KEEPING AI IN ITS ZONE

Sara C. Bronin

ABSTRACT

With zoning-reform fever sweeping local governments, state legislatures, and even Congress, there is an urgent push to understand exactly how zoning works. Unfortunately, zoning is a particularly opaque body of law, and zoning codes are arguably the most complex legal documents in the American legal canon due to their length, inconsistencies, interpretation challenges, and regulatory scope. Given this complexity, it may be tempting to believe that we should use artificial intelligence (AI) to make sense of zoning. Indeed, some recent scholarship has asserted that AI can understand zoning codes as well as, or better than, humans. This Article aims to temper such claims. By honestly acknowledging AI's limitations, we can redirect our energy to harnessing AI's *actual* capacity to help us understand zoning, the hidden power that shapes our lives.

The Article draws from the author's experience building the National Zoning Atlas (NZA), the largest geocoded zoning-data repository ever created, to posit that zoning codes will likely continue to stubbornly defy AI interpretation—even if the models improve and the relevant training corpus expands. Part I identifies relevant regulatory attributes, including inconsistent structure and terminology, the prevalence of alternative regulations, and mismatches between text and map. These attributes confirm that zoning is a dynamic body of law that demands interpretation, contextualization, and other practices of human reasoning, and is not a static dataset waiting to be harvested. Part II goes on to provide a taxonomy of recent, problematic scholarship that uses AI to characterize zoning codes, predict results from zoning codes, and extract information from zoning codes. This scholarship's overestimation of AI's capabilities will only achieve one thing: helping us misunderstand the law faster. Part III expands on the argument that zoning codes challenge AI by revealing the disappointing yet foreseeable results of AI trials conducted by the NZA. Part IV suggests recalibrating our expectations about how AI can be successfully applied to zoning, offering specific characterizing, predicting, and extracting applications that avoid direct code interpretation and require only limited human intervention.

KEEPING AI IN ITS ZONE

Sara C. Bronin

TABLE OF CONTENTS

INTRODUCTION……3

I. INTERPRETING ZONING CODES REQUIRES HUMAN REASONING……7
A. How Zoning Codes Work
B. Interpretation Challenges in Zoning Texts
1. Structure
2. Terminology
3. Alternative Regulations
C. Interpretation Challenges in Zoning Maps
1. Hard-to-Decipher Maps
2. Mismatches Between Texts and Maps

II. CRITIQUING PROBLEMATIC APPLICATIONS OF AI TO ZONING……22
A. Characterizing Applications
B. Predicting Applications
C. Extracting Applications

III. UNRESOLVABLE CHALLENGES: THE NATIONAL ZONING ATLAS AI TRIALS……42
A. Trial 1: Methods, Results, and Lessons
1. The District-Name Extraction Application
2. The District-Classification Characterizing Application
B. Trial 2: Methods, Results, and Lessons
1. Methods
2. Results and Lessons
C. Implications on AI Scholarship About Zoning and Beyond

IV. RECALIBRATING EXPECTATIONS ABOUT AI IN ZONING……54
A. Characterizing Applications
B. Predicting Applications
C. Extracting Applications
1. Deriving Limited Textual Information
2. Transforming Zoning Maps
3. Aggregating Amendments

CONCLUSION……62

KEEPING AI IN ITS ZONE

Sara C. Bronin[*]

INTRODUCTION

In recent years, zoning has received unprecedented levels of public scrutiny, as it is increasingly blamed for high home prices, farmland and forest loss, and constraints on mobility and opportunity. With zoning-reform fever sweeping local governments, state legislatures, and even Congress, there is an urgent push to understand exactly how zoning works. It may be very tempting to believe that artificial intelligence (AI) can fast-track this effort.[1] Indeed, some scholarship has asserted that AI can understand zoning codes as well as, or better than, humans.[2]

---

[*] Sara C. Bronin is the Freda H. Alverson Professor of Law at the George Washington University Law School and the founder of the National Zoning Atlas. She would like to thank Professor Alexander Rush of Cornell Tech for the excellent research partnership described in Part III.B. and the National Science Foundation Human-Environment and Geographical Sciences Program for making our collaboration possible; the American Bar Association Housing and Community Development Section, the George Washington University Law School, and the MIT Leventhal Center for Advanced Urbanism and MIT Joint Program in City Design for hosting workshops; Professors Barak Richman, Patricia Salkin, Jonathan Siegel, and Alicia Solow-Niederman, as well as Scott Markley, Matthew Harris, and Diana Drogaris from the National Zoning Atlas for their review and comment; Denise Ssettimba and Daniel Woods for their research assistance; Judah Axelrod, Lydia Lo, and the Urban Institute for the short but fruitful collaboration mentioned in Part III.A.; National Zoning Atlas analysts who contributed to the manual zoning data collection; Jan Carbonell, Justin Chiu, Max Dumas, Ethan Hao, Junbo Huang, Eesha Khanna, Celine Lee, Jan Llobet, Chenran Ning, Nan Jiang, Ruslana Yurtyn, Shaden Shaar, Angky William, Nathan Yan, Wenting Zhao, and others who contributed to the AI trial described in Part III.B; and Anthony La for pulling key data.

[1] Artificial intelligence (AI) is defined as "software used to perform tasks or produce output previously thought to require human intelligence, esp. by using machine learning to extrapolate from large collections of data." *Artificial Intelligence (n.)*, OXFORD ENGLISH DICTIONARY (Dec. 2023), https://doi.org/10.1093/OED/3194963277. This definition is aligned with the federal statutory definition of AI as a "machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments." 15 U.S.C. § 9401(3).

Similarly, the term "algorithm" has varied definitions depending on its context. An algorithm is defined as a "clearly specified mathematical process for computation; a set of rules that, if followed, will give a prescribed result." *Algorithm*, NAT'L INST. OF STANDARDS & TECH.: COMPUT. SEC. RES. CTR. GLOSSARY, https://csrc.nist.gov/glossary/term/algorithm. *But see* Kristian Lum and Rumman Chowdhury, *What Is an "Algorithm"? It Depends on Who You Ask*, MIT TECH. REV., (Feb. 26, 2021), https://www.technologyreview.com/2021/02/26/1020007/what-is-an-algorithm/ (arguing that current definitions are overbroad and to avoid debates on definition, an algorithm should be evaluated by an impact assessment).

AI includes both machine learning applications, used to characterize and predict outcomes from data, and natural language processing, used to extract information from unstructured data (such as text). In this Article, in the interest of brevity, I refer to AI generally and broadly, and while I do not distinguish or separately evaluate generative AI (which relies on natural language processing and other learning tools to synthesize and generate new output, often using statistics-based predictions) and non-generative AI (expert at pattern matching, counting, and similar tasks).

[2] *See supra* Part II.C. (discussing this scholarship).

This Article seeks to temper such claims, and to clarify how, and to what extent, AI can help us make sense of zoning.

Zoning is a particularly opaque body of law. Adopted and administered locally, zoning operates by regulating geographically-defined districts through a text and accompanying map. More than thirty-three thousand local jurisdictions—cities, towns, townships, counties, villages, borough, and even private associations—have been granted the authority to zone. We do not yet know how many such jurisdictions exercise that authority, or how much land in the United States is subject to zoning. In 2021, I founded the National Zoning Atlas (NZA) to catalogue which jurisdictions exercise zoning authority, and to provide the public with standardized information about key regulatory features of zoning codes.[3] To build the NZA, team members first locate official zoning texts and maps, mostly (but not always) using jurisdiction websites.[4] Then, analysts extract district-specific regulatory information from the texts and digitize the maps, following a detailed methodology laying out interpretive principles and practical guidance.[5] The textual and geospatial analysis is merged, reviewed, and published to the online map for each jurisdiction. Over the last five years, the NZA has become the most comprehensive geocoded zoning data repository ever created,[6] encompassing more

[3] National Zoning Atlas, at www.zoningatlas.org/atlas.

[4] One in seven jurisdictions with zoning fails to post their codes online; we then contact the jurisdiction for copies of these documents. That figure comes from a data pull from the NZA database on December 31, 2025, showing that of 5,828 published jurisdictions, 4,983, or 85.5 percent, had a URL in fields requiring NZA analysts to input the URL for the zoning text and zoning map, where such links are available. Analog records have been more prevalent in places like Alabama, Maine, Montana, West Virginia, and Wyoming, where many jurisdictions have very small populations and make limited investments in information technology infrastructure.

One in forty jurisdictions that we believe have zoning declines (even after freedom-of-information-act requests) to provide the full copy of the code; we annotate these jurisdictions on the public map with the following phrase: "This jurisdiction has zoning, but after diligent efforts, we have been unable to find its zoning text, zoning map, or both." The figure comes from a data pull from the NZA database on June 15, 2026, showing that of 6,385 published zoning jurisdictions, 153, or 2.5 percent, were known to have zoning but did not provide complete information about their codes to the NZA.

[5] Since the date of last public version of the methodology, it has been updated and converted to a wiki available only to NZA staff. *See* Sara C. Bronin, Scott Markley, Aline Fader, & Evan Derickson, *How to Make a Zoning Atlas 2.0: The Official Methodology for the National Zoning Atlas*, Apr. 2024.

[6] Prior efforts have focused on small geographic areas. *See*, *e.g.*, Stephen Menendian, Samir Gambhir, & Chih-Wei Hsu, *Single-Family Zoning in Greater Los Angeles*, OTHERING & BELONGING INST. (Mar. 2, 2022), https://belonging.berkeley.edu/single-family-zoning-greater-los-angeles (greater Los Angeles); Stephen Menendian, Samir Gambhir & Marina Blum, *Single-Family Zoning in the Sacramento Region*, OTHERING & BELONGING INST. (July 13, 2022), https://belonging.berkeley.edu/single-family-zoning-sacramento-region (greater Sacramento); Robert C. Ellickson, *The Zoning Straitjacket: The Freezing of American Neighborhoods of Single-Family Homes*, 96 INDIANA L. J. 395 (2021) (Silicon Valley, greater New Haven, and greater Austin); Emma Nechamkin & Graham MacDonald, Predicting Zoned Density Using Property Records, Urb. Inst. Report, Jan. 2019 (Washington, D.C.); AMY DAIN, RESIDENTIAL LAND-USE REGULATION IN EASTERN MASSACHUSETTS (2005) (eastern Massachusetts); Bengte Evenson & William C. Wheaton, *Local Variation in Land Use Regulations*, BROOKINGS-WHARTON PAPERS ON URB. AFF. 221 (2003) (Massachusetts); *Zoning Atlas*, METRO. AREA PLAN. COUNCIL, https://zoningatlas.mapc.org/; *MassGIS Data: 2001 & 2003 Aerial Imagery*, MASSGIS (Apr. 2001

than 1.36 million pages of zoning texts and more than a hundred and twenty thousand distinctly-regulated zoning districts across more than eleven thousand jurisdictions.[7]

The effort has produced a baseline understanding of zoning codes, which in turn has supported both research and policy-making. NZA data has been used in reports on racial discrimination, homeownership rates, transportation, and climate change, among other topics.[8] NZA data has also informed reforms at the local and state levels, undertaken by advocates and elected officials from across the political spectrum.[9] Our information is only useful to others if it is accurate, complete, and collected at scale. Thus we take great care in initial data collection, undertake multiple levels of review, and implement other quality control measures before we deem our information final and publishable.

The NZA team would enthusiastically welcome AI acceleration of our painstaking task of collecting zoning data, and particularly extracting regulatory information from zoning texts. Yet we have not integrated AI into this workflow, because we have realized that zoning is not a static dataset waiting to be harvested, but a dynamic body of law that demands interpretation, contextualization, and other practices of human reasoning. Zoning texts, averaging nearly two hundred pages in length, are arguably the most complex legal documents in the American legal canon. Across jurisdictions, zoning texts have no common structure or terminology.[10] They offer a surprising number of alternative regulations, applying different rules within the same zoning district depending on the situation.[11] Maps can be hard to read, and frequent mismatches between text and map require resolution, often using information and techniques outside the four corners of the

---

& Apr. 2003), https://www.mass.gov/info-details/massgis-data-2001-2003-aerial-imagery#overview-.

[7] To date, of jurisdictions logged in the National Zoning Atlas database, 35 percent have been confirmed to have zoning authority but not exercise it. To date, published zoning jurisdictions include one-third of the total number of jurisdictions but cover land where 62 percent of Americans live. We are unsure whether this portion of jurisdictions without zoning will be maintained as we analyze the remaining two-thirds of jurisdictions across the country. It is possible that the portion of unzoned jurisdictions will increase, not decrease, as the population-per-jurisdiction of the jurisdictions yet to be analyzed is lower than those already analyzed.

[8] *See, e.g.,* Andrew Rumbach, Will Curran-Groome, Sara McTarnaghan, Annie Rosenow, & Sara C. Bronin, *Zoned Into Risk? Toward a Climate-Resilient Development Index for Housing*, URB. INST. (Feb. 2025), https://www.urban.org/research/publication/zoned-risk; Moses Gates, Marcel Negret, Robert Freudenberg, Ellis Calvin, Christine Garner, Sojung Kim, Maulin Mehta, & Sara C. Bronin, *Averting Crisis: Zoning to Create Resilient Homes for All*, REG'L PLAN. ASS'N (Apr. 2025); Yonah Freemark, Lydia Lo, & Sara C. Bronin, *Bringing Zoning into Focus: A Fine-Grained Analysis of Zoning's Relationships to Housing Affordability, Income Distributions, and Segregation in Connecticut*, URB. INST. (June 2023); Kathryn Blanco & Sara Bronin, *Issue Brief: Get on Board for Equitable Transit-Oriented Communities*, DESEGREGATECT (2021). See also research briefs published by the National Zoning Atlas, most recently *Zoning Report: Metro Kansas City* (2026), which cross-walked Census data and zoning conditions across 212 jurisdictions to find that zoning correlated with racial and income-based segregation.

[9] *See Advocacy*, National Zoning Atlas, www.zoningatlas.org/advocacy.

[10] *See supra* Part I.B.1-I.B.2.

[11] *See supra* Part I.B.3. Variance may depend on occupancy, affordability, topography, proximity to certain amenities (such as transit stations or water and sewer infrastructure), lot size, owner identity (including corporate status), and/or many other factors.

zoning code.[12] Zoning codes demand, and indeed were designed for and anticipate, human involvement as a foundational element of the legal system in which codes operate.

These issues make it hard for humans to master a single code, much less to rely on one code to predict the approach of the next. Further, they confound AI model training, which functions best when interpreting consistent, clear, and non-ambiguous documents. Puzzlingly, several pieces of recent scholarship have asserted that AI can understand zoning codes as well as, or better than, humans, and that AI can quickly scale up efforts to understand zoning. Upon examination, these claims prove hollow. They rest on misconceptions about the basic legal function, structure, and content of zoning codes, and they largely ignore zoning's spatial element. These scholars' overclaims about AI's capabilities only scale up one thing: our ability to misunderstand the law faster.[13]

Even the NZA team, arguably the country's leading experts on zoning code interpretation, has been unable to train AI models to deliver results that match or exceed the accuracy and completeness of our own.[14] Starting in 2022, the NZA partnered with outside researchers on two different efforts to deploy AI to extract information from zoning codes more efficiently, with the goal of building the NZA faster. Despite having highly-refined inputs and a manually-produced training corpus of hundreds of annotated zoning texts, the results were disappointing.[15] The models were particularly challenged by inconsistent terms, alternative regulations, and internal code ambiguities. Our team spent more time confirming or correcting the outputs than it would have taken us to conduct the review without AI assistance in the first place.

Why not keep AI in its zone? There is no need to force AI to pursue hopeless tasks: AI has plenty of capacity to expand our understanding of zoning.[16] For example, AI might be used to extract certain limited types of information from zoning codes. AI-informed geospatial processing could accelerate the digitization of zoning maps. In addition, webscraping AI could monitor and flag the existence, frequency, and content of amendments to zoning codes made by jurisdictions. By honestly acknowledging AI's limitations, we can redirect our energy to harnessing AI's *actual* capacity to help us understand zoning, the hidden power that shapes our lives.

