AI Transparency Statement

Vision

We envision a world where the responsible use of artificial intelligence (AI) empowers educators and students to help every learner succeed.

Teaching and learning are challenging. Like any technology, AI should make those tasks easier, but it must be used responsibly and ethically in order to realize its profound benefit to education. As we continue to discover more research-driven opportunities utilizing AI to address instructor and student challenges, we are committed to upholding our seven guiding principles.

1. We are transparent and accountable in how we ethically use AI.

2. We keep humans at the center of AI oversight and control.

3. We educate people about AI safety and the risks of unethical use of AI.

4. We strictly observe data privacy and security principles.

5. We do not introduce or reinforce biases or inequalities.

6. We continuously learn from experts to improve AI safety and ethics.

7. We use research-driven principles that support the human learning experience in our products and services.

Our Commitment to Using AI in our Educational Products

1. We are transparent and accountable in how we ethically use AI.
Our commitment to transparency is foundational, ensuring clarity about the AI tools in our educational products. We approach the use of AI in our products holistically across the organization bringing in expertise and voices from Product, Data Science, Learning Science, Marketing, Sales, Editorial, and Legal to align our AI practices with industry standards, as well as legal requirements.

2. We educate people about AI safety and the risks of unethical use of AI.
Empowering individuals with the knowledge to use AI safely, ethically, and responsibly is key to how we envision the use of AI in our products. We provide valuable education through public webinars, internal training, and open-access AI communities. These resources help both our team and our customers confidently understand and engage with AI technologies. We will continue to evaluate and update our ethics training to keep pace with the quickly evolving AI landscape.

Please note, all of our product features leveraging AI are in support of student learning, instructor teaching, and are designed to only be used for the specific educational purpose provided, like tutor support. We do not provide any generative AI product features designed to encourage student cheating. We also take measures to mitigate users circumventing our proprietary AI prompts.

3. We keep humans at the center of AI oversight and control.

We are dedicated to keeping human oversight at the forefront of our AI development, reflecting our commitment to ethical technology use. When AI is used in our product, we provide feedback mechanisms and monitoring for users to inform us of AI performance. We leverage in-house subject matter experts in the disciplines we support to regularly review AI responses and performance and we periodically brief our AI advisory board on AI performance.

4. We strictly observe data privacy and security principles.

We protect user information with comprehensive data privacy and security measures. Our comprehensive privacy and security programs reflect industry standards and legal requirements and are assessed and updated regularly.

For details about our data handling practices, please see our Privacy Notice.

See below in AI Features in our Products for additional data handling details for each applicable listed product.

5. We do not introduce or reinforce biases or inequalities.

We have established a robust and rigorous framework to guide our product and technology teams in the ethical implementation of AI tools and mitigate unintended consequences in educational settings. We have and will continue to develop evaluation tools to ensure that AI responses remain accurate and fair. We conduct thorough human reviews of statistically sampled AI outputs to refine our systems continuously. We will thoroughly assess the value and risk of proposed AI features to our product with proofs of concept testing to gather real world feedback from educators and students to inform future product development.

6. We continuously learn from experts to improve AI safety and ethics.

As a learning organization, we are committed to continuously learning about how to improve our use of AI from the voices of instructors, students, institutional representatives, data privacy experts, and AI safety advocates. We have used and will continue to use what we learn to establish best practices in supporting student learning with our products.

7. We use research-driven principles that support the human learning experience in our products and services.

AI Features in our Products

I. AI Tutor

What is the AI tutor?
The AI Tutor available in Macmillan Learning's ACHIEVE platform provides students with assistance on homework problems in ACHIEVE. The AI Tutor uses generative AI to generate responses to questions sent by the student to the Tutor. The tutor is designed to guide students through their questions by stimulating critical thinking, self-reflection, and self-driven learning. Before students attempt a homework question, the AI Tutor will offer limited assistance to help them get started. When students make an incorrect attempt to answer a homework question they can then invoke the chat feature of the AI Tutor, which enables students to chat with the tutor to get help with their specific misconceptions on specific concepts in the homework question.

Is the AI Tutor feature deployed by default?
No, the instructor must turn on the AI Tutor in his or her course in Achieve to activate it and the student can choose whether or not to use the chat feature in the tutor.

