Global position statement on responsible artificial intelligence

Bristol Myers Squibb | November 2026

Our position

At Bristol Myers Squibb (BMS), our mission is to discover, develop and deliver innovative medicines that help patients prevail over serious diseases. Artificial Intelligence (AI) is changing how we pursue that mission through one connected ecosystem. We use AI and machine learning to identify drug targets grounded in causal human biology, to design molecules before they are synthesized, to run more efficient clinical trials, and to enhance our regulatory submissions. Across the company, AI supports how we manufacture and supply medicines, quality and safety surveillance, digital health tools we develop with partners for patients and clinicians, how we deliver information to healthcare professionals and patients, and the everyday work of our colleagues. AI at BMS augments the judgment of qualified professionals. It does not replace it.

We believe responsibly governed AI can shorten the path from discovery to patient and improve the quality and consistency of decisions across our medicines' lifecycle. We recognize its risks: to patient safety, to fairness, to privacy and confidentiality, and to trust. BMS is committed to using AI in a way that captures the benefits while managing those risks with the same patient-centered, evidence-based discipline that guides how we discover, develop and deliver medicines, and to complying with the AI, data protection and sector-specific laws that apply to us wherever we operate.

Our principles

Four principles set out in the BMS Responsible AI Policy govern the use of AI for everyone who works for BMS and every AI system we develop, acquire or use, including generative and agentic AI.

1. **Patient and human safety.** We do not deploy AI in settings that could potentially affect people's health or well-being, including product quality and pharmacovigilance, without testing, monitoring, and establishing human safeguards.
2. **Fairness and non-discrimination.** We design, test and monitor AI to guard against unlawful discrimination and unfair outcomes for individuals, recognizing that unfair bias can be embedded in data and is not always visible in outputs.
3. **Data privacy, security, and confidentiality.** We protect personal information and confidential information when we use AI. Our Privacy Policy permits confidential and personal data to be processed only in approved, secure AI systems, subject to our data minimization, consent, and retention requirements.

4. **Honesty and transparency.** We do not use AI to mislead regulators, healthcare professionals, patients, employees, or anyone else. We meet applicable disclosure requirements.

How we govern AI

BMS applies a risk-based governance model to AI. Under the oversight of the board, our enterprise governance body, the Responsible AI and Data (RAID) Committee, is co-chaired by senior leaders responsible for information security, and AI, data and privacy law and compliance. It approves the highest-risk uses of AI, oversees alignment with AI regulation, sets enterprise standards and tracks AI risks and incidents, drawing on experts from clinical, manufacturing, safety, regulatory, quality, legal, compliance, human resources, sustainability, and commercial functions as needed.

Before deployment, AI systems that BMS develops or acquires, including open-weight models obtained from public repositories, undergo a risk-based assessment that evaluates relevant legal, privacy, cybersecurity, intellectual property, quality, compliance, performance, and other applicable considerations. The assessment establishes binding controls commensurate with the identified risks. We do not adopt new frontier models by default. New models are enabled deliberately, based on risk, and its access is expanded as controls and experience allow. Scrutiny is heightened for uses that could affect patients' health, individuals' employment or rights, regulated quality and safety activities or interactions with regulators. Agentic AI can take actions rather than only generate content and is subject to accountable human ownership, technical limits on what it can access and do, monitoring of its actions and the ability to shut it down. Deployed systems remain under human oversight, with performance monitoring, drift and emerging bias and reassessment when they materially change. Consequential decisions are not based on AI output alone, and the people who use AI remain accountable for what they do with it.

Responsible AI training is required for all colleagues. The cybersecurity program that protects the systems and data behind our AI is independently assessed against the National Institute of Standards and Technology Cybersecurity Framework. Suspected AI incidents, misuse or ethical concerns can be reported through the BMS Integrity Line, which is available to employees and external parties, by name or anonymously where local law permits. BMS prohibits retaliation against anyone who raises a concern in good faith.

Protecting the data behind AI

AI is only as trustworthy as the data behind it. BMS is a steward of data entrusted to us by patients, clinical trial participants, healthcare professionals, employees, and partners. Personal data used with AI is minimized, de-identified or anonymized where that is sufficient, and processed on a lawful basis. Our policy prohibits entering personal information, unpublished research or confidential information into consumer-grade or unapproved AI tools. We are examining the provenance and representativeness of the data behind the AI we use, particularly in tools that may affect patient care or research, and asking our AI providers to do the same.

