for multiparty-based interaction with chatbots; we use affinity diagramming to extract the themes that emerge from the vast number of ideas; and we discuss the final integrated categories and present three considerations for future development of chatbots targeted at multiparty interaction.

In order to inform our work, we build on some of the existing research on the challenges of building and maintaining successful online communities [26]. These challenges include recruiting new users [23], effectively socializing with them, managing misbehavior, [46] and long-term user retention [6]. If chatbots are to contribute as members of communities, these challenges are useful starting points to consider the specific types of contributions they might make. Recent social computing work studying human intra-community bonds has also drawn from Ren et al.'s framework for common identity and common bonds [41], which is similarly applicable to the exploration of bonds between humans and bots in spaces where social interplay between humans and bots is significant (e.g., Wikipedia [12, 13]).

Though there are many possible uses for social chatbots, one particular area within this space where chatbots might be useful is in teaching newcomers norms in online communities. Prior work has found that dealing with newcomers is a significant moderation challenge [48], with this work usually done by moderators after newcomers have (intentionally or accidentally) broken the rules. Social chatbots could assist with this through modeling appropriate behaviors, engaging with newcomers, or behaving in another thought-provoking way.

The next section lays out background on the use and development of chatbots, followed by an academic literature review, and a review of chatbots “in the wild”. We then describe the design process conducted, and the resulted set of seven possible categories for future community-based chatbots that emerged. Finally, we conclude the paper with an attempt to address broader questions about the roles of social chatbots using the insights of this work.

2 A BRIEF HISTORY OF CHATBOT RESEARCH

Performativity in modern chatbots draws from the early tradition of research in artificial intelligence. ELIZA, one of the first chatbots, was created in the mid-1960s as a demonstration of the simplicity of certain human interactions, but participants found it engaging and attributed both intelligence and personality to it [60]. Subsequent chatbots building on this concept have captured public attention repeatedly, often in context of media discussions of the Turing Test [58].

Research on chatbots experienced a resurgence in the late 1990s and early 2000s, as the social internet grew in

scale and influence, and as language processing techniques evolved. For example, Paoillo [37] noted the use of chatbots on Internet-Relay-Chat channels for the purpose of automating certain moderation tasks and sharing information. The development of the Artificial Intelligence Markup Language (AIML) by Richard Wallace and the Alicebot free software community from 1995–2000 facilitated the creation of a wide variety of chatbots, notably including A.L.I.C.E., a one-on-one conversational chatbot designed to compete for the annual Turing-Test-Based Loebner Prize [57, 58].

Following the visible but imperfect conversational achievements of A.L.I.C.E., much research focused on making chatbots “smarter”, with separate but sometimes overlapping “chat-oriented” and “task-oriented” approaches to creating bots [29]. Here, “chat-oriented” refers to chatbots that are designed primarily to converse with users, often as a way of demonstrating advances in underlying technology, while “task-oriented” refers to chatbots that retrieve information or provide services. In early work, Levin, Pieraccini, and Eckert [27] proposed language learning based on the Markov decision process, an approach they refined in subsequent work [28]. A variety of work has made extensive use of models that statistically learn rules for dialogue generation as well as reinforcement learning, including [11], [36], and [40]. In-depth statistically-based approaches to improving conversation beyond the above notable developments are outside the scope of this work.

A complementary line of work, which we draw on more in this paper, has explored the personal and social characteristics of virtual agents, that make them more engaging, easier to use, or more trustworthy. Some of this work has focused on avatar-based conversational agents. For example, Muralidharan, de Visser, and Parasuraman examine the impact of agents’ pitch contour and time-delay on trust [35]. Thompson, Trafton, and McKnight examine how manipulation of movement parameters in walking avatars impact ratings of their humanness, familiarity, and eeriness [52]. Scholars looking at personality of avatars and chatbots found that their personality can also influence perceptions of them. For example, people preferred to interact with an interface with positive personality and a touch of negativity [9], and “sincere” avatars were perceived as more trustworthy than “excited” ones [22]. Other influential factors include dialect [10] and even typeface [4].

3 CLASSIFYING CHATBOTS IN RECENT LITERATURE

In order to understand the current map of chatbots in social contexts, we review literature that has generated, analyzed,

or created concepts for chatbots. Specifically, we are interested to explore dyadic vs. multiparty interactions with chatbots, as well as their intended social roles. Note that much of this literature has been published in the past five years.

We conduct a systematic literature review following methods used in prior work, such as reviews published at CHI and related venues. In prior reviews in domains of HCI literature [7, 18, 47], researchers have used keyword searches of the ACM Guide to Computing Literature or Google Scholar, and summarized the results into “genres” or categories based on a number of criteria. In our review we follow these examples, starting with a keyword-based search of the ACM Guide to Computing Literature. We searched five keywords - “bot”, “chatbot”, “chat bot”, “chatterbot”, and “chatter bot”. These returned 447, 123, 20, 9, and 0 results respectively. We first removed duplicate results, and then chose to remove posters and workshop proposals from the dataset. Next, we manually reviewed these papers to determine whether the primary contribution was about chatbots (with chatbots limited here to text-based interfaces designed for interaction with people, excluding audio-based agents). We removed roughly 75% of the papers from the “bot” keyword, as they focused on botnets and malicious (non-chat) bots intended to attack or exploit systems (e.g., DDos attacks). We also removed a smaller number of papers that briefly mentioned chatbots as a related area of work or as a possible space for future application of the primary contribution. Following these removals, we were left with 104 papers.

In order to understand the social roles that chatbots take in these papers, we summarized the role of the designed chatbot in a sentence, and classified whether chatbots are framed in each paper as *dyadic*, *broadcasting*, or *multiparty-based*.¹

Dyadic chatbots in research

We define *dyadic* chatbots as chatbots that primarily communicate one-on-one with users. ELIZA and A.L.I.C.E. are classic chat-based examples: they communicate with one user at a time (though there could be multiple one-on-one conversations ongoing through use of multiple instances of the bot), responding to users’ statements and asking questions in return. Of the 104 papers in our dataset, 91 framed chatbots as primarily dyadic. Though a small fraction of the dataset focused on chatting as an end unto itself, most dyadic chatbot literature focused on a specific task and context in which the chatbot would be deployed. For example, many dyadic chatbots serve as assistants or provide services, such as a Telegram-based bot for tracking user nutrition [17], a

¹One paper [44] did not explicitly limit chatbots to any one of these categories

Messenger-based bot for communicating with internet-of-things devices [34], or a Twitter-based bot for responding to users' customer service queries [61].

