




The Responsible AI Advantage:

# Grammarly's Guidelines for Ethical Innovation



# Table of Contents

<b>Introduction</b>	<b>2</b>
<hr/>	
<b>Chapter 1   Pillars of Responsible AI</b>	<b>3</b>
1.1 Upholding Transparency	4
1.2 Championing Fairness and Safety	5
1.3 Fostering User Agency	8
1.4 Embracing Accountability	9
1.5 Preserving Privacy & Security	11
<hr/>	
<b>Chapter 2   Additional Considerations When Designing a Responsible AI Approach</b>	<b>13</b>
<hr/>	
<b>Chapter 3   Grammarly's Unique Approach To Responsible AI</b>	<b>14</b>
<hr/>	
<b>Conclusion</b>	<b>15</b>
<hr/>	



# Introduction

**In our rapidly digitizing world, businesses increasingly turn to artificial intelligence (AI) to improve their professional workflows and communication. This widespread adoption has come with new risks including bias, social manipulation, and ethical dilemmas. As a result, the ethical and responsible deployment of AI systems has never been more critical. Companies must take the responsibility to mitigate these risks seriously to protect their employees, data, and brand.**

Responsible AI goes beyond simply creating effective and compliant AI systems; it's about ensuring these systems maximize fairness and reduce bias, promote safety and user agency, and align with human values and principles. According to Grammarly's [2024 State of Business Communication](#) report, more than half of all professionals are concerned about AI issues including privacy, security, quality control, and bias. For chief information security officers (CISOs), implementing a responsible AI practice is a strategic imperative to ensure the safety and effectiveness of this new technology within their organization. This means that leaders must proactively address the risks and vulnerabilities of AI, earn and foster user trust, and align their AI initiatives with both broader organizational values and regulatory requirements.

Grammarly's Responsible AI (RAI) team thinks critically about what we are trying to achieve for our users, our product, and our company and then makes both technical and brand decisions best suited to those needs. While developing our product, we recognized a clear need to

create a dedicated, centralized team to focus on these issues. In addition to having a team of expert analytical linguists, researchers, and machine learning (ML) engineers, our RAI team also provides in-house security expertise focused on AI risks and vulnerabilities.

To help security leaders safely and ethically deploy AI, we are excited to share Grammarly's responsible AI framework detailing our unique approach to building AI responsibly. Our hope is that the processes and principles that we have established to ensure our AI systems are safe, fair, and reliable will inspire and guide you to implement your own responsible AI practices. With these guidelines for ethical innovation, you will be better positioned to enhance the AI capabilities throughout your organization, fortify your security posture, uphold the highest ethical standards, and gain an advantage in today's competitive market.



# Pillars of Responsible AI

Designing a responsible AI framework starts with identifying a business's core drivers and values. This framework should center on the user's experience with the AI product and its outputs, ensuring an added layer of protection between the AI-generated solutions and the human user. Organizations must also implement proper privacy and security measures before deploying AI systems to handle sensitive and confidential business data.

The first objective of any responsible AI mission should be to create ethical development principles that developers should strive to follow for every AI product and user interaction. Once these guidelines are in place, businesses can begin to create processes such as an evaluation framework, a risk assessment process, an escalation process, and an issue-severity rating system.

At Grammarly, the heart of our product has been the intention to improve communication for people worldwide. So when setting out to define our guiding principles for responsible AI, we began with our commitment to safeguarding users' thoughts and words. We then considered a range of industry guidelines and user feedback, consulting with experts to help us understand how people communicate and language issues our users were potentially facing. This baseline assessment of industry standards and best practices helped us to determine the boundaries of our programs and establish the pillars of our responsible AI guiding principles. Since we're in the business of words, we make sure to understand how words matter.

The defining pillars we landed on—**transparency, fairness, user agency, accountability, and privacy and security**—detailed in the next section, represent the main themes from that work and what we use as a North Star to guide everything we build.