This Article proceeds in four parts. Part I explains basic elements of American zoning codes, highlighting the interrelated interpretive challenges they present to both humans and machines alike. Part II goes on to warn against overestimating AI's capabilities. It critiques recent, problematic scholarship that wrongly claims to have used AI to successfully characterize zoning codes, predict results from zoning codes, and extract information from zoning codes. Part III reveals the methods and results of two AI trials conducted by the NZA team to extract information from codes, and discusses the implications of these trials.

[12] *See supra* Part I.C.
[13] *See supra* Part II.
[14] *See supra* Part III.
[15] *See supra* Parts III.A.2, III.B.2.
[16] *See supra* Part IV.A.

Finally, Part IV proposes that we recalibrate our expectations of AI, asking it to achieve tasks that are more viable given both AI's capability and the nature of zoning codes.[17] It identifies characterizing, predicting, and extracting applications that avoid direct code interpretation and can be successfully achieved with only limited human intervention. The Article concludes by observing that delegating the interpretation of zoning law to AI is not just challenging from a practical standpoint, but contrary to zoning's foundational design as a body of law reliant on human cooperation, negotiation, and contestation.

I. INTERPRETING ZONING CODES REQUIRES HUMAN REASONING

It may be tempting to believe that AI can help us better understand zoning codes: we know little about what individual codes actually say, and they are just the kind of tediously long documents humans have little patience for reading. But, as this Part argues, zoning codes—more than most other types of law—demand human engagement for successful interpretation. To support this argument, this Part begins with a brief overview of codes' functional mechanics, explaining how codes' two-part structure, consisting of both text and map, requires cross-referencing and spatial analysis. It also briefly mentions the broader regulatory framework of zoning law, which is structured to resolve applications and approvals through involvement of multiple parties and, often, the general public. It turns next to the interpretation challenges presented by both code parts, drawing heavily from National Zoning Atlas (NZA) findings. On the text side, inconsistencies and complexities in structure, terminology, and content create conflicts that hamper analysis. On the map side, indecipherable images and mismatches between map and text often require multi-step resolution.

The very regulatory characteristics that confound humans, and the framework that necessarily involves humans, present potentially insurmountable challenges to AI interpretation. Humans can hardly be expected to train an AI model to accurately account for codes' myriad illogical, and largely non-replicable, configurations across thousands of jurisdictions. And models cannot be expected to *map* regulatory requirements as required by zoning's two-part structure, much less draw from information beyond the zoning code to resolve ambiguities between text and map in the same way humans might. In explaining how codes are poorly configured for analysis and regulate in unpredictable ways, this Part underscores zoning's fundamental incompatibilities with AI's current interpretative capacity.

A. HOW ZONING CODES WORK

An understanding of how zoning actually works depends on an acknowledgment of four principles: first, the zoning district, which distinctly

[17] *See supra* Part IV.B.

regulates land uses, structures, and lots,[18] is the organizing feature of the code; second, the text laying out regulations for each district governs development in each district; third, the map identifies where zoning districts are assigned, and thus where regulation has effect; and fourth, jurisdictions demonstrate seemingly infinite variety in the way they articulate the rules in the text and delineate districts on the map. Crucially, as discussed below, a zoning code cannot be accurately interpreted by reading the text alone. An interpreter must cross-reference both the text and map to comprehend the scope of each distinctly-regulated district, and the extent to which district overlaps change regulatory outcomes.[19]

The zoning district has been a central organizing element of zoning codes since zoning's inception more than a century ago.[20] Today's districts come in a variety of regulatory flavors, with rules articulated by zoning texts. Districts may allow or prohibit specific uses, mandate particular limitations on structures, and delineate the manner of development of lots. They might be subject to landscaping mandates, parking-lot design guidance, accessory-use restrictions, occupancy requirements, or other rules. These rules come together to articulate how development may proceed in any given district. A residential district, for example, may allow single-family homes or multi-family homes, or both; may limit buildings to two stories; and may require a single acre of land for each unit of housing. A commercial district, meanwhile, may allow offices, restaurants, and apartments; may require buildings to be a minimum of three stories high; and may have no minimum lot size for new development. Local authorities might also create districts allowing industrial, educational, open space, entertainment, or other uses, or mixes of some of all of these. The combination of rules is seemingly endless.

The zoning map delineates the district or combination of districts assigned to a particular parcel. A single lot can be assigned to multiple districts when it is assigned to both base district and one or more overlay districts. When placed atop

---

[18] *See* Restatement (Fourth) of Property: Land Use Volume § 2.1 (Am. Law Inst., Tentative Draft No. 2, 2022). For full disclosure, I was the Associate Reporter solely responsible for drafting the Land Use Volume.

[19] This seems like a pretty obvious point. Unfortunately, many policy interventions do not seem to take zoning maps into account. For example, Congress enacted the bipartisan 21st Century ROAD to Housing Act in 2026 and created the federal PRO Housing program in 2023 to provide funding to government entities to reduce zoning restrictiveness, but it did not require that those entities change some or all of their zoning maps to ensure that additional housing was actually enabled. *See* Sara C. Bronin, *The Road to Producing More Housing Passes Through Meaningful Local Zoning Reforms*, VITAL CITY, July 30, 2026; Consolidated Appropriations Act, 2023, H.R.2617, 117th Cong. (2022). Similarly-written California policies also fail to take the map into account, even when attempting to increase housing supply. *See*, *e.g.*, Letter to California Department of Housing and Community Development, Christopher Elmendorf, Feb. 13, 2022 (on file with author) (advising on a new state regulatory process and noting the "[a]chievement of the [pro-housing] designation does not require cities to adopt prohousing policies that apply citywide or at least to a substantial share of the buildable land in the city").

[20] The zoning district got its start in New York City, which in 1916 adopted a zoning resolution and an accompanying map outlining the boundaries of districts that regulated both uses and building sizes. N.Y.C. ZONING RESOLUTION (1916). In the 1920s, the U.S. Department of Commerce published a model state zoning enabling act solidifying the status of zoning districts as codes' organizing element. U.S. Dep't of Commerce, A Standard State Zoning Enabling Act (1926). All fifty states adopted the Department's model act, with slight variations and amendments across states.

a base district, an overlay district changes the effect of the base zoning in one way or another. For example, a transit-oriented development overlay might be added to an already-zoned area to allow for taller buildings, or a broader mix of uses. Or a floodplain overlay might be overlain atop already-zoned areas to limit or prohibit new development. In dense cities like New York City and San Francisco, many lots have one base district and two, three, four, or even more overlays.[21] Jurisdictions can create, define, and layer zoning districts with virtually no limitation, which has led to incredible diversity across and within counties, regions, and states.

With this background information, the point about the need to consult both text and map to understand a code can be reinforced. Consider the example of a jurisdiction with ten zoning districts. If three of those districts are residential in nature, but seven are commercial (manufacturing, retail, service, office, and so on), it cannot be said, based on the text alone, that the jurisdiction's zoning is primarily commercial. That is because the map could assign 80 percent of land to residential districts, and 20 percent of land to the commercial districts. Similarly, if two of the residential districts allow multi-family housing while the other residential district allows only single-family housing, it cannot be said, based on the text alone, that the jurisdiction is permissive of new housing development. After all, the map may show that each multi-family district is assigned to just a single parcel, while the single-family district is assigned to all others.

Consider another jurisdiction with three residential districts allowing only single-family homes, three commercial districts, and three overlay districts allowing multi-family housing. If one wanted to assess whether the jurisdiction permitted multi-family housing, the text would suggest that it has just as many multi-family as single-family districts. But what if the multi-family districts were each assigned to just one parcel? Jurisdictions often assign multi-family districts to very small areas. Or what if the jurisdiction had not mapped the three overlay districts at all? In that case, the text regarding such unmapped districts would have no legal weight whatsoever.[22] Unmapped districts that only exist in the text should be disregarded in any analysis of what the code actually says.

Similarly, if one wanted to assess whether the jurisdiction just described requires large minimum lot sizes for new housing, one might scan the text of the six residential districts for words like "acres" or "square feet." One might find a 40,000-square-foot minimum lot size for the three single-family housing districts, but no lot-size mandate in the multi-family areas. Where the three single-family districts occupy 90 percent of zoned land in the jurisdiction, the overlay districts occupy 3 percent, and the commercial districts (with no overlay) occupy the remaining, the relative importance of the regulations in the overlay districts pales in comparison to the other residential districts. Yet, with a simple analysis of the text alone, the text of all six districts will be given equal weight. And it may wrongly be seen as a triumph that multi-family housing has no lot-size mandate, rather than a problem that 90 percent of the jurisdiction requires nearly an acre of land for a single home. The calculation would become even more lopsided in a scenario where

[21] One area in Pismo Beach, California, has the largest number of overlays in the NZA dataset: nine.
[22] *See supra* Part I.C.2.

the overlays are unmapped; in those cases, text related to them should not be taken into account when evaluating the jurisdiction's approach to lot sizes.

This brief exposition about the operation of zoning codes has significant implications for the application of AI to zoning. Most significantly, it flags as suspicious any claims that AI can "interpret" zoning codes merely by conducting a textual review. That is because zoning consists of two components, text and map, which *together* constitute the law. One cannot interpret a code without consulting both components, lest the review yield incomplete or highly distorted results. In this sense, zoning differs fundamentally from other areas of law relying solely on text; in short, it requires *more* to be "understood." Moreover, this exposition reveals how the basic structure of codes lends itself to a tremendous amount of variety, even within the common organizing module of the zoning district.

These codes exist within a broader framework of zoning law that expects and facilitates human involvement. The process of adopting, drafting, and amending zoning codes involves multiple bodies (usually an elected officials and an advisory zoning board), as well as the general public. Each zoning application and zoning approval involves, at a minimum staff oversight; even when development is allowed by right, staff review applications to determine compliance with the text and map. Many application and approval processes require additional review by at least one commission (which sometimes holds a public hearing) if not more, and allow for a process of vary from the regulations by appealing to a board of adjustment or appeals. In each of these instances, people determine, advise on, or comment on code implementation. This is not to say that interpreting a code is impossible; it can be done. But ambiguities requiring human intervention should be seen as a feature, not a bug, of zoning law as a whole.[23]

With this overview, we can turn next to the interpretation challenges posed by the two components of zoning codes.

B. Interpretation Challenges in Zoning Texts

Zoning texts use words, charts, and graphics to articulate the rules required by the jurisdiction. Generally speaking, analyzing texts is easier, for either human or machine, when texts are consistent and straightforward. Zoning texts are anything but. Their extended length—around 191 pages each, according to National Zoning Atlas (NZA) statistics to date[24]—and incremental amendments by laypeople naturally generate inconsistencies and complexities. A discussion of the structure and terminology of codes, as well as the prevalence of alternative regulations, illuminate the challenges posed to humans, including humans trying to train AI, and the challenges to AI itself.

1. Structure

[23] Elsewhere, I have critiqued zoning's ambiguities and inconsistencies, urging greater standardization, even while recognizing their prevalence in the current system. *See* Sara C. Bronin, *Simplifying and Modernizing Zoning Codes*, 20 J. of Law Econ. & Pol'y (forthcoming 2026).

[24] This figure comes from a data pull from the NZA database on August 30, 2026, showing that there are 7,136 fully-analyzed jurisdictions with zoning, with 1,362,424 pages of codes between them.

Zoning codes lack standardized, widely-utilized organizational structures across jurisdictions, even within the same region or state. While they are typically divided into numbered sections or chapters, there is no universal method of ordering these sections, nor are there conventions for numbering or formatting them. This lack of standardization hinders streamlined analysis of codes across jurisdictions, and inhibits the ability to train an AI model to search a predictably organized text.

Many zoning texts start or conclude with sections on administration, enforcement, and variances, but not all clearly separate out those elements or put them at the beginning or end. The location of the definitions, further discussed in the next Subpart, is also inconsistent. Some codes offer a single section neatly including every possible definition; others integrate definitions into substantive provisions; still others do both, with a stand-alone definitions section and other definitions sprinkled in between. Definitions may be located either in prose or in tables or charts. Competing definitions, located in different places, can also be problematic. In Boston's code, each neighborhood has its own independent zoning subcode, with different definitions for similar or the same terms.[25] An AI model tasked with understanding how the Boston code defines "apartment" would have to identify and evaluate the myriad instances of the term before answering.[26]

The substantive provisions of a zoning code—namely, the regulation of zoning districts—are also organized in different, at times seemingly random, ways. Some codes organize their texts based on zoning districts, with each district-specific section either describing the purpose of the district or describing its purpose and then outlining rules for new development in the district. Other codes group districts. Bloomfield, Connecticut, has sections articulating general regulations, and then has three sections grouping its nineteen districts by use: six residential districts in the third section of its code, six business districts in the fourth, and seven "special zones" including floodplain, mountain, elderly residential, and design development districts in the fifth.[27] The remaining sections in that code cover development standards, special permits, procedures, and other administrative matters that cut across the various districts. The zoning code of Petaluma, California, includes twenty-eight sections: three sections grouping districts based on their type (base districts, overlay zones, and floodplains), four sections for specific land uses (including alcoholic beverage and cardroom establishments), and separate sections on landscaping, preservation, and hillside protection standards, among other things.[28]

Aside from their organization, the substance of the substantive provisions also takes many forms. Rules may be set forth in prose (paragraph form) or in charts

---

[25] *See* Sara C. Bronin, Reforming the Boston Zoning Code Report, Jan. 2023.

[26] The zoning code provides a definition for an "Apartment hotel," as "containing four or more apartments which do not have kitchens," and a "Dwelling, multifamily" as a "building containing four or more dwelling units." BOSTON, MASS., ZONING CODE art. 2 § 2-1. The term "apartment" and commonly associated residential uses lack definitions, but reappear under different labels across tables and neighborhood-specific articles. *See, e.g.*, *id.* art. 8 tbl. B, art. 46 § 46-9, art. 64 § 64-7 (2025).

[27] BLOOMFIELD, CONN., ZONING REGS. at 1-4 (table of contents).

[28] PETALUMA, CAL., IMPLEMENTING ZONING ORDINANCE.

or graphics. As an example of the latter, a "use table" chart often indicates where uses (listed in, say, row format) are allowed, conditionally allowed, permitted subject to a special permit, or prohibited in particular districts (say, columns). An "area dimensions" chart might cover the minimum lot size, setbacks (the distances from the property line within which no structure may be built), lot-coverage, and density caps required in particular districts. Sometimes, charts are easier to understand than prose; other times, they are more convoluted, running dozens of pages with a perplexing variety of symbols showing different types of permissions. Whether prose or charts, there is no consistency across codes as to the location of district-specific rules.

Making matters even more confusing, some jurisdictions incorporate zoning into what they call "unified development codes," which entangle zoning rules with other development standards derived from separate, non-zoning enabling authority. Cheyenne, Wyoming, embeds zoning in its Unified Development Code, which also covers impact fee, subdivision and design standards.[29] That code includes just seven sections, with the fifth section called "Zoning Regulations" and the seventh called "Form-Based Codes."[30] These two sections, and arguably the section on design regulations, relate to Cheyenne's zoning authority. An AI model would not necessarily know this, and, even if informed that answers may be found in any one of those sections, would likely stumble over the duplicative or overlapping, and often slightly different, approaches contained in each. A search for minimum lot sizes across the entire Unified Development Code might also yield results from the subdivision section, which is adopted under a separate statutory authority than zoning—and is not zoning.[31]

Perhaps the most important consequence of the disorganized and inconsistent structure of zoning codes is the difficulty in readily discerning a complete list of zoning districts, which, as noted in Part I.A., are the very basic building block of codes.[32] A table of contents or an early section of the code might list most or all of the districts. In many codes, though, districts added after the date the code was originally adopted are referenced at random. And not all codes provide consolidated (even if partial) lists. Accordingly, the NZA methodology has an entire chapter providing instructions on how analysts should locate a complete list of zoning districts.[33] It suggests analysts scan the table of contents, search for key phrases (such as "district" and "zone"), and cross-reference the official map and geospatial files. Zoning analysts also contact local jurisdiction officials if a discrepancy cannot be resolved.[34] All of these steps require exercises of human

[29] CHEYENNE, WY., UNIFIED DEVELOPMENT CODE, at TOC-1 to TOC-2.
[30] CHEYENNE, WY., UNIFIED DEVELOPMENT CODE, at arts. 5, 7.
[31] *Compare* WYO. STAT. ANN. § 18-5-301 (subdivision regulations), *with* WYO. STAT. ANN. § 15-1-601 (zoning regulations).
[32] *See supra* Part III.A.1. (describing an AI trial attempting to extract district names from zoning texts and maps).
[33] Sara C. Bronin, Scott Markley, Aline Fader, & Evan Derickson, *How to Make a Zoning Atlas 2.0: The Official Methodology for the National Zoning Atlas*, Apr. 2024, at ch. III.
[34] *See* Jonah Pellecchia & Melissa Hayashida, *Eliminating Map-to-Text Discrepancies*, ZONING PRACTICE, Oct. 2024 (describing the need for humans to resolve conflicts by contacting zoning administrators).

judgment, including the capacity to know when a search for districts may be brought to a close. Given the variables and the potential need to make human contact, it would be challenging to train any AI model to accurately complete this elemental function of zoning analysis.[35]

2. Terminology

As discussed above, an AI model may have difficulties finding definitions in sometimes unexpected or out-of-order locations.[36] Perhaps more confounding, though, are the difficulties posed by the meanings of the terms themselves. Differences in terminology from one code to the next, or even within a single code, mean an interpreter (human or machine) cannot rely on a term at face value, especially AI models that can merely pattern match. They must engage in multiple cross-referencing and comprehension steps: locating the correct definition, understanding the definition, and using context clues where needed to adjust an interpretation. For AI models that rely on pattern matching, confusion about definitions and terminology differences would be highly problematic.

