Malthe Stavning Erslev

I FORCED A BOT TO READ OVER 1,000 PAPERS FROM OPEN ACCESS JOURNALS AND THEN ASKED IT TO WRITE A PAPER OF ITS OWN. HERE IS THE RESULT. OR, A QUASI-MATERIALIST APPROACH TO BOT-MIMICRY

Abstract

The article develops an approach for close reading of auto-generative writing agents (i.e. bots). It introduces the concept of *bot-mimicry* (a practice of writing in a bot-esque style), and argues that bot-mimicry inherently entails that reader and writer alike imagine a conceptual (fictional) bot which *could have* written the text. As such, it investigates the concept as a fruitful way of engaging with cultural, aesthetic and political conceptions and imaginaries surrounding bots. Furthermore, and through an example reading of the "Olive Garden tweet", the paper develops, introduces and applies a *quasi-materialist* approach, where seemingly immaterial elements such as implicit conceptual bots are considered through a framework inspired by materialist media theory from the fields of software studies, media archaeology, and electronic literature.

APRJA Volume 8, Issue 1, 2019
ISSN 2245-7755

Bot or not?

> *my ukulele is not a baby / please do*
> *not reply to this maybe / we did not find*
> *it on the internet / the ukulele*

Consider the poem above: was it written by a human or by an auto-generative writing agent (a bot)? Chances are you will guess 'bot' — at least according to the statistics of the website *bot or not*, an online Turing test for poetry (cf. Laird and Schwartz; 67% guess 'bot'). In fact, it was written by a human, Aaron Koh. There is apparently a somewhat shared feeling that it reads as an auto-generated text — maybe it has to do with the choice of words, the apparent lack of semantic content, the not-quite-right rhythm of the verses. In any case, this poem, *ukulele*, mimics the style of a bot; it is an example of what I will here call *bot-mimicry*: a practice of writing in a bot-esque style. The *bot or not* website contains numerous examples of the blurred line between human-written and bot-generated poetry. This blurred line evidences that a complete distinction between human-written and bot-generated text is difficult if not impossible to uphold in practice. It is not controversial to claim that no written text is the product of pure human creativity, but always already entails 'technical' aspects, including plagiarism, remix, reference to fixed grammars, usage of predefined structures, etc. (Goldsmith). Likewise, there is always at least some human involvement in any bot-generated text, if not in editing/curating the results, then in building and selecting data-sets, and not least in programming the generative software (this holds true even for so-called 'unsupervised' systems).

Still, we may be inclined to maintain an, albeit troubled, difference between texts primarily written by humans and those primarily generated by bots. Indeed, in a time partly defined by continually more advanced text-generation systems, it is increasingly viewed as a democratic concern to do so (Laquintano and Vee; Ferrara et al.). Accordingly, the developers of one of the latest and most advanced text-generation systems, *GPT-2*, highlight its policy implications precisely because of such democratic concerns (Radford et al.). Importantly, the point here is not to scapegoat the bots for our democratic issues, but to recognize that changes in the online textual landscape calls for the development of more nuanced, fine-tuned, and critical reading skills, specifically to navigate an auto-generative situation. One important aspect of contemporary text-generation is a multitude of tech-narratives, reinforcing cultural conceptions of text-generation technologies.

The point of this paper is to develop an approach to investigating and critiquing such cultural conceptions as expressed in narratives. The paper takes as its point of departure a poem which was not written by a bot, but which reads as though it was. At one level, it may seem that the poem *ukulele* is simply a remediation of well-known ways of troubling the idea of human creativity by referencing machinic processes, such as those famously practiced by avant-garde movements including *dada* and *OuLiPo*. While this may in part be true, I nonetheless argue that there is something more at play when humans write texts that are supposed to be read as bot-generated texts without consciously involving any formal logical system. This writing (i.e. bot-mimicry) is necessarily based on cultural conceptions of text-generation technology in general, which are then written into the texts in question. In other words, reading bot-mimicry-texts allows us to study shared conceptions concerning bots precisely because they are not actually written by bots: the writer and reader alike are required to (often implicitly) imagine

a conceptual bot which could have written the text in question.

A quasi-materialist approach

The theoretical grounding of the paper is primarily based on materialist approaches, specifically the field of software studies (Fuller; Cox, McLean, Ward) along with related perspectives situated in media archaeology (Wardrip-Fruin), electronic literature (Cayley), and interface criticism (Andersen and Pold). Seeing that bot-mimicry concerns non-existent, and as such immaterial, bots (as will be elaborated, I read such bots as fictive), it may seem counter-intuitive to apply a materialist approach. Though the actual text which hints at the bot could be studied materially, the imagined bot itself may, at first, seem less appropriate for such inquiry. Nonetheless, I argue that such an approach is not only possible, it is necessary: we need to study the imagined bots present in a multitude of cultural contexts with the same rigor as the actual bots which the imagined ones mimic. To this end, I aim to develop what I call a *quasi-materialist* [1] approach: a framework for applying rigorous materialist theory to imagined (fictive) entities, in this case bot-mimicry. The paper takes on a specific case, and the quasi-materialist approach will be developed in dialectical relation to the case, where various exemplar frameworks from different fields are brought into consideration, while continually referring back to the case as the grounding for the approach.

