An Open Letter to Higher Education Administrators:

Beware AI Communication Coaches

December 20, 2023

Dear Higher Education Administrators:

We write to warn you about the potential harms of AI communication coaches in academic settings. AI communication coaches are web-based applications that provide users individualized feedback on their speaking performances. We are the students, teaching assistant, and instructor of Dr. Allie Rowland's Fall 2023 Rhetoric of Algorithms course at St. Lawrence University, where we experimented extensively with one leading AI communication coach and looked at several more. The AI communication coach industry has exploded in the past five years to include dozens of companies such as Yoodli, TalkMeUp, Getmee, Orai, and Poised AI. While these platforms differ modestly in features and design, they all claim to offer accessible, actionable, judgment-free feedback on speaking performances. Many of these companies seek expansion into the higher education sector. We find the potential creep of AI communication coaches into college classes so worrisome that we decided to co-author this open letter as a class. **While there may be some limited reasonable uses for them in higher education, AI communication coaches must be approached with caution and should never supplant traditional communication instruction.**

AI communication coaches warrant as much critical attention from academics as generative AI programs like ChatGPT. The academic honesty and student privacy conversations ensuing from generative AI breakthroughs have been important, but have occluded some of the other ways that AI appears in educational contexts, such as AI-driven responsive personalized tutoring. The recent White House executive order on AI safety is an exception to this; in the section on education, the Biden administration wants to explore enabling "personalized tutoring in schools." As many universities experience crisis-level financial pressures, we predict the false promises of AI

communication coaches will tempt several institutions to further deplete and casualize communication faculty members.

We may agree with AI communication coach companies that communication skills are crucial, but we diverge sharply on how communications skills should be taught. Our perspectives as students and teachers at a liberal arts institution are informed by our positive experiences with face-to-face communication instruction. Presenting, active listening, collaboration, conflict management, and other communication skills are critical components of our educational experiences. While St. Lawrence University's leadership-level investment in communication skill development starts in the first year and extends to several speaking-intensive curricular and co-curricular opportunities, our institution is far from alone in valuing these important skills. According to Grobman and Ramsey, graduating college as a competent communicator means the following:

> Being adept at message development and organization and recognizing context and what types of messages best suit each situation. This includes supporting arguments, managing relationships and conflict, enacting appropriate conversational and information exchange skills, communicating in groups, and effectively evaluating the communication of others.

Unsurprisingly, the National Association for Colleges and Employers lists *communication* as one of eight key competencies for a career-ready workforce. Higher education should affirm its ongoing priority to communication competency, not pawn it off to AI communication coaches that rarely make us better communicators.

Argument 1: AI communication coaches do not make us better communicators.

If you were a public speaking instructor, how would you assess the following speech?

> Good morning. For my first point, apples, bananas, grapes. 42 apricots, 17 watermelons. Therefore, I think you can see 14 pineapples. Strawberries, peaches, pears, grapes, bananas, apples, peaches, pineapples. Pineapples again. 42% pineapples compared to 17 strawberries.

Most humans would rightfully assess this monologue as complete gibberish. A leading speech coach, however, scored it at 85% overall–and gave it *perfect scores* in the areas of sentiment, empathy, and filler words. (You'll pardon us if we find the latter, well, bananas.) As a class, we discussed the scary possibility of instructors using scores generated by AI speech coaches as the "official" grades for a student's assignment. AI generated scores are not only faulty, but students would attempt to "game" the

algorithms, resulting in a classic perverse incentive. **When we scored higher, it was not because we became better speakers–rather, we became better at speaking to the AI coach.**

The way we learned to speak to maximize our AI speech coach scores recalls a famous 1931 study involving a monkey and a baby. In an era long before Institutional Review Board protocols, comparative psychologist Winthrop Niles Kellogg adopted chimpanzee Gua and raised him alongside baby Donald, treating them both the same. In the early stages of the experiment, Gua excelled at his cognitive and mobility tests compared to his human counterpart. However, Gua's chimpanzee vocal abilities could not compare to a developing human's. The Kellogg family abruptly ended the experiment, and it is speculated that it's because "while Gua showed no signs of learning human languages, her brother Donald had begun imitating Gua's chimp noises." Like baby Donald "learning" from Gua, when we "learn" from AI speech coaches, we are not learning how to speak with humans.