One cannot assume, for example, that the names of zoning districts mean the same things across jurisdictions.[37] A mixed-use district in one jurisdiction may allow a range of uses entirely different than a mixed-use district in the neighboring town. For example, in Dallas, Texas, a standalone single-family detached home is permitted in a mixed-use district, while in Fort Worth, Texas, the mixed-use districts prohibit or restrict single-family detached dwellings.[38] An adult entertainment district will be geared towards cabarets, peep shows, nude dancing, and, in Nevada, brothels[39]—not fine dining and jazz clubs. A human would understand what "adult" adds to the phrase "adult entertainment," but a machine searching for "entertainment" districts may not. And what of the Suburban district of Utqiaġvik, a town of 5,000 people in the northernmost county in the United States (the North Slope Borough of Alaska)?[40] Its stated purpose is to encourage traditional Iñupiat settlement patterns and allow for game processing.[41] Surely the word "suburban," in that context, does not mean what a person would ordinarily think it to mean. A machine searching for "suburban" districts, though, might include this district.

---

[35] *See supra* Part III.A.

[36] *See supra* Part I.C.1.

[37] This observation is relevant to the district-name extraction AI trial conducted by the National Zoning Atlas and described further herein. *See supra* Part III.A.1.

[38] DALL., TEX., CITY CODE § 51A-4.125 (2026) (listing "single family" as a permitted use in MU-1, MU-2, and MU-3 Mixed Use Districts); FORT WORTH, TEX., UNIFIED DEV. CODE § 4.1203 (2026) (listing "one-family detached dwelling" as a prohibited use in Mixed Use District Use Table)

[39] *See*, *e.g.*, Humboldt County, Nev. Zoning Code, § 17.110.020 (Adult Entertainment District allowing "[h]ouses of prostitution" and "[a]dult bookstores and video sales, novelties").

[40] The borough has the largest land area of any county or county-equivalent in the United States.

[41] The Suburban Residential District is designed to encourage "traditional Iñupiat settlement pattern[s]" and large lot sizes with space for "subsistence equipment, such as fish and game drying racks, and individual game processing areas." *See* NORTH SLOPE BOROUGH, AK., ZONING § 19.40.030 (2025).

Now consider an AI model tasked with identifying which zoning jurisdictions allow four-or-more-family multi-family housing; the designer of the experiment would not be satisfied with jurisdictions that only allow duplexes (two units) or triplexes (three units). Simply searching for the term "multi-family" would be insufficient. Adding "multi family," "multifamily," and "apartment" might cast a wider net. But one would also have to review the definition for the term found (be it "multi-family" in one form or another, or "apartment"). The NZA team has found definitions for "multi-family housing" or "apartment" that limit the number of units to two or three. We have found definitions for "multi-family housing" that set a minimum threshold of five or more units. And we have found regulatory provisions, in substantive sections of codes, that cap the number of units—sometimes using the word "multi-family" or "apartment," and sometimes not. Stopping once the term is found, without also reviewing its definition and other context, would yield inaccurate results.

Consider, as another example, an AI experiment with the aim of identifying whether a jurisdiction allows accessory dwelling units (ADUs), defined for the purposes of the experiment as smaller units of housing that can be located on the same lot as a single-family home. The NZA methodology offers the following terms as potential synonyms for an ADU:

- Accessory dwelling/unit/apartment
- In-law dwelling/unit/apartment
- Caretaker(s) dwelling/unit/apartment
- Garage apartment
- Carriage house
- Granny flat
- Servant quarters, servant's quarters, servants' quarters [unfortunately, still used]
- Subsidiary dwelling/unit/apartment
- Business apartment
- Facility manager dwelling/unit/apartment.[42]

Almost certainly a model can locate these terms in the code. After the term is found, however, its meaning must also be understood. For example, a "caretaker's unit" might refer either to a unit occupied by someone taking care of another person living in the main house, or to a unit occupied by a night watchman keeping tabs on a factory. The former would satisfy a query searching for ADUs in single-family homes, while the latter would not. If no definition for the located term is available, then one might review the use table or district-specific regulations to ascertain other uses available in the district where the caretaker's unit is allowed. If a caretaker's unit is allowed in industrial districts, then the query will not be satisfied. If, however, a caretaker's unit is allowed in single-family residential districts, then the

[42] Sara C. Bronin, Scott Markley, Aline Fader, & Evan Derickson, *How to Make a Zoning Atlas 2.0: The Official Methodology for the National Zoning Atlas*, Apr. 2024, at 95. *See also* Judah Axelrod, Lydia Lo, & Sara C. Bronin, *Automating Zoning Data Collection*, Urban Inst. Research Report, Feb. 2023, at appx. B (summarizing the synonyms for a variety of terms, including permitted use, affordable housing, elderly housing, one-family housing, two-family housing, three-family housing, four-family housing, minimum lot size, maximum density, floor-to-area ratio, and minimum unit size).

query will be satisfied. This kind of problem-solving and multi-step analysis is capably done by humans, but may not be so easy to train an AI model to perform.

Most AI queries are not so simple as these two, which focus on the question of whether a condition is met anywhere in the text of the code. As emphasized in Part I.A., any serious inquiry about zoning codes requires understanding information about specific zoning districts—and knowing if and where these districts are mapped within the jurisdiction. But in any event, the above discussion on simple queries illustrates the importance of terminology in understanding the effect of a zoning code, and flags potential issues for an AI application needing to understand words' meaning.

3. Alternative Regulations

Most AI queries run on text assume there will be a single answer: a yes or no, a specific number, or perhaps a string of words located and delivered verbatim from the text.[43] But zoning texts—to a surprising extent—articulate alternative regulations, which apply different rules when different conditions are met. A query of the NZA database reveals that around 58 percent of zoning jurisdictions use alternative regulations in their housing-related provisions.[44]

Alternative regulations can come in myriad, often unpredictable forms. The phenomenon of alternative regulations is so common that the NZA methodology covers it in great detail, and every numerical field maintained in our database can be converted to an "alternative" or "contextual" entry.[45] Take minimum lot sizes, common in nearly all zoning codes and requiring that specific types of development be built on parcels of a certain size. In the case of single-family housing, often the code mandate is quite simple: say, an acre of land for each single-family home. A well-trained model should be able to successfully extract this single piece of information across codes, even if provided in differing alphanumerical formats ("one acre," "1 ac."). It would also be easy to successfully train an AI model to convert units ("43,650 square feet," "43,650 sf"). Alternative regulations, though, present complications. Some codes use formulas for which the *formula* would be the correct answer, rather than a specific number. Such a code might say, for example, that each lot requires 5,000 square feet per dwelling unit, rather than providing a districtwide 20,000 square feet minimum lot size.[46] Yet other codes may offer conditional variables: a text might mandate that an apartment building have a minimum lot size of a half acre, unless the building's occupancy is limited to elderly persons, in which case the minimum lot size may be a quarter acre. In

[43] For example, the pieces of scholarship critiqued in Part II, *supra*, require answers derived from zoning codes be a single binary or numerical answer.

[44] This figure comes from a data pull from the NZA database on January 13, 2026, showing that of 9,200, 5,378, or 58 percent, were had at least one instance of alternative regulations or contextual regulations in fields related to housing regulations analyzed by the NZA.

[45] Sara C. Bronin, Scott Markley, Aline Fader, & Evan Derickson, *How to Make a Zoning Atlas 2.0: The Official Methodology for the National Zoning Atlas*, Apr. 2024, at 61-63.

[46] Another example in the parking context: zoning regulations requiring an office use to have three parking spaces for every 1,000 gross square feet of space, or that a restaurant will need a parking space for every six seats in the dining room.

that case, there are two potential numerical answers. Still other codes establish contextual arrangements, which depend on neighboring conditions. For example, a code might establish building setbacks that depend entirely on existing setbacks of the buildings on adjacent lots.

Summarizing here: even humans can expect to have challenges interpreting zoning texts given the variety of possibilities in the structure, terminology, and regulations presented in typical zoning codes. Consequently, it would be difficult for humans to train a model to expect all of the possible contingencies and options. Not yet mentioned is the situation in which the code offers no answer at all. When reading the zoning text for a jurisdiction that has eliminated all minimum lot mandates, for example, an AI model searching for a minimum lot size might be particularly prone to hallucinate an answer: attempting to find a word or phrase that best fits. Training that algorithm to exercise judgment about when to stop searching is another significant problem.

C. Interpretation Challenges in Zoning Maps

Alongside the inconsistencies and complexities in zoning texts, there are interpretation challenges presented by zoning maps. Two are somewhat common. First, some maps are hard to decipher, either due to cartographical choices or due to issues related to special types of districts. Second, some maps fail to match zoning texts—a legally consequential concern afflicting perhaps one-third of jurisdictions.

These challenges are relevant to a study of AI in zoning because of the critical importance of maps in understanding the scope and impact of zoning regulation. They also underscore the need for human analysts to reconcile discrepancies, either by contacting the jurisdiction's zoning administrators or by using their own best judgment and knowledge of planning practices to ascertain the likely intent of the code. Again, the need for human intervention does not preclude an AI-driven study of zoning codes. It does, however, mean that AI studies cannot ignore the need for reconciliation, and must incorporate human review into an AI workflow.

1. Hard-to-Decipher Maps

On their surface, zoning maps may appear to be easy to understand: a legend filled with colors and hatches, each signifying a different zoning district. But in many cases, maps cannot be easily deciphered. One all-too-common issue is that of color-blurring. As two examples: Meadow, Utah, uses nearly indistinguishable shades of blue to delineate zoning districts, while Cheshire, Michigan, uses a myriad of yellows.[47]

[47] Zoning Map Meadow, Utah (on file with author); Zoning Map Cheshire Twp., Mich. (revised Dec. 31, 2002), https://cheshiretownshipmi.gov/images/zoning%20map.pdf.

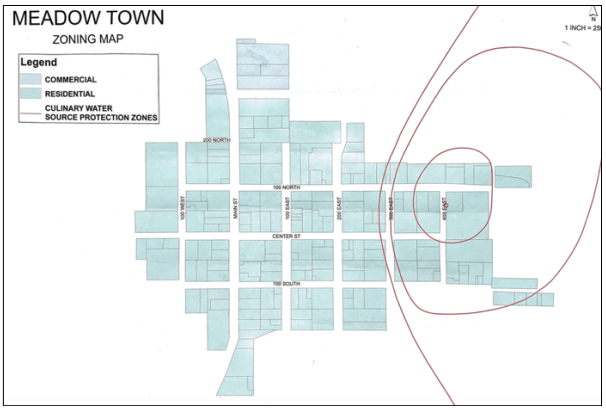

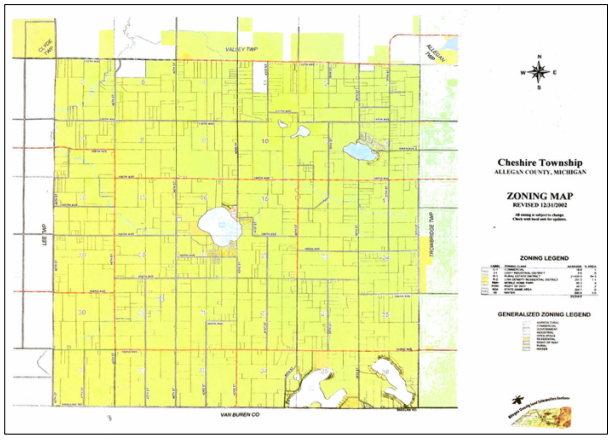

There are also cases, as with the map of South Park, Pennsylvania, where district boundaries are styled like roads, and, as with the map of Beckley, West Virginia, where jumbled labels are disorienting.

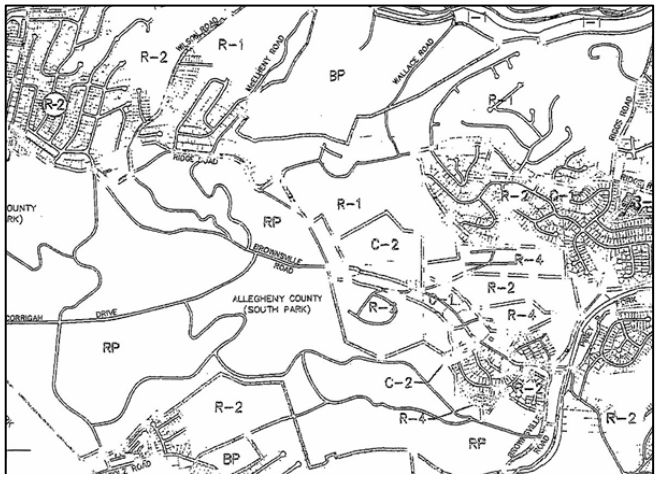

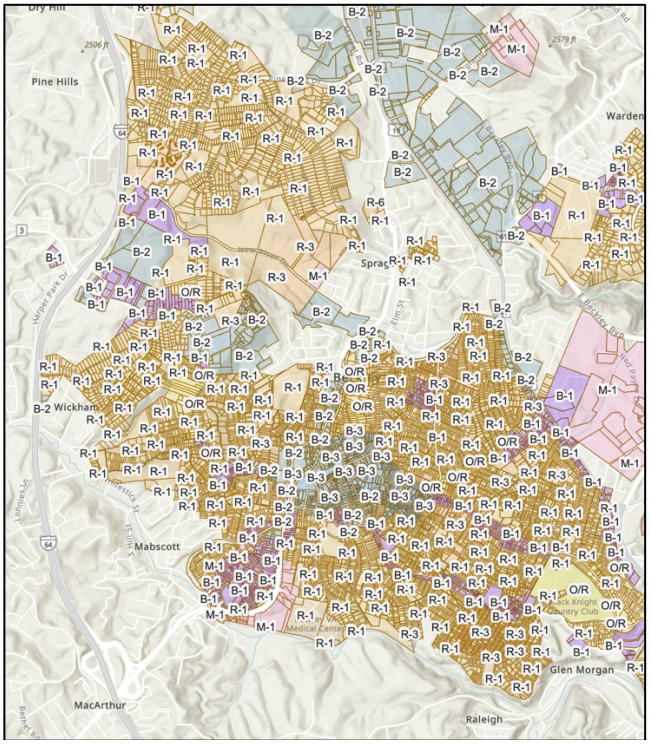

Another common cartographical issue is the difficulty identifying where overlays are located, especially when such overlays are depicted by hatching in small areas. In addition, sometimes special districts, like overlays or planned residential developments, do not appear on the same map as other districts, causing confusion about overlapping requirements.[48] Such is the case with the overlay zones in Walnut Creek, California, like this O-5 district that is identified as a straight line with arrows.[49] But individual parcels assigned to a district should be, at least, two-dimensional on a zoning map—not a line.

