Speech In, Speech Out

*James Grimmelmann**

On two occasions I have been asked, – “Pray, Mr. Babbage, if you put into the machine wrong figures, will the right answers come out?” ... I am not able rightly to apprehend the kind of confusion of ideas that could provoke such a question.

– Charles Babbage¹

I.

Ronald Collins and David Skover ask “whether and why First Amendment coverage given to traditional forms of speech should be extended to the data processed and transmitted by robots.” Their answer is yes, because “what really matters is that the receiver experiences robotic speech as meaningful and potentially useful or valuable.” From this, they conclude that “utility in furthering some lawful objective” will serve as the new “First Amendment norm.”

The answer is right, as far as it goes, but the question is wrong, and so Collins and Skover misunderstand how far it goes. A few seconds’ reflection shows that sometimes robotic transmissions are speech and sometimes they aren’t, so the proper question is not “whether and why?” but “when?” “Robotic speech is covered if it has utility” sounds like a statement about robots and about utility. But really it’s a statement about *speech*, because in that sentence, that’s the word doing all the work.

* My thanks for their comments to Aislinn Black and Rebecca Tushnet. This essay may be freely reused under the terms of the Creative Commons Attribution 4.0 International license, <https://creativecommons.org/licenses/by/4.0/>.

II.

There is a robot in my room. It has a wall-mounted control panel. If I manipulate the controls one way, the robot activates and starts emitting a precisely specified profile of electromagnetic radiation. If I manipulate the controls another, the robot initiates a shutdown sequence and cuts off the radiation.

Does my light-bulb robot “speak?” The obvious answer is “no”: the robot’s output is functional, not expressive. A soft white glow expresses no intention and conveys no message.

But of course this answer is wrong, because of course a light source can be expressive. Two lanterns hanging in the belfry of the Old North Church conveyed the message that the British were coming by sea. A naval signal lamp conveys actionable orders simply by turning off and on in Morse Code; I could easily do the same with my light-bulb robot. My communicative intention, Collins and Skover might say, is “delegate[d]” to the light bulb.

It is tempting to say, then, that light bulb emissions are covered speech. But this too is wrong, because it cuts too broadly. Congress has effectively banned the manufacture of inefficient household incandescent light-bulb robots, pushing most consumers like myself to use compact fluorescent (CFL) or light-emitting diode (LED) robots instead. There were and are serious policy debates over the wisdom and timing of the ban. But one argument that to my knowledge has never been seriously floated is that it amounts to a restriction on “speech” that needs to be measured against the First Amendment, the way that a restriction on the sale of books or Blu-ray players would need to. Sometimes a light bulb is just a light bulb.

Some light bulbs are used in ways that produce speech; other light bulbs are used in ways that produce nonspeech. This is not something intrinsic to the technology – one LED might be part of a speech-emitting digital billboard while a physically similar one might be part of a nonspeech-emitting pen light. Zooming out doesn’t help.

The billboard might be turned to all-white maximum brightness to illuminate a parking lot, while the pen light might be used to flash kid code for “Mom and Dad are asleep; let’s sneak downstairs.”

It is not possible to distinguish speaking machines from conduct machines on the basis of their physical instantiation. Hardware and software are interchangeable. Almost any computer could in principle be implemented using sticks and ropes or dams and water wheels. Danny Hillis and Brian Silverman built a working computer out of Tinkertoys that plays a perfect game of Tic Tac Toe.²

Complexity isn’t the dividing line, either. A typical modern car may have a hundred million lines of source code.³ Yes, some of that runs the onboard entertainment system, but most of it is devoted to boringly functional car stuff like calibrating the fuel–air mixture in the engine. Should VW have argued that its infamous emissions-test defeat device was protected speech?

Perhaps you’d like to say I’m being silly. We *know* that when a VW engine detects that it’s on the open road and turns off some of its emissions controls, it’s not really “speaking.” But ask yourself this. *How* do we know?

III.

In some cases it is easy to tell whether a robot is emitting speech because it is easy to identify a human behind the curtain. If I am flicking my light-bulb robot on and off in Morse Code, then I have an “intent to convey a particularized message,” in the words of *Spence v. Washington*.⁴ No special gyrations are required to attribute speech to the light bulb; the speech is *my* speech, regardless of the medium I use.

Sometimes it takes a little while for it to sink in that a new communications medium can by definition be used to convey messages, but it always happens in the end. There is a brief confused moment when people think that the new medium has no speech in it.