Figure 1: @KeatonPatti's Olive Garden tweet (Patti).

The case of the Olive Garden tweet

Since early 2018, Twitter user @KeatonPatti has popularized a style of tweet in which he claims to have 'forced' a bot to watch over 1,000 hours or episodes of (often pop cultural) video content and then 'asked' it to auto-generate new, similar, content. Though @KeatonPatti is not the only one writing in this style (the style is now recognized as a meme by *KnowYourMeme*; Caldwell), this paper focuses on a specific tweet by @KeatonPatti, posted on June 13th 2018, which parodies commercials for the Italian-themed restaurant chain Olive Garden (see illustrations 1-3). This specific tweet is chosen because it is the (to date) most viral tweet in this style; it has at the time of writing gained ~326.000 likes and ~120.000 retweets (Patti). The tweet also sparked quite a few reactions on and off Twitter, including the online magazines *Futurism* and *Gizmodo*, both focusing on the ability to discriminate between human-written and bot-generated text. These articles referred to a series of tweets by @JanelleCShane, who argued that @KeatonPatti's tweets were "100% human-written with no bot involved," and stating that she "wish people wouldn't present these fakes as bot-written," though she also found at least some aspects of the tweet "pretty darn funny" (Shane; Shane is considered to be an expert on auto-generative writing and is known for her experiments

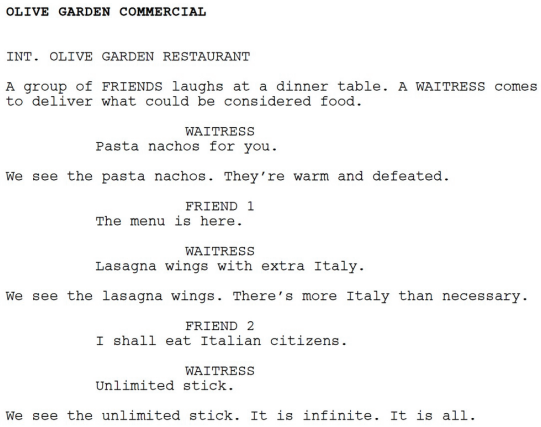

OLIVE GARDEN COMMERCIAL

INT. OLIVE GARDEN RESTAURANT

A group of FRIENDS laughs at a dinner table. A WAITRESS comes to deliver what could be considered food.

WAITRESS
Pasta nachos for you.

We see the pasta nachos. They're warm and defeated.

FRIEND 1
The menu is here.

WAITRESS
Lasagna wings with extra Italy.

We see the lasagna wings. There's more Italy than necessary.

FRIEND 2
I shall eat Italian citizens.

WAITRESS
Unlimited stick.

We see the unlimited stick. It is infinite. It is all.

Figure 2: First half of the screenplay attached to the Olive Garden tweet (Patti).

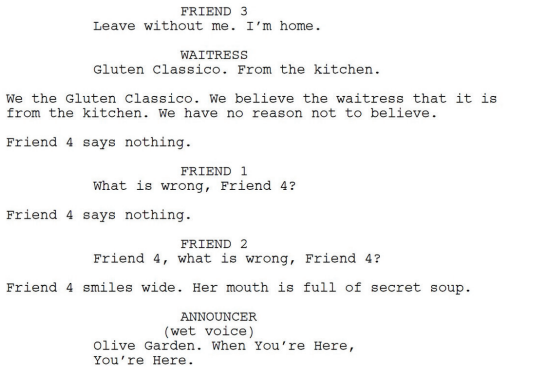

FRIEND 3
Leave without me. I'm home.

WAITRESS
Gluten Classico. From the kitchen.

We the Gluten Classico. We believe the waitress that it is from the kitchen. We have no reason not to believe.

Friend 4 says nothing.

FRIEND 1
What is wrong, Friend 4?

Friend 4 says nothing.

FRIEND 2
Friend 4, what is wrong, Friend 4?

Friend 4 smiles wide. Her mouth is full of secret soup.

ANNOUNCER
(wet voice)
Olive Garden. When You're Here,
You're Here.

Figure 3: Second half of the screenplay attached to the Olive Garden tweet (Patti).

with neural network-driven text-generation on http://aiweirdness.com). Taking the wide range of reactions to the Olive Garden tweet, along with its viral status, into account, the tweet provides an exceptionally fruitful case.