[48] As the NZA research director and engagement director wrote in a recent book chapter, overlays and planned residential developments "are recurring stumbling blocks, but they are just two of many. Conditional-use rules, form-based codes, unmapped districts, expired districts, subdistricts, special districts, extra-territorial jurisdictions (ETJs), and land-use regulations external to zoning (e.g., federal protections) all complicate the translation of zoning codes into an analysis-ready geospatial format." Scott Markley & Diana Drogaris, *Zoning and Municipal Data Brokers: Toward a Political Economy of Data Slop*, The Data Politics of Property & planning (Danielle Hynes, Rob Kitchen, & Sam Mutter, eds.) (forthcoming 2026); preprint available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5188578.

[49] Walnut Creek presents this image in its official map but also presents a digitized version. In this case, NZA staff used the digitized version to determine district boundaries.

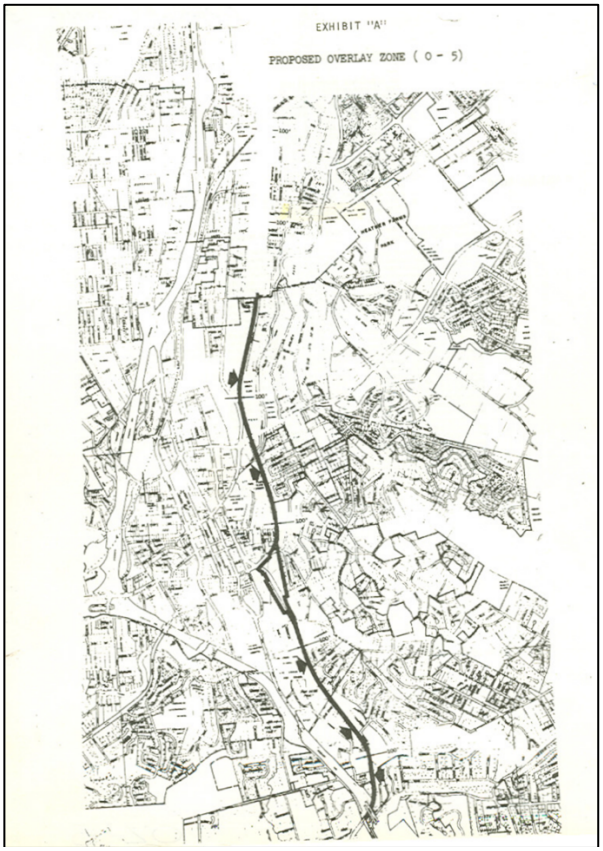

One might think that files provided in geospatial format would avoid some of these issues. But as the NZA methodology explains, such files can vary greatly in quality, often requiring an analyst to remove slivers (unwanted, small polygons), trim zoning district boundaries that wrongly extend past jurisdiction boundaries, and remove roads and protected lands.[50] There may be deeper issues too, such as zoned areas that are entirely unassigned in the geospatial files or zoned areas assigned to multiple base districts.

When human analysts at the NZA encounter such situations, they know how to resolve them: request the official zoning map, consult assessors' maps, scour meeting minutes, and/or pick up the phone and contact jurisdiction staff for guidance. They must use their judgment to determine if they have sufficient proof to map a district, a verification process necessary to ensure the code is appropriately characterized in the NZA database. By contrast, an AI model fed these maps and asked to digitize them might map the districts without further inquiry.

2. Mismatches Between Texts and Maps

The vast majority of zoning districts appear both in the text and on the map: that is, they have rules related to them, and it is clear how and where those rules can be enforced. However, mismatching of zoning district information between text

[50] Sara C. Bronin, Scott Markley, Aline Fader, & Evan Derickson, *How to Make a Zoning Atlas 2.0: The Official Methodology for the National Zoning Atlas*, Apr. 2024, at 140-151.

and map is a significant issue—one whose scope the NZA has helped to reveal. Anecdotally, the team experiences and resolves such mismatches on a daily basis: it is considered part of the job. Empirically, we have one data point so far: two NZA analysts released a small but illuminative in-depth study of 42 jurisdictions across seven states, reporting that 15 (or 36 percent) had major discrepancies between the zoning map and zoning text.[51] The types of discrepancies varied, each raising questions requiring human review and resolution.

Among the minor discrepancies are clerical errors, especially in the map legend naming zoning districts.[52] Spelling errors, errors or omissions of punctuation (such as using a slash instead of a dash), or omission of a word in a multi-word district name (such as using "Transit-Oriented Overlay" on the map instead of the "Transit-Oriented Development Overlay" title in the text) are easily resolved without the need to contact jurisdiction staff. However, some clerical errors require further inquiry. For example, an Arizona town provided a map that gave entirely erroneous names for two of the districts depicted.[53] Staff provided the correct names after an inquiry, but one wonders at how such an error could have been made.

More serious discrepancies raise the same question. In some cases, the map can be more up-to-date than the text, or vice versa. To ensure that the NZA logs the most up-to-date laws, analysts encountering this type of inconsistency contact jurisdiction staff to request meeting minutes, resolutions, and other records of decision-making. Zoning maps sometimes erroneously combine related zoning districts, like districts that all fall into the same general category. A Louisiana city, for example, had zoning text defining the R-1A and R-1B (residential) districts as two distinct districts, but with a zoning map combining them into a single R-1A/R-1B district. Again, such cases require contacting the jurisdiction. For the Louisiana city, the analysts reported that "According to city staff, there is a separate map distinguishing R-1A from R-1B… but this map was never adopted by the city council."[54]

Perhaps most serious of all are those cases where there is a total mismatch, with no ability to reconcile existing text and code. Best exemplifying a total mismatch are zoning districts that in NZA parlance are "extinct." Extinct districts previously existed in, but were deleted from, the text—yet remain on the map. More than 3 percent of zoning districts in the NZA database are extinct.[55] Conversely, unmapped zoning districts are written into the text but not mapped on the map—

---

[51] Jonah Pellecchia & Melissa Hayashida, *Eliminating Map-to-Text Discrepancies*, ZONING PRACTICE, Oct. 2024. *See also* Scott Markley & Diana Drogaris, *Zoning and Municipal Data Brokers: Toward a Political Economy of Data Slop*, The Data Politics of Property & planning (Danielle Hynes, Rob Kitchen, & Sam Mutter, eds.) (forthcoming 2026), at 8-9; preprint available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5188578 (identifying other internal contradictions and ambiguities).

[52] *See* Jonah Pellecchia & Melissa Hayashida, *Eliminating Map-to-Text Discrepancies*, ZONING PRACTICE, Oct. 2024, at 4-5.

[53] Jonah Pellecchia & Melissa Hayashida, *Eliminating Map-to-Text Discrepancies*, ZONING PRACTICE, Oct. 2024, at 4-5.

[54] Jonah Pellecchia & Melissa Hayashida, *Eliminating Map-to-Text Discrepancies*, ZONING PRACTICE, Oct. 2024, at 5.

[55] This figure comes from a data pull from the NZA database on December 31, 2025, showing that of 98,913 published zoning districts, 3,071, or 3.15 percent, were extinct.

and thus the regulations related to them apply nowhere. Around 15 percent of zoning districts in the NZA database are unmapped.[56] The existence of extinct and unmapped districts underscores the need for any analysis of zoning codes, including AI analysis, to consider both text and map, and to reconcile discrepancies between them. In the case of extinct districts, that means locating the prior zoning text that last defined the district, and incorporating those provisions into our database. We do this knowing that the prior zoning text may or may not be considered by a court to be in effect. In the case of unmapped districts, it means flagging their lack of regulatory effect and simply waiting for word of an amendment that maps them.

To deal with these various types of mismatches, analysts use common sense (used to address most clerical errors), contact local officials for interpretation advice, and locate alternative regulatory materials, such as recent decisions or prior copies of codes (to find extinct districts). The reconciliation process depends entirely on human judgment and human intervention, and requires information found outside of the pages of the code.

* * *

Zoning codes present numerous interpretation challenges to anyone seeking to understand what they say, including scholars purporting to use AI to study zoning. Putting the NZA-derived statistics mentioned in this Part in one place underscores this point: More than half of zoning jurisdictions use alternative regulations for even basic fields, meaning that studies claiming to have single answers for every query are either wrong or willfully misleading. A third of codes have map-to-text discrepancies, requiring human resolution before they can be analyzed. And 3 percent of districts are extinct, requiring a treasure hunt for their applicable rules, while 15 percent are unmapped, meaning they have no regulatory weight, even if articulated in the text.

The challenges described in this Part arise not because local officials are incompetent, or wantonly or willfully careless. Rather, they might be explained by the sheer length of codes and breadth of topics codes address. Or by the piecemeal nature of amendments.[57] Or by the fact that laypeople, not lawyers, are often charged with writing and amending the code.[58] Whatever their origin, they must be

[56] This figure comes from a data pull from the NZA database on December 31, 2025, showing that of 98,913 published zoning districts, 14,348, or 14.7 percent, were unmapped. Note that this number of unmapped districts includes some districts the jurisdiction mapped in areas that could not legally be subject to local zoning authority, such as national or state parks. NZA methodology requires excising such protected lands (contained in the U.S. Census Protected Areas Database of the United States) from zoned land, which in turn excises districts located wholly within protected lands, rendering them unmapped.

[57] *See e.g.*, Carol M. Rose, *Planning and Dealing: Piecemeal Land Controls as a Problem of Local Legitimacy*, 71 CALIF. L. REV. 837, 883 (1983); Caleb Hersh, *The NIMBY Filibuster*, 100 N.Y.U. L. REV. 167, 167–206 (2025) (discussing public protest petitions and legislative responses in zoning amendment processes); Matthew J. Lewyn, *The Case Against the Case for Zoning*, 52 GEO. ENVTL. L. REV. 101, 101–38 (2024) (arguing that structural features of zoning laws, including amendment procedures, often reinforce restrictive regulatory outcomes).

[58] *See* Anika Singh Lemar, *Overparticipation: Designing Effective Land Use Public Participation*, 90 FORDHAM L. REV. 1123, 1130-42 (2021) (highlighting how zoning officials' lack of expertise

recognized as inherent to the system in which zoning codes are produced, administered, and updated.[59] And there must be a significant amount of human reasoning to interpret and contextualize them. An AI model might be able to replicate some aspects of human reasoning, but only if the humans training it clearly understand how zoning codes work—and even then, only to a limited extent. The next Part reveals what happens when humans training AI models lack an understanding of, or understand but ignore, zoning's functional mechanics.

II. CRITIQUING PROBLEMATIC APPLICATIONS OF AI TO ZONING

Given the challenges inherent in the way zoning codes are written and amended, it has been surprising to see scholarship claiming to have used AI to successfully interpret them.[60] The fact that this scholarship has been widely cited makes a detailed review imperative and timely.

Here, I draw from the observations made in Part I to show how the authors of AI-related zoning scholarship have, so far, ignored one or all of the four principles outlining how zoning works.[61] They have misunderstood the function of zoning districts, arbitrarily selecting regulations for a single district to characterize the entire jurisdiction. Most have focused on zoning texts, declining to consider the zoning map and thus omitting from analysis the extent to which regulations actually apply within a jurisdiction.[62] And they have refused to accept the heterogeneity and complexity of zoning codes, instead collapsing nuanced regulations into a single answer or number. With such significant human errors clouding inputs, the AI models had no chance of successfully analyzing zoning codes, and authors' claims to the contrary are unfounded. This Part concludes that recent scholarship about AI and zoning has made us do only one thing better and faster: misunderstand zoning.

To make this point, this Part critiques one or two pieces of scholarship on AI and zoning in each of three categories of AI applications to legal contexts: characterizing legal decisions and documents, predicting outcomes of legal

---

can undermine the goals of a public participation process with a lack of accountability and formality).

[59] *See* Alicia Solow-Niederman, YooJung Choi, & Guy Van den Broeck, *The Institutional Life of Algorithmic Risk Assessment,* 34 BERKELEY TECH. L.J. 705 (2019) (arguing that algorithmic tools' interpretation is dependent on the institutional and local decision making in which they function).

[60] This Article does not address private companies making the same types of overclaims, as this topic has already been well-covered in another paper co-authored by National Zoning Atlas participants, calling much of the data collected by these companies "slop." *See* Scott Markley & Diana Drogaris, *Zoning and Municipal Data Brokers: Toward a Political Economy of Data Slop*, The Data Politics of Property & planning (Danielle Hynes, Rob Kitchen, & Sam Mutter, eds.) (forthcoming 2026); preprint available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5188578.

[61] *See supra* Part I.A. ("first, the zoning district, which distinctly regulates land uses, structures, and lots, is the organizing feature of the code; second, the text laying out regulations for each district governs development in each district; third, the map identifies where zoning districts are assigned, and thus where regulation has effect; and fourth, jurisdictions demonstrate seemingly infinite variety in the way they articulate the rules in the text and delineate districts on the map").

[62] One piece of critiqued scholarship, by Margaret Lawrimore et al., does consider zoning maps, as further discussed in Part II.B., *supra*.

decisions or applications, and extracting information from statutes, ordinances, or other legal texts. It may be worth noting that of the different types of AI available, machine learning—the ability of computers (machines) to learn from data and draw inferences and patterns—is typically used for characterizing and predicting applications.[63] Natural language processing—the ability of computers to interpret and structure vast amounts of unstructured data (such as text) using techniques such as text classification and speech recognition—is typically used for extracting applications, and sometimes to support characterizing and predicting applications.[64]

A. Characterizing Applications

1. Generally

Characterizing applications of AI are designed to analyze inputted data to ascertain relationships, properties, and patterns within the data. In the legal context, AI has been used to streamline the arduous task of sifting through large numbers of documents, such as those produced in the discovery phase of litigation, and to summarize judicial opinions, contracts, case files, and other documents. Attorneys have also relied on AI to conduct legal research, asking AI to scan for relevant cases and suggest citations for their briefs, filings, or other documents, though judging by the number of well-documented hallucinations, this AI application needs

[63] *See* David Lehr & Paul Ohm, *Playing with Data: What Legal Scholars Should Learn about Machine Learning*, 653 U.C. DAVIS L. REV. 653 (2017); TOM M. MITCHELL, MACHINE LEARNING, 2–5 (1997); *see also* 15 U.S.C. § 9401(11) (defining machine learning as an application of AI that provides systems the ability to learn and improve from data). For the purposes of this Article, the recognition of ML as a pattern extraction mechanism is critical. The probabilistic approach used by ML models to predict meaning of zoning codes falters where human interpretation is necessary.

[64] *See* CHRISTOPHER D. MANNING & HINRICH SCHÜTZE, FOUNDATIONS OF STATISTICAL NATURAL LANGUAGE PROCESSING, at xix (1999). Again, as the technical distinctions between these processes matter little for this Article, this Part usually uses the term "AI" to cover any of the technologies enabling computers to perform tasks normally completed by humans.

refining.[65] Putting aside the legal research use case, we turn to other characterizing applications of AI that have found their way into academic work.[66]

Scholars have employed AI to examine large volumes of decisions and to understand the reasoning behind and the content of these decisions. For example, Jed Stiglitz and Rosamond Thalkin created a dataset encompassing thousands of United States Supreme Court opinions issued between 1870 and 2023, and asked a generative language model to identify the mode of reasoning used.[67] The opinions, written in prose, were divided into clean paragraphs, which were then fed into the model. They aimed, among other things, to understand how the Supreme Court's mode of reasoning shifted during that period between formalistic reasoning based on rules and "grand" reasoning that is more open-ended.[68] Their results apparently verified established scholarly consensus, and are an example of AI strengthening the empirical basis for a legal hypothesis. Similarly, Jonathan Choi used AI to analyze 182,535 pages of Internal Revenue Bulletins and 470,099 Tax Court opinions to assess how different tools of interpretation were used over time by an agency and by a court considering the same statutes (federal tax laws).[69] Choi trained an AI model to identify keywords or terms associated with four decision-

---

[65] Even legal databases have proven to be unreliable in conducting legal research. *See* Varun Magesh, Faiz Surani, Matthew Dahl, Mirac Suzgun, Christopher D. Manning, & Daniel E. Ho, *Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools*, J. EMP. LEG. STUDIES, Mar. 2025, at 1 (analyzing LexisNexis's Lexis+ AI and Westlaw's AI-Assisted Research and finding hallucinations in between 17 percent and 35 percent of citations). Failure to check results from AI-assisted research has led to attorneys being publicly admonished. *And see* Damien Charlotin, AI Hallucination Cases, at https://www.damiencharlotin.com/hallucinations (documenting nearly 2,000 legal decisions in which generative AI produced hallucinated content); Amol Parikh, *Deception Inspection: Attorney Faces Discipline for Citing Fake Law*, NAT'L L. REV., Feb. 8, 2024, https://natlawreview.com/article/deception-inspection-attorney-faces-discipline-citing-fake-law (reporting Second Circuit referral of an attorney to its grievance panel for investigation and disciplinary action for citing a nonexistent case created by ChatGPT, a Gen AI tool, in a motion to dismiss); Emily Sawicki, *Ex-LA DA's Atty's AI Use 'Affirmatively Misled' Special Master*, Law360, May 15, 2025, https://www-law360-com.gwlaw.idm.oclc.org/articles/2340423/ex-la-da-s-atty-s-ai-use-affirmatively-misled-special-master (reporting prominent law firm's incorporation of AI hallucinations in briefs filed in the Central District of California earned the firm a $31,000 sanction); Evan Gorelick, *Vigilante Lawyers Expose the Rising Tide of A.I. Slop in Court Filings*, N.Y. TIMES, Nov. 7, 2025, https://www-nytimes-com.us1.proxy.openathens.net/2025/11/07/business/lawyers-ai-vigilantes.html?searchResultPosition=1 (reporting a lawyer's use of 32 AI-generated, nonexistent citations in a motion filed in a Texas bankruptcy court, which resulted in the referral of the lawyer to the state bar's disciplinary committee with six hours of AI training).

