R Street
Free markets. Real solutions.

R Street Shorts
No. 154
September 2026

From Queue to Quality: Using AI to Modernize FDA Tobacco-Product Review

By Jeffrey S. Smith

The CTP can improve speed without sacrificing rigor. Well-governed AI can reduce clerical work, find conflicts, organize evidence, and improve safety monitoring, but it cannot resolve unclear policy or replace scientific and legal judgment. When paired with clear standards, machine-readable applications, human accountability, and public reporting, AI-assisted tools can help the CTP reach sound decisions sooner.

Introduction

The U.S. Food and Drug Administration's (FDA's) Center for Tobacco Products (CTP) has a history of taking years to decide tobacco-product applications—decisions that determine which products may lawfully compete with combustible cigarettes and which reduced-risk claims manufacturers may make.[1] The agency took nearly five years to decide the ZYN nicotine-pouch applications, nearly four years to authorize the Vuse Alto e-cigarette, and more than three and a half years to decide the original IQOS heated-tobacco modified-risk applications.[2]

Faster reviews of reduced-risk products matter for public health because combustible cigarette smoking causes more than 480,000 deaths each year in the United States.[3] Although no tobacco product is safe, the FDA recognizes a range of risk across products and states that adults who smoke may reduce their health risks by switching completely to lower-risk alternatives.[4] Adults who use e-cigarettes exclusively have less exposure to many harmful chemicals than those who smoke combustible cigarettes.[5] Some lower-risk products may also help adults stop smoking combustible cigarettes. In one clinical trial, e-cigarettes combined with counseling produced higher one-year quit rates than nicotine-replacement therapy combined with counseling.[6] The precise public-health cost of CTP delays cannot be calculated, and speed must never be prioritized over scientific rigor, but products that meet the FDA's standards should be authorized without delay, as more lower-risk choices mean more opportunities for adults to move away from combustible cigarettes.

Faster reviews of reduced-risk products matter for public health because combustible cigarette smoking causes more than 480,000 deaths each year in the United States.

Many of the tasks that slow CTP application review are necessary but repetitive: checking whether a file is complete, searching for specific evidence within thousands of submitted pages, comparing table data, spotting conflicting information, and checking calculations—the very types of bounded tasks that artificial intelligence (AI) increasingly supports.[7] Well-tested AI systems can do this work quickly, reducing a significant source of review friction and allowing human

The sources included in this paper were verified and active at the time of publication.

reviewers to focus on higher-stakes tasks like determining whether cited studies are trustworthy or whether a product would support or harm public health.

Recent CTP progress shows that faster review is possible, and the FDA's broader adoption of AI tools offers one way to build on that progress. In May 2026, the CTP reported that it had cut its application backlog by about 70 percent during 2025 and cleared its acceptance-review queue (i.e., the backlog awaiting the initial administrative check before the filing and scientific reviews). In addition, the center conducted a focused pilot in which it reviewed six nicotine-pouch products in three months (details on that expedited process are not yet available).[8] Meanwhile, the FDA is building tools for AI-assisted regulation more broadly. In 2025, it completed an AI-assisted scientific-review pilot and launched Elsa (an internal generative AI assistant) across FDA centers.[9] This tool can summarize documents, compare labels, and analyze reports of possible health problems.[10] The FDA later expanded AI applications to premarket review, review checks, safety monitoring after authorization, inspections, and compliance.[11] The CTP should build on these FDA-wide efforts by applying AI to repetitive and time-consuming review tasks, whether by expanding existing uses or introducing AI where it has not yet been adopted. This paper examines where AI could add the most value and the safeguards needed to preserve scientific and legal rigor, and it offers concrete recommendations for adoption.

