

Making Local Law Legible: LLM-Assisted Detection of Discrimination

Dan Bateyko,* Yasmine Mabene,* Derek Ouyang, Daniel E. Ho

The Problem

Local law matters. Across the United States, local law draws lines determining who can vote, who can hold particular jobs, who can attend which schools, and who can be detained or punished. Historically, local lawmakers have often drawn those lines unjustly across protected categories like race, sex, and religion. Ridding codes of these laws removes enforcement risks and public misunderstanding and clears historical stains.

However, making local law legible takes time. Until recently, reformers combed through codes by hand with small armies of researchers and lawyers. In the 1950s, Pauli Murray compiled over 700 pages of segregation laws — a volume Thurgood Marshall called the “bible” of *Brown v. Board of Education*. In the 1970s, then-Professor Ruth Bader Ginsburg led an effort to excise sex bias from federal law, searching the law using just 59 gender-related keywords. In the 2000s, law professor Gabriel “Jack” Chin and his team cataloged vestigial Jim Crow provisions in state statutes, triggering repeals in several states — while acknowledging they likely missed unusually phrased laws. As recently as 2019, a Virginia commission compiled 60 years of discriminatory state laws only by mobilizing law students from three universities, the governor’s staff, practicing attorneys, and a sitting judge. The time and labor required for these painstaking searches are significant. No comparable effort has ever been possible across local codes. There are simply too many.

Key Takeaways

Understanding local law can be overwhelming. Historically, reformers have hired large teams of researchers and lawyers to sift through federal and state codes, but that is particularly infeasible for thousands of cities, counties, and towns across the United States.

Stanford HAI and RegLab researchers gathered 9,623 jurisdictions’ local laws (amounting to 3 billion words of code) and developed an LLM-assisted review pipeline to demonstrate the potential of AI for legal reform. Together, these laws govern in jurisdictions that are home to roughly 252 million Americans.

We identify patently discriminatory laws. Shockingly, we document dozens of jurisdictions that formally still mandate racial segregation, impose poll taxes, and deny women the right to vote, practices that have long been held unconstitutional. A particularly prevalent category is unwarranted citizenship-based discrimination for occupational licenses.

These methods empower lawyers, local governments, and citizens to understand and target reform efforts. LLM-assisted statutory surveys have the potential to update and make the legal code legible again.

In our paper “[Hidden in Plain Text: LLM-Assisted Detection of Discriminatory Local Laws](#),” we present the first effort to develop AI assistance for surveying discriminatory provisions at scale and systematically across local codes. We gathered the largest dataset of local laws to date — spanning 9,623 jurisdictions and 75% of the U.S. population — and used large language models (LLMs) to flag likely instances of discriminatory laws for subsequent human verification.

We found plainly discriminatory laws across the United States in nearly every protected category we searched. These laws can be viewed on our paper’s [companion website](#). We present our approach as a promising solution to a search problem that has hindered code cleanup efforts for decades — and as a foundation for further work on statutory surveys that turn on legal, rather than topical, questions.

*We found plainly
discriminatory laws across
the United States, spanning
nearly every protected
category we searched.*

The Solution

In [previous work](#) from the Stanford RegLab, our colleagues used LLMs for [statutory and regulatory research](#) and found outdated, obsolete, and

cumbersome legal requirements. Yet those statutory surveys did not assess the potential for AI systems to assist in determining the constitutionality of legal provisions, an inquiry whose answers bear on people’s legal rights and equal civic standing.

Here, we introduce a new LLM-assisted method specifically for identifying unconstitutionally discriminatory laws. The method draws on LLM and human judgment, and focuses on laws that are discriminatory in ways that are identifiable from the text itself when read against current constitutional doctrine. (We focus on “disparate treatment” and not facially neutral laws with discriminatory effects.) This necessarily requires developing an approach that is more attuned to antidiscrimination law and legal reasoning than prior work.

We first assembled a corpus of local laws — covering roughly 75% of the U.S. population, spanning 9,623 jurisdictions, and totaling 3 billion words of code — by scraping and parsing codes from two major codification providers. The dataset contains approximately 9 million law sections.

We then used an LLM to label each legal section based on whether it references a protected category — such as race, religion, or gender — and whether the law imposes differential treatment based on that category. We grouped laws by the type of treatment they impose, allowing reviewers to filter out laws where differentiation may be used for accommodations or remedial purposes. Finally, an LLM scored the remaining provisions on five attributes drawn from antidiscrimination doctrine — attributes that signify the presence of discriminatory treatment without a recognized justification. We used multiple LLMs to assign a priority rating based on LLM agreement to direct human review.

Illustration

We identified discriminatory — and likely unconstitutional — laws across the United States. These include actively enforced laws, such as citizenship requirements for occupational licensing, and unenforced, historical vestiges, including Jim Crow-era segregation and disenfranchisement.