[66] *See* Farid Ariai, Joel Mackenzie, & Gianluca Demartini, *Natural Language Processing for the Legal Domain: A Survey of Tasks, Datasets, Models and Challenges*, 58 ACM COMPUTING SURVEYS 1, 15- (2025) (citing research using natural language processing to categorize legal documents based on their content, including whether a document is in a given class or not (a binary classification), whether a document is in one of many classes (multi-class classification) or which set of classes a document belongs to).

[67] Edward H. Stigliz & Rosamond Thalkin, *Historical Trends in Macro-Jurisprudence: A Language Model Assessment, 1870-2023*, 84 MD. L. REV. 46 (2024).

[68] Edward H. Stigliz & Rosamond Thalkin, *Historical Trends in Macro-Jurisprudence: A Language Model Assessment, 1870-2023*, 84 MD. L. REV. 46, 69 (2024).

[69] Jonathan H. Choi, *An Empirical Study of Statutory Interpretation in Tax Law*, 95 NYU L. REV. 363 (2020). The Bulletins dated back to 1919, while the Tax Court decisions dated back to 1942.

making approaches: statutory, normative, textualist, or purposivist.[70] The algorithm measured the frequency that these terms appear in the texts, like when a textualist judge cites to a "dictionary," to characterize each decision.[71] He concluded that while the IRS has become more normative, a phenomenon consistent with increased deference shown to agencies by courts during the period studied, the Tax Court has become more textualist, suggesting it is more like other federal courts than like the IRS. In these representative papers, AI distilled trends within large volume of legal texts written over decades, which were then used to inform theories about the motivations of those texts' authors.

Other scholars have used AI to characterize legal language as similar or dissimilar. Two tax law scholars applied natural language processing to analyze a dataset of 4,052 bilateral, international income tax treaties and twenty-three model treaties to assess how the "controlling legal language" of one treaty compares to the language in other treaties in effect at the same time between different parties.[72] The analysis aimed to measure similarities between each possible pair of treaties during a given year and to calculate the mean similarity in each year.[73] They concluded that language similarity increased between different bilateral tax treaties between 1970 and 2010.[74] In an article published after his tax-law piece, Choi deployed the AI technique of "word embedding," which measures semantic similarity among words using mathematical vectors, to argue that many words used by courts have indeterminate meanings.[75] He referenced a federal court's invalidation of the Biden administration's mask mandate on public transportation during the COVID-19 pandemic, which hinged on whether the term "sanitation" (used in a Center for Disease Control rule authorizing "sanitation" measures) meant "cleaning" or

---

[70] Jonathan H. Choi, *An Empirical Study of Statutory Interpretation in Tax Law*, 95 NYU L. REV. 363, 368 (2020).

[71] Jonathan H. Choi, *An Empirical Study of Statutory Interpretation in Tax Law*, 95 NYU L. REV. 363, 389 (2020).

[72] Elliott Ash & Omri Marian, *The Making of International Tax Law: Empirical Evidence from Tax Treaties Text*, 24 FLA. TAX REV. 155 (2020).

[73] Elliott Ash & Omri Marian, *The Making of International Tax Law: Empirical Evidence from Tax Treaties Text*, 24 FLA. TAX REV. 151, 169-172 (2020). They trained an algorithm to read the clauses of each treaty, split the clauses into four-word phrases (which provide more information than single words), then filter those four-word phrases by part-of-speech and other variables. They then built a frequency distribution for the phrases in each treaty, which they used to identify similarities and differences between the language in the treaties.

[74] Elliott Ash & Omri Marian, *The Making of International Tax Law: Empirical Evidence from Tax Treaties Text*, 24 FLA. TAX REV. 151, 172 (2020).

[75] Jonathan H. Choi, *Measuring Clarity in Legal Text*, 91 U. CHIC. L. REV. 1, 20-28 (2024). For a given set of words, the algorithm creates vectors that originate from the same point, such that the angle between two vectors is smaller the more similar their respective words are. The angle between the vectors is measured as the "cosine similarity" that ranges from 1 for interchangeable words to -1 for antonyms. In this way, this method reflects how appropriate it would be to replace one word with another in certain contexts. Note that the word embedding approach has begun to be superseded by what is known as transformer architecture, essentially a process that builds on word embeddings and refines them by connecting words in different ways. *See* Farid Ariai, Joel Mackenzie, & Gianluca Demartini, *Natural Language Processing for the Legal Domain: A Survey of Tasks, Datasets, Models and Challenges*, 58 ACM COMPUTING SURVEYS 1, 7 (2025). Full treatment of this evolution is beyond the scope of this Article.

"preserving cleanliness."[76] After AI found both meanings plausible given their mathematically-measured similarity of selected terms, Choi presented the judge as legal interpreter that must look beyond the text to make their decisions.[77] As shown in these examples, the use of AI to understand the meanings of words can reveal legal trends and enhance interpretive nuance, while also reconfirming the need for humans to ultimately arbitrate ambiguity.

2. Zoning-Related Scholarship

No legal scholars have yet used AI to characterize zoning-related legal decisions and documents.[78] However, two planning scholars, Arianna Salazar-Miranda and Emily Talen, attempted to use AI to characterize zoning codes, including a particular type of zoning code known as a "form-based code" (FBC). They downloaded 2,723 municipal ordinances, including ordinance chapters on zoning, subdivision, and streets.[79] From this group of texts, they set aside 70 codes that had been independently identified as exemplary FBCs.[80] They then used a word embeddings approach (as Choi did) and compared the codes with each other, using a scatterplot to represent codes' similarities. Realizing that the known FBCs clustered together, they then classified the remaining codes as FBC-like and non-FBC-like. From fifty FBC-like and fifty non-FBC-like codes, they extracted paragraphs about setbacks and floor-to-area ratios (FAR), and created word clouds to highlight that the setback and FAR provisions in FBC-like codes were more open (using words such as "landscape" and "architectural"), while the same provisions in non-FBC-like codes were more closed (using words such as "adverse," "performance," and "hazard"). Finally, they reviewed on-the-ground conditions—including actual building footprints, building heights, and parcel sizes—to understand lot sizes, building placement within lot boundaries (built setbacks), and square footage of buildings (built FARs). From this analysis, they concluded that "FBCs are associated with higher floor to area ratios, narrower and more consistent street setbacks, and smaller plots" than non-FBC-like codes.[81] They presented their approach as an alternative to manual data collection like that used by the National

[76] *Health Freedom Def. Fund, Inc. v. Biden*, No. 8:21-CV-1693-KKMAEP, 2022 WL 1134138, at *7 (M.D. Fla. Apr. 18, 2022).
[77] *See* Jonathan H. Choi, *Measuring Clarity in Legal Text*, 91 U. CHIC. L. REV. 1, 35 (2024).
[78] *See supra* Part IV.B. (offering ideas about future scholarship in this arena).
[79] Arianna Salazar-Miranda & Emily Talen, *An AI-Based Analysis of Zoning Reforms in US Cities*, 2025 NATURE CITIES 304, 305 (chapter titles included "'zoning', 'land division', 'subdivision', 'land use', 'land development', 'streets', 'master plan', 'development' and 'neighborhood design'"). All codes were downloaded from municode.com.
[80] Arianna Salazar-Miranda & Emily Talen, *An AI-Based Analysis of Zoning Reforms in US Cities*, 2025 NATURE CITIES 304, 305 (pulling from a list of FBCs created by a planning and design firm and a list of FBC award winners created by a nonprofit program called the Form-Based Codes Institute).
[81] Arianna Salazar-Miranda & Emily Talen, *An AI-Based Analysis of Zoning Reforms in US Cities*, 2025 NATURE CITIES 304, 304. *See also id.* at 305.

Zoning Atlas, and asserted that they had created a "comprehensive and systematic assessment of FBC [form-based code] adoption and its impacts."[82]

Unfortunately, misguided inputs may have hindered achievement of the "comprehensive and systematic assessment."[83] First, the authors admitted that their corpus included a significant amount of text not normally considered within the legal scope of local zoning authority. The text of ordinance chapters on subdivision and streets should have been excluded from their analysis. Compromised corpus aside, the authors also made a puzzling choice to focus their AI analysis on two regulatory features—FAR and setbacks—considered antithetical to FBCs.[84] As

---

[82] Arianna Salazar-Miranda & Emily Talen, *An AI-Based Analysis of Zoning Reforms in US Cities*, 2025 NATURE CITIES 304, 305 (adding that "efforts, such as the National Zoning Atlas, rely on the manual input of zoning information. We address these issues by employing natural language processing…").

[83] On a cosmetic note, the title, "An AI-Based Analysis of Zoning Reform in U.S. Cities," may be the first clue of the authors' propensity to overclaim. Despite the centrality of "AI-Based Analysis" in the title, the authors only use AI in a limited respect—the similarity analysis and the word clouds—while devoting much more attention to statistical and spatial methods common in planning academia. Discussion of the application of these methods is omitted from the summary in the main text, but measurably more prominent in the article text than the AI methods. For example, they conducted regressions, created scatter plots, conducted a manual validation, engaged in spatial mapping, and drew and measured built patterns. The title's use of the phrase "zoning reform" also misleads readers. Despite the authors' feeble attempt to equate FBCs with zoning reform, no significant zoning reform movement advocates make FBCs a key part of their platform. The "yes-in-my-backyard" (YIMBY) groups that have led the American zoning reform movement for years instead advocate almost exclusively for zoning reform that facilitates the construction of more housing. They rarely mention form-based codes. *But see To HCD: Santa Cruz YIMBY Comments on Santa Cruz County Draft Housing Element (July 2023)*, SANTA CRUZ YIMBY (Sept. 12, 2023), https://santacruzyimby.org/news/to-hcd-santa-cruz-yimby-comments-on-santa-cruz-county-draft-housing-element-july-2023/?_gl=1*1yjsfc8*_gcl_au*OTI5OTUwMDg0LjE3Njc4NTQ1NDc.*_ga*MTI1NzgyMDY1Ny4xNzY3ODU0NTQ3*_ga_YW3PRLFK0Y*czE3Njc4NTc1NDgkbzIkZzEkdDE3Njc4NTc4MjMkajM0JGwwJGgw (advocating for the inclusion of form-based zoning in affordable housing development). Moreover, the article does not document *changes* in zoning codes, a phenomenon implied in the word "reform"; it merely documents the status quo at a single point in time. Perhaps the authors incorporated these phrases in the title to attract more readers than "mixed methods" and "form-based codes" would have attracted. But they confuse a reader given that the article itself centers on neither concept. While a cosmetic point, the choice of title hints at the authors' inclinations.

[84] *See*, *e.g.*, William J. DeCoursey, Tyler S.W. Berl, & Jeremy J. Rothwell, Office of State Planning Coordination, Form-Based Codes: A Primer for Delaware Municipalities 9 (contrasting zoning, which uses "floor-area ratio (FAR), setbacks" and other strategies with form-based codes, which use "build-to-lines, number of floors" and other strategies); Kaizer Rangwala, *Form-Based Codes*, 11 ECON. DEV. J. 35, 36 (2012) (differentiating "conventional use-based codes and FBCs" because "FBCs do not determine permitted square footage entitlements through FAR"); DANIEL G. PAROLEK, KAREN PAROLEK & PAUL C. CRAWFORD, FORM-BASED CODES: A GUIDE FOR PLANNERS, URBAN DESIGNERS, MUNICIPALITIES, AND DEVELOPERS 318 (2008) (holding that "[FAR] has absolutely no role in FBCs"); *see also* AM. SOC'Y OF PLAN. OFFS., FLOOR AREA RATIO, INFORMATION REPORT NO. 111 3–4 (1958) (demonstrating the basic formula used to compute FAR and concluding that although FAR "affects, volume, shape, and spacing of buildings on the land, it does not determine a particular shape or spacing"); Noal Leonetti & Abby Newbold, *Floor Area Ratio: What Is It And How Local Governments Can Use It*, MSRC (Dec. 3, 2025), https://mrsc.org/stay-informed/mrsc-insight/december-2025/floor-area-ratio (explaining that

their name suggests, FBCs focus on form of the building above all else.[85] Thus typical FBC regulatory provisions include height minima and maxima, building-style requirements (like Victorian, Cape Cod, or Federal style), roof design standards, and fenestration (door and window) standards. Many FBCs eliminate FARs altogether and eschew setbacks in favor of "build-to" lines.[86] Even if the AI model accurately extracted information about FAR and setbacks, it should not have been instructed to do so in the first place.

The authors also embedded in their analysis incorrect assumptions that about zoning's functional mechanics. Chief among them is their stated "assum[ption] that what is written in the zoning code applies to the entire municipality."[87] As noted in Part I, jurisdictions with zoning divide their land into distinctly-regulated districts, and one cannot assume that a rule in one district applies to the entire jurisdiction.[88] In fact, one should assume the opposite. In this case, the FBC-like provisions of interest to the authors may have governed only a small portion of land, or even lacked regulatory weight because they were unmapped (a possibility the article does not consider).[89] By misapplying district-specific regulations to the whole jurisdiction, and by eschewing map-related analysis, the authors could not have trained a model to "comprehensively" assess FBCs or their impacts.[90]

This brief critique reveals issues common in scholarship involving zoning and AI.[91] First, it illustrates the tendency to offer misleading or unsubstantiated

---

"[u]nlike other density measurements, such as units/acre and lot coverage, FAR does not address building form or shape").

[85] *See*, *e.g.*, Emily Talen, *Design by the Rules: The Historical Underpinnings of Form-Based Codes*, 75 J. AM. PLAN. ASS'N 144, 144 (quoting *Definition of a Form-Based Code*, FORM-BASED CODE INST. (2008): "Form-based codes create a predictable public realm primarily by controlling physical form, with a lesser focus on land use").

[86] For example, in a major code overhaul producing an award-winning FBC, Hartford, Connecticut, eliminated FARs and setbacks entirely, as inconsistent with the purpose of the FBC. *See* Sara C. Bronin, *Comprehensive Rezonings*, 2019 BYU L. REV. 725 (2020).

[87] Arianna Salazar-Miranda & Emily Talen, *An AI-Based Analysis of Zoning Reforms in US Cities*, 2025 NATURE CITIES 304, 311.

[88] *See supra* Part I.A. (describing the necessity of understanding both text and map to understand the actual impact of codes).

[89] *See supra* Part I.D. (noting that 15 percent of districts currently incorporated into the National Zoning Atlas dataset are unmapped).

[90] *See supra* Part I.D.1.