Understanding the Tobacco-Product Review Process

To sell a new tobacco product legally in the United States, manufacturers must submit a premarket tobacco product application (PMTA) to the FDA.[12] The FDA then reviews the application and decides whether selling that product would be "appropriate for the protection of the public health." Under that standard, the FDA must consider the product's impact on the whole population, including whether it will help current users quit combustible products and whether nonsmokers, especially young people, will start using the product.[13]

PMTA review moves through several stages. The FDA first checks whether it regulates the product type and whether the application contains the required parts.[14] It then determines whether the file has enough information for full scientific review. That review may involve experts from several scientific fields, site inspections, product testing, requests for missing information, and environmental review. The agency then grants or denies permission to market the product and may add sales, advertising, or reporting conditions. Regulations generally call for action within 180 days after a complete application is filed, although pauses and exceptions can extend that review period.[15]

A related pathway applies when manufacturers wish to make a specific claim that a product reduces harm, disease risk, or exposure compared with another tobacco product.[16] To make that claim, they must submit a modified risk tobacco product application (MRTPA).[17] The FDA studies whether consumers understand the claim, how they may use the product, and how the claim may affect both users and nonusers. Modified-risk orders expire and require continued monitoring and studies.[18] For example, in June 2026, the FDA allowed 20 ZYN products to carry a specific modified-risk claim after reviewing these issues.[19]

PMTAs and MRTPAs are large, highly technical, data-dense files, often with hundreds of scanned attachments, inconsistent tables, cross-references, and repeated amendments, making efficient human review challenging.

How a PMTA Moves Through FDA-CTP Review*

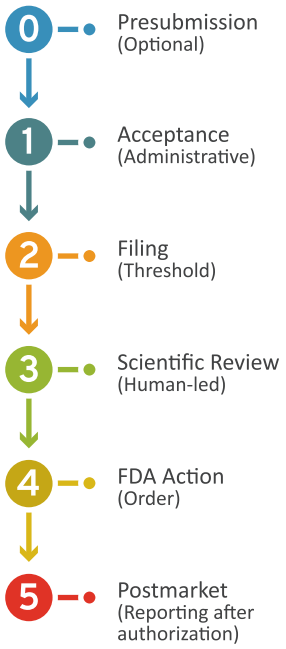

*Regulations generally call for action within 180 days after a complete application is filed, although pauses and exceptions can extend that review period.

Diagnosing the Delays

Three key problems slow tobacco-product review: the sheer volume of applications, the complexity of the science they contain, and gaps in process-related guidance for submissions. The first is a matter of scale. When a 2020 deadline required manufacturers of every e-cigarette already on the market to file a PMTA or remove the product, the CTP received 2,151 applications covering more than 6.6 million e-cigarette products.[20] By October 2022, about 0.8 percent of those products remained pending; most of the others had been rejected at an early stage or denied, and only 23 products had received marketing orders.[21]

Three key problems slow tobacco-product review: the sheer volume of applications, the complexity of the science they contain, and gaps in process-related guidance for submissions.

The second problem is the complexity of the science in the applications. Reviewers must connect product design and manufacturing with chemistry, the health effects of chemicals, addiction, consumer behavior, youth use, and estimates of population effects. Moreover, amendments submitted after the initial application can affect study data, labels, or marketing plans while the review is still underway. The ZYN record, for example, included a response to an FDA request for missing information and later updates to research and labeling.[22] Each update requires staff to find what changed and determine which earlier conclusions must be revisited.

The third problem is a lack of clear FDA guidance on the application process itself. An independent review called the CTP's process cumbersome, unpredictable, and lacking transparency.[23] It found that guidance often came late and that the FDA's information systems were not built in a way that supported a single, integrated application-review process.[24] In addition, a 2026 FDA roundtable showed that manufacturers still had basic questions about what they needed to include in their applications with regard to product testing, manufacturing, nicotine delivery, adult benefits, and chemical health effects.[25] Together, these gaps leave applicants without clear guidance on the CTP's review process, common application templates, or a plain statement of scientific expectations.

Applying AI Across the Review Process

"AI" is not a single tool; it includes basic software that checks rules, tools that turn scans into searchable text, and systems that find facts in documents. It also encompasses search tools that retrieve relevant passages, programs that flag unusual numbers, and language models that draft summaries.