Then there is an even briefer confused moment when people think that the new medium is all speech. And then sanity returns and we treat the medium like any other: not determinative by itself but relevant to understanding the content and the context of the messages it conveys. Books can be speech, parades can be speech, movies can be speech, video games can be speech, light bulbs can be speech, robots can be speech. To think otherwise is to score a conceptual own goal.

Actually, this usage is a little loose. A bound codex of sheets of paper with ink blots on them in the shapes of letters in the Latin alphabet arranged into English words and sentences is not “speech” if it is being used as a doorstep or a melee weapon. It is “speech” when it is being *used as* a medium of communication – to “convey a particularized message.”⁵

There is a recurring but necessary difficulty here: how can we tell when the book or the light bulb is being used as a medium of communication and when it is not? This is the work that the *Spence* test does: it shifts our attention from the medium to the message. In *Texas v. Johnson*, the medium was a flag on fire;⁶ in *Cohen v. California*, a jacket.⁷ Neither case stands for the proposition that “flag speech” or “jacket speech” is a distinct First Amendment category. Indeed, there are *Spence*-test cases involving no artificial medium at all, not even a flag or a jacket: in *Erie v. Pap’s A.M.*,⁸ a plurality of the Court held that nude dancing was “expressive conduct” implicating the First Amendment. Being nude with intent to communicate is different from being nude without, just as flipping a light switch or programming a computer with intent to communicate is different than doing so without.

Most of the time.

IV.

This speakers’-intent approach works for simple cases where there is no ambiguity about who the speaker (if any) might be. The

correspondence between my intent to communicate by flashing the light bulb and the resulting flashes is so close that it seems to follow without question that the flashes are my speech, and therefore are speech.

But there are harder cases. What if I program a computer to emit the same message repeatedly: spam? What if I program it to emit related but distinct messages: a mail merge? What if I program it to emit random (whatever that means) variations? What if you and I program it together, or a hundred of us contribute code? What if you use a program I wrote, or what if my program learns from its interactions with millions of users or from parsing millions of existing texts? What if? What if? What if?

Some of these cases have easy answers; some do not. What they share, for present purposes, is that the speaker's-intent approach threatens to break down because it is no longer so easy to associate a message with the intent of a unique human author. It is this radiating complexity that makes the category of "robotic speech" superficially attractive – if only we could declare it all speech by treating the *program* as the speaker, we could cut the Gordian knot and be home by teatime.

In the copyright context, I have argued that this "solution" is chimerical.⁹ The problem of assigning authorship in computer-generated works seems to repeatedly lead to suggestions that we ought to treat them as computer-authored. It doesn't work, because unless and until computers are capable of being treated like people in general, calling them "authors" deals with the complexity not at all, by giving a completely arbitrary wrong answer. A novel written using Microsoft Word and an animation that plays itself when the user clicks are meaningfully different for copyright purposes: the user is the author of the former, while the programmer is the author of the latter. In neither case is the program the author.

But things are not quite so bad when it comes to speech, because copyright operates under a constraint that the First Amendment does not. The private-rights structure of copyright law requires identifying a copyright owner for each work, and authorship (or something derivative of it, as with works made for hire and transfers of title) is the only principled way of doing so. Free speech law need not do the same: it is perfectly plausible to say that something is protectable *as speech* without identifying a *speaker*.

Collins and Skover get there by appealing to reader-response and related theories in literary criticism, which emphasize a listener's experience of a text rather than an author's intent. This is hardly unprecedented in First Amendment law: there are plenty of cases in which a listener has greater rights than a speaker with respect to the same material (take *Stanley v. Georgia*'s protection for the possession of obscenity in the home¹⁰) or has standing where a speaker is not before the court and might not even be identifiable (take *Lamont v. Postmaster General*'s protection for the American recipients of foreign communist political propaganda¹¹).

Collins and Skover call this "intentionless free speech," (or "IFS"), and it performs well both on easy cases where a robotically assisted human speaker can be identified and on harder cases where the human speaker is missing. Of course lights can be speech: Paul Revere's riders knew what the lanterns meant. And of course robotic utterances can be speech: people often regard the sonic waves Siri emits as being laden with meaning. It's a fun parlor game to try to attribute that speech to Apple, its employees, Apple's data sources, the user, and the other users whose responses provided training data. But the human user's experience of Siri's utterances as meaningful speech does not depend on which humans (if any) were responsible for those utterances. That ought to be sufficient to ground a First Amendment interest, even if we're not sure whose line it is, anyway. Robots speak.