[91] I have other concerns with this article, but these are not relevant to their AI analysis. For example, I noted the authors' willingness to conflate analysis of the law with analysis of as-built conditions. In the first paragraph of the article, they asserted that FBCs "are associated with higher floor to area ratios, narrower and more consistent street setbacks, and smaller plots." Arianna Salazar-Miranda & Emily Talen, *An AI-Based Analysis of Zoning Reforms in US Cities*, 2025 NATURE CITIES 304, 304-305. Pages later, the article revealed that they explored these correlations only in their analysis of as-built conditions, not their analysis of legal texts. Either way, as-built conditions do not reveal what zoning codes actually say. In Boston, for example, as-built conditions differ significantly from code requirements. Using two measures of interest to Salazar-Miranda and Talen (FAR and lot sizes), a recent analysis revealed that 58 percent of parcels would be illegal if current zoning regarding FARs were applied to them, and 42 percent of parcels would be illegal if current minimum lot-size mandates were applied. Sara C. Bronin, Reforming the Boston Zoning Code Report, Jan. 2023, at 8. The extent to which as-built conditions in communities reviewed by the authors actually

claims. Second, it reveals the importance of properly framing an AI experiment: the inputs here—including a corpus compromised with non-zoning texts and a prompt misdirected to focus on regulatory features antithetical to the type of zoning studied—raise questions about the validity of the results. And the critique exposes how ignoring zoning maps led to unsupportable claims about the adoption rates and impacts of a particular type of zoning code. Zoning can be a difficult and complex topic to write about, but for secondary research to be successful, zoning fundamentals must be more universally understood. We turn next to another type of application of AI, which scholars have used not to explain what is, but to imagine what could be.

B. Predicting Applications

1. Generally

Predicting applications of AI aim to identify patterns and relationships within inputted data to make informed forecasts about future action. Practitioners have used predicting applications of AI to analyze past judicial decisions and sentencing decisions, including those made by particular judges, to predict the outcome of new cases. They have been used to predict future trends in regulation and legislation, based on past trends. And they have been used to analyze legal documents, including contracts, to predict which phrases or sections may lead to litigation and to adapt those documents accordingly. Scholars, too, have analyzed historical legal outcomes, opinions, laws, and documents to identify trends and foresee potential outcomes.[92] Recent studies highlight this potential.

Studies in the United States and elsewhere have used AI to predict the way legal decisions will be made in the future. For example, a 2017 study aimed to predict decisions of the Supreme Court of the United States based only on Court cases previously decided. It generated a database of 240,000 votes (attributed to individual Justices) and 28,000 cases between 1816 and 2015, with database

---

fail to conform with zoning texts is unknown, as is the timing of the construction of these buildings relative to the adoption of particular provisions deemed FBC-like. Moreover, as noted above, the authors did not incorporate zoning maps into their analysis, so they cannot say whether relevant regulations applied to areas whose as-built conditions were studied. In addition to being more rigorous about their methods, the authors should be clearer about their claims, lest they be taken out of context.

The authors' reliance on the Wharton Residential Land Use Restrictive Index (WRLURI) to verify results is also highly misguided, for reasons fully articulated in a related article co-authored by this author. *See* Scott Markley, Sara C. Bronin, and Matthew Harris, *It's Time to Retire the Wharton Index*, at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=7391479.

[92] *See* Farid Ariai, Joel Mackenzie, & Gianluca Demartini, *Natural Language Processing for the Legal Domain: A Survey of Tasks, Datasets, Models and Challenges*, 58 ACM COMPUTING SURVEYS 1, 12-15 (2025) (citing numerous studies from around the world, including the Medvedeva article cited below, in which scholars used natural language processing to predict legal judgments); Harry Surden, *Machine Learning and Law*, 89 WASH. L. REV. 87, 102-107 (2014) (describing how machine learning might be used to predict information about the likelihood of cases having certain outcomes depending on sets of facts and past cases and noting that whether the model can be successfully trained on past cases often depends on future cases being similar to past cases in terms of their facts).

variables including the Justice, month of argument, and court of origin.[93] The authors initially trained their model using data from cases from the Court's first twenty-five years (1791 to 1816), but extended the training set period after the initial data yielded poor results.[94] Once the model received training on a sufficient number of cases, it achieved a successful prediction rate of about 70 percent.[95]

Curiosity about whether machines can help inform our understanding of legal decisions is not limited to the United States. Canadian scholars have used machine learning to analyze Canadian tax law, compiling six hundred cases over twenty years, and using 70 percent of the cases as training data and 30 percent for testing.[96] They created a portal in which facts about a case (say, a case involving whether a worker is an employee or an independent contractor) may be entered, and which predicts a case outcome with, they say, 90 percent accuracy.[97] Across the pond, European scholars have used natural language processing to analyze decisions made by the European Court of Human Rights, composed of many judges (currently numbering forty-seven) who oversee cases involving European Convention on Human Rights.[98] The authors downloaded all English-language case texts available on the court's website until 2017 and used machine learning to predict decisions between 2014 and 2017 based on cases decided the decade prior and based on the judge assigned. Assigning each case to one of three specific Convention articles that appeared frequently in the caselaw, they found accuracy rates between 58 and 68 percent (depending on the year) for predictions based on past decisions,[99] and accuracy rates of for 65 percent for predictions based on the assigned judge.[100] Perhaps this means that individual judges are more consistent than the court as a whole.

This cited scholarship uses predicting applications of AI to make educated guesses about the way humans will make decisions in the future. Because humans will make inconsistent decisions, or even evolve in their ideology, these applications would never be expected to be 100 percent accurate. Greater accuracy, though, might be expected in applications of AI designed to predict what the law actually says, a topic we turn to next.

---

[93] Daniel Katz, Michael Bommarito, & Josh Blackman, *A General Approach for Predicting the Behavior of the Supreme Court of the United States*, 12 PLoS ONE 1, 1 (2017).
[94] Daniel Katz, Michael Bommarito, & Josh Blackman, *A General Approach for Predicting the Behavior of the Supreme Court of the United States*, 12 PLoS ONE 1, 7 (2017).
[95] Daniel Katz, Michael Bommarito, & Josh Blackman, *A General Approach for Predicting the Behavior of the Supreme Court of the United States*, 12 PLoS ONE 1, 1 (2017).
[96] Benjamin Alarie, Anthony Niblett, & Albert H. Yoon, *Using Machine Learning to Predict Outcomes in Tax Law*, 58 CANADIAN BUS. L. J. 231 (2016).
[97] Benjamin Alarie, Anthony Niblett, & Albert H. Yoon, *Using Machine Learning to Predict Outcomes in Tax Law*, 58 CANADIAN BUS. L. J. 231, 242, 244 (2016).
[98] Marsha Medvedeva, Michael Vols, & Martijn Wieling, *Using Machine Learning to Predict Decisions of the European Court of Human Rights*, 28 ARTIFICIAL INTELLIGENCE & L. 237 (2020).
[99] Marsha Medvedeva, Michael Vols, & Martijn Wieling, *Using Machine Learning to Predict Decisions of the European Court of Human Rights*, 28 ARTIFICIAL INTELLIGENCE & L. 237, tbl. 10 (2020). Accuracy rates were 72 percent (for article 3), 64 percent (article 6), and 69 percent (article 8).
[100] Marsha Medvedeva, Michael Vols, & Martijn Wieling, *Using Machine Learning to Predict Decisions of the European Court of Human Rights*, 28 ARTIFICIAL INTELLIGENCE & L. 237, tbl 11.

2. Zoning-Related Scholarship

Trying to use AI to predict zoning codes may be a fruitless task. Unlike judges, whose lengthy, reasoned opinions have been mined by legal scholars in predicting applications, local zoning officials often fail to fully explain their logic and are not bound to a stare decisis-level of consistency across their decisions. Also unlike judges, local officials lack life tenure and may (in the case of elected officials involved with zoning) rotate out as regularly as each local election cycle. While it is difficult to predict future zoning rules within a single jurisdiction, it is even more difficult to guess at what zoning rules say by simply digesting what neighboring jurisdictions say. NZA analysis has confirmed that individual jurisdictions can be unpredictable, with a great degree of variety between adjacent localities.

Perhaps as a result, just as they have so far avoided publishing articles deploying AI to characterize zoning law, legal scholars have yet to use AI to predict outcomes of legal decisions or applications related to zoning. But we have an example of this type of analysis from another field. A 2024 article by a Ph.D. student in geospatial analytics, Margaret Lawrimore, used a corpus of information derived from North Carolina zoning codes to predict zoning conditions across the state, and to create a map showing these characteristics.[101] First, her research team manually collected regulatory and geospatial information about zoning districts from some or all jurisdictions within 39 counties, and classified these districts into three broad categories—residential, nonresidential, and mixed use—and thirteen

[101] Margaret A. Lawrimore, Georgina M. Sanchez, Cayla Cothron, Mirela G. Tulbure, Todd K. BenDor, & Ross K. Meenemeyer, *Creating Spatially Complete Zoning Maps Using Machine Learning*, 112 COMPUTERS, ENV'T & URB. SYSTEMS (2024), at 3–4. While the paper had many co-authors, in the text, I refer primarily to Lawrimore for convenience and for what I believe to be the appropriate attribution.

In another article, economist Jae-Hee Song built an algorithm purporting to detect zoning regulations from observed property characteristics. Jaehee Song, *The Effects of Residential Zoning in U.S. Housing Markets*, 149 J. URB. ECON. 1 (2025). She did not explicitly discuss AI processes in her article, so I do not include it here. But her methods provide an interesting counterpoint to Lawrimore's. Specifically, she theorized that she should design the algorithm to detect the lot sizes associated with new construction, assuming that new construction tends to be more frequent on lots that are just above the local minimum lot area. By detecting the point of "bunching" for new constructions, Song identified the minimum lot areas for nearly all municipalities in the forty-eight contiguous states and Washington, D.C. Song used the algorithm to create a dataset of 75 million single-family homes in 909 metropolitan areas across the United States for her study. Using this dataset, Song performed her analysis the relationship between minimum lot sizes and different housing outcomes and concluded that higher minimum lot areas decrease housing density, increase housing prices, and disproportionately attract Whiter and wealthier households. *Id.* at 23. While Song's study is more limited in scope (selecting one type of zoning regulation, minimum lot size) and in that way potentially more precise, her study shares several things in common with Lawrimore's study and can be similarly critiqued. Both studies rely on non-legal characteristics to predict the law. Both studies neglect the possibility that some jurisdictions decline to zone. Both studies claim victory in identifying zoning law from observed characteristics, a virtual impossibility given what we know about zoning (though Song does check her results against manually coded information in the greater Boston area and 55 other municipalities—still a small verification set relative to her created dataset of over 16,000 communities). And both studies, in misleading rhetoric, wrongly conflate the data they produced (a guess at the law) with the law itself.

sub-categories.[102] Then, for jurisdictions in all one hundred counties in the state, they collected information about 39 "predictor variables" documenting population, community characteristics, the built environment, and the natural environment across several dimensions.[103] They inputted the zoning information and the predictor variables into an AI model that then projected zoning in other areas.[104] These areas were located both in counties partially analyzed already (interpolation) and in the 61 counties not analyzed at all (extrapolation). Lawrimore claimed an astonishingly high 98.8 percent accuracy rate for information produced in interpolated counties, stating that of the sub-categories the office uses presented most challenges for the model.[105] She also claimed a 28 percent (for sub-categories) to 85 percent (for the three broad categories) accuracy rates for information produced in extrapolated counties, noting that nonresidential uses were least successfully predicted.[106] Several predicted-zoning maps of North Carolina illustrated the results.

As the only published article using AI to predict the content of American zoning law, the article has some laudable elements.[107] Notably, it recognizes three of the four principles articulating how zoning actually works: that codes are organized by zoning districts, that text regulates the district, and that the map identifies where regulation has effect.[108] However, the core assumption of her experiment—that one can predict zoning in one place based on zoning nearby—is

---

[102] Margaret A. Lawrimore, Georgina M. Sanchez, Cayla Cothron, Mirela G. Tulbure, Todd K. BenDor, & Ross K. Meenemeyer, *Creating Spatially Complete Zoning Maps Using Machine Learning*, 112 COMPUTERS, ENV'T & URB. SYSTEMS (2024), at 4 (classifying them into "three core districts (residential, non-residential, and mixed use) and thirteen sub-districts (open space, industrial, commercial, office, planned use, high-density residential, medium-high-density residential, medium-density residential, medium-low-density residential, low-density residential, agricultural residential, mixed use, and downtown")) (omitting emphases).

[103] Margaret A. Lawrimore, Georgina M. Sanchez, Cayla Cothron, Mirela G. Tulbure, Todd K. BenDor, & Ross K. Meenemeyer, *Creating Spatially Complete Zoning Maps Using Machine Learning*, 112 COMPUTERS, ENV'T & URB. SYSTEMS (2024), at 6 (tbl. 2) (listing the thirty-nine predictor variables in these four categories, including population, jurisdictional boundaries, parcel area, distances to amenities such as hospitals and schools, parks, voter registrations, social vulnerability as measured by indices or by housing type, transportation infrastructure, building footprints, building heights, elevation, hydrography, permanently protected areas, and land cover).

[104] Margaret A. Lawrimore, Georgina M. Sanchez, Cayla Cothron, Mirela G. Tulbure, Todd K. BenDor, & Ross K. Meenemeyer, *Creating Spatially Complete Zoning Maps Using Machine Learning*, 112 COMPUTERS, ENV'T & URB. SYSTEMS (2024), at 7–8, 11 fig. 8.

[105] Margaret A. Lawrimore, Georgina M. Sanchez, Cayla Cothron, Mirela G. Tulbure, Todd K. BenDor, & Ross K. Meenemeyer, *Creating Spatially Complete Zoning Maps Using Machine Learning*, 112 COMPUTERS, ENV'T & URB. SYSTEMS (2024), at 10.

[106] Margaret A. Lawrimore, Georgina M. Sanchez, Cayla Cothron, Mirela G. Tulbure, Todd K. BenDor, & Ross K. Meenemeyer, *Creating Spatially Complete Zoning Maps Using Machine Learning*, 112 COMPUTERS, ENV'T & URB. SYSTEMS (2024), at 11, 13.

[107] I also appreciated how Lawrimore explicitly noted that zoning codes had changed since her analysis, then made the case for her approach. Margaret A. Lawrimore, Georgina M. Sanchez, Cayla Cothron, Mirela G. Tulbure, Todd K. BenDor, & Ross K. Meenemeyer, *Creating Spatially Complete Zoning Maps Using Machine Learning*, 112 COMPUTERS, ENV'T & URB. SYSTEMS (2024), at 14 ("most zoning updates are targeted and specific to certain areas or properties, [so]… the majority of the zoning data is likely to remain relevant and informative for analysis and modeling purposes").

[108] *See supra* Part I.A.

at odds with our fourth basic principle of zoning: that zoning codes have incredible variety. Lawrimore attempts to resolve this contradiction by using 39 selected predictor variables to reverse-engineer zoning codes from built conditions, demographic characteristics, and environmental and other conditions. But her results confirm the difficulty of her task.

A visual comparison between her predicted-zoning maps and the manually-coded National Zoning Atlas (NZA)[109] roughly bear out the dismal accuracy rates she reports. When compared to the NZA, Lawrimore's maps show inaccurate residential densities and misplaced commercial areas across the state.[110] Another discrepancy is the extent to which the one-third of land logged as unzoned by the NZA[111] shows up as zoned in Lawrimore's maps. Her implicit assumption that nearly every inch of the state had zoning unfortunately resulted in a much higher inaccuracy rate than claimed in the article.[112] And yet the article trumpets its production of "the first spatially complete zoning map depicting core and sub-districts across the State of North Carolina."[113] When taken out of context, this statement might be seen as making a claim about the law itself.[114] In reality, this project did not produce a "zoning map," but a *guess* at what zoning maps might say. Further, the statement implies that the "zoning map" is accurate: if it is not accurate, then how could it purport to "depict" rather than "predict" the nature of

[109] Verification that could have informed refinement could have been partly done while the paper was being written; the NZA had completed and published to its public map full information about zoning in 34 of the 39 counties in the interpolated dataset, and full or partial information about zoning in about a dozen counties in the extrapolated dataset by 2024, when the paper was published. In total, by 2024, the NZA had published about 55 percent of North Carolina's 652 zoning jurisdictions, representing areas where 80 percent of the population lived.