The case as meta-parody

Reading through the responses to both @KeatonPatti's and @JanelleCShane's tweets, one gets the sense that only relatively few people are actually tricked into thinking that the Olive Garden tweet was written by a bot. Many reference the fact that @KeatonPatti is a known comedy writer, who e.g. writes for the parody newsmedium *The Onion*. It seems, then, that the comical aspects of the tweet are not at all reliant on the reader believing that the bot is real. Rather, I argue, the tweet contains two closely connected jokes — it is a two-fold parody: both a parody of Olive Garden commercials and a kind of meta-parody of text-generation bots in general, specifically those common on Twitter.

I argue that the implied bot is not inspired by a single text-generation technique (such as Tracery grammars, markov chains, predictive text keyboards, word2vec, or recurrent neural networks). Rather, the implied bot relates to auto-generated text in general, an amalgamation of a multitude of text-generation techniques and the style they generally write in. In this case, reading bot-mimicry does not rely on cultural conceptions relating to a single technique (though it might in other cases), but rather on cultural conceptions of artificial intelligence/machine learning (AI/ML), and auto-generative text in general. As such, the tweets also become somewhat platform-specific to Twitter.

The tweets are situated in a context where generative text is commonplace, often in the form of so-called *Twitter-bots* (cf. Flores), but also represented in the popular predictive keyboard-based narratives by e.g. Botnik Studios (Botnik Studios). Twitter-bots are in fact so common that they are viewed as a problem by some, and action has been taken towards limiting the presence of automated bots on Twitter, or at least to make it possible to locate bot-driven accounts automatically (Siddiqui, Healy, Olmsted; Davis et al.). Most users of Twitter are used to seeing auto-generated content, and many of them have a somewhat technical understanding of how Twitter-bots (and auto-generative text in general) works (which also shows in many of the responses calling @KeatonPatti out

for not actually involving a bot in the writing process). Indeed, it seems likely that @KeatonPatti's many bot-mimicry-texts would not have been successful outside of Twitter. As such, @KeatonPatti relies on his readers being used to reading these kinds of texts — this is virtually necessary in order for them to appreciate the tweet's meta-parody.

Though one would arguably still be able to find the screenplays funny without appreciation of the meta-parody, the parody of Olive Garden commercials changes when the reader is aware of the meta-parody: the relation between the text, the platform, and the output becomes negotiable, and the reader engages in a creative act of combining the reading of the parody and that of the meta-parody. My approach is primarily concerned with the meta-parody, which relates to the implied bot and its alleged generative process.

Reading the implied bot

The implied bot as diegetic prototype

I read the Olive Garden tweet as fiction, maybe even a kind of science fiction. It contains two stories – the story present in the screenplay and a meta-story of its generation. When reading the tweet as fiction, I in part follow Paul Dourish and Genevieve Bell who have studied the interesting relation between science fiction and ubiquitous computing (ubicomp) research by reading ubicomp research alongside science fiction (Dourish and Bell). They inquire into the collective imagining that shapes much of ubicomp research. Their argument is not, however, that ubicomp research is science fiction, but that their reading of it *alongside* science fiction provides an opportunity to "point to a series of themes that illuminate contemporary imaginings of the relationship between science, technology, and society" (ibid. 773). Though @KeatonPatti's tweet hardly illuminates ubicomp research, this approach inspired by Dourish and Bell is equally fruitful when applied here.

In the present quasi-materialist study of @KeatonPatti's implied bot, the concept of the *diegetic prototype* provides a valuable perspective. The term *diegetic* refers to that which is part of a story. A diegetic prototype is a prototypical technology embedded in a story as a way to communicate or explore possibilities and dangers connected to widespread implementation of these (yet fictional) technologies, as is fairly common within science fiction (Kirby). With the perspective of *design fiction*, Julian Bleecker has shown how the line between 'science fiction' and 'science fact' is blurred — how the diegetic prototypes known from e.g. Stanley Kubrick's *2001: A Space Odyssey* or William Gibson's *Neuromancer* have played major roles in both technology development and discourse (Bleecker).

At this point, it is not entirely clear how @KeatonPatti's implied bot can be viewed as a diegetic prototype. As underlined by Joshua Tanenbaum, diegetic prototypes only work when they are embedded as part of a story (Tanenbaum) — and to what extend does that apply to @KeatonPatti's implied bot? On one level, the Olive Garden tweet is obviously a little story (it takes form as a screenplay), but in addition to this, I argue, it contains a meta-story which relates to the generation of the screenplay. The screenplay relies on the readers' understanding that it was written by a (fictional) bot. This understanding is of course derived from @KeatonPatti's brief intro, written in the tweet, where the screenplay is attached as two images.