Table 1. Examples of discriminatory provisions appearing in local codes and charters available online

PROVISION TYPE	JURISDICTIONS AFFECTED (LOWER BOUND)	RESIDENTS IN AFFECTED JURISDICTIONS	EXAMPLE
Racial segregation	15	1.8M	“The town council shall have the power to enact ordinances ... providing separate cemeteries for white and black ” (<i>Smithfield, NC, Charter § 37</i>)
Race- or gender-based voter eligibility	29	101K	“Voters ... shall consist of all male persons of the age of twenty-one and upwards.” (<i>Opp, AL, Charter § 5</i>)
Race-based voting restrictions	18	189K	“No person shall be permitted to vote [without] pay[ing] ... poll taxes. ” (<i>Valparaiso, FL, Charter § 4</i>)
Gender-based indecent language restrictions	19	400K	“It shall be unlawful ... to use indecent or obscene language in the presence ... of any woman or child. ” (<i>Walled Lake, MI, Code § 50-157</i>)
Women-only police roles †	50	2.6M	“The police department shall consist of ... patrolmen ... a matron of police, two police women. ” (<i>Holyoke, MA, Code § 17-16</i>)
Citizenship-based occupational licensing	1.6K	49M	“ No [junk dealer] license shall be issued ... if the applicant is ... an alien. ” (<i>Alameda County, CA, Code § 3.48.060</i>)

† Although now antiquated, lawmakers originally enacted many gender-specific police roles as remedial measures to expand women’s participation.

Race and ethnicity. More than 70 years after *Brown v. Board of Education* (1954), we documented dozens of jurisdictions that still mandate racial segregation. The charter for Memphis, Tennessee, provides: “Said board of education shall provide and maintain separate schools for the use and accommodation of the white and colored youths of the City.” (Charter § 1000). Segregation is not the only Jim Crow holdover. Poll taxes have been unconstitutional in all elections since the 1960s, yet the charter for Valparaiso, Florida, still conditions voting on having “pa[id] ... poll taxes” (Charter § 4).

Sex and gender. We identified nearly 2,000 laws discriminating on the basis of sex or gender. These include provisions restricting voting eligibility to men, such as an Alabama city’s charter providing that “voters . . . shall consist of all male persons of the age of twenty-one.” (City of Opp, AL, Charter § 5).

Alienage. Notably, we found over 2,000 laws nationwide barring noncitizens from obtaining occupational licenses. These restrictions formally prevent noncitizens from operating businesses ranging from go-cart tracks and bowling alleys to hotels and bars.

Classifications based on alienage are subject to strict legal scrutiny.¹ Under this standard, the Supreme Court has struck down citizenship requirements for occupational licenses — ranging from notaries to attorneys² — carving out only a narrow exception for political functions, such as police officers.³ Licensing rules for bowling alleys and hotels would fall outside that exception.

Some of these laws are nearly a hundred years old. In the 1930s, most states and major cities passed laws excluding noncitizens from dozens of occupations.

More than 70 years after Brown v. Board of Education (1954), we still found provisions mandating racial segregation.

Other provisions are more recent. A 1993 ordinance in Madison County, Illinois, imposes costs on residents with limited English proficiency by requiring them to pay for their own interpreters in administrative hearings. Such requirements are at odds with federal law: In 2010, the Department of Justice’s Civil Rights Division advised that requiring litigants with limited English proficiency to pay for their own interpreters is inconsistent with the language-access obligations that Title VI of the Civil Rights Act imposes on jurisdictions receiving federal funds.⁴

Offensive language. Separately from discriminatory laws, we flagged thousands of provisions for derogatory or offensive language, a broad net that includes explicit slurs, outdated clinical terminology, and references to bygone practices like minstrel and “freak” shows. One example, from Euclid, Ohio: “Every person . . . housing . . . invalids, children, aged persons or individuals adjudged . . . as feeble-minded, insane, lunatics, imbeciles or idiots, shall . . . comply with all the laws of the State” (Code § 753.13).

While many discriminatory laws are historical vestiges, many continue to be enforced. Take liquor licensing. The Nunes family — Brazilian American green card holders — were recently barred from applying for a liquor license under Massachusetts state law. The Boston Licensing Board similarly denied a liquor license to a bakery owner, requiring his wife, a U.S. citizen, to apply instead. Cases like the Nunes family’s prompted

Massachusetts lawmakers to amend the state's 1933 citizenship-restriction liquor statute in 2025. Nearly a year later, however, we found that similar restrictions persist in Massachusetts towns including Lexington, Topsfield, Wakefield, and Orleans, underscoring the need for more efficient methods to detect outdated rules and keep municipalities aligned with state law.

Recommendations

Local codes are sticky. While reform efforts have focused on federal and state law, local codes offer no guarantee of following suit.

Capacity and sludge. Policy sludge — outdated, obsolete, or unworkable legal requirements — accrues when no one is responsible for removing it. Most Americans are governed by part-time local lawmakers working with little or no dedicated legal staff. Absent a triggering event, codes go decades without systematic review. LLM-assisted methods may dramatically lower the cost of systematic reform efforts.