[110] For example, the predicted-zoning maps depicting the thirteen sub-categories show a large amount of land predicted to be residentially zoned land with medium or high densities. Yet according to the NZA, much of the state's actual zoning is rural-residential or residential-agricultural, allowing residential at low or very low densities only. Lower densities occur when zoning codes impose large lot-size mandates, and the NZA has measured a quarter of single-family land in the state requiring an acre or more per home, with another third of such land having varying lot-size mandates (including many with large minimum lot sizes). National Zoning Atlas, Zoning Snapshot: North Carolina, at https://www.zoningatlas.org/snapshots?state=59. The patches of red on these maps predicting commercial zoning are also largely off. A large red section in central part of the state is, according to the NZA, actually residential and agricultural zoning; the northern part of Onslow County, also shown in red, is a Rural Agricultural district; the Alligator River National Wildlife Refuge in Dare County is also not only not commercial, but not zoneable by local authorities because it is federal land.

[111] The NZA shows that eighty-nine of the six hundred and fifty-two jurisdictions (cities, towns, and counties) across North Carolina decline not exercise zoning; still other jurisdictions decline to zone all of their land. National Zoning Atlas, Zoning Snapshot: North Carolina, at https://www.zoningatlas.org/snapshots?state=59.

[112] Perhaps a later study could explain why the algorithm would think that unzoned areas had zoning at all. It might well be possible, for example, that people living in unzoned areas organize their land uses even in the absence of law, as the work of legal scholar Robert Ellickson suggests may be possible. *See* ROBERT C. ELLICKSON, ORDER WITHOUT LAW: HOW NEIGHBORS SETTLE DISPUTES (1994).

[113] Margaret A. Lawrimore, Georgina M. Sanchez, Cayla Cothron, Mirela G. Tulbure, Todd K. BenDor, & Ross K. Meenemeyer, *Creating Spatially Complete Zoning Maps Using Machine Learning*, 112 COMPUTERS, ENV'T & URB. SYSTEMS (2024), at 34.

[114] Talen and Salazar-Miranda had the same issue. *See supra* n. __.

zoning districts across the state? The author should have been more careful about these claims, which again speak to the tendency of those using AI in the zoning context to say they achieved more than they did.

Many have critiqued the field of predictive analytics for, among other things, being prone to prior errors.[115] But here, AI was not to blame for the failure of the experiment to achieve its goal of using use AI to predict the content of zoning regulations. Even though the author did many things right, the assumption that zoning codes are predictable created an insurmountable problem.[116] Future scholarship should narrow the use of predicting applications of AI to those situations for which partly-right information about zoning at scale is good enough, and temper claims about what predicting applications of AI can do for our understanding of zoning. For example, predicted-zoning maps could be used by companies seeking a broad-brush understanding of which areas likely have more regulations than others, particularly if they need a place to site undesirable land uses, such as data centers, warehouses, or solar farms discussed below.[117] On the other hand, a predicted-zoning map could only help those companies so much. People who want to understand zoning most often want to know what zoning codes actually say. To build this knowledge, extracting applications of AI hold the most promise. We turn to those next.

C. Extracting Applications

1. Generally

AI extracting applications aim to pull information from large datasets, particularly unstructured ones, and turn that information into structured data—like turning a lengthy text into a spreadsheet with key attributes from the text organized into categories. In the legal context, extracting applications have been used to parse regulatory information or to take information from multi-part contracts or case files and then put that information into software, case management systems, or databases. Of the three AI applications described in this Part, extracting

[115] Critics of algorithmic predictions have noted its challenges from data input to output that doesn't reflect legal reasoning. *See* Sandra G. Mayson, *Bias In, Bias Out*, 128 YALE L.J. 2218, 2226 (2019) (arguing that predictive analytics used by algorithms function like a "mirror to the past" and project past racial disparities present in past-crime data); Hideyuki Matsumi & Daniel J. Solove, *The Prediction Society*, 2025 U. ILL. L. REV. 1, 11-15, 25 (2025) (discussing problems with "algorithmic predictions" that speculate the future, unlike "algorithmic inferences" which uses of verifiable present information, more akin to legal practice). Moreover, the complexity of zoning codes requires an understanding of the countervailing interests which decision-makers are considering. *See* Niva Elkin-Koren, *Contesting Algorithms: Restoring the Public Interest in Content Filtering by Artificial Intelligence.*, 7 BIG DATA & SOC. (2020), at 6 (arguing that machine learning algorithms output decisions without articulating the types of social tradeoffs made in balancing legal norms such as freedom of speech and national security).

[116] See Alicia Solow-Niederman, YooJung Choi, & Guy Van den Broeck, *The Institutional Life of Algorithmic Risk Assessment*, 34 BERKELEY TECH. L.J. 705, 724-725 (2019) (explaining how the same algorithmic tool may operate differently across local jurisdictions due to varying legal and institutional structures).

[117] *See supra* Part IV.A.1.

applications deal with the most challenging problems in comprehending legal texts.[118]

That may be why a 2025 article surveying the "tasks, datasets, models, and challenges" of natural language processing applications to legal contexts was replete with many examples of scholarship employing characterizing and predicting applications of AI, but very few examples of extracting-focused scholarship. The authors of that survey explained that, to perform the task of gathering information from legal texts, AI models must be "adapt[ed] to legal texts and further train[ed]… on specialised legal documents" and that "due to the challenges present in the legal domain, there are many legal LLMs [large language models] that have not undergone peer review."[119]

The authors went on to highlight, among other challenges, two relevant to zoning. First, they recognized the need to conduct extensive processing of legal texts to prepare them for AI use. They explained that "legal documents often consist of raw texts that require[] extensive cleaning and transformation before use in ML models. Their complex, nested structures, such as clauses within clauses and cross-references to other cases, statutes or provisions, make it hard to segment them into coherent units for analysis."[120] Zoning codes, in all their complexity, certainly exhibit these characteristics, perhaps more so than other areas of law.[121] The authors also identify a second challenge: scholars' inability to use "reinforcement learning from human feedback" to guide an AI model to the correct result through an iterative approach, due to a lack of credible, structured information about legal fields.[122] In the form of the National Zoning Atlas, the zoning field now has credible, structured information about legal fields and is the rare dataset manually assembled by experts that could help train an AI model on interpreting and applying complex concepts in law. (However, as Part III will discuss, AI has not necessarily advanced to the point where it can be trained on it.)

The scarcity of non-zoning-related legal research employing extracting applications of AI means we should turn immediately to articles purporting to use AI to extract information from zoning codes.

---

[118] A 2023 study of scholarship since 2013 involving natural language processing and law found information extraction (including labeling, text extraction, and event extraction) yielded the second highest number of papers among seven AI engineering tasks studied. Daniel Martin Katz, Dirk Hartung, Lauritz Gerlach, Abhik Jana, & Michael J. Bommarito, Natural Language Processing in the Legal Domain, Feb. 2023, at https://arxiv.org/pdf/2302.12039, at fig. 1, tbl. 1.

[119] *See* Farid Ariai, Joel Mackenzie, & Gianluca Demartini, *Natural Language Processing for the Legal Domain: A Survey of Tasks, Datasets, Models and Challenges*, 58 ACM COMPUTING SURVEYS 1, 25 (2025).

[120] *See* Farid Ariai, Joel Mackenzie, & Gianluca Demartini, *Natural Language Processing for the Legal Domain: A Survey of Tasks, Datasets, Models and Challenges*, 58 ACM COMPUTING SURVEYS 1, 25 (2025).

[121] *See supra* Part I.B. (zoning codes are ill-suited—in format and structure—for AI application).

[122] *See* Farid Ariai, Joel Mackenzie, & Gianluca Demartini, *Natural Language Processing for the Legal Domain: A Survey of Tasks, Datasets, Models and Challenges*, 58 ACM COMPUTING SURVEYS 1, 28 (2025) ("legal experts such as lawyers and judges must provide guidance to ensure the AI models accurately interpret and apply complex legal concepts, which is a barrier due to the significant cost of employing such professionals"). The authors stated that only one peer-reviewed work explores the approach.

2. Zoning-Related Scholarship

The exercise of collecting information from regulatory documents is complicated for humans.[123] But two recent, popular[124] papers from non-legal scholars that have applied AI extracting processes to zoning codes assert that machines can easily and accurately extract information from them.[125] Both a 2023 paper by Matthew Mleczko and Matthew Desmond, and a 2024 paper by Alexander Bartik, Arpit Gupta, and Daniel Milo claim to have used large language models to extract information from zoning codes to deliver accurate, comprehensive databases of regulatory information and accompanying restrictiveness indices for thousands of jurisdictions.[126] Unfortunately, the authors of both papers undermined their claims by failing to account for zoning's fundamental elements.

The earlier article, by Mleczko and Desmond, claimed to use natural language processing of zoning code text to build the "National Zoning and Land

---

[123] *See supra* Part II.

[124] Their popularity has been demonstrated by thousands of downloads from online repositories. *See* Alexander Bartik, Arpit Gupta & Daniel Milo, *The Costs of Housing Regulation: Evidence From Generative Regulatory Measurement*, SSRN, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4627587 (last visited Aug. 30, 2026) (downloaded 4,981 times); Matthew Mleczko & Matthew Desmond*, Using Natural Language Processing to Construct a National Zoning and Land Use Database*, SAGE J.: URB. STUD., https://journals.sagepub.com/doi/full/10.1177/00420980231156352 (choose "Metrics and citations" to view "Publication usage") (last visited Aug. 30, 2026) (downloaded and viewed 8,309 times).

[125] Another unpublished paper from 2021 also dealt with this issue, but it has not been as widely cited. Brendan Shanks used natural language processing to construct an index of zoning restrictiveness from 40,000 pages of Massachusetts zoning laws across 341 of the state's 352 zoning jurisdictions. Brendan Shanks, Land Use Regulations and Housing Development: Evidence from Tax Parcels and Zoning Bylaws in Massachusetts, Working Paper, Feb. 10, 2021. (The paper says that there are 351 towns, but there are 352 zoning jurisdictions, as, in addition to towns, the legislature granted the Devens Regional Enterprise Zone zoning authority.) He used what is called a "latent dirichlet allocation" model, which estimates the probability that a specific word was drawn from certain categories of topics. *Id.* at 3. Then, matching certain categories to another index for stringency, he created a "Natural Language Processing Zoning Stringency Index" to measure the restrictiveness of the zoning regulations. Using information about the built environment, and controlling for lot sizes, school districts, or property tax rates, he found that the more restrictive his index, the more expensive and the older the housing. *Id.* at 8. One could critique the text-centric nature of this analysis, as well as the issues related to the index development, in a similar manner as the other papers analyzed in the main text in this Subpart.

[126] Matthew Mleczko & Matthew Desmond, *Using Natural Language Processing to Construct a National Zoning and Land Use Database*, 60 URB. STUDIES 2564, 2565 (2023) ("As an *efficient and accurate* alternative… we access publicly available zoning and land use information and utilize natural language processing (NLP) to construct the National Zoning and Land Use Database (NZLUD)."); Alexander Bartik, Arpit Gupta, & Daniel Milo, The Costs of Housing Regulation: Evidence From Generative Regulatory Measurement, Working Paper, at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4627587, at 6-7 ("The central contribution of our project is the creation of a standardized dataset of zoning across the United States… [through] modern AI LLMs, which provide both the comprehensiveness and granularity" missing in prior methods).

Use Database," an "efficient and accurate alternative" to manual data collection.[127] Among other variables, the database includes information about the extent to which selected municipalities permit multi-family housing by right, cap the number of permits or units, require open space, mandate a minimum lot size, impose inclusionary-zoning mandates (such as fees or unit set-asides for affordable housing), and permit a certain number of dwellings per acre.[128] To build this database, the authors downloaded 2,639 municipal codes.[129] They cleaned these texts to enable machine readability.[130] Then to extract desired information about the variables, they identified an initial list of search keywords, and refined the list further based on an analysis of the first round of outputs.[131] After running the algorithm, the authors organized information extracted from the zoning texts into a database containing single answers for each variable. They offered this database for free public download. The authors also generated a new measure, the "Zoning Restrictiveness Index," that they claim illuminates the relative restrictiveness of zoning conditions uniformly across thousands of jurisdictions.

The later piece, by Bartik et al., took a similar approach.[132] The authors downloaded municipal ordinances for around 4,900 municipalities.[133] They cleaned and then fed the text of those ordinances into a large language model and asked the model to answer eleven binary questions, including whether multi-family housing is allowed, ADUs are allowed, mandates or incentives for affordable housing exist,

---

[127] Matthew Mleczko & Matthew Desmond, *Using Natural Language Processing to Construct a National Zoning and Land Use Database*, 60 URB. STUDIES 2564, 2565 (2023).

[128] Matthew Mleczko & Matthew Desmond, *Using Natural Language Processing to Construct a National Zoning and Land Use Database*, 60 URB. STUDIES 2564, 2569-70 (2023).

[129] Matthew Mleczko & Matthew Desmond, *Using Natural Language Processing to Construct a National Zoning and Land Use Database*, 60 URB. STUDIES 2564 (2023). They noted 3 percent of jurisdictions in the hoped-for sample of 2,660 municipalities did not put their zoning codes online. This is a lower number than the percentage documented by the NZA, likely because the selected municipalities were located in urbanized areas with larger jurisdictions more likely to have the administrative capacity to maintain a public website.

[130] Matthew Mleczko & Matthew Desmond, *Using Natural Language Processing to Construct a National Zoning and Land Use Database*, 60 URB. STUDIES 2564, 2570 (2023) (explaining how they filtered out unnecessary words, punctuation marks, or symbols and standardized numeric information, among other cleaning steps).

[131] For example, they asked the machine to exclude phrases that had led to false positives in the initial round of outputs. They also added words that matched, such as the various names for inclusionary housing programs used by the selected municipalities.

[132] The article is very lengthy, and this critique focuses only on the parts relevant to AI extraction processes.

[133] Alexander Bartik, Arpit Gupta, & Daniel Milo, The Costs of Housing Regulation: Evidence From Generative Regulatory Measurement, Working Paper, at 40, tbl. 1, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4627587. Without evidence, the authors claim that "We obtain municipal codes for 63% of the population governed by local zoning ordinances." Alexander Bartik, Arpit Gupta, & Daniel Milo, The Costs of Housing Regulation: Evidence From Generative Regulatory Measurement, Working Paper, at 1, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4627587. The authors have no way of knowing the *denominator*—i.e., the total population of people subject to zoning codes—because we still do not have a final count of the jurisdictions that exercise and do not exercise zoning authority.

and any housing has age restrictions.[134] They also asked the model to extract information about certain numerical variables, such as the size of the smallest minimum lot size and the size of the largest minimum frontage requirement across all single-family districts in the jurisdiction.[135] They acknowledged that the algorithm must "make judgment calls when answering questions (i.e., choosing the more common type of single-family home minimum lot size requirements within a district)."[136] Like Mleczko and Diamond, the authors made their collected data and technical coding publicly available, while asserting that their "AI-driven approach is scalable, auditable, and allows for refinement."[137]

In both of these papers, the authors boasted that they used AI to accurately extract regulatory information from zoning texts, and that they have put the extracted information into definitive databases about zoning conditions and corresponding indices zoning restrictiveness at the level of the jurisdiction. Yet in both articles, the authors could not have possibly produced the claimed results, because their methods failed to account for the way zoning codes actually work.

Critically, these authors again ignore all four principles outlined in Part I and describing how zoning actually works.[138] Of these four principles, most egregiously given their claim is that they declined to account for the zoning maps that are essential to understanding codes' restrictiveness.[139] No study can assess a jurisdiction's regulatory restrictiveness without accounting for the geographic extent to which specific regulations apply.[140] In identifying jurisdictional permissions for multi-family housing, for example, Mleczko and Diamond required the algorithm to track such permissions for each zoning district; then, the authors simply divided the number of districts allowing multi-family housing by the total

---

[134] Alexander Bartik, Arpit Gupta, & Daniel Milo, The Costs of Housing Regulation: Evidence From Generative Regulatory Measurement, Working Paper, at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4627587, at 42.

[135] Alexander Bartik, Arpit Gupta, & Daniel Milo, The Costs of Housing Regulation: Evidence From Generative Regulatory Measurement, Working Paper, at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4627587, at 12. For this and other numerical variables, the LLM was also required to name all of the districts allowing single-family housing.

[136] Alexander Bartik, Arpit Gupta, & Daniel Milo, The Costs of Housing Regulation: Evidence From Generative Regulatory Measurement, Working Paper, at 6, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4627587.