A second category of literature on dyadic chatbots included user research to identify perceptions, expectations, and desired interactions with dyadic chatbots systems. Zamora [63] grouped common words used to describe desired chatbots into four categories- "High-performing," "Smart," "Seamless," and "Personable," and also found that users expressed various preferences for text vs. audio input for different types of tasks. In subsequent work, Zamora also noted deficiencies in current chatbots' ability to handle the blending of multiple languages ("code-mixing"). Though focused primarily on conversational agents (e.g., Amazon Alexa) rather than chatbots, recent work by Druga et al. [8] analyzed how children interact with these types of conversational agents.

Several papers did not discuss specific bots that had been created or studied, but rather provided additional tools or systems to support chatbots. Examples include platforms for automated testing of chatbots [31, 55], tools for automatically creating chatbots from apps [30] or websites [56], and methods to support their creation [62].

The final subcategory within the literature on dyadic chatbots dealt with frameworks for designing chatbot features. Though these frameworks aim to be comprehensive, their foci remain at least implicitly on dyadic forms of chatbots. For example, Valério et al. [53] examine the features available to designers for conveying the chatbots' capabilities to users for Facebook Messenger bots (a platform that is mostly intended for dyadic interaction). Pereira and Díaz [38] also examined Messenger chatbots to identify four types of features chatbots might possess - "support of a minimal set of common commands", "foresee language variations in both inputs and output", "human-assistance provision" and "timeliness".

Broadcasting chatbots in research

We define *broadcasting* chatbots as those that have the capability to send messages that are seen by many users at once, but are not situated as part of a multiparty conversation or interaction.² We identified 6 papers in our dataset that describe this type of chatbots. These bots tended to be situated in large online communities. For instance, Long et al. [33] study the processes for requesting customized Reddit bots, many of which are designed to post in public threads. This type of communication sometimes comes in the form of regularly-scheduled administrative posts, but might also take the form of humorous responses to keywords or phrases detected in thread text. Other work has examined anti-social

bot behaviors on Twitter, where bots can broadcast to a large audience. Abokhodair, Yoo, and McDonald analyzed a Twitter botnet that was active during the Syrian uprising [1], a phenomenon distinct from cybersecurity research on "bot-nets" due to the bots' significant social engagement with and deception of users.

Addressing similarly malicious bots, Wang, Angarita, and Renna [59] conceptualize a scenario where social media becomes completely overrun by false information spread by bots. A final strain of cybersecurity-adjacent work builds methods to detect whether Twitter accounts are bots or humans based on behavior patterns and the content of their messages [15, 49].

Multiparty-based chatbots in research

We define *multiparty-based* chatbots as those that engage in back-and-forth conversations that involve multiple users. In our dataset, 6 papers used multiparty framings for chatbots. While dyadic chatbots may have many individual conversations with users, and broadcasting chatbots reach many separate users at once, multiparty-based chatbots are recognized as participants in a group conversation. As such, multiparty-based chatbots are most likely to be deployed to platforms with defined groups (e.g., Facebook Groups, Subreddits, Twitch channels, and Slack groups), rather than to networks like Twitter (though the latter did exist in our dataset). Gianvecchio et al. [14] present a classification system to support the detection of potentially malicious chatbots that attempt to deceive users in chatrooms, listing six types of such chatbots - *periodic*, *random*, *responder*, *replay*, *replay-responder*, and *advanced-responder*. These types of chatbots sit in chatrooms and post messages in one of three ways: at intervals, in response to other users' messages, and by imitating other users with the intent of seeming more human. Though these types of chatbots are not intended to meaningfully contribute to the group conversation, their modes of engagement mirror those of chatbots developed in literature and elsewhere, (e.g., [38, 53].)

The clearest example of a chatbot that meaningfully contributes to a multiparty interaction comes in Candello et al.'s recent installation of a group of chatbots having an open, flowing conversation around a table with a user [5]. As they note, standard human group conversation often flows from one person to another without a predetermined order but with cues from each speaker indicating who might speak next. Though Candello et al.'s installation was not formally situated within a broader community of users, it did illustrate the challenges of identifying when it is a chatbot's "turn" to contribute to casual social conversation, without explicit cues provided by keywords or commands. Other work has also explored chatbot participation in multiparty interactions, but with strict pre-determined rules about when

²Note that, as we define them, some broadcasting chatbots may also be designed to have separate dyadic conversations with users.

and what chatbots should say: Avlva et al.’s “SearchBots” [2] presents a Wizard of Oz system where a chatbot joins a discussion between a pair of users to help them find a place to eat; Savage, Monroy-Hernandez, and Höllerer’s “Botivist” messages up to three Twitter users in a group to facilitate collaboration between them. In the latter, the bot’s goal was to promote a discussion between the three rather than lean towards dyadic conversations between each user and the bot. [43]

4 CLASSIFYING “IN THE WILD” CHATBOTS

Given the close ties between chatbot research and chatbots deployed in the wild, we elected to supplement our literature review with a review of active or recently active chatbots outside of academia. As noted above, some recent literature has classified types of features of chatbots [38] or patterns of deceptive communication [53], but no work has performed a review of social chatbot engagement across multiple platforms.

To establish a baseline for the current state of chatbot design, we aggregated a list of 130 chatbots from a wide variety of domains. In order to qualify as a chatbot for this list, a system was required to be able to engage in conversation with one or more users at a time via text; responses could be pre-determined or learned, but the chatbots were required to react to the content of users’ typed messages in some way. We did not document every single chatbot we found—for example, there are perhaps tens of thousands of Twitter bots when various botnets are included (e.g., [1]). We therefore selected representative examples of Twitter bot concepts that we observed. This list was not intended to be a comprehensive list of all chatbots on all platforms, but rather a comprehensive list of the types of chatbots that have been developed to this date.

As there is no verifiable central repository that documents chatbots³, we proceeded by exploring based on topic and platform. We began by searching by platform, using terms such as “Twitch bot” and “Twitch chatbot”, searching for Twitch, Reddit, Facebook, Discord, Twitter, Messenger, Kik, SMS, AOL Instant Messenger, Slack, Tinder, WeChat, and Telegram chatbots. We also searched for “chatbot app” for both iOS and Android to capture standalone bots, and “web-based chatbot” to capture in-browser chatbots. Platforms were selected according to the presence in the literature we had previously reviewed (notably [25]), but additional platforms emerged as we explored the space.

Next, we searched by topic, such as “healthcare bot” and “healthcare chatbot”. We explored a variety of topics, including healthcare, entertainment, finances, relationships, customer service, humor, romance, sex, shopping, education, news, fitness, nutrition, and productivity. We stopped our search when a team of four researchers could not find any meaningfully different chatbots from those already documented within the span of one hour. We then classified these bots using the same schema as used above - dyadic, broadcasting, and multiparty-based bots. Note that the specific proportions of chatbots in the corpus in each category should be taken with caution, as this search cannot be verified to cover the space in the same manner as a review of the ACM Guide to Computing Literature.

