


Designing AI for Meta-Expertise in Regulatory Drafting

Dan Bateyko¹

¹Cornell University



Overview

This paper asks how AI writing assistants can assist federal rule writers. Draft rules must ground authority in statute, reflect current science, and persuade diverse audiences. Drafting such rules demands what Frank Pasquale calls *meta-expertise*: the ability to arbitrate among competing expert domains [10], in this case, between legal precision, scientific accuracy, and persuasiveness. I argue that designing for this meta-expertise is key to building systems that support rule writers' expert judgment.

This paper synthesizes HCI scholarship on legal, scientific, and persuasive writing assistance, drawing on its insights to outline the task of rule writing. I argue that rule writing's blended tasks come with tradeoffs that have consequences for model accuracy and user trust in AI systems.

Rule Drafting as a Site of HAI Study

HCI scholarship has focused little on rule drafting. The more popular site of study is the *public comment period* in rule-making, with scholarship for detecting fake comments [12], machine-assisted comment sorting [11], and measuring public input's influence [1]. Legal literature tends to be optimistic about tech's role in rulemaking [7], but rarely tests actual tools or designs empirically [4]. We know little about how rule writers collaborate, share drafts, or manage revisions within agencies [4].

Drafting decisions can have real public consequences. Delays and mistakes can be costly. From OSHA to the FDA, uncontroversial rules often stall for years, delaying needed protections [3, 6, 9]. Old rules remain on the books even when outdated; OSHA admits its exposure limits are “outdated and inadequate,” yet revising them is rarely prioritized [8]. Updating old rules requires setting aside drafting time, money, and political capital, which are resources leaders prefer to spend elsewhere.

