American Consumer Institute
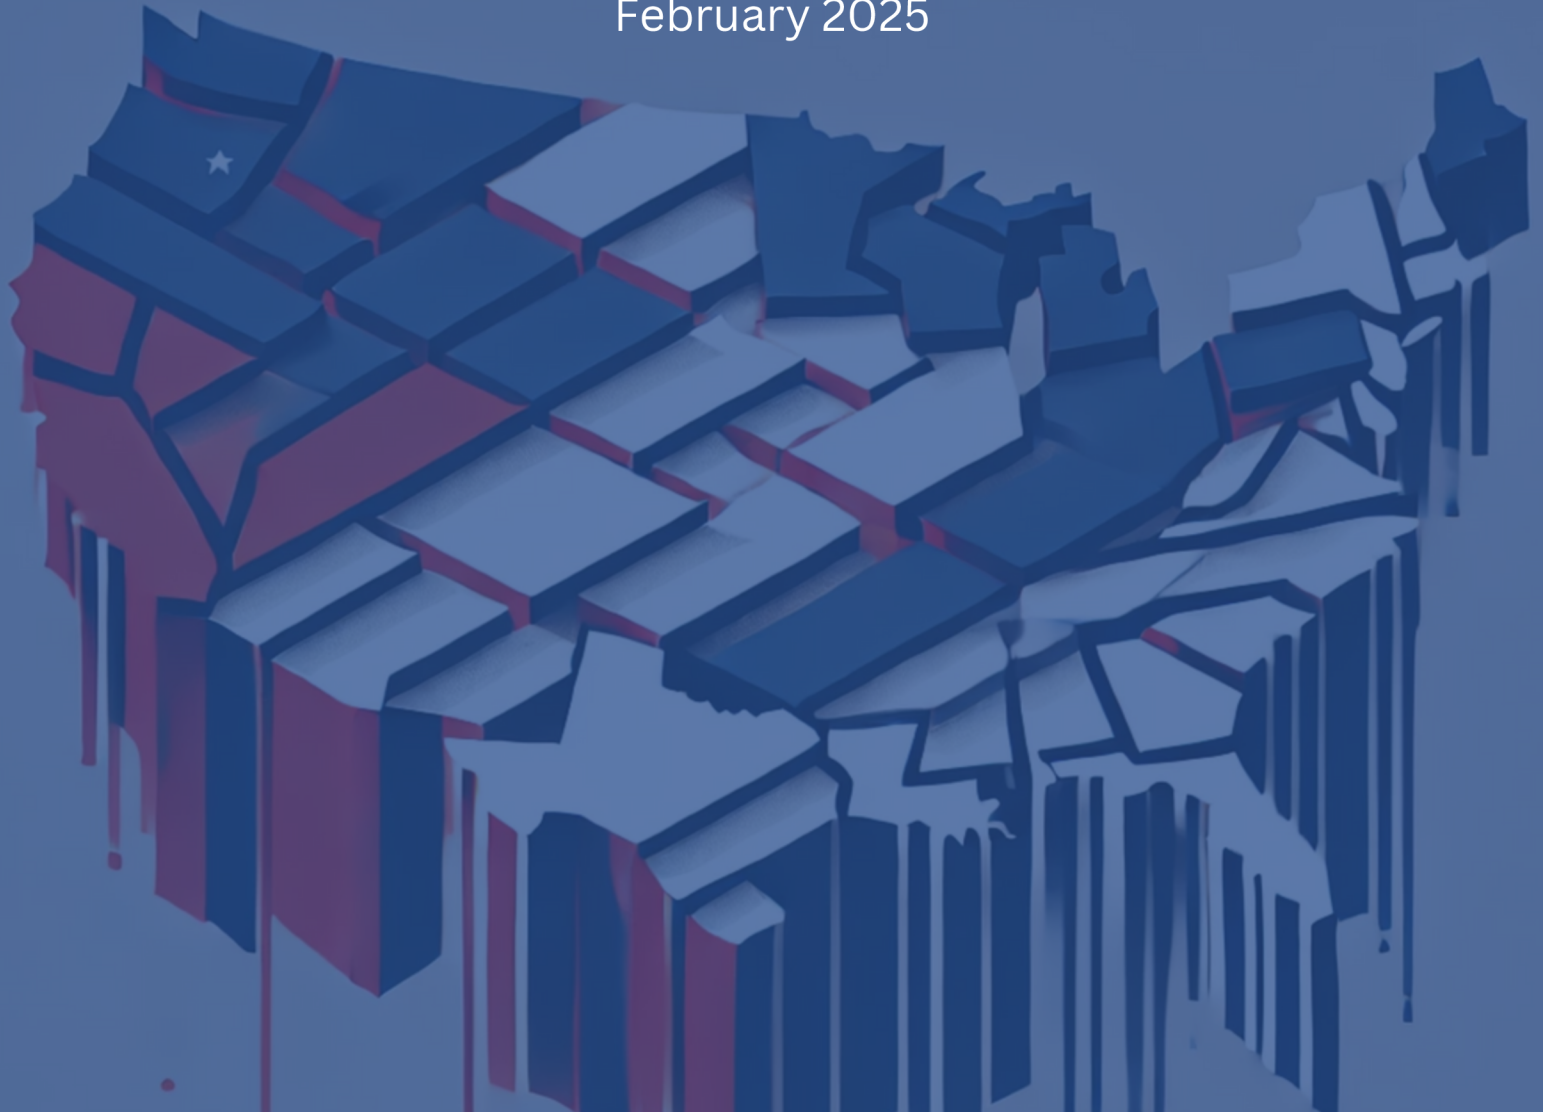
ConsumerGram