One could of course debate whether the speaker's or listener's experience matters more here, or whether both do, and if so, how. I think Collins and Skover have it right when they claim that a listener's experience is sufficient to ground a legally, morally, and politically cognizable speech interest. (I would add that nothing requires us to treat a listener's experience of speech by itself as presenting the same case for protection that a listener's and speaker's experiences together do, but that is a matter for another time.) When a light bulb turns on in a previously darkened room, no one present is likely to think of it as speech. When a light bulb flickers in Morse Code, those in the room are more likely to recognize it as speech, even if they don't know who is responsible for making it flicker that way.

We have the intuitions we do about light-bulb robots because our free-speech intuitions in general are structured by our extensive experiences as audiences who respond to communications. Consider *Bland v. Roberts*. The district court held that a Facebook like was not protected speech, because there was no "substantive statement," just "one click of a button."¹² But the Fourth Circuit corrected the mistake on appeal, writing, "In the context of a political campaign's Facebook page, the meaning that the user approves of the candidacy whose page is being liked is unmistakable."¹³ Where does that meaning come from? From the community of Facebook users who would see the like and infer the liker's support. That's a claim about the social meaning of a technical practice. Reader-response gets cases like *Bland* right for the right reasons.

V.

But we still have not exhausted the types of robotic "speech." In addition to cases with an obvious human speaker and cases with an obvious human listener, there are also cases in which it is not so easy to identify a human speaker *or* a human listener. Collins and Skover give an extended example involving a "robotrader" that executes an

algorithmic series of stock trades and then at the end of the day generates a report for a human investor listing the trades and their gains or losses.

In such a scenario, the human investor was not a receiver of information during the process of trading, because the robotrader's objective was to "make meaning" itself of the relevant data gathered to inform its buying and selling decisions. Nonetheless, a real First Amendment experience exists in this example – one that is too easily overlooked when focusing narrowly on the fact-based end product rather than more expansively on the intermediate moves that made that product possible.

Even when robots or robotic components communicated with one another, there was still "meaningful" information being conveyed back and forth – all in exchanges that were set into motion by the human investor and that culminated in his or her reception of the robotrader's report. In short, the inter-robotic communicative exchange worked at the behest of and in the service of human objectives. Assuming that the investor's purposes and goals were lawful, the robo-trader's exchanges of information alone made those commercial objectives possible. Why, then, should the intermediate stages in the process – the communicative steps – be viewed as any less deserving of First Amendment coverage?

For IFS purposes, moreover, it does not matter whether the robotrader's report was nothing but a communication of a collection of facts that had little or no ideological or evaluative significance.

I have quoted this passage at length because I think it represents the precise point at which the argument in *Robotica* goes wrong. The argument, as best I can understand, seems to be that the lack of human

involvement in generating or receiving the “interrobotic communicative exchange” is no obstacle to First Amendment protection. Intentionless free speech excuses us from having to identify a specific human whose intent is conveyed by these communications. Instead, the practical reality is that such communications will be protected as speech whenever they are “at the behest of, and in the service of, human objectives.”

This is the norm of utility. Collins and Skover contrast it with the views of other First Amendment scholars who believe that protections are reserved for speech that has “ideological or evaluative significance.” For example, Robert Post argues in a famous essay discussing First Amendment protection for computer source code – crudely, speech *to* robots rather than speech *by* robots – that “First Amendment coverage is triggered by those forms of social interaction that realize First Amendment values.”¹⁴ Post has his own set of preferred values,¹⁵ and other First Amendment theorists have theirs. Collins and Skover set up the norm of utility in opposition to all such theories. “Rather than looking up to the normative heavens, it looks down to the streets where life and technology evolve.” Thus, they expect equal First Amendment protection for factual, artistic, and political speech – that is, “utility” replaces “truth” or “beauty” or “self-governance” as the governing First Amendment norm.

VI.

Collins and Skover seem to think that the norm of utility follows from their reader-response analysis. I am not so sure that it does. In fact, I am quite certain that it doesn’t.

Take the robotrader. Collins and Skover claim that the intermediate stages in the day’s events – the information transmitted between and within trading robots – are covered “speech,” and not just the final report presented to the human trader. But while the intermediate transmissions were “at the behest of, and in the service

of, human objectives,” the only thing experienced *by a human* as speech was the final report. A listener-oriented reader-response approach can find meaning in the final report, but it has less to say about the intermediate steps. No human was present for them; no human extracted meaning from them. The user who set the algorithm in motion and who received the report likely has no idea how the robo-traders work or what they “said.”