Tradeoffs in Rule Writing

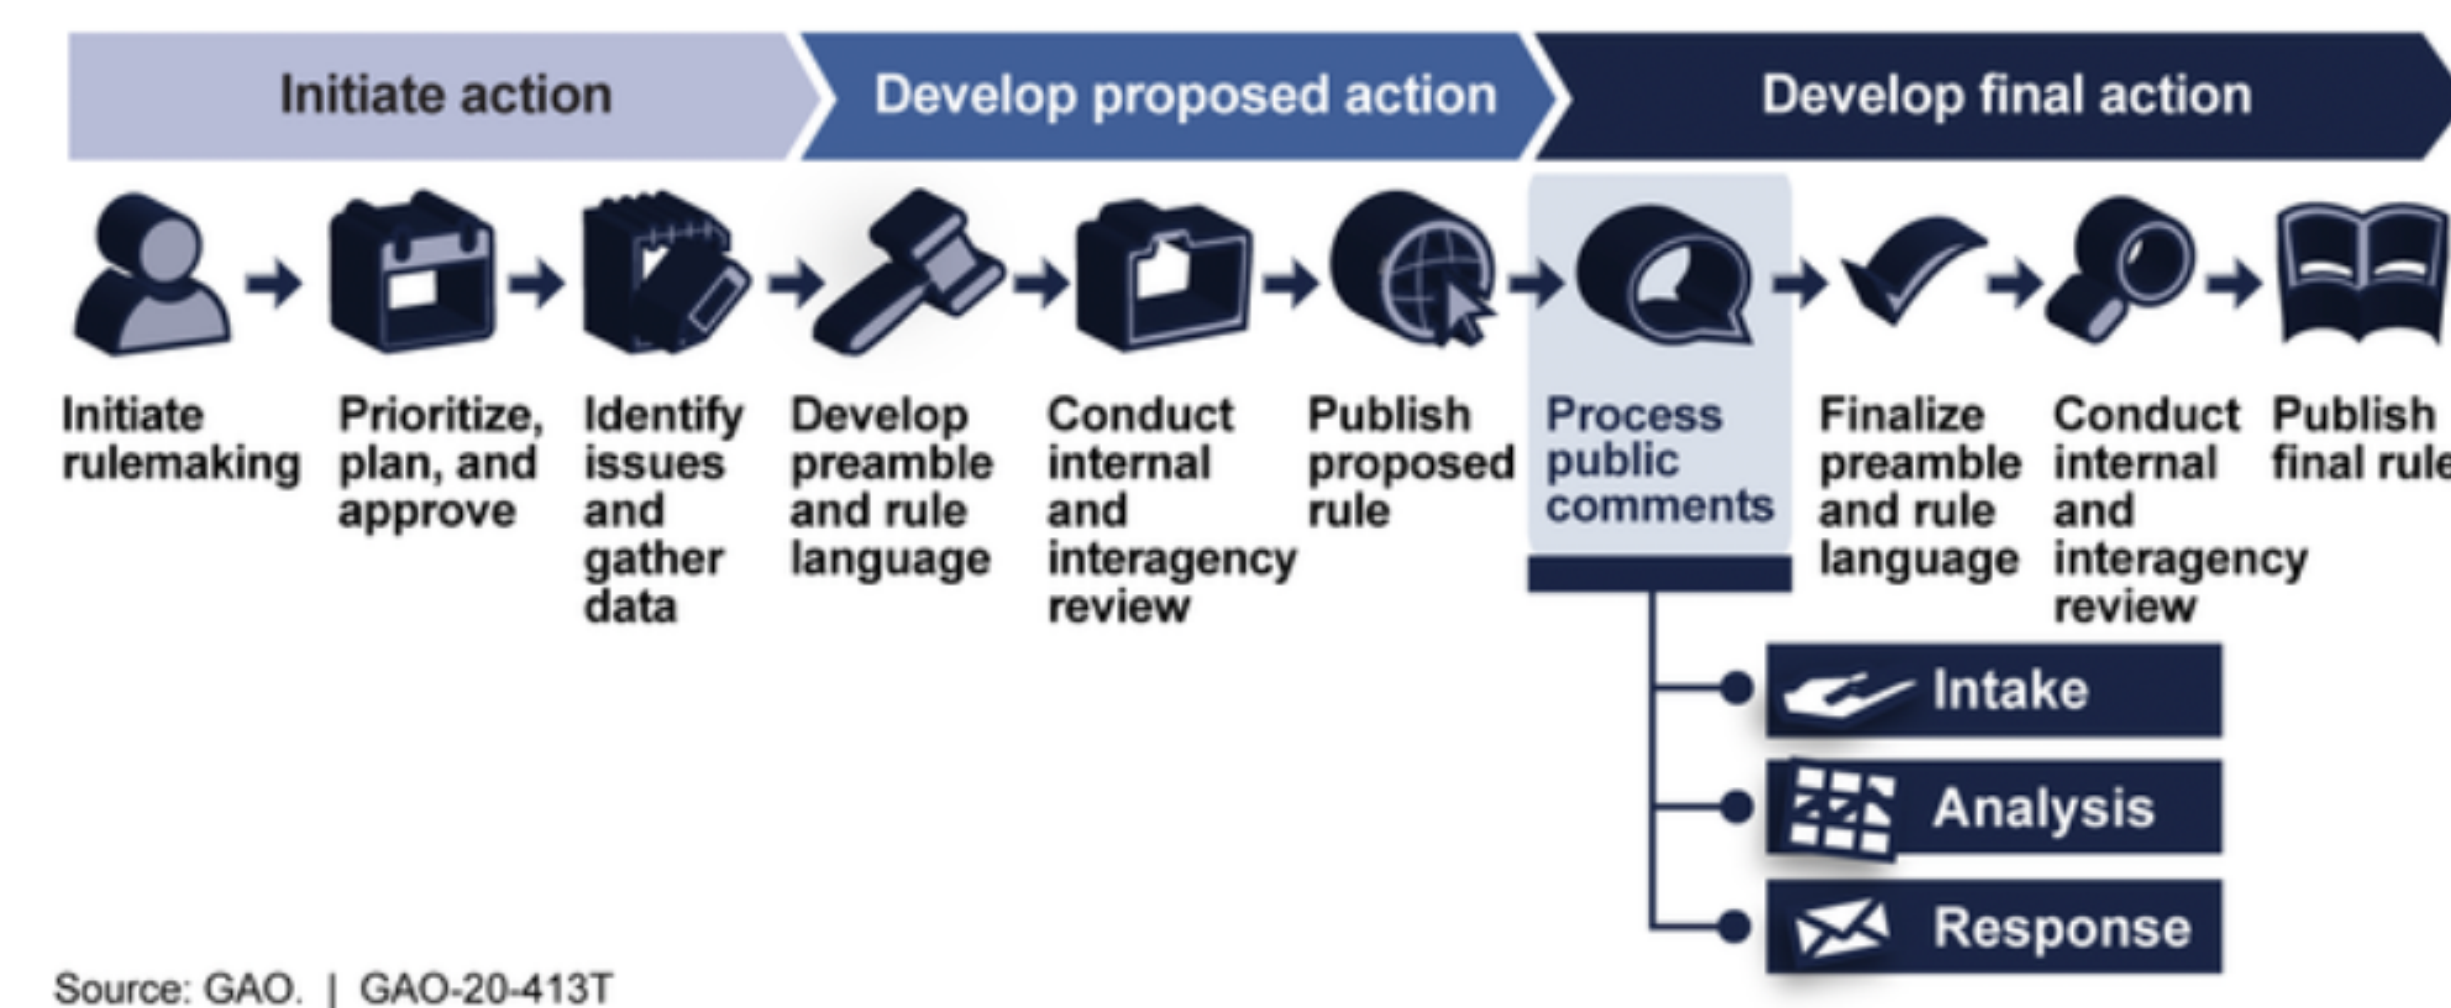


Figure 1. The figure shows the three stages of rulemaking: initiating a rulemaking, drafting the rule, and revising the rule based on public input. HCI scholarship to date has focused on the final step, specifically processing public comments.

Balancing legal precision, scientific validity, and persuasive clarity affects both user-interface and recommendation decisions.

- **Clarity:** There is an intelligibility tradeoff between legal writing and persuasive writing that the AI tool might model. Legal documents are often unclear to readers: when should the tool preference rule clarity over precision in language? Simplifying syntax can improve readability but strip necessary scientific or technical jargon.
- **Nuanced recommendations:** While an AI-supported legal writing task might focus on precise and accurate legal reference retrieval, an AI-supported persuasive writing task might be open-ended idea-generation. This kind of efficiency-ideation tradeoff is discussed variously in the literature; for example, should an assistant give multiple suggestions for phrases to use in a draft? Assistants might indicate whether a suggestion is legal, scientific, or persuasive in nature.
- **Bias & hallucination risks:** Summaries or corrections can misstate legal or scientific facts. Users may benefit from clear, task-specific rationales for each recommendation to calibrate trust and avoid over-reliance.

Future Directions

In future work, I will explore tradeoffs in rule writing by conducting a walkthrough and audit of a federal government-procured, open-source policy platform GAMECHANGER. The tool, suggested for use in policy drafting, draws on a BERT-based model for open-ended question-asking [2]. By investigating their implementation, I aim to surface opportunities for improving AI writing assistants in policy settings.