Dyadic chatbots in the wild

Of the 130 chatbots we identified, we classified 103 as dyadic. These fell into a number of categories, though most could be grouped as either chat-oriented or task-oriented per the categories in Li et al. [29]. Chat-based chatbots included Steve Worswick’s *Mitsuku*, the self-described “world’s best conversational chatbot” and four-time Loebner Prize winner; *Roman Mazurenko*, a virtual avatar designed in memory of a friend of one of the co-founders of Luka Inc., which was trained to speak in the same style as Roman through analysis of his communications; and SmarterChild, a now-defunct chatbot, originally designed for humorous and sometimes informative chats on AOL Instant Messenger. Other bots explored the space of providing personal companionship, and even romance and sex (e.g., the web-based *Sexbot Ciran* and *Love Droids*).

A significantly wider variety of chatbots fell into the task-oriented category. Medically-related chatbots included the web-based *Dr. AI* and *Symptomate* for diagnosis of symptoms, *Joy*, a Messenger-based mental-health check-in bot, and various nutrition-related chatbots for meal-tracking. We identified numerous personal assistant bots that aimed to support productivity or assist with basic tasks, from Facebook’s now-defunct *M* to the *Skyscanner* Skype bot. Many service-oriented bots were commercial in nature. For example, *SephoraBot*, a Messenger and Kik-based service allows users to book appointments at a nearby Sephora store and “try on” Sephora makeup in-app. General-purpose customizable customer service chatbot platforms include *Morph.ai*, *Twyla*, and *Clare.AI*.

Broadcasting chatbots in the wild

Broadcasting chatbots were found primarily in network-style platforms like Twitter. We identified 14 of these bots in our search. Microsoft’s infamous *Tay* bot and its successor *Zo* both have had Twitter presences, though *Zo* also appears in Kik. Both used a combination of one-on-one conversations

³Though <https://botlist.co/> provides a wide variety of examples with brief descriptions

and broadcasted tweets. Other types of Twitter chatbots have been developed simply to periodically post humorous or interesting content. Allison Parrish’s *@the_ephemerides* regularly tweets images from NASA’s OPUS database accompanied by often-nonsensical computer-generated text. Jia Zhang’s *@censusAmericans* tweets short, anonymous biographies of Americans based on previously submitted census data. *Angry Insulting Donald Trump Bot* is a Slack-based bot that listens for mentions of Donald Trump and posts “horrible, offensive comments” in response. Although this chatbot lives within a group, we classify it here as a broadcasting chatbot because its function is only to post public messages intended to be viewed by any present users, rather than to engage in a multiparty conversation.

Multiparty-based chatbots in the wild

We identified 13 multiparty-based chatbots, with all but two found on online community platforms that host chatroom-style conversations, including Twitch, Discord, and Slack.⁴ Prior research has identified bots as prominent social actors in the Twitch space [45], with functions of major bots ranging from engagement with users to moderation of user posts. *TranscriberBot*, for example, is a Twitch bot designed as an assistive technology by developer Myles Proulx-Vanevery, which also became a prominent member of several Twitch communities as its successes and failures in attempting to transcribe domain-specific language endeared it to community members.⁵

The group-based Slack bots we identified leaned more toward task-oriented functions. *Suttna* is a Slack bot for conducting remote check-ins with team members, while *Howdy* surveys team members and helps conduct stand-up meetings. *Niles* is a Slack bot that acts as an interactive Wiki for a group, though its social features are limited. Discord, a more socially-oriented channel-based discussion platform, has social variants on Slack bots. For instance, *Guilded* is a Discord bot for organizing gaming sessions, coordinating calendars, and recruiting new members for a guild or team. *Gnar* supports music playing in Discord voice chats, provides a suite of memetic content and allows users to organize votes.

Overall, we found that a very small fraction of chatbots in both research literature and “in the wild” communicate in multiparty interactions. We use this as evidence to suggest that this paradigm has been under-explored, with very little attention given to conversational opportunities within a group [5]. Furthermore, none of the chatbots described in the research literature were designed to be *members* of

a community, but rather they were all designed as tools to support their communities.

5 NEW DIRECTIONS FOR SOCIAL CHATBOTS

In order to explore the potential of chatbots as community members in future research and development, we took a research-through-design approach [64]. We elected to use a design approach as we do not have a clear understanding of the social roles chatbots might serve in a community or a group—humans’ social roles and the technical and social features of previously developed chatbots can serve as a starting point, but an exploration of chatbots as community members engaging in multiparty interactions could benefit from consideration of a broader space of possibilities that have not yet been defined.

We used multiple design strategies, beginning with generating as many ideas of scenarios for multiparty interaction with a chatbot as possible. We analyzed these new ideas to come up with recurrent themes and roles that chatbots can serve in a community, using an Affinity Diagramming method [3], and present the seven resulting chatbot role categories here. We conclude with a discussion about the challenges and considerations that arose in the course of our design process.

In this work, we did not strive to come up with a single-point solution or implementation. We instead attempted to present a broad range of possibilities for multiparty chatbot interactions within communities. We hope that this work inspires further research that expands, critiques, or challenges these categories.

Ideation

We used several ideation methodologies to generate a substantial pool of ideas: We first brainstormed within the existing design space, with each team-member rapidly generating about 100-150 micro-concepts of one sentence or less. Next, we worked through a custom card-based ideation method, in which we randomly selected a set of cards as a prompt for ideation [16]. We used card categories of chatbot social roles, goals, relationships to users (e.g., friendly, supportive, etc.), and platform. Finally, we used New Metaphors [32], a method designed to facilitate using metaphors as part of an ideation process [42]. These combined methods yielded roughly 600 ideas for community- or group-based chatbots. Note that in our ideation we focused on spaces like Facebook Groups, Subreddits, Twitch channels, and Discord servers, all of which are primarily social spaces where users spend time, make friends, and learn and develop as people. However, similar concepts could apply to more professional spaces as well.