Irresponsible Collaboration: Evidence of a Growing AI Fairness Patchwork

Logan Kolas and Nate Karren

February 2025

Executive Summary

More than 200 state lawmakers from over 45 states have been meeting as a part of the Multistate AI Policymaker Working Group to coordinate AI policy across the country.[1] Despite widespread participation, states are still embarking on what will likely turn out to be a patchwork approach to AI governance, especially in AI fairness law. Colorado took the lead in passing legislation of this kind in May 2024, despite reservations from Governor Jared Polis, who warned that the Colorado law could hurt "an industry that is fueling critical technological advancements across [the] state for consumers and enterprises alike." Governor Polis rightly feared that a state level "patchwork across the country can have the effect to tamper innovation and deter competition in an open market."[2] He wrongly signed the bill anyway: igniting the beginning of an innovation destroying AI patchwork just as the next technological revolution sweeps the planet.

Other states are now following Colorado's lead, making their own modifications to the law, thereby spinning a web of vaguely similar but not the same algorithmic fairness AI laws—each with its own unique scope applications, reporting requirements, and enforcement provisions. Legislation has already been introduced in Arkansas, California, Connecticut, Georgia, Hawaii, Illinois, Maryland, Massachusetts, Nebraska, New Mexico, New York, Oklahoma, Texas, Vermont, and Viriginia [Figure 1].[3] Utilizing standard tools of text analysis, we find significant evidence of a growing AI patchwork threat within state AI fairness legislation. An overall assessment of word overlap of the proposals finds significant differences. Even accounting for word frequency and rarity, legislative proposals differ significantly. To avoid a new patchwork of administrative rules, federal lawmakers should pre-empt a growing labyrinth that threatens to smother AI innovation before it can get off the ground.

