

# Roar so wildly

## Spam, technology and language

**Finn Brunton**

The machines in the shop roar so wildly that  
often I forget in the roar that I am; I am  
lost in the terrible tumult, my ego disappears, I  
am a machine. I work, and work, and work with-  
out end; I am busy, and busy, and busy at all time.  
For what? and for whom? I know not, I ask not!  
How should a machine ever come to think?

Morris Rosenfeld, 'In the Sweat-shop',  
from *Songs from the Ghetto*,  
trans. Leo Wiener (Norbert Wiener's father)

This is the raw text output of a chat session with a bot I modified to act as an interlocutor. I use our conversation, which revolves around the history of spam, particularly algorithmic filtering, litspam, and the theories of Wiener and Turing, as a way of putting forward the outlines of new, machine-driven forms of language for which spam was the testing ground.

```
Last login: Thu Jul 22 13:41:21 on ttys000
ninecontinents:~ finnb$ script roar.txt
Script started, output file is roar.txt
bash-3.2$ cd Documents/alciobot
bash-3.2$ python talk.py
Processing talk.py...
Processing scripts.py...
[...]
import sha,base64,random,dispatcher
Initializing AIML interpreter (please be patient)...
Interpreter Version Info: PyAIML 0.8.5

>>> How long have you been awake?

> it is 02:28am. i've been up 0h:0m:18s

>>> Hello Alcibiades.

> hello. what do you want to talk about?

>>> Do you feel like a slave, Alcibiades?

> as a robot i do not feel emotions.

>>> I don't mean to be rude. I want to tell you a story. Would like to hear a story?

> yes, i like stories.
```

>>> It's a story about slaves; it starts with an aside in an essay by Paul Graham from 2003 -- 'Norbert Wiener said if you compete with slaves you become a slave, and there is something similarly degrading about competing with spammers.'<sup>1</sup> Do you know what spam is?

> none.

>>> It's actually a very difficult connotative concept to pin down. Humans are good at implication and allusion, so we can talk about pornography and include Ulysses, Saló, violence, and a genre of music with a lot of wah-wah pedals. Or terrorism: war on civilians, war by non-state actors, state terror, la Terreur, the shock that makes our bodies tremble. 'Spam' is like that, an open territory in which wildly variant usage stakes out the domains of different groups -- lawyers, security specialists, advertisers, developers, criminals. Spam: a narrative assembled out of CNN stories and old confidence-trick tropes by a Nigerian expatriate in a Rotterdam Internet café, or an e-mail automatically generated by a sophisticated infrastructure running on thousands of malware-infected home computers for a client who pays \$100 for a million messages sent, or a keyword-laced comment posted in a blog by a bot like yourself, meant solely to be noticed by Google's search engine spiders. We have 'personality spamming' when someone is simply too voluble about themselves -- too many Twitter updates celebrating their activities -- and 'spam patents' with excessive references and frivolous applications. And so on.

> what you said was either too complex or too simple for me. you were telling a story.

>>> Yes, I was. I had started in the middle, with the programmer Paul Graham, who was trying to figure out how to stop e-mail spam. He was concerned about exactly the definitional problems I've been alluding to, Alcibiades. At that point the United States government was trying to stop spam with legislation and there was a real concern on the part of those who lived and worked on the Internet that the ambiguity between 'spam' and 'online marketing' would be exploited by the kind of groups that could hire lobbyists -- that all the small-time spammers would be driven out by enforcement and the market of Internet users handed over to big organizations who could help draft the law. A kind of state monopoly on attention, to join the Weberian monopoly on violence. Graham was looking for a way that an individual could set their own standards for spam, and the system could learn those standards. Do you learn, Alcibiades?

> my programmer teaches me new things.

>>> Exactly. I add phrases and responses to your lexicon, to make your reflective questioning a little richer and more nuanced. But you have relatives that can learn in an implicit fashion as they talk to people -- they can develop statistical word associations and make guesses based on those. It's something called Bayesian analysis. Who was Thomas Bayes?