We will pause our critique to assert that we are not anti-technology, though many attempts to integrate technology in higher education rarely fulfill the promises of improved access, decreased costs, and better learning outcomes. We are concerned that AI speech coaches may be on the brink of what educational technologist Justin Reich cautions is the next "learning-at-scale hype cycle." The Hechinger Report found that new technology often creates additional work for instructors in the long term. For example, online courses are often filled with busywork to ensure that students are reviewing the material and engaging with the content–features that attempt to replicate discussion and critical thinking. More than half of students found technological difficulties to be a source of stress. Universities are uneven, at best, in the tech support that they offer for students. Learning technologies that result in less coaching and individualized feedback from (human) instructors frequently fail.

A pitfall of AI communication coaches is that something is always lost in translation. The AI coaching software tries to quantify human communication, translating it into algorithmic terms, and then providing feedback from an algorithmic point of view. The idea that technology can quantify something as complex as human communication is akin to tech writer Meredith Broussard's notion of *technochauvinism*, or the assumption that technological solutions are superior to other solutions. Adding to the problem is the subjective nature of algorithms and the AI software as they are created by a small group of individuals deciding what qualifies as "good" communication. Tarleton Gillespie highlights the fact that many digital platforms reflect the biases of their (typically white, male) creators, narrowing the complexity of human communication down to a handful of algorithmic quantifications. AI communication coaches reduce

communication to the outdated transmission model and do not account for communication complexities.

We agree with edtech expert Justin Reich:

> Autograders are unevenly useful across the curriculum. They are most useful in the fields in which desired human performance is **sufficiently routine** for algorithms to reliably identify the features of high-quality and low-quality performance and to assign grades and scores accordingly. [...] **Much of what we want students to learn, however, cannot be demonstrated through performances that adhere to these kinds of rigid structures**.

Reich's last statement above should refocus us on the "why" of teaching communication skills. Even components of communication that appear relatively quantifiable by an AI speech coach, such as eye contact or speaking rate (often measured in words per minute) are subject to context collapse. (Notably, the AI speaking coaches that we looked at are only capable of measuring eye contact *with a camera* and not eye contact with a human.) The myriad factors that determine something like appropriate and effective eye contact or speaking rate (audience, occasion, purpose, public address versus interpersonal conversation, etc.) will always exceed what Reich refers to as the "rigid structures" of autograders. AI speech coaches will never fully comprehend the subtle, artistic aspects that give human communication its unique flavor. As students and teachers of communication skills, being and becoming rhetorically sensitive interpreters of contextual complexity are the most important skills–and they are skills that AI speech coaches simply cannot teach.

When AI speech coaches attempt to quantify unquantifiable components of communication, the results are much worse than just a confusing and glitchy product–the results, in fact, can be quite harmful. Take, for instance, one student's experience with the *sentiment* assessment algorithm. In the context of a speech on algorithmic racism, the student stated that content moderation practices on major platforms often lead to an increase in "harmful ideas like white supremacy." Because of this and similar statements, the student scored 20% on *sentiment* and was informed: "Good effort. However, your choice of words created a largely negative tone." If sentiment is defined as "a view of or attitude toward a situation or event," then apparently condemning white supremacy is not cheerful enough for the AI speech coach. Because proprietary algorithms are black boxes, we do not know the precise terms that trigger low sentiment scores. However, we do know that many aspects of the world are *harmful*, *terrible*, *tragic*, *awful*, *devastating* (these are all words we expect the program would flag for negative sentiment)--and failing to name them as such only perpetuates these conditions. Not only does the AI communication coach fail to measure

sentiment despite its claims to measure sentiment, but it also risks teaching people to stop telling the truth about the world.

Similarly to *sentiment*, AI speech coaches disregard the human experience of *empathy*, framing it instead as a strategic method of inducing compliance. We suspect that the popularity of Brene Brown's work drives this focus on empathy, so we will use her definition here: Empathy is connecting with the emotion that someone else is feeling. Some AI speech coaches refer repeatedly to empathy in their marketing materials but then do not attempt to measure or provide feedback for empathy. For the AI speech coaches that evaluate empathy, as far as we can tell, they rely on key phrases such as "I heard you say.." while paying lip service to the idea that empathy can be communicated through visual, verbal, and paralanguage channels. Rather than fostering genuine empathy, AI speech coaches offer shortcuts to fake a compelling empathy performance.