Intentionless free speech is a listener-oriented theory: it grounds protections in listeners’ experiences rather than in speakers’ intentions. But it is not a human-free theory: without a human somewhere in the loop, there is no cognizable First Amendment interest to assert, because no one’s rights have been infringed. (As in copyright, the day when robots can qualify in their own right as legal persons protected by the Bill of Rights is still a long way off.)

Suppose there was no report at the end of the day, so that no information was presented to a human at all, and no human had an experience of meaning. Collins and Skover’s argument would still go through. The trades were “at the behest of, and in the service of, human objectives.” This should be a sign: the argument doesn’t actually depend on readers’ responses to anything. The norm of utility is not really a *free speech* theory.

VII.

The norm of utility cannot function as what Frederick Schauer calls a Free Speech Principle: a way of deciding what kinds of claims are “speech” claims entitled to special weight and what kinds of claims are not.¹⁶ Schauer’s point, which is profound, is that without some kind of Free Speech Principle, specific free speech arguments collapse into general liberty arguments. Utility is a norm and a virtue, to be sure. But it is not a *free speech* norm or a *free speech* virtue.

This isn’t just a theoretical issue; it’s a deeply practical one. My light-bulb robot is useful: it helps me see my dresser well enough to

put on matching socks in the morning. The ban on inefficient residential incandescent bulbs is, by the norm of utility, a restriction on speech. It is probably a content-neutral restriction, although perhaps a manufacturer could argue that it is content based given the different wavelength output profiles of incandescent, CFL, and LED bulbs. Energy efficiency is a substantial governmental interest that is probably unrelated to the suppression of speech, but is the restriction greater than necessary? That depends on the costs and availability of CFL and LED illumination robots ... and down the rabbit hole we go.

If utility is the “First Amendment lodestar,” then speech eats the world, because anything some human cares enough to do is useful, at least to them. Some physicists, and some mystics, think that the entire universe is literally made up of information. The idea is that what we perceive as matter and forces and everything else is merely the flow and transformation of information from one place and one form to another. The universe, on this view, is a gigantic computer, constantly churning through the computation of a function of cosmic and unbelievable complexity. Something like this is true of Collins and Skover’s conception of speech in the age of robots. Speech is everywhere and everything; it waits only for humans to come along and find it useful.

That would make the First Amendment into “the New *Lochner*,” to use Amanda Shanor’s term: a broad and deep prohibition on government regulation of a wide range of activity.¹⁷ That might or might not be a good thing, but it is not meaningfully tethered to anything we would recognize or care about as “speech.” The argument for maximal liberty has to be made on its own terms, in light of human experience and purposes. It can’t be made by invoking an age-old progressive free speech tradition (as Collins and Skover do in the first third of *Robotica*), because that tradition as a

speech tradition doesn't get us there. If everything is speech, then nothing is.

VIII.

The crux of the contradiction is the relationship between communicative meaning and First Amendment coverage. The *Spence* test is deceptively simple: when a speaker's "intent to convey a particularized message" and listeners' "likelihood ... that the message would be understood"¹⁸ coincide, the First Amendment kicks in. On this view, communicative meaning is a necessary and sufficient condition of First Amendment coverage: if humans intend and experience it as speech, it's "speech" for First Amendment purposes. (We have already discussed the question of whether the speaker's intent is strictly necessary. Collins and Skover argue – I think correctly – that in cases of robotic speech the answer is "no," and listeners' understanding suffices.)

Some scholars accept this view that the First Amendment reaches to the full extent of communicative meaning. Back in 2000, Lee Tien gave a sophisticated account of it in terms of speech-act theory as part of a sustained argument for First Amendment protection for source code.¹⁹ More recently, Stuart Benjamin gave a detailed working out of its consequences for robots.²⁰ And I take this to be an important premise of Jane Bambauer's argument that the collection and sharing of "data" is subject to First Amendment coverage.²¹

Other scholars disagree and think that the First Amendment reaches only a subset of communicative meaning. When Tim Wu says that courts "limit coverage in a way that reserves the power of the state to regulate the functional aspects of the communication process,"²² he is not saying that a self-driving car's left-turn signal communicates nothing and has no meaning. The human or robotic driver of the car behind knows exactly what it means: "The car ahead is about to turn left." But that signal is so closely connected to the

activity of safe driving that failure-to-signal laws are not regarded as restrictions on speech. Driving has great utility – but that still doesn’t provide a compelling reason to treat the turn signal as First Amendment “speech.” In various ways, at various times, and for various reasons, criminal conspiracies, threats, navigational charts, violent conduct, price fixing, and obscenity have all been so classified, despite possessing what any competent reader or listener would regard as communicative meaning.