Does the AI Tutor rely on third party systems?
Yes, the AI Tutor currently leverages OpenAI's GPT family of models. OpenAI describes its data privacy and security practices in the OpenAI Security Portal at https://trust.openai.com/ and Enterprise Privacy at OpenAI at https://openai.com/enterprise-privacy. We may consider other models over time, which will conform to similar behavior and safeguards.

How is the AI Tutor initially trained?
See https://openai.com/index/gpt-4o-system-card/#model-data-and-training. Note: Macmillan Learning's data has not been used to train the LLM.

What customer personal data is input into the AI Tutor?
The AI Tutor is designed to work without any personal data from the user. We provide notice to users not to input personal information into the chat.

Is customer data from the use of the AI Tutor used to improve the third party large language model (LLM)?
No, the third party AI system is not trained on our customer data for its own purposes.

What data does the third party LLM have access to?
To provide an effective tutoring experience, the AI Tutor sends the content of the question and student submission data to a third party LLM provider, none of which contains personally identifiable information.

How will Macmillan Learning use the data from the AI Tutor?
In addition to using the data to provide the feature to the user, Macmillan Learning may analyze the student messages in de-identified form to improve the AI Tutor and other Macmillan Learning products. Additionally, we maintain audit capabilities that allow us to monitor responses input into the tutor for accuracy and bias. For additional details about our data privacy practices, please see our Privacy Notice.

Where is the data processed by the AI system stored?
All OpenAI servers are US based.

How is the AI Tutor evaluated for accuracy, reliability and fairness?
Our internal reviews by subject matter experts indicate that the AI Tutor provides helpful responses and learning assistance to students. We ensure accuracy of the responses by grounding the tutor in our own assessment content, which serves as the primary source for the AI Tutor's responses to the students. However, due to the nature of LLM technology, the tutor may occasionally produce inaccurate or biased information. To reduce inaccuracies and bias, we have implemented both human and automated evaluation methods to randomly review a sampling of the AI Tutor's interactions with students, and we carefully evaluate external benchmark evaluations of model accuracy, focus on instructions, and bias.

II. iClicker Question Creator

What is the AI Question Creator?
iClicker's AI Question Creator utilizes generative AI technology to assist educators in creating engaging and effective questions for their students.

Is the AI Question Creator deployed by default?
The feature is available in all iClicker instructor accounts, but its use is optional and must be actively selected by instructors.

Does the AI Question Creator rely on third party systems?
Yes, the Question currently leverages OpenAI's GPT family of models. OpenAI describes its data privacy and security practices in the OpenAI Security Portal at https://trust.openai.com/ and Enterprise Privacy at OpenAI at https://openai.com/enterprise-privacy/. We may consider other models over time, which will conform to the similar behavior and safeguards.

How is the AI Question Creator initially trained?
See https://openai.com/index/gpt-4o-system-card/#model-data-and-training. Note: Macmillan Learning's data has not been used to train the LLM.

What customer personal data is input into the AI Question Creator?
The Question Creator is intended not to require any personal data from the user.

Is customer data from the use of the AI Question Creator used to improve the third party LLM?
No, it is only used to improve Macmillan Learning's product or AI features. We do not train any third party AI system with our customer data.

What does the third party LLM have access to?
To provide an effective question creation experience, the Question Creator sends the following pieces of data to the third party LLM provider, none of which contains personal data:

☐ A proprietary set of instructions and rules to guide the LLM in how it should respond to the request.

☐ The content of the instructors' requested topic and question style preferences for question creation.

How will Macmillan Learning use the data from the AI Question Creator?
In addition to using the data to provide the feature to the user, Macmillan Learning may analyze question creator submissions and questions created in de-identified form to improve the AI Question Creator and other Macmillan Learning products. Additionally, we maintain audit capabilities that allow us to monitor responses input for accuracy, prevention of malicious use, and bias. For additional details about our data privacy practices, please see our Privacy Notice.

Where is the data processed by the AI system stored?
All OpenAI servers are US based.

How is the AI Question Creator evaluated for accuracy, reliability and fairness?
Our internal reviews indicate that the AI Question Creator provides helpful question and answer content for instructors to utilize for engaging students. However, due to the nature of LLM technology, it is possible the Question Creator may occasionally produce inaccurate or biased information. To reduce inaccuracies and bias, we have implemented evaluation methods to randomly review a sampling of the questions created and review any questions flagged by instructors for possible tool improvement. We carefully evaluate external benchmark evaluations of model accuracy, focus on instructions, and bias. We recommend all questions created be evaluated by the instructor to ensure fit for their instructional needs.