⁴These platforms have low barriers to entry for novice bot development. See <https://help.twitch.tv/customer/portal/articles/1302780-twitch-irc>, <https://discordpy.readthedocs.io/en/rewrite/discord.html>, and <https://api.slack.com/tutorials>

⁵<https://lexbot.ca/TranscriberBot/learn/>

Categorization

Following ideation, we used Affinity Diagramming to make sense of the large volume of generated ideas [3]. In a team of 4 designers, we identified the characteristics of each scenario and grouped them based on their behavior and social role. The process resulted in eight preliminary categories: *Organizer*, *Host*, *Adviser*, *Archivist*, *Assistant*, *Antagonist*, *Novice*, and *Regular user*. This first draft of categories aimed to cover the different ways chatbots could socialize within their communities; for example, an *Antagonist* might bully certain users or speak in an offensive manner to gather the community against the bot as a whole. A *Host* would welcome new users and help in onboarding them, perhaps introducing them to more established members and supporting them in their early attempts to join conversations. The *Regular user* chatbot is closest to the conceptualizations of chatbots in [14], where chatbots are designed to act as regular members of a community (albeit not necessarily maliciously). Though we do not further explore this category here, it may be interesting in the future as a development challenge. As noted in [5], enabling chatbots to engage in natural turn-taking in group conversation (e.g., interjecting at the most appropriate time without having been summoned) is a significant challenge left for future work.

As a second iteration in the categorization analysis, we went through each category, reflecting and discussing agreements and disagreements among the team of researchers. For example, some of the discussion was focused on whether categories should be classified according to the chatbot's action (e.g. shared data), or based on their social role in the group (e.g. archivist). While some categories were easily understood and agreed upon in the team, like having a *dependent* chatbot, other categories demanded more debate in order to categorize them. The process also lead the team to discuss the possible outcomes of each chatbot behavior—how may it influence users' overall experiences in an online community.

We concluded with seven final categories: *Antagonist*, *Archivist*, *Authority figure*, *Dependent*, *Clown*, *Social Organizer*, and *Storyteller*. We present each of the categories in detail below (see Fig 1).

Antagonist. Both offline and online, members of communities sometimes misbehave. As noted in [14], spambots are a consistent problem, but we propose a different take on malicious bots. How might communities benefit from bots that are rude or crude or even harass users? An *Offensive Joke Bot*, which targets certain groups with crude humor or snide attacks, might initially seem of questionable value. However, if deployed within the correct context, such a bot might push a community to reconsider its values, think about what types of behaviors should be permitted, or even come together to support vulnerable members of the community. In classic

work on misbehavior, Herring et al. [19] describe the impact of a “troll” on the community within a feminist forum. Though rude, argumentative, and probably even hurtful toward some of the members, the troll did lead the community to grow through discussions of how to handle such situations in the future. A carefully designed bot might inspire such conversations as a result of its misbehavior, but might also be able to be “taught” how to be a better citizen in the community. While previous iterations of chatbots that “learn how to behave” have not always turned out as intended (e.g., Tay), the structure offered by a bounded community with a strong identity might be more amenable than the open reaches of Twitter to teaching standards for behavior to a bot.

Archivist. Previously-discussed bots have already been employed to assist in documentation, retrieval, and presentation of shared knowledge [45]. We conceive of *archivist* bots as a more social, spontaneous, and participatory variant of these assistants. In our ideation, we imagined an *RIPBot* for example, who catalogs and occasionally shares humorous “obituaries” of users who have been banned from the channel. This behavior can be entertaining, make light of previous users' misbehavior, and perhaps allow the users offended some satisfaction. Prior work has shown the importance of visible examples of what behaviors are permitted and what behaviors are not [24, 46]; while some platforms like Reddit, and now Twitch, have shifted to allow content to be removed before it can cause harm, this also may remove visibility of examples of how not to behave. A bot that can accumulate these examples over time while injecting some measure of humor could help a community retain a sense of continuity and maybe even evolution over time while keeping clear signals of acceptable behaviors.

Authority figure. The concept of a chatbot as a leader in an online community beyond simple automated rule-enforcement is one that has not yet been explored in the literature. Per the Social Identity Theory of Leadership [20], leaders emerge as prototypical members of a community; they are selected officially or unofficially because other members of the community like them, and this gives them authority. We consider *LawmakerBot*, a bot designed to enforce rules for behavior as a typical moderation bot but also to decide what the rules should be, changing them over time as the community evolves and reacts. Ideally, *LawmakerBot* would come to “understand” the community and its values better over time, but this process would work very differently depending on whether the bot is accepted as a legitimate authority figure by the community; designers might explore in the future what it means for chatbots to have a “social identity” in the classic sense [51], and what is required for other users to see this identity as legitimate. A longer-term

Social Role	Bot Example
Antagonist	OffensiveJokeBot - A bot that tells offensive jokes about users or in general.
Archivist	RIPBot - A bot that presents memories of those who came into the chat once and never came back again.
Authority Figure	LawMakerBot - A bot that makes a new rule every morning. If someone breaks it, they will be punished.
Dependent	NoviceBot - A bot who makes all of the “beginner” mistakes available. This bot allows for other users to learn from their mistakes while inserting humor into the community.
Clown	SuperlativesBot - A bot that gives out superlatives for group members, based on analysis of prior participation.
Social Organizer	AmbassadorBot - The bot pairs viewers with other viewers from other channels, based on needs or interests.
Storyteller	CoupleBots - Two bots that interactively tell the story of their secret relationship with each other. The users need to interact with them to advance the plot.

Table 1: Emergent categories of chatbots and examples of each

variant on *LawmakerBot* might be a bot that begins as a regular user (though one that is honest about being a bot) and works over time to try to gain the trust, respect, and favor of members of the community so they gradually elevate it to a position of authority.

Dependent. Many online communities have to deal with onboarding newcomers [26] and teaching them how to behave appropriately. In communities of expertise, novices may need training or support in learning how to complete tasks. In socially-based communities, newcomers may not know how to behave or who other users are, making it difficult to know who to listen to or who to trust. The process for onboarding new users can thus be compared with the process of taking care of someone who is a dependent. A *NoviceBot* could be a way to practice and clarify this process; for example, in a professional software development community, a *NoviceBot* could be developed to identify many of the most common questions asked in a community and address random members with these question (who might or might not know the answers). The community would be pushed to formalize procedures for onboarding, including designating who might answer which types of questions and what a reasonable progression of contributions might be. Broadly, a *NoviceBot* could help users rethink the newcomer experience in their communities.

Clown. Entertainment is a core part of human social experiences, from board games to stand-up comedy to horror films. It is therefore unsurprising that many chatbots across a number of online social spaces from Reddit to Slack to Discord and Twitch have already been designed to let users play in-chat minigames. Popular Twitch bots can be set to let users play “roulette” with fake currency [45]; Reddit bots have been made to play games like Tic-tac-toe [33]. The Discord bots described above allow for minigames and sharing of memes. Though these games tend to be very simplistic, they do engage multiple members of the community simultaneously, sometimes in group efforts. However, the chatbots running these games are usually based on a scripted series of prompts. In nuanced social interactions, entertainment can come from a variety of sources including shared references, humorous comparisons, and subversion of expectations. We build on these concept to consider a *SuperlativesBot*, a chatbot that keeps ongoing “rankings” of community members within yearbook-style superlative categories, occasionally stopping to engage in humorous debate with targets about the validity of their ranking. This serves both to poke fun at established users and to give clues to new users about the identities of community members.