Importantly, @KeatonPatti's intro is arguably not diegetic to the screenplay. It is not part of that story, but is rather a *paratext*, i.e. a text that is part of the work at hand, but is not part of the story *per se*, like the text written on the back of a book. Still, I argue that the implied bot is diegetic to a different story, the meta-story in which the person Keaton Patti (played by @KeatonPatti) developed, trained, and initiated a generative bot, which then outputted the screenplay in question. Seeing that all these steps are entirely fictional, I read them as a meta-story in which the implied bot functions as a diegetic prototype. The relation between the story (the screenplay) and the meta-story (the generative process) is the same as the relation between the parody of Olive Garden and the meta-parody of text-generation bots which was outlined above. To reiterate, my focus lies on the meta-story and the meta-parody, which is where the study of the implied bot can be conducted.

In my approach to reading the implied bot, this view of the bot as a diegetic prototype is the point of departure. The reading of @KeatonPatti's tweet as fictional and the bot as a diegetic prototype is fruitful in that it allows for an approach to the bot as a concrete entity to be studied through a reading of the story in which it occurs.

The implied bot is, of course, only implied. This makes it particularly difficult to study thoroughly, even when viewed as a diegetic prototype. In the following, I review selected perspectives from the fields of media archaeology, software studies, electronic literature, and interface criticism, in order to clarify how one might study *any* generative bot. In reviewing these perspectives, I seek to identify concrete methods and techniques to apply in the reading of @KeatonPatti's implied bot.

The implied bot as imagined generative system

In Noah Wardrip-Fruin's reading of Christopher Strachey's 1952 *Love Letter Generator*, situated as media archaeology, Wardrip-Fruin views the generator, "not as a process for generating parodies, but as itself a parody of a process" (Wardrip-Fruin 316). The love letter generator uses simple prewritten grammars and a fairly small database sampled from a thesaurus to generate almost rambling expressions that somehow mimic love letters, but in Wardrip-Fruin's reading, the letters themselves are "not really the interesting part of the project" (306). What is interesting to Wardrip-Fruin is the generator's data and processes, and in his reading of these elements, the love letter generator is viewed as a parody of mainstream love letter-writing activities. This relation between the outputted love letters and the parody of the letter-writing process is similar to that of the screenplay and the meta-story (and thus the meta-parody) within the Olive Garden tweet. The generative process of the love letter generator is in itself the parody in question, which maps onto the meta-story of the Olive Garden tweet — the meta-story and the meta-parody both also refer to the implied bot's generative process. Thus, we turn to the study of generative software in general in order to further develop a quasi-materialist approach to our case of bot-mimicry.

One immediate issue is, though, that the software per se is arguably an integral aspect of generative art (Cox). As such, the reading of generative art entails a close reading of the system's generative process (see Wardrip-Fruin's reading). A close reading of data, processes, and code is of course impossible in our case – these simply do not exist. Still, we may be able to study aspects

of the (implied) generator without having access to its (imaginative) technical elements.

Consider Alex McLean's generative work *forkbomb.pl* (cf. Cox; Cox, McLean, Ward; see figure 4). A 'forkbomb' is a computer program which forks (copies itself) continuously until the system crashes, and are usually very simple programs (they often require only a single line of code). McLean's work consists of a few more lines than that: in addition to being a simple forkbomb, it also generates a visual output while the program is executing; the visual output is a binary pattern which glitches as the program forks and the system crashes. This visual output has been read by generative art practitioners and scholars Geoff Cox, Alex McLean and Adrian Ward as "a 'watermark' of the processor and operating system" (Cox, McLean, Ward n.p.). While Cox, McLean, and Ward maintain the importance of considering the code of generative works, they include a focus on execution as equally important. The aesthetic appreciation of a generative work is here considered partly dependent on an appreciation of what the work actually generates. There is a clear notion that the output bears an imprint of its generator's technical elements — even those beyond the source code (e.g. the operating system). In the case of the Olive Garden tweet, is it possible to get a sense of the 'watermark' of @KeatonPatti's implied bot by considering its alleged output?

The 'watermark' that Cox, McLean, and Ward discern from *forkbomb.pl* is enlightened by a knowledge of what the system is doing 'behind the scene' — i.e. what kind of process is executing. In our case, @KeatonPatti's short explanation of how the screenplay was generated hints at some aspects of the processing which is allegedly happening in the Olive Garden tweet. Firstly, the relation between inputting video material and the bot outputting screenplays hints at an extremely