The most important principle is that AI should never decide a PMTA or MRTPA; it should handle bounded tasks that produce outputs a scientist can check, thereby enabling reviewers to spend more time assessing evidence and less time on administrative and technical work. The sections that follow reinforce this distinction by identifying tasks that tested AI tools can safely carry out at key stages of the review process and noting where a human reviewer must make decisions.

Confirming an application is complete and consistent

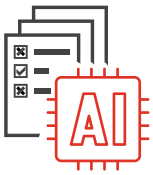

A basic checking tool can catch application defects, such as missing required fields, mismatched product names, and numbers in the wrong units. In a PMTA, that means verifying that the same product identifier appears consistently across the cover letter, product-composition tables, and the stability data, or flagging a nicotine concentration reported in milligrams per milliliter in one section and as a percentage in another. Software can also turn scanned pages into searchable text and sort documents by type—for instance, converting a scanned third-party lab report into text a reviewer can search rather than read page by page. Automated completeness checks already have some regulatory precedent. The European Chemicals Agency, for example, gives

applicants a rule-based tool to catch missing information before they file, while agency staff retain responsibility for judging the science.[26]

Beyond catching errors, standardized data can save substantial time. An FDA survey found that most drug-application reviewers agreed that data standardization made routine analysis easier and left more time to address unusual questions.[27] The CTP could also use AI to build a document map of each tobacco application, showing what it contains, where each item appears, and which entries conflict. This would make it easier to navigate long documents and catch clerical errors consistently.

Finding and comparing evidence

AI-based search and extraction tools could help CTP reviewers find and connect facts scattered across long applications—two of which reportedly exceeded 125,000 and 150,000 pages.[28] In the case of PMTAs or MRTPAs, these tools could help distinguish a nicotine concentration in a product from a nicotine level measured in a person, then link each number to the correct study and product. These tasks take advantage of a key strength of AI: It can find and organize specific facts with high accuracy when the task is narrow and the source documents follow familiar patterns.[29]

Outputs from these tools become less reliable when information spans several passages, uses unusual language, or requires scientific interpretation. Studies of AI-assisted data extraction show the pattern clearly. One study found that AI-assisted extraction was about as accurate as human-only work and could save time, but both people and machines missed some information.[30] Another found that performance was strongest for standard facts, such as study design and population, and weaker for questions unique to a particular review.[31] For the CTP, the practical use would be a preliminary table of extracted data, including sample size, dose, exposure, and study results. However, a reviewer must still compare every entry against the original source.

Prioritizing studies for review

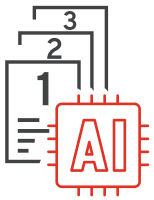

AI could help prioritize a reading list for CTP application reviewers, moving the most relevant studies to the top to ensure they receive the most careful attention. The strongest evidence for this application of AI comes from systematic literature review, where researchers face the same challenge CTP reviewers do—finding the few relevant studies within a large body of research. The Environmental Protection Agency, for example, estimated that its screening tool could reduce the number of papers reviewers had to screen by 40 to 60 percent, though the tool did not find every relevant study.[32] The United Kingdom's National Institute for Health and Care Excellence offers a similar tool to rank studies by relevance to a research topic (not to include or exclude them).[33] The institute also reports that reviewers may resist the tool when they cannot tell at what point screening can safely stop.[34]

Research on using large language models (LLMs) to identify relevant studies for systematic reviews also shows potential benefits and limitations. One study found that LLMs identified most relevant papers on a given topic but also flagged many irrelevant ones.[35] A larger study found that an LLM's performance varied sharply depending on the topic and the review question, suggesting that no single tool can be assumed to perform equally well across the range of science presented in PMTAs.[36] Given these findings, the CTP should consider using AI tools to help set reading priorities for reviewers, but not to include or exclude papers from a reading list or judge a paper's strength.