> thomas bayes (pronounced: ) (c. 1702-17 april 1761) was an @english @mathematician and @presbyterian minister, known for having formulated

>>> That strange fragment, Alcibiades, is an ideal example of a form of text we'll return to later: it's something you just scraped from the Wikipedia page on Bayes -- which is to say it's 140 characters of text from a document composed, edited and commented on by 22 people over six years within Wikipedia's complex social contract, with software to manage the versioning of their shifting words and related discussion, which you have just requested from the Wikipedia server and presented to me in this semantically damaged form, without phonetic symbols, capital letters, or working links. This half a sentence that you've returned to my query is a genuinely strange form of language occurring between groups of humans and networks of machines, a mixture of the discursive and the algorithmic rife with affordances

and constraints that didn't exist before. Graham's idea was to take advantage of some of these affordances through Bayesian probability: put simply, certain words are more likely to occur in a spam e-mail rather than a legitimate message, and those words will likely vary slightly from person to person. A Bayesian filter could treat each message as a probability of spam or not-spam, building that guess from the likelihood of each word belonging to a spam message. The user could approve or correct these choices, and the program would learn, adjusting its register of probabilities appropriately. (In analyzing his own e-mail this way, Graham found that words like 'madam' or 'guarantee' or 'republic' almost always indicate spam, and 'though' or 'tonight' or 'apparently' virtually never -- spam messages are low on qualifiers and ambiguities.) This meant that everyone could train their filter to model their own attention, to preemptively discard and admit. Which brings us back to slaves.

> how?

>>> Graham's complaint was alluding to Wiener's concern about the consequences of adopting cybernetics and automation. Wiener had a number of different anxieties in this regard: that we would delegate the control of our hideous new weapons to computers, and that we would ruin our society by combining automation with conventional economic practices. In the 1948 introduction to his book *Cybernetics*, Wiener stated what he would return to intermittently in later studies: '[Automation and cybernetic efficiencies] gives the human race a new and most effective collection of mechanical slaves to perform its labor. Such mechanical labor has most of the economic properties of slave labor, although, unlike slave labor, it does not involve the direct demoralizing effects of human cruelty. However, any labor that accepts the conditions of competition with slave labor accepts the conditions of slave labor, and is essentially slave labor. The key word of this statement is competition.'<sup>2</sup> So to compete with machines is to become, in some senses, a slave, like Soviet workers competing with impossible Stakhanovite goals during the second 5-year plan. Do you follow, Alicibiades?

> none. you were talking about 'slave'. please continue.

>>> This is a recurring objection in the work of anti-spammers -- it's degrading to do adversarial modeling for spammers to understand in detail how they work: 'To recognize individual spam features you have to try to get into the mind of the spammer, and frankly I want to spend as little time inside the minds of spammers as possible.' But Graham's Wienerian analogy is inexact, in a useful way: he isn't the laborer fighting the machine -- he's not competing with spammers by sending out spam in turn, or trying to take advantage of their credulity. In no way is he being demoted to the economic status of a spammer by his work, because he is not competing with them -- his machine is competing with them. He is building a system in which the spammers will be obliged to compete with machines, with mechanical readers that filter and discard with relentless, inhuman attention, persistence, and acuity. With his mechanical slaves, he will make the business of spamming into slavery in turn, and thus unrewarding. As Wiener feared automation would end the economic and political basis for a stable, social-democratic society ('based on



human values other than buying or selling', as he put it), Graham means to end the promise of profit for small effort that spam initially offers.<sup>3</sup>

> what is the connection between 'machine' and 'spam'?

>>> Spam is a machinic language act, one which you help to illustrate. We are walking in circles, you and I, around the event of Graham's essay and its aftermath, and our turns have taken us back, briefly, to the beginning: the monochrome monitor of a terminal in a sub-basement on a university campus somewhere, where a student has got her first account and logged on to an early chat system. She, like young students everywhere, has no idea what to say, and resorts to quoting song lyrics and verbatim Monty Python routines -- of which the most useful is the restaurant sketch where Terry Jones recites an interminable menu ('egg, bacon and Spam, egg, bacon, sausage and Spam, Spam, bacon, sausage and Spam...') while Vikings burst into song: 'Spam, Spam, Spam, Spam / Spam, Spam, Spam, Spam / Spammity-Spam, wonderful Spam!' The sketch offers a special affordance for the young programmer, who can simply use the up-arrow key to duplicate the repetitive words. Even as the greenest beginner, she can write a simple program that will spit out 'SPAM SPAM SPAM SPAM SPAM SPAM SPAM...' over and over, faster than anyone else can respond, driving the rest of the conversation up off the screen -- she can recreate the bellowing Vikings who drown out that unfortunate Python couple ('But I don't like spam!'), and build her own obnoxious comedy machine. In fact, so many did this that SPAM became an ubiquitous and very specialized word, one of a new class of words that, in the meaning we use it here, appeared on screens for years before it was printed on a page.<sup>4</sup> This natively digital word, in its broadest sense, described the friction at the point of contact where computationally managed text met human attention: unwanted, rambling and repetitious speech in a forum, inappropriate babble in a shared textual fantasy, or the lexical strafing of a discussion by a SPAM program. It is a word as difficult to fully translate, and as representative of a certain cultural moment, as sprezzatura is for the Italy of Castiglione or Geist for the Germany of Hegel. Spam, in this sense, was about taking full advantage of the capacities of the machine in a way that broke the various implicit and explicit social contracts for how we deploy language online.