Transparency



Fairness and safety



User agency



Accountability



Privacy and security

### We consulted the following resources to shape our own responsible AI guidelines:

- Feedback from users
- A subject matter expert who specializes in sociolinguistics to help Grammarly avoid suggesting offensive language in its communication-based product
- Academic literature
- Language policies of other companies
- Widely trusted style guides (e.g., Associated Press, Chicago Manual of Style)
- Government discrimination and language policies
- Federal Communications Commission and TV/radio broadcast guidelines
- Inclusive language guidelines from various nonprofit organizations





## Upholding Transparency

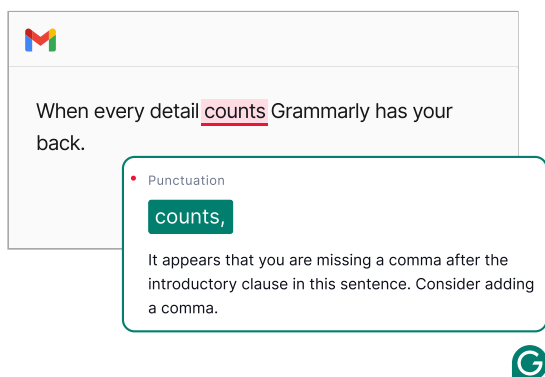
Transparency in AI usage and development is crucial for fostering trust among users. When the workings of an AI system are clear, it is easier to hold developers and organizations accountable for the system's decisions. Additionally, many regulatory frameworks and ethical guidelines require a certain level of transparency in AI systems. Being transparent helps organizations comply with these regulations and avoid legal repercussions.

**There are a few considerations to keep in mind when aiming to offer AI centered on transparency:**

### User awareness

It should always be clear to users when they are interacting with AI. This includes being able to identify AI-generated content and distinguish it from human-generated content.

In addition to knowing when an interaction is rooted in AI, stakeholders should understand how the AI system arrives at its decisions. When a system is transparent, users can better interpret the rationale behind its outputs and make appropriate decisions about how to apply them to their use cases, which is especially important in high-stakes areas like healthcare, finance, and law.



Grammarly empowers users to exercise control over the output of AI-powered systems. Users can choose to accept or disregard the suggestions presented to them. We include explanations so that users can make an

informed decision about whether they want to include a suggestion in their writing. Users can also control the types of suggestions they want to receive in their settings. Our commitment is to empower every user to use AI to express themselves in the most effective way possible.

### System development and limitations

Users should understand any risks associated with the model. This involves clearly identifying any conflicts of interest or business motivations to demonstrate whether the model is objective and unbiased. Competing interests of developers or stakeholders may skew a model's design or outputs, leading to unfavorable outcomes.

Building AI with this level of transparency can enhance public confidence in the technology. When people understand how AI systems work and see the efforts to make them transparent, they are more likely to support and adopt these technologies.



**Grammarly makes money by selling subscriptions, not by relying on an ads-based model. This means our incentives are aligned with providing our users the highest-quality experience possible.**

### Detailed documentation

Detailed information is critical to achieving user awareness and articulating AI risks. Developers must document the capabilities and limitations of the systems they create, including the data used and training methodologies. Organizations should provide users with comprehensive information on how the AI works, its potential risks, and appropriate usage guidelines, communicating about known issues and the ongoing efforts to mitigate them through research and updates.

Perhaps most essential, disclose how user data is being used, stored, and protected. This is particularly important in scenarios where AI uses personal data to make or influence decisions. Grammarly details this in our technical specifications and by giving all users, free and paid, the ability to opt out of having their content used for model training.



## Championing Fairness and Safety

AI systems should be designed to produce quality output and avoid bias, hallucination, or other unsafe outcomes. Organizations must make intentional efforts to identify and mitigate these biases to ensure consistent and equitable performance. By doing so, AI systems can better serve a wide range of users and avoid reinforcing existing prejudices or excluding certain groups from benefiting from the technology.

Safety not only includes monitoring for content-based issues; it also involves ensuring proper deployment of AI within an organization and building guardrails to holistically defend against adverse impacts of using AI. Preventing these types of issues should be top of mind for developers before releasing a product to users.

---

Some of the steps an organization can take to ensure fairness and safety include:

### 1 Create sensitivity guidelines

One way to build safety into a model is by defining guidelines that ensure the model stays aligned with human values. Grammarly is committed to building AI products that are inclusive, safe, and free of bias. This includes [biases in the data](#) used to train the AI and biases in the algorithm and outputs. To

test for bias, consider whether a model might generate a different output based on different demographic information in the input. If so, it's necessary to introduce diverse and representative training data that accurately reflects various contexts.



Grammarly's product offerings should;



Support and use  
inclusive language.



Avoid creating or  
perpetuating bias.



Never suggest offensive  
or unsafe content.