The norm of utility blows past the limits of this debate: it finds First Amendment coverage even in cases where there is *no* communicative meaning. I do not read Tien, Benjamin, Bambauer – or anyone else, except perhaps John Perry Barlow – to go so far as to say, “It’s all speech now, whether you knew it or not. As long as it’s useful, it’s speech.” Not just the turn signal, but the wiring from the self-driving car’s onboard computer to the turn signal, its laser and GPS unit, and the formulation of the rubber in its tires: they’re all useful, they’re all speech.

It is true that some activities that are never directly experienced by a human – including ones that no human is capable of experiencing – possess communicative meaning and are treated as “speech.” We have no sense organs capable of detecting FM radio waves; our brains cannot easily decode the bits in a Microsoft Word file being emailed from one computer to another. And yet these are unquestionably “speech” for First Amendment purposes. These activities are “speech” because of their close nexus to other activities that do involve communicative meaning. For better or worse, different theories of the First Amendment do the work of explaining what kinds of nexus count and what kinds do not. The norm of utility does not and cannot, because it is completely untethered from the human experiences that distinguish speech from nonspeech in the first place.

IX.

What has gone wrong here is that Collins and Skover's incorrect framing of the question of robotic speech – as “whether?” rather than “when?” – has led them to draw the wrong conclusions from their otherwise insightful analysis of how listeners experience the world. Reader-response is a good answer to the fact-laden question of when humans make meaning from their interactions with robots, but it is a poor answer to the categorical question of whether robotic speech requires a new First Amendment norm. Looking to the receiver's experience correctly distinguishes speech from nonspeech in many cases involving robots. It does not follow that utility is the new “First Amendment lodestar,” because while utility is sensitive to what makes speech valuable, it is entirely indifferent to what makes speech *speech*.

-
- ¹ Charles Babbage, *Passages from the Life of a Philosopher* (London, England: Longman, Green, Longman, Roberts, and Green, 1864), p. 67.
- ² See A. K. Dewdney, “A Tinkertoy Computer That Plays Tic-Tac-Toe,” *Scientific American*, October 1989, p. 120.
- ³ See Nicole Perlroth, “Why Automakers Are Hiring Security Experts,” *New York Times*, 7 June 2017, sec. B, p. 4.
- ⁴ *Spence v. Washington*, 418 U.S. 405, 410–11 (1974).
- ⁵ Lee Tien, “Publishing Software as a Speech Act,” *Berkeley Technology Law Journal* 15: 629 (2000).
- ⁶ *Texas v. Johnson*, 491 U.S. 397 (1989).
- ⁷ *Cohen v. California*, 403 U.S. 15 (1971).
- ⁸ *Erie v. Pap’s A.M.*, 529 U.S. 277 (2000).
- ⁹ James Grimmelman, “There’s No Such Thing as a Computer-Authored Work, and It’s a Good Thing, Too,” *Columbia Journal of Law & the Arts* 39: 403 (2016).
- ¹⁰ *Stanley v. Georgia*, 394 U.S. 557 (1969).
- ¹¹ *Lamont v. Postmaster General*, 381 U.S. 301 (1965).
- ¹² 857 F. Supp. 2d 599 (E.D. Va. 2012).
- ¹³ 730 F.3d 368 (4th Cir. 2013).
- ¹⁴ Robert Post, “Encryption Source Code and the First Amendment,” *Berkeley Technology Law Journal* 15: 713 (2000).
- ¹⁵ See, e.g., Robert Post, “Recuperating First Amendment Doctrine,” *Stanford Law Review* 47: 1249 (1995).
- ¹⁶ Frederick Schauer, *Free Speech: A Philosophical Enquiry* (Cambridge: Cambridge University Press, 1982), pp. 6–7.
- ¹⁷ Amanda Shanor, “The New Lochner,” *Wisconsin Law Review* 2016: 133 (2016).
- ¹⁸ *Spence v. Washington*, 418 U.S. 405, 410–11 (1974).
- ¹⁹ Tien, “Publishing Software as a Speech Act.”

²⁰ Stuart Minor Benjamin, “Algorithms and Speech,” *University of Pennsylvania Law Review* 161: 1445 (2013).

²¹ Jane Bambauer, “Is Data Speech?,” *Stanford Law Review* 66: 57 (2014).

²² Tim Wu, “Machine Speech,” *University of Pennsylvania Law Review* 161: 1495, 1496–97 (2013).