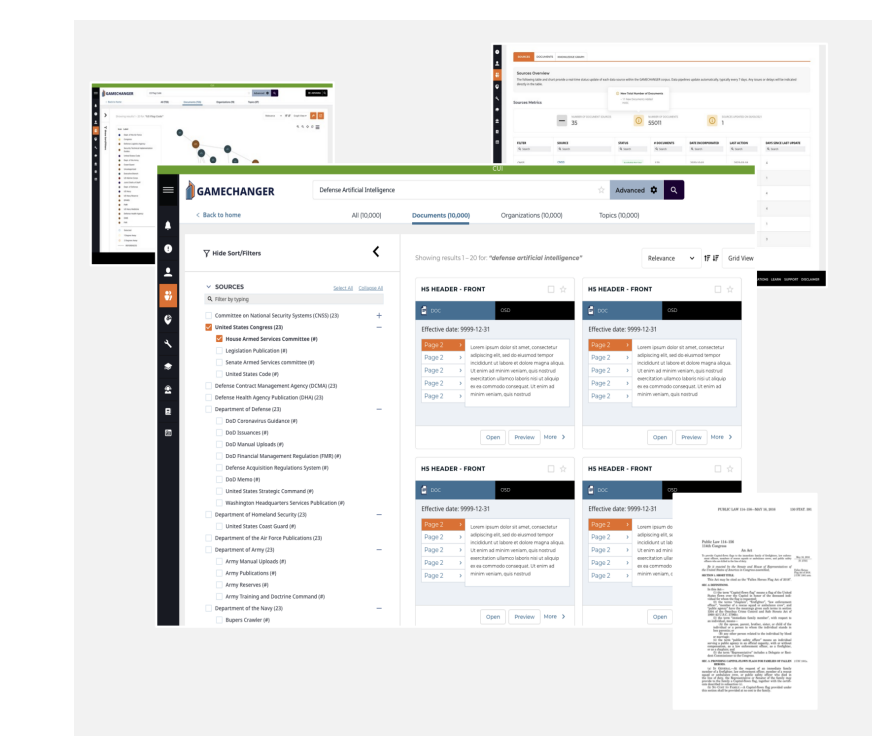


Figure 2. The user interface of the GAMECHANGER platform.

Future interviews with rule writers will map when and how AI should surface sources, balance recommendations across domains, and signal their relevance. By attending to this design space [5], I aim to develop AI systems that meaningfully support expert judgment.

References

- [1] Anastasia Deligiaouri and Jane Suiter. A policy impact tool: Measuring the policy impact of public participation in deliberative e-rulemaking. *Policy & Internet*, 13(3):349–365, September 2021.
- [2] dod-advana. GAMECHANGER aspires to be the department's trusted solution for evidence-based, data-driven decision-making across the universe of DoD requirements, 2024. Retrieved December 13, 2024 from <https://github.com/dod-advana/gamechanger-web/tree/dev>.
- [3] Thomas J. Hwang, Jerry Avorn, Daniel Carpenter, and Aaron S. Kesselheim. Quantifying the food and drug administration's rulemaking delays highlights the need for transparency. *Health Affairs*, 33(2):309–315, February 2014.
- [4] Cornelius M. Kerwin and Scott R. Furlong. *Rulemaking: How Government Agencies Write Law and Make Policy*. CQ Press (Sage Publications), Thousand Oaks, CA, 5 edition, 2019.
- [5] Mina Lee, Katy Ilonka Gero, John Joon Young Chung, Simon Buckingham Shum, Vipul Raheja, Hua Shen, Subhashini Venugopalan, Thiemo Wambgsanss, David Zhou, Emad A. Alghamdi, Tal August, Avinash Bhat, Madiha Zahrah Choksi, Senjuti Dutta, Jin L.C. Guo, Md Naimul Hoque, Yewon Kim, Simon Knight, Seyed Parsa Neshaei, Antonette Shibani, Disha Shrivastava, Lila Shroff, Agnia Sergeyuk, Jessi Stark, Sarah Sterman, Sitong Wang, Antoine Bosselut, Daniel Buschek, Joseph Chee Chang, Sherol Chen, Max Kreminski, Joonsuk Park, Roy Pea, Eugenia Ha Rim Rho, Zejiang Shen, and Pao Siangliulue. A design space for intelligent and interactive writing assistants. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pages 1–35, Honolulu HI USA, May 2024. ACM.
- [6] Taylor Lincoln and Negah Mouzoon. Cranes and derricks: The prolonged creation of a key public safety rule. Technical report, Public Citizen, 2011.
- [7] Lauren Moxley. E-rulemaking and democracy. *Administrative Law Review*, 68:661, 2016.
- [8] Occupational Safety and Health Administration (OSHA). Permissible exposure limits – annotated tables. Retrieved November 1, 2024 from <https://www.osha.gov/annotated-pels>.
- [9] Catherine O'Neill, Amy Sinden, Rena I. Steinzor, James Goodwin, and Yee Huang. The hidden human and environmental costs of regulatory delay. *SSRN Journal*, 2009.
- [10] Frank Pasquale. Battle of the experts: The strange career of meta-expertise, March 2021.
- [11] Julia Romberg and Tobias Escher. Making sense of citizens' input through artificial intelligence: A review of methods for computational text analysis to support the evaluation of contributions in public participation. *Digital Government: Research and Practice*, 5(1):1–30, March 2024.
- [12] Max Weiss. Investigation plan: Deepfake submissions on federal public comment servers. Technical Report