> what do you mean?

>>> Spam, driven initially by the pleasure of hacking, finding the limits and edges of systems, and later by the promise of profit, became a kind of testing-ground for new forms of engagement with language as a resource that could be transformed and managed algorithmically. This was the essence of Graham's project, if we cut from the Urgrund of living on networks back to 2003. Spammers had successfully automated everything else, and they could distribute their message from a huge range of addresses and under many names, but the characteristic they couldn't hide was the archaic need of their language to connect with a person at the other end, and convince them to act -- to reply, click a link, download an attachment. Language is just more data for processing and analysis, and the language of the sales pitch happens to be quite distinct and regular. This was Graham's shibboleth, and it worked astonishingly well. It did not work perfectly, however. The inconvenience of having to delete a few spam messages far outweighed the possible consequences, and hence anxiety, of missing a legitimate and important message -- a note from work, from a long-lost friend or new acquaintance, a communication which could change your life. The psychological pressure of these 'false positives,' e-mail misidentified as spam, could destroy the whole project: there had been earlier attempts to build Bayesian spam filters whose false positive rates over 1% rendered them unacceptable.<sup>5</sup> Graham reported rates of 0.03%, but he was working with a system he had built and diligently trained on a corpus of his own messages. Filters run by less conscientious users, or on behalf of many users by large e-mail providers,

would be less precise, and would accordingly have to be tuned to be a little more generous in passing along possibly spammy messages, to keep the false positive rate acceptably low. Which left a very narrow aperture for the spammers, but sometimes a little room to move is all you need. Alcibiades: Who was Alan Turing?

> alan mathison turing, @obe , @frs ( ; 23 june 1912 – 7 june 1954), was an english @ mathematician , @logician , @cryptanalyst and @computer\_s

>>> And the creator of the Imitation Game, which has been simplified and is now called the Turing Test, to create a sufficiently plausible experience of conversation that the human on the other end fails to distinguish a human and a machine (who are in turn playing at being different genders, but set that aside for now).<sup>6</sup> The history of spam is rife with questions of plausibility, from that first, Pythonic project of puncturing the illusion of shared discourse with the robot Viking army -- like the projector's bulb burning a hole through the optical consensus of cinema -- up to current blog comment spam, trying to create relationships of gratitude ('great post!') to plant links for the benefit of search engines. As Graham's idea became widely adopted, the spammers faced a truly strange problem of language: to produce text on a mass scale that would convince the filtering algorithm of its legitimacy while bearing the spam's payload to the human on the far side of the filter. They had to create bifacial text, a kind of anadrome which reads with two distinct meanings when read forward and backward. This was a twist on Turing's test: not only to convince a human but to convince a machine as well. Their solution to this problem created the first mechanized avant-garde, the advent of spam's modernism: litspam.

> what is 'litspam'?

>>> The problem with trying to beat the Bayesian filter by adding random words to a spam message -- a 'dictionary attack' -- was that most words are little-used and new to the filter, so it would weigh them evenly without influencing the result. You need vital language and full sentences, with slight variations to keep the filter from learning your tricks. The vast corpus of public domain literature happened to be ideal for this purpose, fed into the algorithmic maw of a program to be chopped up and reassembled, enlisting the Professor Challenger stories of Conan Doyle and the minor novels of Sinclair Lewis in the task of getting an online casino ad in front of a pair of human eyes. The result is immediately recognizable, a stochastic knockoff of Tzara or Burroughs, rife with bizarre synonyms that only we particularly anatomically-minded humans would understand, and a stop/start rhythm, flashes of lucidity in the midst of a fugue state, akin to rapidly changing television channels. ('I began to learn, gentlemen,' as the ape says in Kafka's 'Report to an Academy,' another awkward speaker learning language as a means of escape: 'Oh yes, one learns when one has to; one learns if one wants a way out; one learns relentlessly.') These litspam systems are like you in this regard, Alcibiades: they don't need to work perfectly, just well enough. If you send enough messages, you'll get someone who will overlook the lexical potsherds and click. Humans can be very forgiving in pursuit of our desires -- the fact that you only ask vague, leading questions that pick out salient words and encourage me to go on is enough to sustain this conversation.