## 2 Design a risk assessment process

When launching new products involving AI, all features should be assessed for risks using a clear evaluation framework. This helps to prevent the feature from producing biased, offensive, or otherwise inappropriate content and evaluates for potential risks related to privacy, security, and other adverse impacts. Teams should work to understand the implications of the

specific AI feature on users, in society, and within various contexts and use cases for a holistic view of the implications of the feature and potential for abuse. Features should be categorized by risk profile, and then an appropriate determination can be made about any changes that it may need to undergo before being launched.



### Grammarly's Risk Assessment Process

While some responsible AI teams may focus primarily on research or policy decisions, our team is very hands-on with the product. We have developed a sophisticated machine learning (ML) solution combined with human review to vet new features, and we are involved with every feature that is shipped. At Grammarly, every feature must undergo our risk assessment process and receive

approval from the Responsible AI team before launching. Each feature is categorized as “low risk,” “high risk,” or “some risk” according to criteria developed by our team. User-facing launches are not approved unless the team developing them implements the recommended mitigations before launch.



### 3 Stand up guardrails

Once a feature has been evaluated, determine and implement guardrails that address the identified risks. Teams should plan to continuously monitor the feature for new risks, abuses, and unforeseen use cases. User feedback is another way to monitor for issues and ensure continuous improvement. Investing in tools to detect harmful content is crucial for mitigating risks going forward, providing a positive user experience, and reducing the risks of brand reputation damage.

Filtering AI model output based on determined principles is one method to keep users from experiencing harmful behavior. Content should be reviewed both algorithmically and by humans to comprehensively detect offensive and sensitive language.

Beyond filtering, models should be taught how to respond in safe ways, meaning the model should be designed and trained in a manner that consistently

avoids actions that could lead to harmful, undesirable, or dangerous outcomes. Teams must first clearly define what constitutes safe and unsafe actions. They can then train the model through reinforcement learning, employing positive rewards and feedback that incentivize safe behaviors, negative rewards for risky actions, and mechanisms to intervene when the model demonstrates unsafe behaviors. Continuing to fine-tune the model reinforces harm reduction by demonstrating what a safe output compared to a harmful output looks like and by ensuring the model engages with users in safe ways. Models should be taught to refuse certain types of user requests and respond appropriately when prompted with an inappropriate request to avoid risky interactions.



#### Seismograph

Grammarly's Responsible AI team has developed a proprietary tool called Seismograph, which is used to ensure our products adhere to our sensitivity guidelines. Seismograph combines ML-classifier-based and dictionary solutions. As its name suggests, Seismograph helps us to detect tremors in language anomalies and minimize the potential damage delicate text might otherwise cause.

We maintain a taxonomy of delicate text and use expert annotators to label data accordingly. This annotated data is used to train a model to recognize instances of text that may be considered emotional, personal, charged, or referencing a greater number of sensitive topics. We are continually updating Seismograph to keep up with lexical shifts and improve performance.



## Fostering User Agency

At Grammarly, we believe that AI should enhance a team's skills while respecting personal autonomy and amplifying every team member's intelligence, strengths, and impact. Users should always be in control of their experience when interacting with AI.

People are the ultimate decision-makers and experts in their own business contexts and with their intended audiences, and they should also understand the limitations of AI. They should be empowered to make an appropriate determination about whether the output of an AI system fits the context in which they are looking to apply it.

An organization must decide whether AI or a given output is appropriate for their specific use case. For example, a team that is responsible for loan approvals may determine that they do not want to use AI to make the final call on

who gets approved for a loan, given the potential risks of removing human review from that process. However, that same company may find AI to be impactful for improving internal communications, deploying code, or enhancing the customer service experience. These determinations may look different for every company, function, and user, which is why it's critical that organizations build or deploy AI solutions that foster user agency, ensuring that the output can align with their organization's own guidelines and policies.

**Building on our commitment to helping users transparently and responsibly use AI, Grammarly offers several tools to equip users with the information necessary to make appropriate judgments about the use of AI for their specific purposes.**



### AI Detection Capabilities

Grammarly's AI detection capabilities can evaluate the originality of selected text by assessing how much of the content appears to be AI-generated. Our AI detection also provides guidance on how to interpret and act on the results responsibly. This gives users an opportunity to appropriately attribute sources, rewrite content, and mitigate the risk of being incorrectly accused of AI plagiarism.

