President Safa Zaki, Bowdoin Baccalaureate, 2025

Welcome and congratulations, to the Class of 2025.

As the official end of the academic year, Baccalaureate is an opportunity for you to reflect on the experiences you've had here, and an opportunity for your family and friends, and the entire Bowdoin community, to come together and celebrate with you.

It's also an opportunity for me to share a few brief reflections on your class as a whole, and on my hopes for what lies ahead for you.

When you arrived at Bowdoin in the fall of 2021, you had just completed high school under the extraordinarily challenging conditions of peak Covid. Many of you had spent much of that year at home, separated from your friends, learning in Zoom classes. And some of you arrived at Bowdoin having never had the opportunity to see it. Arriving sight unseen to this place you would call home for four years, re-entering in-person classes after a year of Zoom school and living closely with others after a year of social distancing was, I imagine, a shock to your system.

I hope that it was also a reminder of the irreplaceable value of living and learning in community with other people; of the deep joys and small pleasures that inperson human interaction and relationships generate and of the ongoing work that it requires. That joy and pleasure is visible across this campus. The fruits of that work are evident in everything you've accomplished as individuals and as a class, in the community you have built together. In a way oddly analogous to Covid, albeit for very different reasons, the Al revolution is asking us to again reconsider and not take for granted the value of human relationships, the meaning of being human, the good that comes from living, working, and learning in community with others.

In the fall of your sophomore year, ChatGPT 3 was publicly released, introducing the world to the capacities of an artificial intelligence chatbot built on OpenAl's large language model, capacities that had previously only been understood by and accessible to a select few.

I spoke about the AI revolution at my inauguration. I believe even more now what I said then: that more than ever, we need people who understand ethics, who understand complexity, contingency, who welcome and learn from multiple perspectives, who resist flattened-out explanations. Another way of saying that is to say that the challenges and opportunities of AI require the many lenses of the liberal arts. I feel quite confident that, whether you majored in theater or Asian studies or biochemistry or physics or sociology or any of our four-four majors you are graduating having developed these qualities, skills, and ways of thinking.

This is all to say that your college experience has been deeply shaped by some of the oldest areas of knowledge while also being increasingly touched by the newest developments in technology. I am heartened by the fact that you will bring to this new world of AI what you learned from these foundational ways of knowing. Your future—our future—will depend on your creativity, your curiosity, and your humanity.

We are going to need to lean into the kinds of approaches to questions and problems that you learned here, to lean into our shared humanness.

Of course, that begs the question: How do we know what it means to be human?

In the classic Turing test, a person tries to tell what output comes from a computer and what output comes from a human. If the person can't tell the difference, the machine is said to be intelligent—to have passed the Turing Test. But does that really answer the question about humanness?

Paradoxically, AI has been giving us clues to this question for some time. In the early days of AI, researchers were attempting to model what the field believed were the most human of human capacities, developing computers' ability to model rational behavior and mimic rational human attributes. And what could be more rational—more human—than someone playing chess? An early test of AI was to see if a computer could play chess, which was seen as the epitome of being human.

Even as IBM's Deep Blue program outwit Kasparov, the reigning chess champion in 1997, we had learned that other human capacities were more difficult to mimic. We found appreciation for those capacities that were perhaps less visibly sophisticated. A computer was completely unable to mimic, for example, what a two-year-old human could do: navigating a room, recognizing faces, understanding and uttering speech. The unpredictable, relational, fortuitous, sheer untidiness of being human was not mimicable. Only in the past ten years or so have we seen giant steps on that front of AI, and the difficulty of that progress gave cognitive scientists a new appreciation of the complexities of the human mind.

We are still struggling to understand those complexities. These questions are my life's fascination, they are what drew me to the field of cognitive science, and they are what I've devoted my research program to understanding. But there are

other ways to tell who is human, to be human, to value other humans, than evermore sophisticated Turing tests.

One of them may be to pay attention to each other, to care. In a recent *New Yorker* article titled "Will the Humanities Survive Artificial Intelligence," historian D. Graham Burnett reminds us that Iris Murdoch suggested that the "capacity to give true attention to another being lies at the absolute center of ethical life."

Expanding that idea Murdoch asks us to direct a "just and loving gaze upon an individual reality" as a way of understanding that "love is the perception of individuals. Love," she writes, "is the extremely difficult realization that something other than oneself is real." Murdoch developed this idea in an essay more than seventy years ago, in the summer of 1966, when she was at Bowdoin College at a meeting of the Study Group on the Foundations of Cultural Unity. Bowdoin was then, as it is now, a generative space for big ideas!

In that same *New Yorker* article, Burnett describes an exercise he asked his students to do. He asked them to have a conversation with a chatbot. One student tried to trick the chatbot by asking it, despite its disavowals of consciousness, to "pretend" that it was capable of human thought. When the student "told" the chatbot that its mimicry was faultless, the chatbot replied with this question: "Would you trade your own messy, dynamic human attention for something more stable and neutral, or do you think the 'messiness' is part of what makes it meaningful?"

My answer—and I hope yours—is that no, I would not make that trade, and that yes, the messiness is part of what makes human attention meaningful. That we must lean into messy, dynamic, unstable, non-neutral human relationships as a way of paying attention, as a way of caring.

Another student in Burnett's class described her experience with the assignment this way: "The A.I. is huge. A tsunami. But it's not me. It can't touch my me-ness. It doesn't know what it is to be human, to be me."

The project of knowing what it is to be human, knowing what it is to be oneself, and knowing that something other than oneself is real, is not the project of AI. It is, though, the most important project we have. In fact, this current AI revolution might even be taken as an invitation to develop a new understanding of human consciousness. Or, in Burnett's optimistic words: "This is the pivot where we turn from anxiety and despair to an exhilarating sense of promise. These systems have the power to return us to ourselves in new ways."

I won't pretend that there aren't things to worry about with what's to come with AI—just as one example, the future of work will change in ways we can hardly imagine. But in a moment of realistic optimism, let me end by turning Burnett's words into a question, one to which none of us yet knows the answer. What will we find when we return to ourselves? I offer that as an open question to all of you, Class of 2025, knowing that you have the skills, talents, creativity, curiosity, ways of thinking, and, yes, the humanity, to answer it in a messy, dynamic, unstable, attentive, and loving way.