The data we submit to regulators and share with partners contains trade secrets and confidential commercial information. When we entrust that data to a regulator, a vendor, or a collaborator, it is for a defined purpose. BMS expects that it will not be used to train, fine-tune or otherwise improve non-BMS AI systems, directly or indirectly, without our explicit permission and binding, durable safeguards, including segregation of data belonging to different companies. Permission alone is not sufficient. We expect our technology partners and the developers of the models we use to meet the same standards of safety, security, data protection and bias mitigation that we apply to ourselves, to test their systems for safety and security before deployment, and to disclose significant safety and security incidents that affect us. We build these expectations into our contracts.

Our public policy positions

BMS engages constructively with policymakers, regulators, standards bodies and industry associations to help shape AI frameworks that protect patients, enable innovation, and help patients and health systems benefit more quickly from advances in science. The governance we apply internally is what allows us to advocate credibly for the following.

1. **Regulation should be risk-based, focused on use, and establish clear pathways.** Requirements should be clear, predictable, and proportionate to the risk of the context in which AI is used, concentrating where AI may affect patient safety, the integrity of evidence or regulatory decisions, and should not burden low-risk internal and administrative uses. For sponsors, regulators should provide clear, predicable pathways for the use of AI in generating and evaluating evidence across development, submission and review, without compromising regulatory standards.
2. **Global alignment, not duplication.** New AI rules should build on the frameworks that already govern medicines and be harmonized across borders. The research, development, manufacture and safety surveillance of medicines are subject to extensive regulatory oversight and data protection law, and health authorities are best placed to assess AI risk in these settings. New AI obligations should not duplicate or conflict with these frameworks, and compliance with them should be recognized. The same holds across jurisdictions: a patchwork of differing requirements adds cost and delay without adding patient protection. BMS welcomes the joint FDA and EMA Guiding Principles of Good AI Practice in Drug Development, and the accompanying work on harmonized terminology, as a foundation for regulatory convergence. BMS also supports international standards, regulatory sandboxes, and frameworks that preserve secure, lawful cross-border data flows and responsible access to high-quality data for research, with protection for personal data, trade secrets and confidential commercial information.
3. **Research and development should be protected.** AI used in drug discovery, biomarker identification, clinical trial design, monitoring and analysis, companion diagnostic development and safety surveillance should fall clearly within research exemptions, and AI used in regulated clinical investigations should not be subject to obligations that conflict with clinical trial regulation. In addition, AI used for operational efficiency, including internal workflows, resource planning and drafting of documentation, that

does not affect patient safety, product quality or the reliability of study results should not attract AI-specific obligations.

4. **Meaningful human oversight should be recognized.** AI that informs a decision made by a qualified professional should be treated as decision support, not as automated decision-making, and notice requirements designed for consumer credit or employment screening should not be transposed to research, clinical development and patient support settings where established professional and regulatory standards already apply.
5. **Transparency should protect trust and innovation alike.** BMS supports disclosure when AI interacts directly with patients. Transparency requirements should not compel disclosure of trade secrets, confidential commercial information, or proprietary algorithmic methods. Where full access to training data and code is not feasible, robust quality management and demonstrated performance monitoring are the appropriate basis for trust.
6. **Regulators' use of AI should protect sponsor data.** BMS supports regulators' own responsible use of AI, which can improve the consistency and efficiency of review and speed patient access to medicines. Data that sponsors submit should not be used to train or improve regulator or vendor AI systems, directly or indirectly, without the explicit permission of the data owner and binding, durable safeguards, including segregation of data from different sponsors and independent security audits. Regulatory decisions should remain in the accountable judgment of human reviewers, regulators should be transparent about where AI informs them, and any expansion of permissible data use should proceed through formal rulemaking with public comment.

Our broader commitments

Our approach to AI is consistent with our Human Rights Global Position Statement and our Environmental Stewardship position statement, and we are extending both to AI, including the representativeness of the data behind AI, the oversight of patient-facing tools, human rights due diligence, and the associated environmental impact of AI deployment and usage. We will continue to manage AI as part of our broader sustainability and social impact strategies to ensure we hold ourselves and our partners to high standards on environmental stewardship and human rights.

Continuing the work

AI is evolving faster than the frameworks that govern it. We will keep this statement, our Responsible AI Policy and our governance current as technology, regulation and stakeholder expectations change, and we will continue to engage patients, healthcare professionals, regulators, policymakers, peers and technology partners to advance responsible AI in life sciences.

Concerns about the use of AI at BMS can be raised through the BMS Integrity Line at bms.integrity.ethicspoint.com.