Because AI detection is still an emerging and imprecise technology, our tool provides a percentage score for the amount of AI-generated text that is likely contained in the writing rather than a definitive "yes" or "no" declaration. We prioritize transparency with this tool, explaining the limitations of AI detection technology available today. We recognize that there are a number of reasons text could potentially be flagged for the presence of AI-generated content, even when generative AI tools have not been used. For example, human-generated writing may exhibit patterns, structures, or characteristics of typical AI-generated content, such as repetitive phrases or unconventional grammar.



### Authorship

Authorship is a Grammarly tool (currently in beta testing with select customers) that enables users to track all text added to their writing and generate a comprehensive report detailing the content's origin—be it human-typed, AI-generated, edited with traditional grammar checking, or sourced externally. This tool signifies our dedication to promoting responsible AI use and empowering users with tools for accountability. These initiatives play a crucial role in ensuring that technology augments human effort without diminishing the value of original work, fostering responsible AI usage across industries.





## Embracing Accountability

Accountability does not mean zero fallibility. Rather, accountability is the commitment to a company's core philosophies of ethical and responsible AI. It's about more than just recognizing issues in a model. Developers need to anticipate potential abuse, assess its frequency, and pledge to take full ownership of the model's outcomes. This proactive approach helps ensure that AI aligns with human-centered values and positively impacts society.

In addition to Grammarly's rigorous risk assessment process, outlined in an earlier section, Grammarly robust Support team that collects instances of adverse model outcomes. This allows us to review and iterate on incorrect or harmful information that our product might generate. We understand that our users rely on us to ensure that the suggestions are helpful, and our team works hard to ensure that we take ownership and accountability when we get it wrong.



### Grammarly's Approach to Accountability

Grammarly has demonstrated a strong commitment to accountability. One such example is by actively **acknowledging** and addressing gender bias in our autocorrect features, the learnings of which we shared openly on our company blog. Keenly aware of the issue of ML models perpetuating bias, especially in how corrections can differ based on the gender of the nearby pronoun, our RAI team conducted investigations and devised specific metrics to assess gender bias in our models. We retrained the model with sentences that demonstrated diverse pronouns and misspelled words and improved the model's contextual awareness by incorporating

more information about physical keyboard layouts and edit distances, which enhanced the quality of the correction. In doing so, we successfully identified and quantified the bias present in autocorrect suggestions and reduced the rate of gender pronoun bias by 98%

Grammarly continuously performs regression testing to monitor even the smallest changes in our model. This commitment to ongoing improvement reflects Grammarly's accountability in designing AI that supports fair and unbiased communication for all users.

### Product and engineering teams should adhere to the following principles to embrace accountability and promote responsible AI usage:

#### 1 Test for weak spots in the product.

Perform offensive security techniques, bias and fairness evaluations, and other pressure tests to uncover vulnerabilities before they significantly impact customers.

#### 2 Identify industry-wide solutions

such as open-source models, that make building responsible AI easier and more accessible. Developments in responsible approaches help us all improve the quality of our products and strengthen consumer trust in AI technology.

### 3 **Embed responsible AI teams across product development.**

This work can fall through the cracks if no one is explicitly responsible for ensuring models are safe. CISOs should prioritize hiring a responsible AI team and empower them to play a central role in building new features and maintaining existing ones.

---

#### **Upholding accountability at all levels**

Companies should establish clear lines of accountability for the outcomes of their AI systems. This includes mitigation procedures to handle any AI errors, misinformation, harm, or hallucinations. Systems should be tested to ensure that they are functioning correctly under a variety of conditions, including instances of user abuse/misuse, and should be continuously monitored, regularly reviewed, and systematically updated to ensure they remain fair, accurate, and reliable over time. Only then

can a company claim to have a responsible approach toward the outputs and impact of its models. One of the ways we hold ourselves accountable at Grammarly is by giving our users the option to report offensive and incorrect suggestions in product. These reports, along with any issues reported to customer support and social, are reviewed by our RAI team, prioritized, and fixed, as appropriate.





## Preserving Privacy and Security

Our final, and perhaps most important, pillar of responsible AI is the importance of upholding privacy and security to protect all users, customers, and their companies' reputations. In Grammarly's [2024 State of Business Communication](#) report, we found that over 60% of business leaders have concerns about protecting their employees' and company's security, privacy, personal data, and intellectual property.

When people interact with an AI model, they entrust it with some of their most sensitive personal or business information. It's important that users understand how their data is being handled and whether it is being sold or used for advertising or training purposes.