[1] Future of Privacy Forum, "Future of Privacy Forum Convenes Over 200 State Lawmakers in AI Policy Working Group," FPF.org, October 21, 2024, https://fpf.org/press-releases/future-of-privacy-forum-convenes-over-200-state-lawmakers-in-ai-policy-working-group/; Adam Thierer, "There are Two Different AI Policy Debates, and You Need to Start Paying More Attention to the One about "'Woke AI,'" *Medium*, November 11, 2024, https://medium.com/@AdamThierer/there-are-two-different-ai-policy-debates-and-you-need-to-start-paying-more-attention-to-the-one-41f2c473a5f9; and Dean W. Ball, "AAmerica is Sleepwalking Into a Permanent DEI Bureaucracy That Regulates AI," Pirate Wires, November 11, 2024, https://www.piratewires.com/p/america-is-sleepwalking-into-a-permanent-dei-bureaucracy-regulating-ai.
[2] David Wright Tremaine, "SB 24205 Signing Statement," May 2024, Future of Privacy Forum, "FPF Multistate AI Policymaker Working Group: Summary of October 2024 Meeting," Google Drive, May 2024, https://drive.google.com/file/d/1i2cA3IG93VViNbzXu9LPgbTrZGqhyRgM/view.
[3] Author Research; and Adam Thierer, "Updated Compendium of Bills Pushed by the Multistate AI Policymaker Working Group," *Medium*, January 23, 2025, https://medium.com/@AdamThierer/updated-compendium-of-bills-pushed-by-the-future-of-privacy-forum-fpf-multistate-ai-policymaker-40cb0566cb2f.

Figure 1: States Pursuing AI Fairness Legislation

A map of each of the states have passed or introduced AI Fairness Bills
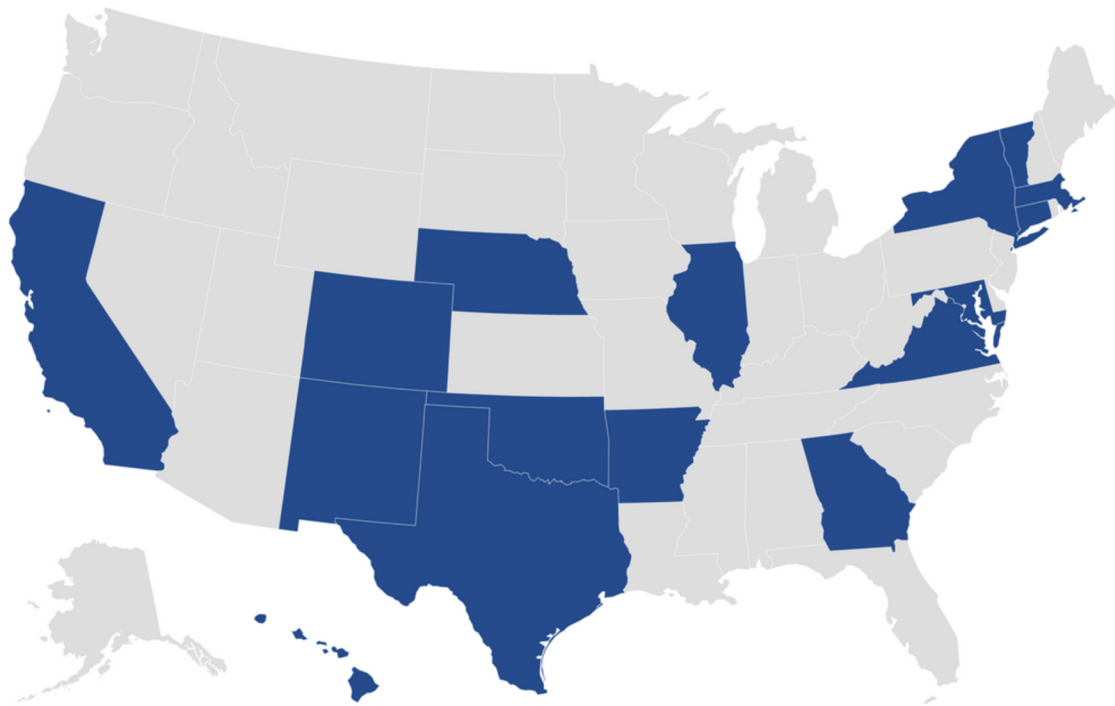
Dangers of an AI Patchwork