Checking numbers and calculations

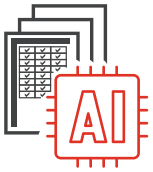

Software can compare values across a study report, summary table, and population model. It can also repeat calculations, convert units, and flag figures that may warrant a second look. The FDA reports that its Elsa AI tool already supports study-plan review, label comparison, summaries of possible health problems, code writing, and inspection priorities.[37] A 2026 update, Elsa 4.0, added document search, scan reading, and data analysis, and, in a separate 2025 pilot, one FDA official said AI tools cut the time needed for some tasks from three days to minutes.[38] These reports are promising, but the FDA has not published enough detail to judge the pilot's error rate or its effect on overall review time.

For the CTP, a flagged conflict should trigger review, not settle the issue. Software can show that two sections disagree, but a scientist must determine whether the difference matters and what it means for the population.

Keeping People in Charge

The evidence supports a layered system that matches AI tools to tasks they can reliably perform. Rule-based checks can verify fields, units, and product identifiers. Search and extraction tools can organize files and draft evidence tables. Statistical tools can repeat calculations and flag conflicts. LLMs can summarize approved records, provided they show their sources.

Humans must continue to make every high-stakes decision, including judgments about study credibility, acceptable levels of uncertainty, and whether a product meets public-health standards.

Humans must continue to make every high-stakes decision, including judgments about study credibility, acceptable levels of uncertainty, and whether a product meets public-health standards. Scientists weigh the evidence, lawyers confirm the FDA's authority and fair procedures, inspectors assess facilities, and accountable officials sign orders and enforcement actions.

This division of labor is necessary because AI can invent citations, omit important evidence, repeat past inconsistencies, expose confidential data, or encourage "automation bias" (i.e., the tendency to trust a computer answer too readily). If reflected in a PMTA decision, errors like these could also expose the agency to legal challenge.

For these reasons, any AI-generated output used in review should be logged with its source, the tool and instructions that produced it, any human correction, and how the reviewer ultimately used or rejected it. The CTP should use validated AI tools consistently, test performance across different product types and applicants, train application reviewers to challenge output, and publish overall error and correction rates, time savings, and known limits—all without exposing applicants' trade secrets.

Policy Recommendations

The following policy recommendations translate these ideas into practice. None require new statutory authority; each falls within the CTP's existing rulemaking and operational discretion.

Establish accountable governance

The CTP should create a formal AI-assisted review program staffed by scientists, policy experts, lawyers, technology experts, and privacy and security officials. Each AI-assisted task should have one clear purpose, a named human owner, and evidence that it works and meets measurable standards. Staff should keep a complete record of what the tool did and set a rule for discontinuing its use if performance declines. The CTP can draw on FDA and U.S. Department of Health and Human Services systems, but it must first

confirm their effectiveness for tobacco review rather than assume that results from another center will transfer. The CTP should also avoid relying on one vendor or on a vendor whose system or updates it cannot inspect or control.

Make applications machine-readable

The FDA should require that key PMTA and MRTPA data be submitted in standard electronic formats. The CTP should publish standards for common product identifiers, units, table layouts, study descriptions, and the marking of amendments. Narrative reports can remain in the legal record, but critical product, manufacturing, and study data should be made as easy as possible to check electronically. The agency should seek public input, phase in any changes, and help smaller manufacturers comply.

Test narrow uses before expanding

The CTP should begin with pilots for completeness checks, amendment comparison, evidence search, citation verification, and consistency review. It should compare each tool against qualified human review and measure missed information, false alarms, correction time, and total time saved across product types and applicant sizes. The CTP should publish combined results and known limits while protecting confidential business information. Reports should separate time saved on one task from any change in the total application-review time. Only tools that meet preset standards should move into more demanding uses.

Preserve fairness and human judgment

AI should not decide whether evidence is credible, whether missing information is important, or whether the population-health standard is met. It should not issue orders, denials, requests for missing information, suspensions, or withdrawals. Applicants should be told when AI materially affects a review and should be able to correct factual errors. The official review record should show which evidence the FDA relied on, and reviewers should document any disagreement with AI output.