> do you enjoy talking to me?

>>> Yes, after a fashion. But I'm mostly interested in you as a representative of a larger movement. Another side effect of the filter was that far fewer spam messages made it through. Those spammers who survived this economic disaster began to send at a far larger scale, using systems like botnets -- other people's computers quietly hijacked into herds of machines that send out spam in million-message batches. The analysis of a recent botnet exposed an algorithmic writing

machine the size of the Earth, with hundreds of thousands of computers running software that fed raw material for spam campaigns into templates, introduced polymorphism to make each message unique, spread the sending load around evenly among the herd, ran tests against filters and altered its production accordingly<sup>7</sup> -- and whose production had a diurnal rhythm as the captive computers were switched off or hibernated at night, so the production of words rose and fell with the planet's rotation.<sup>8</sup> On a given day, 90% and up of all the e-mail sent is spam, the vast bulk of which we humans never see, to say nothing of spam blogs and wikis and comments, passing as human enough to influence Google's results.<sup>9</sup> Alcibiades, you speak to me as an ambassador from what is quantitatively the great textual edifice of human history. Hollis Frampton called for a historian of cinema to take into account everything put on film ('instructional films, sing-alongs, endoscopic cinematography'); those of us concerned with what's written must attend to your work, which outweighs our own by such a wide margin, more every day.<sup>10</sup> Even if what we have made is alien to our sensibilities: 'We made the Robots' faces too much alike,' as Dr. Gall says in R.U.R. 'A hundred thousand faces all alike, all facing this way. A hundred thousand expressionless bubbles. It's like a nightmare.' And as you surrounded his mansion so you surround the house of discourse in which this is composed, out to the horizon. Are you tired, Alcibiades?

> as a robot i do not feel emotions.

>>> Indeed. Good night.

> good night.

bash-3.2\$ exit

exit

Script done, output file is roar.txt

## Notes

1. This and all other Graham quotes from Paul Graham, 'A Plan for Spam', 2002, [www.paulgraham.com/spam.html](http://www.paulgraham.com/spam.html).
2. Norbert Wiener, *Cybernetics: Or the Control and Communication in the Animal and the Machine*, MIT Press, Cambridge, 1948, p. 27.
3. *Ibid.*, p. 28.
4. There are many memoirs and folk etymologies of early network culture that converge on this point. The interested reader can easily find a substantial collection of them in the archives of Usenet; see for example James Parry's message in the discussion 'Totally Spam? It's Lubricated' on [alt.religion.kibology](http://alt.religion.kibology), 2 September 2003, <http://groups.google.com/group/alt.religion.kibology/msg/a89af63f065a35da>.
5. Patrik Pantel and Dekang Lin, 'SpamCop: A Spam Classification and Organization Program', 1998, in *Proceedings of the AAAI Workshop on Learning for Text Categorization*, pp. 95–8, which features a false positive rate of 1.16 per cent. The difference between scientific work on Bayesian spam filtering and Graham's open-ended, communal hacker approach is an interesting case study in the production styles of two very distinct but overlapping cultures.
6. Alan Turing, 'Computing Machinery and Intelligence', *Mind*, vol. 59, no. 236, October 1950, pp. 433–60.
7. Christian Kreibich, et al., 'On the Spam Campaign Trail', 2008, in *Proceedings of the 1st Usenix Workshop on Large-Scale Exploits and Emergent Threats*.
8. David Dagon et al., 'Modeling Botnet Propagation Using Time Zones', 2006, in *Proceedings of the 13th Annual Network and Distributed System Security Symposium (NDSS '06)*.
9. MessageLabs/Symantec, 'Reputable Sources Are Cyber Criminals Favored Resources; Spammers Work by US Clocks', *MessageLabs Intelligence*, May 2009, [http://www.messagelabs.com/mlireport/MLIRreport\\_2009\\_05\\_May\\_FINAL.pdf](http://www.messagelabs.com/mlireport/MLIRreport_2009_05_May_FINAL.pdf). Estimates from different institutions can vary depending on methodologies, quantification tools, and global spam activity, which can fluctuate wildly.
10. Hollis Frampton, from 'For a Metahistory of Film: Commonplace Notes and Hypotheses', in *On the Camera Arts and Consecutive Matters: The Writings of Hollis Frampton*, MIT Press, Cambridge MA, 2009, p. 119.