Benefits of uniformity in AI law are hotly debated. Laboratories of democracy have historically generated competition between state governments to create the best laws, which could then be exported and adopted by other states. Patchworks of different rules create hurdles and compliance costs that are burdensome to business. In the context of advanced algorithmic computation and AI—which inherently involves free expression and interstate commerce—such an outcome is often not desirable, especially as compliance costs and regulatory burdens proliferate under cumbersome frameworks like what has been proposed. To lower patchwork and compliance costs, states theoretically could approximate a federal standard in AI fairness law by passing identical laws in statehouses across the country. Such a framework would avoid messy compliance and administrative burdens that have been shown to concentrate markets in other technology policy arenas, such as privacy,[4] but raises concerns about collusion, often referred to as "cartel federalism," where states collaborate on keeping regulatory burdens elevated. Collusion foists the costs on everyone by insulating states from intergovernmental competition. More importantly, uniformity is only valuable insofar as it promotes a light-touch, regulatory structure that is limited in scope and allows innovation and competition to flourish. In other words, uniformity is only as wise as the foundational idea that it spreads.

[4] Will Rinehart, As If We Needed More Evidence There's a Privacy-Competition Tradeoff, Exformation, March 3, 2022, https://exformation.williamrinehart.com/p/as-if-we-needed-more-evidence-theres.

Picking a model and sticking to it is, of course, politically and practically tricky. Proposals must comply with political requirements in different regions of the country. Additionally, state lawmakers naturally want to improve their versions of the bill—but gathering support to pass the law through the state legislature often involves compromise, further increasing legislative differences and patchwork costs. Texas, for example, proposed one of the most strict and cumbersome versions of an AI fairness bill—but wisely modified it to address complaints from experts.[5] Texas lawmakers were right to amend the bill, but even those changes differentiated it from other states that copied their approach, thus increasing patchwork costs. As more states debate, amend, and then pass their own independent versions of AI algorithmic fairness doctrines, regulated businesses will face an increasingly complex and burdensome mishmash of rules and regulations that significantly raise the cost of business, particularly among smaller firms that lack the resources of bigger companies.

Methodology

To assess the relative similarity between algorithmic fairness proposals in different states, we deploy two commonly used methods in text analysis to assess similarity in Python: Jaccard score (J-Score) and Cosine score (C-Score). The raw bill is extracted, normalized for formatting and spacing, and then leveraged to calculate the similarity through the following methods, both scored on a scale of 0 to 1:

- Jaccard Similarity Score (J-Score): Shared Words / Total Unique Words [6]

- Cosine Similarity Score (C-Score): Vector-based method, paired with TF-IDF to weight frequency and rarity [7]

The Jaccard similarity score is a conservative measure deployed to assess word overlap. Meanwhile, the Cosine similarity score infers importance from word frequency and rarity. The Cosine similarity score is normalized to account for bill text length differences and is considered a more aggressive approach, which down weights common words. Since every word in bill text could shift the legal interpretation of the law, down weighting seemingly innocuous terms could mask nuance in legal interpretation. The Jaccard score is more likely to understate similarity while the Cosine score is likely to overstate similarity. Together, the scores provide a range of bill text similarity.

[5] Dean W. Ball, "Texas Plows Ahead," Hyperdimensional, January 2, 2025, https://www.hyperdimensional.co/p/texas-plows-ahead.
[6] See Appendix for Example from ChatGPT.
[7] See Appendix for Example from ChatGPT.

Discussion of Results

American Consumer Institute analysis finds evidence of a growing AI fairness patchwork [Figure 2]. Comparing all other AI fairness proposals to Colorado—the only bill currently signed into law—Jaccard scores range from 12 percent (Maryland and Oklahoma) to 60 percent (Nebraska), and average 35 percent. Cosine scores, which are likely to overstate similarity, range from 12 percent (California) similar to 91 percent (Massachusetts) similar, with an average Cosine score of a meager 51 percent.

Virginia has passed their iteration along party lines and the proposal now awaits a decision from Governor Glenn Youngkin. Virginia does a better job than their peers in sticking to the framework set out by Colorado, but their proposal is still less than 90 percent similar, even under the more forgiving Cosine score.