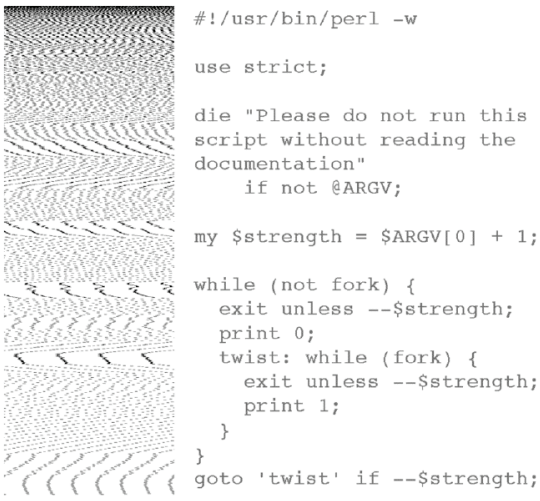

```
#!/usr/bin/perl -w

use strict;

die "Please do not run this
script without reading the
documentation"
    if not @ARGV;

my $strength = $ARGV[0] + 1;

while (not fork) {
  exit unless --$strength;
  print 0;
  twist: while (fork) {
    exit unless --$strength;
    print 1;
  }
}
goto 'twist' if --$strength;
```

Figure 4: "Output and program script from Alex McLean's forkbomb.pl" (Cox 2). The output differs on different machines, hence the idea of a 'watermark'.

sophisticated algorithm which is seemingly able to discern what is to be considered stage directions, which lines are being spoken by whom, and so on. This is impossible (or at least very unlikely), and is one aspect which hints that @KeatonPatti's tweet should not be taken seriously, but rather be read as a joke. But it also hints at certain tech-narratives that computers are able to extract any kind of information from any datatype. Indeed, there seems to be a narrative that computers are able to extract a kind of essence from given input data, and this narrative is clearly reflected in @KeatonPatti's tweet. Another 'watermark' of the generator is the ability to maintain a structurally cohesive narrative throughout the screenplay, including keeping track of characters. Again, though this is not entirely impossible, it is a distinctively sophisticated aspect of the screenplay. Both these aspects were also brought up by @JanelleCShane in her critical reading of the tweet, but where she read them as fallacies, I read them as parts of the meta-parody, and as hints at the tech-narratives at play in the implied bot. These fairly technical aspects of the implied bot refer to how the bot processed the inputted Olive Garden commercials.

The implied bot's transformation of corpus into output

So how can we scrutinize the data, or the 'corpus' in natural language processing terms, with a quasi-materialist approach? Let us consider a concrete example of a reading of a corpus by considering the output, situated within the field of electronic literature. In his essay, *Writing to be found and writing readers*, John Cayley writes with(in) the Google search engine, and through his writing practice, he investigates the system's corpus (Cayley). His writing technique is relatively simple: he uses the Google search engine to look up sequences of words from a source text, and re-writes the text based on the results — or rather, the new text is based on various concepts relating to which sequences of words are *not* found in Google's database. In one example, Cayley searches for the longest sequences within a predefined string of words which do not get any results, letting the poem take form accordingly (see figure 5). These sequences of words are, then, technically original, at least in comparison to what has been indexed by Google's 'spiders'. Cayley's technique is interesting to our case in that it is a striking example of how engaging with output text can inform an understanding of a system's corpus text. Cayley's writing engages directly with a corpus, exploring it by querying into it. Furthermore, Cayley's reading demonstrates an engagement with the way this corpus is organized and processed: his results vary depending on which Google server he arbitrarily accesses — something which he discovers through his writing practice.

Returning to my first process, with the supply text just quoted, for example:

"The purpose of this writing is to address an"
"is to address an edge of"
"address an edge of chaos."

completed with Google at 9:17 EST on Oct 1, 2009, became:

"The purpose of this writing is to address an edge"
"is to address an edge of chaos."

a little over two hours later at 11:30 on the same day.

Figure 5: An example of how Cayley inquires into Google's database. In the example above, it is shown how the resulting 'originality' changes over time, and depends on which mirror of Google's database he arbitrarily accesses (Cayley n.p.).

In relation to the Olive Garden tweet, this is important as it allows us to scrutinize @KeatonPatti's idea of a corpus text and how this is processed by his implied bot by reading the output. One striking insight into the relation between corpus and output is the bot's ability to be creative — i.e. to create something new that did not exist in the corpus text. This is, as of yet, impossible to do with any text-generation software (or, indeed, at all using AI/ML [2]). Accordingly, this feature was among the most frequently highlighted by critical readers of @KeatonPatti's tweet, where e.g. the ideas of 'Gluten Classico' or to 'eat Italian citizens' were taken as concepts which arguably could not have originated in any Olive Garden commercial. Thus, the bot seems to have the ability to not only extract suspiciously exact data from the video content fed to it, it is also seemingly able to synthesize new and highly creative concepts from this data. So, taken together, these 'watermarks', along with the relation between corpus and output, hint that the implied bot is extremely sophisticated. In the following section, I move beyond this fairly simple reading and aim to consider a more nuanced understanding of the tweet, as well as to provide a framework for conducting political critique on the cultural conceptions and technology imaginings written into the implied bot.