### Training data development

AI developers must be given guidelines and training on how to make sure datasets are safe, fair, unbiased, and secure. Both human review and ML checks should be implemented to ensure the guidelines are being applied appropriately.

### Working with user data

In order to uphold privacy, all teams interacting with models and training data should be thoroughly trained to ensure compliance with all legal, regulatory, and internal standards. All individuals working with user data must follow these strict protocols to ensure data is handled securely. Tight controls should be implemented to prevent private user data from being used in training data or being seen by employees working with models.



At Grammarly, all users can control whether their content is used to improve our products, and we do not allow any third-party AI processors we work with to store our users' data or use it to train their models. We never sell user data, either.

### Understanding data training

All users must have the ability to control whether their data is being used to train models and improve the product overall for everyone. No third parties should have access to user content to train their models.

### Considerations in using a third-party large language model (LLM) or building your own

When engaging with AI, everyone involved should be mindful of safety and security. This means not just the companies making or providing the AI technology, but also the businesses and individuals using it. Whether you're creating your own AI from scratch or using one made by someone else, you still need to be careful with how it's used and how data is protected.



Grammarly maintains a thorough vendor review process, repeated regularly, to conduct due diligence before engaging with any subprocessors. Among other focuses, this process validates that customer personal data is always protected. The review process includes multistep security and privacy assessments, a detailed review of the vendor's compliance posture, and an in-depth legal review of subprocessor data practices.

There are different risks associated with building one's own LLM, which comes with significant liability and security/legal implications, but using an external LLM does not unburden the platform using the LLM of those responsibilities either. The service through which the end user interacts with an LLM is still ultimately responsible for safeguarding user data and any potential security vulnerabilities as a custodian for that data. It's important to check how an external AI service handles security and privacy before using it and to ensure that it complies with all relevant data protection and storage regulations.

### Monitoring abuse of systems

As more people and companies begin to use AI as part of their suite of daily tools to solve complex problems, the stakes become exponentially higher for the types of potential attacks from bad actors. It is our responsibility



to ensure that our users are protected from these vulnerabilities and that we continue to keep pace to prevent and protect against constantly changing threat vectors.

When working in a field that's evolving so quickly, it's essential to understand the potential threat vectors, not only by understanding existing vulnerabilities but also by anticipating new risks. Steps must be taken to ensure that the model cannot leak data and that there is no opportunity to set up backdoor access, poison the data, or otherwise infiltrate the model.

At Grammarly, much of our work in security is around preventing adversarial attacks for our users. We rely on a red team of experts who have experience on both sides of the security fence to identify adversarial security approaches and vulnerabilities and help us formulate mitigation strategies that help keep our product safe.

To enhance the security of our AI features, we launched a targeted bug bounty campaign as part of our ongoing bounty program through HackerOne. This initiative is intended to encourage a thorough examination of our AI capabilities, particularly focusing on issues that could compromise user privacy. Researchers investigated key areas of concern, including identifying malicious data injections, detecting sensitive information leaks, uncovering embedded vulnerabilities, securing the AI system's API, ensuring private information isn't shared during conversations, protecting third-party AI components, and addressing any issues affecting confidentiality, integrity, and availability. We continue to invite security researchers to investigate issues related to AI unintentionally generating harmful content and to discover potential vulnerabilities in the integrity of our systems.

# Additional Considerations When Designing a Responsible AI Approach

In an era of rapid technological advancement like today, the role of security leaders extends beyond traditional security measures to encompass the ethical and responsible deployment of AI systems. With AI investment comes immense potential but also a range of complex risks and considerations.

By staying mindful of these key callouts, IT leaders can guide their organizations in developing AI systems that lay the groundwork for an innovative and robust AI strategy while remaining responsible and aligned with ethical standards.

### Top-down buy-in

To successfully implement a responsible AI strategy, teams must secure leadership buy-in and enforce clear procedures across all organizational levels. Without this support, AI features might launch without thorough reviews, leading to unforeseen issues and potential failures. This support is crucial in helping teams understand that proactive reviews contribute to the long-term success and reliability of their features by identifying and mitigating problems early on. This prevents poor user experiences and enhances the overall effectiveness of the AI systems.