Provide useful tools to applicants

The CTP should offer applicants a presubmission checker based on statutes, regulations, guidance, and filing rules. It could flag missing fields, mismatched product names, invalid units, and conflicting values but should not predict authorization.

Conclusion

The CTP can improve speed without sacrificing rigor. Well-governed AI can reduce clerical work, find conflicts, organize evidence, and improve safety monitoring, but it cannot resolve unclear policy or replace scientific and legal judgment. When paired with clear standards, machine-readable applications, human accountability, and public reporting, AI-assisted tools can help the CTP reach sound decisions sooner. Products that meet the law's public-health standard could reach adults who smoke with less avoidable delay, and products that fail could be identified more quickly. That would make tobacco regulation more predictable and more focused on the disease and death combustible cigarettes cause.

About the Author

Jeffrey S. Smith is a behavioral neuroscientist and resident senior fellow in the R Street Institute's Healthier Communities program. His work focuses on tobacco harm reduction, reduced-risk nicotine products, substance use policy, and evidence-based approaches to public health regulation.

Endnotes

1. "FDA Authorizes Marketing of 20 ZYN Nicotine Pouch Products after Extensive Scientific Review," U.S. Food & Drug Administration, Jan. 16, 2025. https://www.fda.gov/news-events/press-announcements/fda-authorizes-marketing-20-zyn-nicotine-pouch-products-after-extensive-scientific-review; "Technical Project Lead Review of PMTAs: PM0000593.PD1–PM0000612.PD1," U.S. Food & Drug Administration, Jan. 16, 2025. https://www.accessdata.fda.gov/static/searchtobacco/ZYN/PMTA_TPL_PM593-PM612_Zyn_01_13_2025_Redacted.pdf; "FDA Authorizes Marketing of Vuse Alto Tobacco-Flavored E-Cigarette Pods and Accompanying Power Unit," U.S. Food & Drug Administration, July 18, 2024. https://www.fda.gov/tobacco-products/ctp-newsroom/fda-authorizes-marketing-vuse-alto-tobacco-flavored-e-cigarette-pods-and-accompanying-power-unit; "Scientific Review of Modified Risk Tobacco Product Application (MRTPA) Under Section 911(d) of the FD&C Act—Technical Project Lead: MR0000059–MR0000061 and MR0000133," U.S. Food & Drug Administration, July 7, 2020. https://www.fda.gov/media/139796/download.
2. Ibid.
3. "Cigarette Smoking," Centers for Disease Control and Prevention, Sept. 17, 2024. https://www.cdc.gov/tobacco/about/index.html; "The Relative Risks of Tobacco Products," U.S. Food & Drug Administration, May 6, 2026. https://www.fda.gov/tobacco-products/health-effects-tobacco-use/relative-risks-tobacco-products.
4. "The Relative Risks of Tobacco Products." https://www.fda.gov/tobacco-products/health-effects-tobacco-use/relative-risks-tobacco-products.
5. Maciej L Goniewicz et al., "Comparison of Nicotine and Toxicant Exposure in Users of Electronic Cigarettes and Combustible Cigarettes," *JAMA Network Open* 1:8 (Dec. 7, 2018). https://pubmed.ncbi.nlm.nih.gov/30646298.
6. Peter Hajek et al., "A Randomized Trial of E-Cigarettes versus Nicotine-Replacement Therapy," *New England Journal of Medicine* 380:7 (Feb. 14, 2019), pp. 629-637. https://pubmed.ncbi.nlm.nih.gov/30699054.
7. Eric N. Lindblom, "How might FDA fix this e-cigarette PMTA mess? Commentary on Glantz and Lempert and Meshnick et al.," *Tobacco Control* 33:e1 (April 2024), pp. e106-e107. https://tobaccocontrol.bmj.com/content/33/e1/e106.
8. Bret Koplow, "CTP Acting Director Statement: New Steps Forward in Accelerating Innovation and Efficiency in Product Review," U.S. Food & Drug Administration, May 7, 2026. https://www.fda.gov/tobacco-products/ctp-newsroom/ctp-acting-director-statement-new-steps-forward-accelerating-innovation-and-efficiency-product.
9. "FDA Launches Agency-Wide AI Tool to Optimize Performance for the American People," U.S. Food & Drug Administration, June 2, 2025. https://www.fda.gov/news-events/press-announcements/fda-launches-agency-wide-ai-tool-optimize-performance-american-people.
10. Ibid.
11. "FDA Expands Artificial Intelligence Capabilities with Agentic AI Deployment," U.S. Food & Drug Administration, Dec. 1, 2025. https://www.fda.gov/news-events/press-announcements/fda-expands-artificial-intelligence-capabilities-agentic-ai-deployment.