Social Organizer. No online community exists completely in isolation; platforms are host to a wide variety of communities that often share similar interests. However, many of

these communities never find each other. For example, more than 25,000 livestream-based communities may be active at a given time on Twitch [45], but most of these “communities” have few or no members because they have not been discovered by users. In order to facilitate connections between these communities, we imagine an *AmbassadorBot* that moves between communities on Twitch and attempts to form friendships and introduce users to each other. Ultimately, its goal would be to grow its host community by identifying users in similar communities who might be interested in joining, ideally leading to a network of related communities optimized to share a large group of users with common interests. This is a particular challenge on Twitch, as advertising one’s community in another community is generally seen as poor etiquette and may result in a ban from the community [46]. An *AmbassadorBot* would have to navigate the line between advertising and connecting, perhaps by initially making “friends” before moving on to try to connect individual users from different communities.

Storyteller. Narrative storytelling as a social practice has a long history in communities both online and offline. An engaging story, particularly when there are opportunities to comment or even participate, can draw a group of people together and provide a meaningful experience. Rather than having chatbots explicitly tell stories, we propose chatbots that are regular members of a community who happen to be involved with other chatbots in engaging ways, creating a more interactive, almost “live” narrative experience. Consider a pair of *CoupleBots*, where each chatbot lives in a separate channel. Over time, one chatbot drops hints that it might like to get to know the other chatbot better, maybe in a romantic way. Once the community begins to see the possible routes for the story to take, they can decide whether to talk with the other chatbot, to encourage the first to make a grand statement, or to simply sit back and watch events unfold.

6 CHALLENGES IN CHATBOTS AS COMMUNITY MEMBERS

Through the research-through-design practice of ideation, scenario building, classification, reflection and discussion, we revealed a variety of challenges that need to be considered for the design and development of future chatbots that meaningfully engage within a community.

The first of these is a technical challenge. As noted above, this would be a significant barrier to full automation of any of the above examples; turn-taking in group conversations is difficult even for humans to accomplish smoothly, and humans make use of signals from a variety of complex cues including social status, body language, and intonation of voice. We hope that more work will continue in the line

described in [5] to develop methods for smooth group conversations with chatbots.

Beyond this, we identify three broader questions that designers and developers might consider in the process of creating new community-based chatbots:

- (1) Does the chatbot become recognized as a legitimate participant within the community?
- (2) Does the chatbot contribute meaningfully to the development of the community?
- (3) Does the chatbot’s role in the community evolve over time?

Does the chatbot become recognized as a legitimate participant within the community?

To date, chatbots have been seen as tools (task-oriented chatbots), curiosities, or companions (chat-oriented chatbots) [39, 50]. If a chatbot really is to become part of a community, it must achieve some form of *legitimacy*, as determined by members of the community. Given that standards for behavior and group identities [51] will be different even across very similar communities, there is unlikely to be a one-size-fits-all solution for gaining legitimacy beyond imitating and hoping for good results. A more thoughtful approach to designing a community-specific chatbot would be to perform one of a variety of ethnographic or ethnography-adjacent methods to understand the community in depth, perhaps even with researchers going through the process of gaining legitimacy themselves and taking notes to inspire design concepts.

Gaining legitimacy might be a multiple step process, with different behaviors expected at each step. For example, a *NoviceBot* would be expected to be respectful in its early questions about how to contribute, and to contribute without taking too much of others people’s time. As it improves its ability to contribute, it might be expected to collaborate with other users or even to teach future newcomers. A failure to adhere to these expectations could result in a loss of perceived legitimacy, or could even lead community members to treat it more like a broken tool than part of the community. Though at first this might seem like a purely technical challenge, where the bot is only legitimate if it performs its tasks perfectly, it is just as much a question of design. Expectations will need to be managed, and the bot will need to be able to fail gracefully. These chatbots’ actions and personalities must be carefully considered in the planned design, as well as adjusted in-the-moment.

Does the chatbot contribute meaningfully to the development of the community?

Current chatbots in Twitch channels, Slack teams, and Discord servers all provide additional functionalities, from humorous meme-interjections to running minigames to providing information. In this sense, they could be thought of as tools that augment users' capabilities. However, they do not meaningfully contribute to the ways in which a community and its identity evolves over time. In order to achieve this level of impact, chatbots must be engaged in the social interactions that drive community development. Typically in online text-based communities this would mean some engagement in conversation; the story of the *CoupleBots* depends on chat messages posted by the chatbots in separate channels, though well-designed messages might lead users to discuss the story outside the platform. If this is the case, one of the bots' presence alone might be enough to drive the story forward - if users are used to one of the bots being particularly present and talkative, they might sense a shift in the storyline if the bot were quiet or even missing.

Meaningful development can also happen through forms beyond text conversation. *AmbassadorBot*'s attempts to build networks and connections between communities depend on text conversation, but in the end the value comes from new relationships and connections that are made with the bot's help. Similarly, the value to the community of *Superlatives Bot* is the experience of having to negotiate and come up with plans in a group that contains a member who thinks quite differently, and may not always understand what is going on. This frustration, though ideally followed by the chatbot's learning and improving, would be important if only to help community members learn about each other and how they deal with challenging social situations.

Does the chatbot's role in the community evolve over time?

Beyond bots like *Tay* and *Zo*, which attempt to improve their conversational skills over time, most chatbots change over time when new features are added or new services are required. We suggest that a future challenge may be to design bots that change over time not only in their appearance, features, personality, etc., but in the role that the community accepts them in. Just as a human member of a community may gain more responsibilities or more trust over time as they learn and grow, chatbots could become more or less central to the community as they adapt to it. As with their human counterparts, they might be "asked" to take on a different role from their original one. For example, an *OffensiveJokeBot* might initially be seen as peripheral or annoying, but might come to the center of attention if its jokes worsen or if it begins harassing users. If it were able to "reform,"

it would make sense for it to move to a more central position in the social structure of the community and in the future might even be tasked with teaching newcomers how to behave because of its experiences with misbehaving.

Similarly, for the *LawmakerBot* described above, it is likely that the chatbot would not be able to join a community and immediately be given the legitimacy to start making rules. Early in its life-cycle it would need a significantly different design, focused much more on ingratiating itself with various users and finding the right way to engage in discussions so as to bring favor to itself.