The implied bot's political tendency

As a final aspect of the present reading of @KeatonPatti's Olive Garden tweet, the field of interface criticism will frame a political critique as well as a general reading of implied bot. In their latest book, *The Metainterface*, Christian Ulrik Andersen and Søren Pold explore how the concept of 'the interface' has changed from being something located in a specific place (e.g. desktop computers) to being something ever-present, ever-connected, and *seemingly* immaterial – as shown by Andersen and Pold, the material reappears gradually as aspects of a metainterface (Andersen and Pold). In their numerous analyses of exemplar artworks, one common aspect is that the artworks in question are viewed as self-exploratory; they are "a material exploration of [their] own technological means of production" (24). Andersen and Pold conduct political criticism of these technological circumstances by applying a focus on the Benjaminian concept of *tendency*. Tendency here refers to a deeper political tendency as materially embedded in the technological conditions of production, which is revealed in and can be leveraged by artistic production.

In Andersen and Pold's work, such political tendency is explored through focusing on various types of interface-critical artworks, and their approach is to analyze various (artistic) interfaces. Their approach is in part inspired by Espen Aarseth's concept of *cybertext* (Aarseth), and his model of what he calls the *texual machine*, which includes three aspects: *operator*, *medium*, and *verbal sign* which are brought together in the *text/machine*, any of these aspects and can only be defined in its relation to the other two (21; see figure 6). Note that 'the operator' denotes what is typically called 'the reader', who in Aarseth's approach is situated as an integrated, constructive, part of the text/machine, and not as purely receptive, situated outside of it. Put briefly, the operator may read into (and in doing so, also reconstruct) the text/machine (the implied bot) by considering the medium and verbal signs (the tweet and screenplay). Though the tweet, the screenplay, and even the Twitter platform are also parts of the text/machine, my argument is that the implied bot is integral to the text/machine in the operator's engagement with the it, since the tweet and screenplay are somewhat nonsensical without the addition of the implied bot to the text/machine. Thus, I argue that we, situating our reading within Andersen and Pold's interface criticism, may approach the tendency of the imaginary bot by considering it as part of the text/machine.

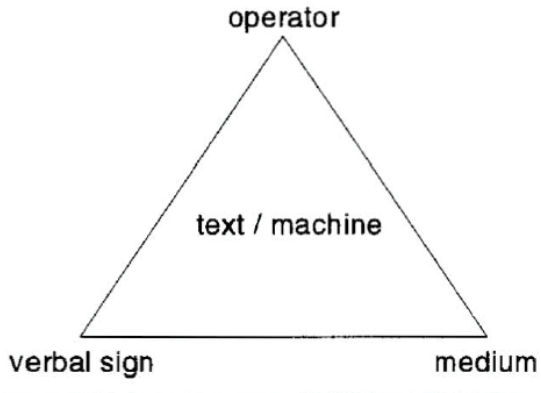

Figure 6: Aarseth's model of "The Textual Machine" (Aarseth 21).

In the case of the Olive Garden tweet, I argue that the tendency of the tweet revolves around a seemingly counter-intuitive dynamic between docility and autonomy. In @KeatonPatti's words: "I *forced* a bot to watch [...] then *asked* it to write" (Patti; my emphasis). These words give insight into the kind of generative process that @KeatonPatti imagines. The bot may be *forced* to repeat a somewhat typical computer-task (i.e. processing large amounts of data), but it has to be *asked* to

perform a typical human-task (writing being a traditionally 'creative' act). The content of the screenplay reflects this dynamic as well, as the screenplay seems to be *both* seemingly random *and* at the same time strikingly accurate in its depiction of Olive Garden (and their commercials). Two examples of this are the concepts of 'lasagna wings with extra Italy' and 'unlimited stick'. Both these concepts seem somewhat randomly generated as a result of the computer 'forgetting' what it was writing and thus combining elements that do not usually belong together, as is typical for much auto-generative text. At the same time they both parody the menu selection at Olive Garden, one referring to the highly Americanized version of Italian food served at Olive Garden (lasagna wings being a mix of the Italian dish lasagna and the American hot wings-concept), while 'extra Italy' is added to make the Olive Garden experience appear more authentic — though the attempt fails, as there is 'more Italy than necessary'. The other refers to the option to get unlimited bread sticks at Olive Garden, which is then taken to the absurd in claiming that 'it is infinite, it is all'.