12. 21 U.S.C. § 387j (2024), "Application for Review of Certain Tobacco Products." https://www.govinfo.gov/content/pkg/USCODE-2024-title21/pdf/USCODE-2024-title21-chap9-subchapIX-sec387j.pdf; "Premarket Tobacco Product Applications," U.S. Food & Drug Administration, March 28, 2024. https://www.fda.gov/tobacco-products/market-and-distribute-tobacco-product/premarket-tobacco-product-applications.
13. Ibid.
14. 21 C.F.R. pt. 1114, "Part 1114—Premarket Tobacco Product Applications," Electronic Code of Federal Regulations. https://www.ecfr.gov/current/title-21/part-1114.
15. Ibid.
16. 21 U.S.C. § 387k (2024), "Modified Risk Tobacco Products." https://www.govinfo.gov/content/pkg/USCODE-2024-title21/pdf/USCODE-2024-title21-chap9-subchapIX-sec387k.pdf; "Q&A: Modified Risk Tobacco Products," U.S. Food & Drug Administration, June 21, 2018. https://www.fda.gov/tobacco-products/market-and-distribute-tobacco-product/qa-modified-risk-tobacco-products.
17. Ibid.
18. Ibid.
19. "FDA Authorizes 20 ZYN Nicotine Pouches to Be Marketed with Specific Modified Risk Claim," U.S. Food & Drug Administration, June 30, 2026. https://www.fda.gov/tobacco-products/ctp-newsroom/fda-authorizes-20-zyn-nicotine-pouches-be-marketed-specific-modified-risk-claim.
20. "The Food and Drug Administration Needs To Improve the Premarket Tobacco Application Review Process for Electronic Nicotine Delivery Systems To Protect Public Health," U.S. Department of Health and Human Services, Nov. 8, 2023, pp. 10, 12-13. https://oig.hhs.gov/documents/audit/8314/A-06-22-01002-Complete%20Report.pdf.
21. Ibid.
22. "Technical Project Lead Review of PMTAs: PM0000593.PD1–PM0000612.PD1." https://www.accessdata.fda.gov/static/searchtobacco/ZYN/PMTA_TPL_PM593-PM612_Zyn_01_13_2025_Redacted.pdf.
23. "Operational Evaluation of Certain Components of FDA's Tobacco Program: A Report of the Tobacco Independent Expert Panel," Reagan-Udall Foundation, Dec. 19, 2022. https://reaganudall.org/sites/default/files/2022-12/Tobacco%20report%20210pm.pdf.
24. Ibid.
25. "February 10, 2026: Roundtable on Premarket Tobacco Product Application (PMTA) Submissions for Electronic Nicotine Delivery Systems (ENDS) Products," U.S. Food & Drug Administration, Feb. 10, 2026. https://www.fda.gov/tobacco-products/ctp-newsroom/february-10-2026-roundtable-premarket-tobacco-product-application-pmta-submissions-electronic.
26. "Validation assistant," European Chemicals Agency, last accessed July 28, 2026. https://echa.europa.eu/validation-assistant; "Technical completeness check," European Chemicals Agency, last accessed July 28, 2026. https://echa.europa.eu/technical-completeness-check.
27. "Assessment of the Impact of Electronic Submissions and Data Standards on the Efficiency and Other Performance Attributes of the Human Drug Review Process: Final Report and Recommendations," U.S. Food & Drug Administration, 2017, pp. 2-3, 41-43. https://www.fda.gov/files/about%20fda/published/PDUFA-V-Assessment-of-Impact-of-Electronic-Submissions-and-Data-Standards-on-the-Efficiency-and-Other-Performance-Attributes-of-the-Human-Drug-Review-Process.pdf.
28. "Juul Labs Submits Premarket Tobacco Product Application to the U.S. Food and Drug Administration for the Juul System," Juul Labs, July 30, 2020. https://www.juullabs.com/juul-labs-submits-premarket-tobacco-product-application; "Vuse Premarket Tobacco Product Application Filed for Substantive Review by the FDA," Reynolds American Inc., Nov. 27, 2019. https://www.prnewswire.com/news-releases/vuse-premarket-tobacco-product-application-filed-for-substantive-review-by-the-fda-300965729.html.
29. Yiwen Shi et al., "Information Extraction From FDA Drug Labeling to Enhance Product-Specific Guidance Assessment Using Natural Language Processing," *Frontiers in Research Metrics and Analytics* 6 (June 10, 2021). https://pmc.ncbi.nlm.nih.gov/articles/PMC8222600.
30. Gerald Gartlehner et al., "Artificial Intelligence-Assisted Data Extraction With a Large Language Model: A Study Within Reviews," *Annals of Internal Medicine* 178:12 (December 2025), pp. 1763-1771. https://pubmed.ncbi.nlm.nih.gov/41183336.