7 CONCLUSION

In this work we present a review of prior work in the chatbots space, finding an underexplored and important opportunity to consider chatbots that engage within communities to contribute meaningfully to their development. We take insights from a research-through-design process to present conceptual categories that we hope will encourage exploration of new types of chatbots in this space. Lastly, we present a set of design considerations that emerged by reflection on the process through a social computing framework.

None of the chatbot concepts that we present here are of a type that would be complete when first deployed. As communities and their members are constantly growing and evolving, so too should these chatbots, and designers should expect chatbots to make mistakes and sometimes engage awkwardly just as human newcomers do in online communities. Rather than being problems to try to avoid entirely, these mishaps can provide opportunities for development of a narrative and a history for chatbots over time as they "realize" their mistakes and learn how to connect better with community members. If true meaningful engagement within a community is the end-goal, design and development cycles would be important throughout the bot's life cycle.

Broadly, it is valuable to consider the extent to which future chatbots should emulate human social roles. Because our ideation process was designed with the intention to explore participation specifically in online communities, our concepts currently relate to roles that exist in these communities or in broader human social settings. Future work could explore ways that chatbots might perform meaningful "inhuman" social roles, rather than building off of the ways in which humans socially interact. Moreover, as we found through development of the categories we present above, the extent to which chatbots can active participants within or distanced from their communities and still contribute to their development. As with the *CoupleBots*, even absence of a chatbot can be meaningful in the right context.

Through this work we hope to provide insight into this new design space. We also hope to inspire future work to push further on the challenges we identify and the potential

for new categories of chatbots that engage meaningfully within online groups and communities.

ACKNOWLEDGEMENTS

We would like to thank Amazon for their support of our work. We would also like to thank Nathan McKenzie for feedback throughout the ideation processes.

REFERENCES

- [1] Norah Abokhodair, Daisy Yoo, and David W. McDonald. 2015. Dissecting a Social Botnet: Growth, Content and Influence in Twitter. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '15)*. ACM, New York, NY, USA, 839–851. <https://doi.org/10.1145/2675133.2675208>
- [2] Sandeep Avula, Gordon Chadwick, Jaime Arguello, and Robert Capra. 2018. SearchBots: User Engagement with ChatBots During Collaborative Search. In *Proceedings of the 2018 Conference on Human Information Interaction & Retrieval (CHIIR '18)*. ACM, New York, NY, USA, 52–61. <https://doi.org/10.1145/3176349.3176380>
- [3] Hugh Beyer and Karen Holtzblatt. 2016. *Contextual Design: Design for Life* (2nd ed.). Morgan Kaufmann, Burlington, MA, USA.
- [4] Heloisa Candello, Claudio Pinhanez, and Flavio Figueiredo. 2017. Type-faces and the Perception of Humanness in Natural Language Chatbots. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 3476–3487. <https://doi.org/10.1145/3025453.3025919>
- [5] Heloisa Candello, Claudio Pinhanez, Mauro Carlos Pichiliani, Melina Alberio Guerra, and Maira Gatti de Bayser. 2018. Having an Animated Coffee with a Group of Chatbots from the 19th Century. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems (CHI EA '18)*. ACM, New York, NY, USA, Article D206, 4 pages. <https://doi.org/10.1145/3170427.3186519>
- [6] Boreum Choi, Kira Alexander, Robert E. Kraut, and John M. Levine. 2010. Socialization Tactics in Wikipedia and Their Effects. In *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work (CSCW '10)*. ACM, New York, NY, USA, 107–116. <https://doi.org/10.1145/1718918.1718940>
- [7] Carl DiSalvo, Phoebe Sengers, and Hrönn Brynjarsdóttir. 2010. Mapping the Landscape of Sustainable HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. ACM, New York, NY, USA, 1975–1984. <https://doi.org/10.1145/1753326.1753625>
- [8] Stefania Druga, Randi Williams, Cynthia Breazeal, and Mitchel Resnick. 2017. "Hey Google is It OK if I Eat You?": Initial Explorations in Child-Agent Interaction. In *Proceedings of the 2017 Conference on Interaction Design and Children (IDC '17)*. ACM, New York, NY, USA, 595–600. <https://doi.org/10.1145/3078072.3084330>
- [9] D Christopher Dryer. 1997. Ghosts in the machine: Personalities for socially adroit software agents. In *AAAI Fall Symposium 1997*.
- [10] Samantha Finkelstein, Evelyn Yarzebinski, Callie Vaughn, Amy Ogan, and Justine Cassell. 2013. The effects of culturally congruent educational technologies on student achievement. In *International Conference on Artificial Intelligence in Education*. Springer, 493–502.
- [11] Milica Gašić, Catherine Breslin, Matthew Henderson, Dongho Kim, Martin Szummer, Blaise Thomson, Pirros Tsiakoulis, and Steve Young. 2013. On-line policy optimisation of bayesian spoken dialogue systems via human interaction. In *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*. IEEE, 8367–8371.
- [12] R STUART GEIGER and AARON HALFAKER. 2017. Operationalizing Conflict and Cooperation between Automated Software Agents in Wikipedia: A Replication and Expansion of "Even Good Bots Fight". *Proceedings of the ACM on Human-Computer Interaction* 1, CSCW (2017).
- [13] R Stuart Geiger and David Ribes. 2010. The work of sustaining order in wikipedia: the banning of a vandal. In *Proceedings of the 2010 ACM conference on Computer Supported Cooperative Work (CSCW '10)*. ACM, 117–126.
- [14] Steven Gianvecchio, Mengjun Xie, Zhenyu Wu, and Haining Wang. 2011. Humans and Bots in Internet Chat: Measurement, Analysis, and Automated Classification. *IEEE/ACM Trans. Netw.* 19, 5 (Oct. 2011), 1557–1571. <https://doi.org/10.1109/TNET.2011.2126591>
- [15] Zafar Gilani, Ekaterina Kochmar, and Jon Crowcroft. 2017. Classification of Twitter Accounts into Automated Agents and Human Users. In *Proceedings of the 2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2017 (ASONAM '17)*. ACM, New York, NY, USA, 489–496. <https://doi.org/10.1145/3110025.3110091>
- [16] Michael Golembewski and Mark Selby. 2010. Ideation decks: a card-based design ideation tool. In *Proceedings of the 8th ACM Conference on Designing Interactive Systems*. ACM, 89–92.
- [17] Bettina Graf, Maike Krüger, Felix Müller, Alexander Ruhland, and Andrea Zech. 2015. Nombot: Simplify Food Tracking. In *Proceedings of the 14th International Conference on Mobile and Ubiquitous Multimedia (MUM '15)*. ACM, New York, NY, USA, 360–363. <https://doi.org/10.1145/2836041.2841208>
- [18] Andrea Grimes and Richard Harper. 2008. Celebratory Technology: New Directions for Food Research in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*. ACM, New York, NY, USA, 467–476. <https://doi.org/10.1145/1357054.1357130>
- [19] Susan Herring, Kirk Job-Sluder, Rebecca Scheckler, and Sasha Barab. 2002. Searching for Safety Online: Managing "Trolling" in a Feminist Forum. *The Information Society* 18, 5 (2002), 371–384. <https://doi.org/10.1080/01972240290108186>
- [20] Michael A Hogg. 2001. A social identity theory of leadership. *Personality and social psychology review* 5, 3 (2001), 184–200.
- [21] Paris (Pei-Ting) Hsu, Jingshu Zhao, Kehan Liao, Tianyi Liu, and Chen Wang. 2017. AllergyBot: A Chatbot Technology Intervention for Young Adults with Food Allergies Dining Out. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '17)*. ACM, New York, NY, USA, 74–79. <https://doi.org/10.1145/3027063.3049270>
- [22] Seung-A Annie Jin and Yongjun Sung. 2010. The roles of spokes-avatars' personalities in brand communication in 3D virtual environments. *Journal of Brand Management* 17, 5 (2010), 317–327.
- [23] Raghav Pavan Karumur, Bowen Yu, Haiyi Zhu, and Joseph A. Konstan. 2018. Content is King, Leadership Lags: Effects of Prior Experience on Newcomer Retention and Productivity in Online Production Groups. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA, Article 506, 13 pages. <https://doi.org/10.1145/3173574.3174080>
- [24] Sara Kiesler, Robert Kraut, Paul Resnick, and Aniket Kittur. 2012. Regulating behavior in online communities. *Building Successful Online Communities: Evidence-Based Social Design* (2012).
- [25] Lorenz Cuno Klopfenstein, Saverio Delpriori, Silvia Malatini, and Alessandro Bogliolo. 2017. The Rise of Bots: A Survey of Conversational Interfaces, Patterns, and Paradigms. In *Proceedings of the 2017 Conference on Designing Interactive Systems (DIS '17)*. ACM, New York, NY, USA, 555–565. <https://doi.org/10.1145/3064663.3064672>
- [26] Robert E Kraut and Paul Resnick. 2012. *Building successful online communities: Evidence-based social design*. MIT Press, Cambridge, MA, USA.
- [27] Esther Levin, Roberto Pieraccini, and Wieland Eckert. 1997. Learning dialogue strategies within the markov decision process framework. In