These two examples demonstrate the dual nature of the bot as both a docile machine randomly stitching together unrelated concepts from a source text *and* at the same time a potent comedic parody of the inputted data, referencing concepts far beyond those that would be present in the alleged data-set. This situates the bot as harmless while simultaneously having almost mystic powers to extract an essence of a given input and synthesize it into a condensed form. As mentioned, I read this dynamic as at least one aspect of the tweet's tendency, which reflects more broad ideological conceptions surrounding the development of AI/ML. The bot's dual position as both docile and mystic echoes Wendy Chun's reading of the fetish-like ideas of source code as 'sourcery', which both gives the computer magic abilities while at the same time reinforcing an idea of complete user control (Chun). Thus, @KeatonPatti's tweet exhibits a tendency which is not reserved for AI/ML, but which applies to cultural conceptions of computers in general, yet this narrative is arguably only amplified when relating to AI/ML, as one criterion for successful AI/ML is that the system exhibits relatively high degrees of autonomy while still remaining controllable.

A quasi-materialist approach to bot-mimicry

In order to briefly sum up the quasi-materialist approach applied here, the individual aspects of it, outlined above, are here put in context to one another. The approach is considered relevant for practices of bot-mimicry, i.e. situations where humans write in a bot-esque fashion. This practice entails the (implicit or explicit) imagining of a conceptual bot which *could have* produced the written text. When analyzing this non-existent (fictional) bot, viewing it as a diegetic prototype allows for studying it by reading the story in which it is situated. Following the idea of a 'watermark' of technical aspects of a generative work, such 'watermarks' can be located in the story and hint at the imagined technical aspects of the fictional bot. By considering how one can read into a corpus by considering the output, we can then also analyze the imagined processing of corpus into output. Finally, by focusing on the political tendency inscribed into the text in question, we are able to conduct critique of the work's fictional conditions of production, and relate these to contemporary conceptions which dominate tech-narratives.

Implications and future work

The quasi-materialist framework explored here points to several interesting aspects relating to contemporary developments within natural language processing. In an era where people continually worry about bots posing as humans, one way of coping is to imitate and parody these suspected malicious bots by exaggerating particular aspects of computational writing. Such imitation can either be rather convincing (the poem *ukulele* by Aaron Koh) or openly fake (the Olive Garden tweet). The relation of these imitations to computational (real) bots is dialectic as the imitations are based on encounters with (and conceptions of) real bots, while they may themselves take part in exploring bot writing, potentially discovering blind spots. As such, these imitations may then influence the development of computational bots, likely making them yet more difficult to recognize. What makes @KeatonPatti's tweet interesting in this context is that it is not concerned with tricking the reader, but rely on the reader noticing its being fake to conduct comedic critique of both AI/ML discourse and Olive Garden commercials.

With continuing developments within natural language processing to make computationally generated text indistinguishable from humanly written text, the 'style' of the Olive Garden tweet — what I call bot-mimicry — is increasingly interesting. This style is not inherent to the ability to computationally generate language, but feeds into techno-cultural conceptions of bots, including imaginaries surrounding AI/ML as well as robots in general. What is interesting here is not so much if @KeatonPatti and others represent bot-writing accurately, but rather how these writing experiments exemplify, inquire into, and communicate shared cultural conceptions of bots and AI/ML.

Far from claiming that the readings carried out here can enable people to distinguish bots from humans online, the paper provides a different proposition: That bot-mimicry can be employed as a creative and critical way to inquire into technological conceptions, narratives, and imaginaries. I propose that facilitating writing experiments and conducting readings in bot-mimicry-texts may be fruitful ways of engaging directly with these phenomena that are otherwise difficult to articulate concretely by encapsulating them in narratives and considering them as though they were material entities. Such quasi-materialist experiments may, then, provide an opportunity to further examine and critique these cultural conceptions. A lens through which such conceptions can be negotiated, explored, and potentially challenged: through a practice of bot-mimicry.

Notes

[1] My usage of the term *quasi-materialist* is not related to its meaning within philosophy of mind, where the term relates to the question of mind-body dualism. Rather, the prefix 'quasi-' simply refers to the fact that the material is not there in a traditional sense, yet my approach is to consider the cases as though it was.

[2] This claim that it is impossible to generate something new with AI/ML technology is in part based on a lecture given by Professor Matteo Pasquinelli. The lecture was given at Cambridge Digital Humanities Learning Programme, University of Cambridge, on January 14, 2019, as part of the *Machine Feeling* research workshop. Pasquinelli argued that one integral aspect which defines the capabilities of machine learning software is "the undetection of the new" (which was also the title of the lecture). The undetection of the new refers to the way machine learning algorithms 'learn' — only by statistically aggregating the existing data of the data-set, and thus nothing entirely 'new' (which did not already exist in the data-set) can emerge in a machine learning-based system. Cf. http://matteopasquinelli.com/cambridge-ai/

Works cited

Aarseth, Espen J. *Cybertext: Perspectives on Ergodic Literature.* Johns Hopkins University Press, 1997.