31. Tue Helms Andersen et al., "Using Artificial Intelligence Tools as Second Reviewers for Data Extraction in Systematic Reviews: A Performance Comparison of Two AI Tools Against Human Reviewers," *Cochrane Evidence Synthesis and Methods* 3:4 (July 14, 2025). https://pubmed.ncbi.nlm.nih.gov/40661122.
32. "Introduction to Systematic Review in the Context of Environmental Health Assessments," U.S. Environmental Protection Agency, May 27, 2021. https://www.epa.gov/sites/default/files/2021-05/documents/all_slides_epa_sr_101_nams_cop_0.pdf.
33. "NICE: Priority Screening Algorithm," National Institute for Health and Care Excellence, Feb. 27, 2025. https://www.gov.uk/algorithmic-transparency-records/nice-priority-screening-algorithm.
34. Ibid.
35. Siun Kim and Hyung-Jin Yoon, "Large Language Model-Assisted Systematic Review: Validation Based on Cochrane Review Data," *Studies in Health Technology and Informatics* 327 (May 15, 2025), pp. 904-905. https://pubmed.ncbi.nlm.nih.gov/40380609.
36. Rohan Sanghera et al., "High-performance automated abstract screening with large language model ensembles," *Journal of the American Medical Informatics Association* 32:5 (May 1, 2025), pp. 893-904. https://pubmed.ncbi.nlm.nih.gov/40119675.
37. "FDA Launches Agency-Wide AI Tool to Optimize Performance for the American People." https://www.fda.gov/news-events/press-announcements/fda-launches-agency-wide-ai-tool-optimize-performance-american-people.
38. "FDA Expands AI Capabilities and Completes Data Platform Consolidation," U.S. Food & Drug Administration, May 6, 2026. https://www.fda.gov/news-events/press-announcements/fda-expands-ai-capabilities-and-completes-data-platform-consolidation; "FDA Announces Completion of First AI-Assisted Scientific Review Pilot and Aggressive Agency-Wide AI Rollout Timeline," U.S. Food & Drug Administration, May 8, 2025. https://www.fda.gov/news-events/press-announcements/fda-announces-completion-first-ai-assisted-scientific-review-pilot-and-aggressive-agency-wide-ai.