- Proceedings of the 1997 IEEE Workshop on Automatic Speech Recognition and Understanding*. IEEE, 72–79.
- [28] Esther Levin, Roberto Pieraccini, and Wieland Eckert. 2000. A stochastic model of human-machine interaction for learning dialog strategies. *IEEE Transactions on speech and audio processing* 8, 1 (2000), 11–23.
- [29] Jiwei Li, Will Monroe, Alan Ritter, Michel Galley, Jianfeng Gao, and Dan Jurafsky. 2016. Deep Reinforcement Learning for Dialogue Generation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. ACL, 1192–1202.
- [30] Toby Jia-Jun Li and Oriana Riva. 2018. Kite: Building Conversational Bots from Mobile Apps. In *Proceedings of the 16th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys '18)*. ACM, New York, NY, USA, 96–109. <https://doi.org/10.1145/3210240.3210339>
- [31] Lue Lin, Luis Fernando D'Haro, and Rafael Banchs. 2016. A Web-based Platform for Collection of Human-Chatbot Interactions. In *Proceedings of the Fourth International Conference on Human Agent Interaction (HAI '16)*. ACM, New York, NY, USA, 363–366. <https://doi.org/10.1145/2974804.2980500>
- [32] Dan Lockton. 2018. New Metaphors. <http://imaginari.es/new-metaphors/>
- [33] Kiel Long, John Vines, Selina Sutton, Phillip Brooker, Tom Felthwell, Ben Kirman, Julie Barnett, and Shaun Lawson. 2017. "Could You Define That in Bot Terms?": Requesting, Creating and Using Bots on Reddit. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 3488–3500. <https://doi.org/10.1145/3025453.3025830>
- [34] Wail Mardini, Yaser Khamayseh, and Ashraf Smadi. 2017. Messenger Bot for IoT Devices. In *Proceedings of the 9th International Conference on Information Management and Engineering (ICIME 2017)*. ACM, New York, NY, USA, 182–186. <https://doi.org/10.1145/3149572.3149611>
- [35] Laya Muralidharan, Ewart J. de Visser, and Raja Parasuraman. 2014. The Effects of Pitch Contour and Flanging on Trust in Speaking Cognitive Agents. In *CHI '14 Extended Abstracts on Human Factors in Computing Systems (CHI EA '14)*. ACM, New York, NY, USA, 2167–2172. <https://doi.org/10.1145/2559206.2581231>
- [36] Lasguido Nio, Sakriani Sakti, Graham Neubig, Tomoki Toda, Mirna Adriani, and Satoshi Nakamura. 2014. Developing non-goal dialog system based on examples of drama television. In *Natural Interaction with Robots, Knowbots and Smartphones*. Springer, 355–361.
- [37] J. C. Paolillo. 1999. The virtual speech community: social network and language variation on IRC. In *Proceedings of the 32nd Annual Hawaii International Conference on Systems Sciences*. 1999. HICSS-32. Abstracts and CD-ROM of Full Papers, Vol. Track2. 10 pp.–. <https://doi.org/10.1109/HICSS.1999.772680>
- [38] Juanan Pereira and Oscar Diaz. 2018. A Quality Analysis of Facebook Messenger's Most Popular Chatbots. In *Proceedings of the 33rd Annual ACM Symposium on Applied Computing (SAC '18)*. ACM, New York, NY, USA, 2144–2150. <https://doi.org/10.1145/3167132.3167362>
- [39] David M. W. Powers, Martin H. Luerksen, Trent W. Lewis, Richard E. Leibbrandt, Marissa Milne, John Pashalis, and Kenneth Treharne. 2010. MANA for the Ageing. In *Proceedings of the 2010 Workshop on Companionable Dialogue Systems (CDS '10)*. Association for Computational Linguistics, Stroudsburg, PA, USA, 7–12. <http://dl.acm.org/citation.cfm?id=1870559.1870561>
- [40] Adwait Ratnaparkhi. 2002. Trainable approaches to surface natural language generation and their application to conversational dialog systems. *Computer Speech & Language* 16, 3–4 (2002), 435–455.
- [41] Yuqing Ren, F Maxwell Harper, Sara Drenner, Loren Terveen, Sara Kiesler, John Riedl, and Robert E Kraut. 2012. Building member attachment in online communities: Applying theories of group identity and interpersonal bonds. *Mis Quarterly* (2012), 841–864.
- [42] Dan Saffer. 2005. The role of metaphor in interaction design. *Information Architecture Summit* 6 (2005).
- [43] Saiph Savage, Andres Monroy-Hernandez, and Tobias Höllerer. 2016. Botivist: Calling Volunteers to Action Using Online Bots. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW '16)*. ACM, New York, NY, USA, 813–822. <https://doi.org/10.1145/2818048.2819985>
- [44] Ari Schlesinger, Kenton P. O'Hara, and Alex S. Taylor. 2018. Let's Talk About Race: Identity, Chatbots, and AI. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA, Article 315, 14 pages. <https://doi.org/10.1145/3173574.3173889>
- [45] Joseph Seering, Juan Pablo Flores, Saiph Savage, and Jessica Hammer. 2018. The Social Roles of Bots: Situating Bots in Discussions in Online Communities. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW '18)*. ACM, New York, NY, USA.
- [46] Joseph Seering, Robert Kraut, and Laura Dabbish. 2017. Shaping Pro and Anti-Social Behavior on Twitch Through Moderation and Example-Setting. In *Proceedings of the 2017 ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW '17)*. ACM, New York, NY, USA, 111–125. <https://doi.org/10.1145/2998181.2998277>
- [47] Joseph Seering, Felicia Ng, Zheng Yao, and Geoff Kaufman. 2018. Applications of Social Identity Theory to Research and Design in Social Computing. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW '18)*. ACM, New York, NY, USA.
- [48] Joseph Seering, Tony Wang, Jina Yoon, and Geoff Kaufman. 2019. Moderator Engagement and Community Development in the Age of Algorithms. *New Media & Society* (2019).
- [49] Sohil L. Shrestha, Saroj Panda, and Christoph Csallner. 2018. Complementing Machine Learning Classifiers via Dynamic Symbolic Execution: "Human vs. Bot Generated" Tweets. In *Proceedings of the 6th International Workshop on Realizing Artificial Intelligence Synergies in Software Engineering (RAISE '18)*. ACM, New York, NY, USA, 15–20. <https://doi.org/10.1145/3194104.3194111>
- [50] Ulrike Spierling. 2005. Beyond Virtual Tutors: Semi-autonomous Characters As Learning Companions. In *ACM SIGGRAPH 2005 Educators Program (SIGGRAPH '05)*. ACM, New York, NY, USA, Article 5. <https://doi.org/10.1145/1187358.1187365>
- [51] H Tajfel. 1982. Social Psychology of Intergroup Relations. *Annual Review of Psychology* 33, 1 (1982), 1–39. <https://doi.org/10.1146/annurev.ps.33.020182.000245>
- [52] James C Thompson, J Gregory Trafton, and Patrick McKnight. 2011. The perception of humanness from the movements of synthetic agents. *Perception* 40, 6 (2011), 695–704.
- [53] Francisco A. M. Valério, Tatiane G. Guimarães, Raquel O. Prates, and Heloisa Candello. 2017. Here's What I Can Do: Chatbots' Strategies to Convey Their Features to Users. In *Proceedings of the XVI Brazilian Symposium on Human Factors in Computing Systems (IHC 2017)*. ACM, New York, NY, USA, Article 28, 10 pages. <https://doi.org/10.1145/3160504.3160544>
- [54] Teun A Van Dijk. 1997. *Discourse as structure and process*. Vol. 1. SAGE Publications.
- [55] Marisa Vasconcelos, Heloisa Candello, Claudio Pinhanez, and Thiago dos Santos. 2017. Bottester: Testing Conversational Systems with Simulated Users. In *Proceedings of the XVI Brazilian Symposium on Human Factors in Computing Systems (IHC 2017)*. ACM, New York, NY, USA, Article 73, 4 pages. <https://doi.org/10.1145/3160504.3160584>
- [56] Mandana Vaziri, Louis Mandel, Avraham Shinnar, Jérôme Siméon, and Martin Hirzel. 2017. Generating Chat Bots from Web API Specifications. In *Proceedings of the 2017 ACM SIGPLAN International Symposium*