Given the aggressiveness of the approach in Texas, we also score state AI fairness proposals according to how similar they align to the Texas Responsible AI Governance Act (TRAIGA). Virginia ranks most similarly, with a Cosine score of 84 percent similarity, with Colorado ranking 78 percent similar —suggesting that Virginia has blended the approaches of Colorado with other state features. The variance in state-level scores—and the fact that no state has eclipsed more than 91 percent similarity in comparison to Colorado or Texas—we conclude that there is significant evidence of a growing AI fairness patchwork.

Figure 2: Similarity Score Comparison Results

State	Bill No.	Overall Similarity Scores			
		Texas Baseline		Colorado Baseline	
		J-Score	C-Score	J-Score	C-Score
Arkansas	SB258	36%	61%	37%	50%
California	AB1018	27%	28%	28%	30%
California	AB-2930	27%	11%	32%	12%
Colorado	SB24-205	32%	78%	100%	100%
Connecticut	SB2	36%	79%	39%	78%
Georgia	SB167	32%	43%	50%	46%
Hawaii	SB59	22%	23%	22%	19%
Hawaii	SB2572	21%	66%	21%	49%
Illinois	SB2203	22%	30%	28%	34%
Maryland	SB936	33%	74%	48%	78%
Maryland	SB957	11%	10%	12%	13%
Maryland	HB1255	10%	10%	12%	13%
Massachusetts	HD396	19%	40%	23%	38%
Massachusetts	HD4053	27%	80%	54%	91%
Nebraska	LB642	32%	75%	60%	84%
New Mexico	HB60	31%	77%	54%	83%
New York	A768	33%	68%	52%	75%
New York	SB1962	33%	68%	52%	75%
Oklahoma	HB3835	21%	24%	25%	27%
Oklahoma	HB3453	11%	30%	12%	26%
Oklahoma	HB1916	20%	33%	19%	25%
Texas	HB1709	100%	100%	32%	78%
Vermont	H710	31%	65%	37%	65%
Virginia	HB2094	37%	84%	50%	88%

Conclusion

No matter what state policymakers, experts, or other nonprofit organizations may say on the matter, state governments have not delivered a system of state proposals that approximates a uniform national algorithmic fairness framework on artificial intelligence. Legislative proposals vary drastically, driving up compliance costs for businesses innovating to bring new and different AI applications and tools to consumers. State lawmakers should avoid this approach to AI governance entirely or at least resist the temptation to make additional tweaks in the legislative process, but that is unlikely. As state policymakers build support coalitions and expend political capital to amend these proposals in line with perceived needs in their states, frameworks will continue to evolve, making the problem even worse. Federal policymakers should pre-empt that outcome by asserting supremacy on the issue with a clear, concise, simple, and light-touch national standard.

American Consumer Institute

Appendix

To understand how Jaccard and Cosine similarity are calculated, we asked ChatGPT to give a hypothetical example of each. Two examples of the calculation are listed below.

Jaccard Similarity Calculation:

1. Unique words in State A = {consumer, means, person, engages, transactions, artificial, intelligence, system, directly, affected, use, system}

2. Unique words in State B = {consumer, means, individual, Colorado, resident}

3. Intersection (Shared words) = {consumer, means}

4. Union (All words combined) = {consumer, means, person, engages, transactions, artificial, intelligence, system, directly, affected, use, individual, Colorado, resident}

5. Jaccard Score = Intersection / Union = 2 / 14 = 0.14

Cosine Similarity Calculation:

1. Convert definitions into word frequency vectors.
   - State A: {consumer: 1, means: 1, person: 1, engages: 1, transactions: 1, artificial: 1, intelligence: 1, system: 2, directly: 1, affected: 1, use: 1}
   - State B: {consumer: 1, means: 1, individual: 1, Colorado: 1, resident: 1}

2. Compute dot product of vectors.

3. Normalize by the magnitude of each vector.

4. Cosine Score ≈ 0.22

Disclaimer: Appendix generated by ChatGPT with minor tweaks by authors.