Andersen, Christian U., and Søren B. Pold. *The Metainterface – The Art of Platforms, Cities and Clouds*. MIT Press, 2019.

Bleecker, Julian. *Design Fiction: A Short Essay on Design, Science, Fact and Fiction*. Near Future Laboratory, 2009.

Botnik Studios. Botnik. https://www.botnik.org/. Accessed 4 Mar. 2019.

Caldwell, Don. *"I Forced a Bot." Know Your Meme*, https://knowyourmeme.com/memes/i-forced-a-bot. Accessed 04 Mar. 2019.

Cayley, John. "Writing to Be Found and Writing Readers." *Digital Humanities Quarterly*, vol. 005, no. 3, Nov. 2011.

Chun, Wendy H.K. "On 'Sourcery,' or Code as Fetish." *Configurations*, vol. 16, no. 3, 2008, pp. 299–324. Project MUSE, doi:10.1353/con.0.0064.

Cox, Geoff. *Generator: About Generative Art and/or Software Art.* 2006. *Self-published*, http://www.academia.edu/download/36518561/generator.pdf. Accessed 04 Mar. 2019.

Cox, Geoff, Alex McLean, and Adrian Ward. "The Aesthetics of Generative Code." *Proc. of Generative Art*, vol. 03, 2000.

Davis, Clayton A., Onur Varol, Emilio Ferrara, Alessandro Flammini, and Filippo Menczer. "BotOrNot: A System to Evaluate Social Bots." *Proceedings of the 25th International Conference Companion on World Wide Web* – WWW '16 Companion, ACM Press, 2016, pp. 273–74. Crossref, doi:10.1145/2872518.2889302.

Dourish, Paul, and Genevieve Bell. "'Resistance Is Futile': Reading Science Fiction alongside Ubiquitous Computing." *Personal and Ubiquitous Computing*, vol. 18, no. 4, Apr. 2014, pp. 769–78. Crossref, doi:10.1007/s00779-013-0678-7.

Ferrara, Emilio, Onur Varol, Clayton A. Davis, Filippo Menczer, and Alessandro Flammini. "The Rise of Social Bots." *Communications of the ACM*, vol. 59, no. 7, June 2016, pp. 96–104. Crossref, doi:10.1145/2818717.

Flores, Leonardo. Genre: *Bot*. http://iloveepoetry.org/?p=5427. Accessed 4 Mar. 2019.

Fuller, Matthew, editor. *Software Studies: A Lexicon*. MIT Press, 2008.

Goldsmith, Kenneth. *Uncreative Writing: Managing Language in the Digital Age*. Columbia University Press, 2011.

Kirby, David. "The Future Is Now: Diegetic Prototypes and the Role of Popular Films in Generating Real-World Technological Development." *Social Studies of Science*, vol. 40, no. 1, Feb. 2010, pp. 41–70. Crossref, doi:10.1177/0306312709338325.

Koh, Aaron. "Ukulele." *Bot or Not*, http://botpoet.com/poem/ukulele/. Accessed 4 Mar. 2019.

Laird, Benjamin, and Oscar Schwartz. "Bot or Not." Bot or Not, http://botpoet.com/. Accessed 04 Mar. 2019.

Laquintano, Timothy, and Annette Vee. *How Automated Writing Systems Affect the Circulation of Political Information Online*. 2017. Crossref, http://dx.doi.org/10.21623%2F1.5.2.4.

Patti, Keaton. "@KeatonPatti's Olive Garden Tweet." *Twitter*, https://twitter.com/KeatonPatti/status/1006961202998726665. Accessed 04 Mar. 2019.

Radford, Alec, Jeff Wu, Dario Amodei, Daniela Amodei, Jack Clark, Miles Brundage, and Ilya Sutskever. "Better Language Models and Their Implications." *OpenAI*, 14 Feb. 2019, https://blog.openai.com/better-language-models/.

Shane, Janelle. "@JanelleCShane's Response." *Twitter*, https://twitter.com/JanelleCShane/status/1007061610005794817. Accessed 04 Mar. 2019.

Siddiqui, Husna, Elizabeth Healy, and Aspen Olmsted, "Bot or Not." *12th International Conference for Internet Technology and Secured Transactions* (ICITST), 2017, pp. 462–63. IEEE Xplore, doi:10.23919/ICITST.2017.8356448.

Tanenbaum, Joshua. "Design Fictional Interactions: Why HCI Should Care about Stories." *Interactions*, vol. 21, no. 5, Sept. 2014, pp. 22–23. Crossref, doi:10.1145/2648414.

Wardrip-Fruin, Noah. "Digital Media Archaeology." *Media Archaeology: Approaches, Applications, and Implications* (Ed. Jussi Parikka and Erkki Huhtamo), University of Minnesota Press, 2011, pp. 302–22. Print.