- on *New Ideas, New Paradigms, and Reflections on Programming and Software (Onward! 2017)*. ACM, New York, NY, USA, 44–57. <https://doi.org/10.1145/3133850.3133864>
- [57] Richard Wallace. 2003. The elements of AIML style. *Alice AI Foundation* (2003).
- [58] Richard S Wallace. 2009. The anatomy of ALICE. In *Parsing the Turing Test*. Springer, 181–210.
- [59] Patrick Wang, Rafael Angarita, and Ilaria Renna. 2018. Is This the Era of Misinformation Yet: Combining Social Bots and Fake News to Deceive the Masses. In *Companion Proceedings of the The Web Conference 2018 (WWW '18)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 1557–1561. <https://doi.org/10.1145/3184558.3191610>
- [60] Joseph Weizenbaum. 1966. ELIZA—a Computer Program for the Study of Natural Language Communication Between Man and Machine. *Commun. ACM* 9, 1 (Jan. 1966), 36–45. <https://doi.org/10.1145/365153.365168>
- [61] Anbang Xu, Zhe Liu, Yufan Guo, Vibha Sinha, and Rama Akkiraju. 2017. A New Chatbot for Customer Service on Social Media. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 3506–3510. <https://doi.org/10.1145/3025453.3025496>
- [62] Mengting Yan, Paul Castro, Perry Cheng, and Vatche Ishakian. 2016. Building a Chatbot with Serverless Computing. In *Proceedings of the 1st International Workshop on Mashups of Things and APIs (MOTA '16)*. ACM, New York, NY, USA, Article 5, 4 pages. <https://doi.org/10.1145/3007203.3007217>
- [63] Jennifer Zamora. 2017. Rise of the Chatbots: Finding A Place for Artificial Intelligence in India and US. In *Proceedings of the 22Nd International Conference on Intelligent User Interfaces Companion (IUI '17 Companion)*. ACM, New York, NY, USA, 109–112. <https://doi.org/10.1145/3030024.3040201>
- [64] John Zimmerman, Jodi Forlizzi, and Shelley Evenson. 2007. Research Through Design As a Method for Interaction Design Research in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*. ACM, New York, NY, USA, 493–502. <https://doi.org/10.1145/1240624.1240704>