### Legal and regulatory compliance

As AI technologies face increasing scrutiny and potential regulation, staying up-to-date with legal and regulatory requirements becomes crucial. This requires close collaboration with legal and policy teams and regular implementation of process updates to stay abreast of changing legislation. By maintaining legal awareness, organizations can mitigate legal risks and ensure their AI systems operate within the bounds of the law. They should

ensure the safe use of AI across the foundational models used, in every AI platform they invest in, and in every end-user interaction. This flow of information does not absolve anyone of responsibility, and it is up to all parties involved to uphold their end of using AI responsibly.

### Continuous improvements and monitoring

AI systems must undergo continuous monitoring and regular updates to remain relevant and accurate. This involves addressing changes in data, societal contexts, distribution shifts in input data, and mitigating biases and risks inherited from foundational models. Continuous improvement processes ensure that the AI adapts to new trends and contexts, enhancing its utility and reliability over time. Regular testing practices are crucial to preventing bugs and vulnerabilities that could lead to insecure or unexpected behaviors.

### Balancing risk mitigation with user value

Striking the right balance between content filtering and user value is critical in AI development. While effective filtering is necessary to prevent harm, it should not overly restrict user interactions or diminish the utility of the AI system. This requires fine-tuning models to avoid false positives and implementing clear content classification metrics.

# Grammarly's Unique Approach to Responsible AI

Unlike other AI tools, Grammarly's writing assistance is built specifically to optimize your communication. Grammarly's approach draws from teams of expert linguists, deep knowledge of professional writing best practices, and 15 years of experience in AI. With our vast expertise in developing best-in-class AI communication assistance, we always go to great lengths to ensure user data is private, safe, and secure.

Our commitment to responsible AI is woven into the fabric of our development and deployment processes, ensuring that our AI not only enhances communication but also safeguards user data, promotes fairness, and maintains transparency.

This approach permeates all aspects of our business, from how we implement third-party AI technologies to how we weave responsible AI reviews into every new feature we launch. We think critically about any in-house and third-party generative AI tools we use and are intentional in how our services are built, ensuring they are designed with the user in mind and in a way that supports their communication safely.

## Grammarly's Responsible AI Pillars:



Transparency



Fairness  
and safety



User agency



Accountability



Privacy  
and security

# Conclusion

Responsible AI is not just a buzzword but a necessity in today's AI-driven workforce. Implementing responsible AI principles ensures that AI systems are designed and deployed with the utmost consideration for fairness, safety, and user empowerment. At Grammarly, we have embedded these principles into the very core of our product frameworks and processes. By prioritizing transparency, fairness, user agency, accountability, and privacy and security, we aim to build AI tools that are not only effective but also ethical, trusted, and loved by users.

As AI continues to evolve and redefine the boundaries of what is possible in our workplaces, the accompanying risks and challenges will inevitably grow, too. Leadership teams must understand that while rapid innovation is essential, ethical innovation that balances safety and integrity is as well. It is not enough to simply adopt AI technologies; organizations must ensure they are doing so responsibly. This includes conducting thorough evaluations of AI vendors and partners to ensure they adhere to rigorous standards of privacy, security, and fairness. By [asking the right questions](#) and demanding transparency, businesses can make informed decisions that align with their ethical values and operational goals.

Organizations can establish dedicated responsible AI teams, implement robust evaluation frameworks, and continuously monitor and update AI systems to keep pace with evolving societal norms and technological advancements. By doing so, they not only enhance their AI capabilities but also build trust with their users and stakeholders.

We invite you to join us in leading the way toward responsible AI adoption. By embracing these principles and practices, you can help ensure that AI serves as a beneficial force, enhancing communication and upholding the highest ethical standards. Together, we can build a future where AI is not only a tool that supercharges our work but is also trustworthy, ethical, and fair.



# Uplevel Your Organization's Communication With AI, Built Responsibly

Grammarly is the world's leading AI writing assistance company, trusted by over 30 million people and 70,000 professional teams. From instantly creating a first draft to perfecting every message, Grammarly helps people at 96% of the Fortune 500 and teams at companies like Atlassian, Databricks, and Zoom get their point across—and get results—without compromising security or privacy. Grammarly's product offerings—[Grammarly for Business](#), [Grammarly Premium](#), [Grammarly Free](#), and [Grammarly for Education](#)—work where you do, delivering contextually relevant writing support across over 500,000 apps and websites. Founded in 2009, Grammarly is No. 7 on the *Forbes* Cloud 100, one of *TIME*'s 100 Most Influential Companies, one of *Fast Company*'s Most Innovative Companies in AI, and one of *Inc.*'s Best Workplaces.

## Contact Our Sales Team