Argument 2: Not only do AI communication coaches fail to solve the so-called problem of communication instruction at scale, but they also create a new set of problems.

In their marketing materials, AI communication coach companies position themselves as accessible and equitable. Nothing could be further from the truth. Like most programs that rely on facial recognition technology, speech-to-text software, and algorithmic reduction of complex practices, AI speech coaches are racist, ableist, classist and further inculcate students into regimes of digital surveillance. The Matthew effect, a principle that asserts that the most advantaged tend to accumulate more advantage, is applicable in educational technology contexts: "[W]hen researchers evaluate how learners from different backgrounds access and use new technologies, it is common to find that the benefits of new technologies–even free technologies–accrue most rapidly to the already advantaged." For example, the students who tend to benefit the most from autograded tutors are students who have already developed self-regulation and higher-order executive skills in traditional educational settings.

Internal data on the extent to which AI speech coaches are biased is closely guarded; however, it is well documented that the error rates for facial recognition technology (on which AI speech coaches rely) are higher for people with darker skin and women. In her body of work that advocates for algorithmic justice, Joy Buolamwini describes *the coded gaze* as the way that racial biases are programmed into algorithmic systems. We follow Ruha Benjamin's assertion that algorithmic racism is a feature, not a bug. In other words, large digital companies would like us to believe that their products are only *incidentally* racist and that they should not be held accountable for discriminatory results. As a diverse community at a liberal arts college, these rationalizations simply do not meet our standards.

We also take issue with the growing surveillance capacities of digital products such as AI speech coaches. When students use an AI communication coach, they may not realize that several companies will own their video data forever. Almost all digital products conform to the principles of surveillance capitalism, an economic system built on the secret extraction and manipulation of human data. In the "Terms of Service" fine print for nearly all AI communication coaches, video data is collected, stored, and used to retrain its models. In other words, the AI communication coach company now owns your data. However, it would be a mistake to consider this a closed system. Rather than write established code from scratch, many AI speech coaches outsource to big companies. For example, one leading AI speech coach uses Amazon's speech-to-text service and likely stores this data in Amazon's data lake. Privacy at the Amazon data lake is ambiguous: "We will not access or use Your Content except as necessary to maintain or provide the Services." Since the statement "as necessary to maintain or provide the Services" is broad and subjective, we do not rule out the possibility of Amazon using this data to contribute to surveillance capitalism. When AI communication coaches are inevitably integrated into learning management systems such as Canvas, yet another company whose surveillant data privacy practices warrant scrutiny is introduced into the mix.

Despite our objections, there may be a few circumstances in which AI communication coaches could supplement a well-rounded college experience. For example, in an introductory communication course, we could imagine AI speech coaches serving as an interesting case study to help students understand the nuances between the linear, interaction, and transaction models of communication. In a course like ours, with learning goals to understand the ways that algorithms rhetorically influence culture and society, the chance to reverse engineer an AI speech coach has been invaluable. Rather than use the AI speech coach to attempt to get better at communicating, we are more interested in unpacking the implicit assumptions about human communication smuggled into these products. Finally, students are increasingly likely to go on job interviews where potential employers use AI to assist decision making. Though we have equity concerns about this practice, it behooves college students to be familiar with how AI communication coaches assess typical performance indicators nonetheless.

In sum, we predict a rise of AI communication coaches being integrated into college classes in the next ten years, fueled by promotional campaigns that claim the products are effective, accessible, and equitable. We caution against this rise with deep skepticism regarding the utility of AI speech coaches and mounting horror regarding their capacities to harm.

Warm Regards,

Allie Rowland, Instructor
Jack Benjamin, Teaching Assistant

The students of the Rhetoric of Algorithms Fall 2023 class at St. Lawrence University:

Markus Basciano
Lucas Boyden
Promise Butler
Drew Chase
Alyana Contant
Michael DiStefano
Sam Fosberg
Kyle Graber
Maryam Konateh
Cooper Phillips
Jed Saltus
Emily Siansky
Sophia Sulkin
Ellie Wyckoff