LLM-assisted methods may dramatically lower the cost of systematic reform efforts.

The structured workflow we developed can replace prior, painstaking manual enumeration that impedes reform. Our pipeline makes human review tractable for scholars, policymakers, legal clinics, or advocates. For instance, San Francisco undertook a code cleanup effort in collaboration with Stanford RegLab to cut over a third of its outdated reporting requirements.

Jurisdictions should follow suit and utilize these tools to engage in code cleanup.

Public codes, privately hosted. Access to municipal codes is also a data infrastructure challenge. Third-party, private code publishers host public versions of the code. Such code is often not available on typical legal search platforms, and date-stamped versions are inaccessible to the public. The difficulty of conducting sophisticated searches on municipal codes (with some lawyers simply using Control+F on local PDF copies) poses another impediment to reform. Making such codes available in bulk data (as the U.S. Congress has done) could improve innovation, understanding, and accessibility to law.

Repository of reforms. We found nearly identical discriminatory laws in adjacent jurisdictions, often because municipalities borrow language from one another. Repositories of reforms — such as Massachusetts removing the citizenship requirement to operate a liquor store — could also power reform efforts, as identifying sister provisions across Massachusetts municipalities becomes very cheap with AI-assisted search. Similarly, while our pipeline focused on operative codes, some jurisdictions keep historical provisions for archival or transparency purposes, which can also become seed instances for superseded or unenforceable legal provisions. In the same way that *model acts* serve as templates for legislatures to enact statutes, *model reforms* could consolidate provisions for omnibus cleanup.

Doctrinal frontier. While our approach shows great promise, we also document that it is no substitute for human judgment. The pipeline, designed with basic antidiscrimination law principles in mind, radically reduces the search cost, but candidate provisions still require human validation. For instance, LLMs do not exhibit reliable judgment in applying rationales across

types of occupational licenses. Much productive work remains to train and develop systems capable of reasoning about complicated, conflicting, and dynamic bodies of law.

Nationwide, discriminatory laws affect people's livelihoods in ways that seldom make headlines. A resident reading their local code may not be able to distinguish between dead-letter statutes and those that can be enforced. This is the tragedy of legal legibility: Codes were written to make the law scrutable, but policy sludge guarantees that they are not. Those denied opportunities under discriminatory laws may never learn that the rule itself is unlawful. Even when unenforced, discriminatory laws are relics of the past that should no longer haunt the legal code.

Endnotes

- 1 See *Graham v. Richardson*, 403 U.S. 365 (1971).
- 2 See, e.g., *Bernal v. Fainter*, 467 U.S. 216 (1984) (notary public); *In re Griffiths*, 413 U.S. 717 (1973) (attorneys), *Examining Board v. Flores de Otero*, 426 U.S. 572 (1976) (engineers).
- 3 *Foley v. Connelie*, 435 U.S. 291 (1978).
- 4 Moreover, the Supreme Court has recognized that failure to provide adequate language assistance can constitute national origin discrimination. See *Lau v. Nichols*, 414 U.S. 563 (1974).

Reference: The original article is accessible at Dan Bateyko,* Yasmine Mabene,* Derek Ouyang, Daniel E. Ho, “**Hidden in Plain Text: LLM-Assisted Detection of Discriminatory Local Laws,**” International Conference on Artificial Intelligence and Law (ICAIL) 2026, <https://hidden-in-plain-text.reglabapp.com/hidden-in-plain-text.pdf>.

Stanford University’s Institute for Human-Centered Artificial Intelligence (HAI) applies rigorous analysis and research to pressing policy questions on artificial intelligence. A pillar of HAI is to inform policymakers, industry leaders, and civil society by disseminating scholarship to a wide audience. HAI is a nonpartisan research institute, representing a range of voices.

Stanford University’s Regulation, Evaluation, and Governance Lab (RegLab) partners with government agencies to design and evaluate programs, policies, and technologies that modernize government.

The views expressed in this policy brief reflect the views of the authors. For further information, please contact HAI-Policy@stanford.edu.



Dan Bateyko is a PhD student in information science at Cornell University and was a Summer Institute Fellow at the Regulation, Evaluation, and Governance Lab (RegLab) at Stanford University.



Yasmine Mabene is a research fellow at the RegLab at Stanford University.



Derek Ouyang is a research director at the RegLab at Stanford University.



Daniel E. Ho is the William Benjamin Scott and Luna M. Scott Professor of Law, Professor of Political Science and Computer Science (by courtesy), Associate Director and Senior Fellow at the Stanford HAI, Senior Fellow at the Institute for Economic Policy Research (SIEPR), and director at the RegLab.



Stanford University
Human-Centered
Artificial Intelligence

Stanford HAI: 353 Jane Stanford Way, Stanford CA 94305-5008

T 650.725.4537 **F** 650.123.4567 **E** HAI-Policy@stanford.edu hai.stanford.edu