[137] Alexander Bartik, Arpit Gupta, & Daniel Milo, The Costs of Housing Regulation: Evidence From Generative Regulatory Measurement, Working Paper, at 35, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4627587.

[138] *See supra* Part I.A. ("first, the zoning district, which distinctly regulates land uses, structures, and lots, is the organizing feature of the code; second, the text laying out regulations for each district governs development in each district; third, the map identifies where zoning districts are assigned, and thus where regulation has effect; and fourth, jurisdictions demonstrate seemingly infinite variety in the way they articulate the rules in the text and delineate districts on the map").

[139] One of the papers recognizes at the end of their article that "[i]ncorporating zoning map data into our procedure could be a worthwhile avenue for future research." Matthew Mleczko & Matthew Desmond, *Using Natural Language Processing to Construct a National Zoning and Land Use Database*, 60 URB. STUDIES 2564, 2581 (2023).

[140] *See supra* Part I.A. (describing the necessity of understanding both text and map to understand the actual impact of codes). *See also* Scott Markley, Sara C. Bronin, and Matthew Harris, *It's Time to Retire the Wharton Index*, at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=7391479.

number of districts.[141] But percentages of zoning districts satisfying a criteria cannot possibly measure of the extent to which regulations may apply in a jurisdiction, because districts differ greatly in size. A jurisdiction with ten zoning districts, among them six multi-family districts that collectively cover only a few blocks within the jurisdiction, cannot be said to be permissive of multi-family housing.[142] Yet applying the authors' approach would suggest that this jurisdiction is permissive of multi-family housing because 60 percent of the districts permit multi-family housing.

In the Bartik et al. paper, similar issues arose. Illustrating this were maps of four metropolitan areas depicting the minimum lot size variable incorporated into the restrictiveness index.[143] The maps color code each jurisdiction based on their *minimum* minimum lot size (the lowest lot size offered by the jurisdiction).[144] South Fulton, a town of 100,000 adjacent to Atlanta, is depicted as having a minimum minimum lot size of between 0 to 5,000. This figure is correct but misleading, as according to the NZA, only 11 percent of single-family land meets the criteria.[145] In other words, 89 percent of land has a higher minimum lot size, which means that South Fulton is not as permissive of small lots as the map suggests. The authors' Atlanta map logs a minimum minimum lot size of between 7,000 and 8,000 square feet. But the NZA indicates that fully 15 percent of single-family land has *no* minimum lot size.[146] In this case, the answer was incorrectly extracted by the algorithm, possibly because it did not know how to accept a null answer.[147] Results from the San Francisco area are similarly distorted. For example, the NZA shows a *single acre* of Walnut Creek (which has a total of 7,930 zoned acres) having a minimum lot mandate of between 0 and 5,000 square feet.[148] Yet Bartik et al. attribute the entire jurisdiction to that range. (I hasten to note that .) Meanwhile, Berkeley appears to have been bumped into the category of cities with a minimum minimum lot mandate of 5,000 to 7,000 square feet, despite having a mixed-use

---

[141] Matthew Mleczko & Matthew Desmond, *Using Natural Language Processing to Construct a National Zoning and Land Use Database*, 60 URB. STUDIES 2564, 2571 (2023).

[142] *See supra* Part I.A. (explaining an example jurisdiction with ten districts of different geographic areas).

[143] Alexander Bartik, Arpit Gupta, & Daniel Milo, The Costs of Housing Regulation: Evidence From Generative Regulatory Measurement, Working Paper, at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4627587, at 63, fig. 11.

[144] Alexander Bartik, Arpit Gupta, & Daniel Milo, The Costs of Housing Regulation: Evidence From Generative Regulatory Measurement, Working Paper, at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4627587, at 14, 22. The discussion here is confusing; the authors indicate that they have estimated "the minimum of all lot size requirements," but they also refer to an average of lot size requirements. Because the "minimum minimum" appears twice, I believe that this is what their minimum lot size figure calculates.

[145] National Zoning Atlas, Zoning Snapshot: South Fulton, at https://www.zoningatlas.org/snapshots/?jurisdiction=3569.

[146] National Zoning Atlas, Zoning Snapshot: Atlanta, at https://www.zoningatlas.org/snapshots/?jurisdiction=3556.

[147] Had the algorithm gotten the answer correct, depicting it on the map would have been misleading for the same reasons that South Fulton's depiction was misleading.

[148] National Zoning Atlas, Zoning Snapshot: Walnut Creek, at https://www.zoningatlas.org/snapshots/?jurisdiction=10763.

district with a smaller mandate.[149] In some cases, the algorithm did not work very well. In others, the algorithm functioned as instructed, but the instructions themselves skewed results.

Relatedly, the authors of both papers trained their algorithms to select single answers (whether binary yes-no answers or numerical answers) for inclusion in the database for each variable, using criteria that seem arbitrary. By attempting to collapse a very high amount of regulatory complexity, the authors failed to respect the role of the zoning district or the variety of zoning regulations. For example, Mleczko and Diamond incorporated into the database only a single number for minimum lot size: the largest minimum lot size required by the jurisdiction.[150] The figure selected was highly misleading, for the same reasons the minimum minimum lot size was misleading in the Bartik et al. paper. For example, consider a jurisdiction that has a single, large-lot neighborhood of a few blocks with a lot size of 1.5 acres, while the vast majority of the jurisdiction has 0.25-acre lot sizes. That jurisdiction would be represented in the database at the misleading, and higher, 1.5-acre figure. Meanwhile, Bartik et al. instructed the algorithm to identify the minimum lot size for single-family housing in each residential district, and to choose "the most common standard size, excluding sizes for historic properties or special cases."[151] However, as noted above, often minimum lot sizes are presented as alternative regulations, and forcing the choice of a single number belies the way zoning works on the ground. Moreover, the question of what constitutes "the most common standard size" would be difficult to answer without broader context. For example, in a jurisdiction known to be a retirement community, senior, age-restricted housing may be more common than other types of housing; an algorithm would have no way of knowing that it should select as "the most common standard size" the lot size mandate for senior housing over the mandate applying more generally to other types of housing. The authors similarly attempted to oversimplify parking mandates,[152] heights,[153] and other types of regulations.

---

[149] National Zoning Atlas, at https://www.zoningatlas.org/atlas.

[150] Matthew Mleczko & Matthew Desmond, *Using Natural Language Processing to Construct a National Zoning and Land Use Database*, 60 URB. STUDIES 2564, 2568-69 (2023). They did make efforts to filter out minimum lot sizes for planned residential developments of multiple houses, typically much larger than the minimum lot sizes for individual houses. *Id.* at 2571.

[151] Alexander Bartik, Arpit Gupta, & Daniel Milo, The Costs of Housing Regulation: Evidence From Generative Regulatory Measurement, Working Paper, at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4627587, at 96.

[152] The authors took the average of the median and mode of all parking requirements for residential uses, a nonsensical choice given the variety of regulations in different districts and for different residential uses, the frequent incidence of alternative parking regulations, and the failure to map mandates' applicability. Matthew Mleczko & Matthew Desmond, *Using Natural Language Processing to Construct a National Zoning and Land Use Database*, 60 URB. STUDIES 2564, 2570 (2023).

[153] The authors arbitrarily average the median and mode building heights across all residential zoning districts in the jurisdiction in feet and in stories, and it's not clear whether they differentiate between single-family housing heights and multi-family housing heights. Matthew Mleczko & Matthew Desmond, *Using Natural Language Processing to Construct a National Zoning and Land Use Database*, 60 URB. STUDIES 2564, 2570 (2023). They convert a single story to fourteen feet, *id.* at n. 5, an unusually high number.

The authors of both papers might have salvaged their research by being more modest and precise in their claims—at a minimum jettisoning their assertion that their particular mix of arbitrary answers could have been combined to measure a jurisdiction's restrictiveness in any way. In the scholarship studying land use, we have seen the troubling rise of scholars eager to have *some* measure of zoning incorporating other flawed, publicly available indices of jurisdictional zoning restrictiveness into their research.[154] For the reasons articulated above, other researchers should avoid using the Mlezcko and Diamond and Bartik databases and indices for any study that needs a measure of zoning conditions, including zoning restrictiveness, in the covered jurisdictions.

* * *

In the scholarship purporting to use AI to help us understand zoning and discussed in this Part, scholars have tried but generally failed to harness AI to achieve their characterizing, predicting, or extracting purposes. That is because they have either misunderstood or ignored how zoning actually works. The most egregious omission is zoning's spatial element, which is core to understanding the extent to which regulatory conditions apply, and thus to understanding the extent to which a jurisdiction is more or less restrictive.[155] Moreover, these scholars generally took a reductionist approach, attempting to simplify complex regulations into binary answers or arbitrary numbers, and wrongly obscuring necessary nuance.[156] They also made false assumptions—including that all land is zoned, or that all regulations in text apply throughout a jurisdiction. These methodological errors and incorrect assumptions significantly detracted from the accuracy claims made by scholars about their work. Often made without explanation, these accuracy claims led to grand pronouncements of the "first-ever" or "comprehensive" map, database, or study. Unfortunately, these scholars embedded so many errors into their AI models that the only thing we know for sure is that the algorithms never had a chance of succeeding at their intended tasks.

The irony is that scholars' overclaiming, oversimplifying, misunderstanding, and miscalculating is diminishing—not improving—our understanding of zoning. To use AI to generate new knowledge about zoning, scholars must commit to better comprehending the subject matter, and building skills needed to equip AI models with the information they need to fulfill their intended tasks. Even if they do everything right, though, I contend that AI's current capability is incapable of fully replacing human reasoning in this particular legal sphere. I base that contention primarily on the AI experiments described next.

---

[154] Chief among these is the Wharton Residential Land Use Restrictive Index (WRLURI) to verify results. *See* Scott Markley, Sara C. Bronin, and Matthew Harris, *It's Time to Retire the Wharton Index*, at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=7391479.

[155] As explained above, Lawrimore did not ignore zoning's spatial element.

[156] They do not, for example, log alternative regulations which are prevalent in zoning codes. *See supra* Part I.B.3.

III. UNRESOLVABLE CHALLENGES: THE NATIONAL ZONING ATLAS AI TRIALS

Critiques of scholarship on AI applications to zoning reveal that, so far, scholars have only programmed AI models to fail at illuminating zoning. Because Part II focused on human error, it did not evaluate AI's technological capacity to parse zoning codes. This Part now turns to the capacity question, using one consequential application of AI to read and interpret zoning codes.

As noted above, I lead the National Zoning Atlas (NZA), the country's largest geocoded data repository about zoning laws. The NZA uses a rigorous methodology to parse zoning codes: first, finding and verifying texts and maps; second, extracting the names of zoning districts; third, classifying these districts in broad categories based on their allowed uses; and then fourth, delving into the texts and maps independently, with text analysts manually searching for, and logging, up to 200 regulatory characteristics per district and geospatial analysts drawing individual zoning districts. The fifth and final step in our process—supervisory review and approval leading to merging of the textual and geospatial analysis—results in jurisdictions being published to the public, online map. The NZA team has repeated and refined these steps for more than five years, as it manually produces the most robust dataset possible given zoning's inconsistencies and complexities.

In theory, a properly trained AI model could accelerate completion of each of these laborious steps. Thus I jumped at two opportunities to explore the application of AI to three of these steps. The NZA's first AI trial focused on the second and third steps in the NZA workflow: district-name extraction and district classification. Our second AI trial focused on the text-analysis component of the fourth step: i.e., the district-specific data extraction. (Neither trial focused on predicting applications, because, for reasons explained above, imputing zoning data from existing zoning data is not a reliable way to achieve our specific task, to convert information in codes to a structured, standardized dataset.) To both of these trials, the NZA team brought two major assets. First, the team's considerable expertise helped us to avoid the design errors enumerated in Part II: we strove to instruct the model to perform exactly the same nuanced analysis we conducted as humans. Second, the NZA dataset formed an up-to-date, accurate corpus capable of training models. Still, our AI trials still left us wanting—not because of human error, but because the technology has not yet sufficiently advanced to replicate the human reasoning of knowledgeable experts.

After discussing the methods, results, and lessons of each trial, this Part concludes with some thoughts about the implications of the NZA AI trial results on AI scholarship on zoning specifically and legal text analysis more generally.

A. TRIAL 1: METHODS, RESULTS, AND LESSONS

In 2022, we launched a small AI experiment with the Urban Institute, aiming to achieve two tasks: to extract the names of zoning districts from a corpus of zoning texts and maps (an extracting application) in 180 Connecticut

jurisdictions and then to classify them as primarily residential, mixed with residential, or nonresidential (a characterizing application).[157] These two steps both come at the beginning of the manual workflow for building the NZA, and both typically involve human labor and judgment. An explanation of methods and results follows.[158]

1. The District-Name Extraction Application

The district-name task charged the AI model with scanning the zoning text and maps for keywords and providing, as output, the names of found districts. The NZA had already assembled the relevant codes, having completed review of the full state of Connecticut the prior year. As a necessary first step, optical character recognition software was used to convert these codes' unreadable features, such as scanned images or hand-drawn maps, to readable text.[159] After conversion, the model was asked to scan both texts and maps for words such as "district" or "zone," as well as phrases like "divided into" or "are hereby established," which often precede a list of districts.[160] It was asked to only extract names that appeared in both the zoning map and zoning text. The hypothesis was that developing a process to extract mapped, non-extinct[161] districts would be more straightforward than having to search for unmapped or extinct district names appearing in only one of the text and map. We knew this would mean that the model would not spot the 11 percent of Connecticut zoning districts we knew to be either extinct or unmapped, because by definition these districts only appeared in either the text or map, but not both.[162] Thus the maximum success rate was capped by design at 89 percent.

The team addressed many potential challenges in its instructions. For example, the model was told that zoning codes typically use a full name and an abbreviated name to refer to the same district; and it was told that these pairs of names are often found together in the table of contents, map legend, or section titles. It was also told that maps sometimes use abbreviated names, rather than full names, in the map legend. With those instructions, the algorithm was asked to put the names into a database, and de-duplicate the list. The output included 2,256 district

---

[157] Judah Axelrod, Lydia Lo, & Sara C. Bronin, *Automating Zoning Data Collection*, Urban Inst. Research Report, Feb. 2023.

[158] The tasks were divided as follows: NZA assembled and prepared zoning codes, provided a glossary of terms, and provided and cleaned the training data, while the Urban Institute authors developed the AI prompts, processes, and coding, and both sides collaborated on experiment design, hypothesis development, and verification of outputs against the training data.

[159] Judah Axelrod, Lydia Lo, & Sara C. Bronin, *Automating Zoning Data Collection*, Urban Inst. Research Report, Feb. 2023, at 9.

[160] Judah Axelrod, Lydia Lo, & Sara C. Bronin, *Automating Zoning Data Collection*, Urban Inst. Research Report, Feb. 2023, at 9-10.

[161] *See supra* Part I (explaining that some districts may be unmapped, because they are referenced in the text but not located on the map, and that some districts may be extinct, being named in the map but not in the code).

[162] This figure comes from a data pull from the NZA database on December 31, 2025, showing that of 2,776 zoning districts in Connecticut, 33, or 1.2 percent, of districts are extinct, and 274, or 10 percent, are unmapped. The overall number of districts increased since the time of this AI trial, but the proportion of extinct and unmapped districts remained the same.

names, which were then manually checked against the 2,385 districts listed in the database, leaving 1,317 district names that roughly matched NZA results, for a success rate of 55 percent.[163]

Despite all appropriate instructions, the AI model could not match the success of human coders. It often failed to find districts listed outside of the table of contents, map legend, or section titles. It also failed to find districts where texts and maps remained illegible even after attempted conversion. And zoning district names could vary between the map and text, with abbreviations used in one document (typically the map) but not the other, making matching impossible for an algorithm that could not make a leap between "Industrial I" in the text and "IND-1" in the map.

The research team concluded that an AI model might assist with finding the page of the zoning text on which a list of districts could appear, while a human reviewer could then review the selected pages and then review the zoning map.[164] That said, this two-step process would not likely save time for the typical human analyst—who can easily search for "table of contents," "divided into" and "are hereby established" themselves, and then can pick up the scattered district names after. (Instructions for undertaking this search are already included in the NZA methodology.[165]) Moreover, it might take a typical human analyst longer to identify which of the districts were correct, and locate missing districts, than if the analyst started from scratch.

2. The District-Classification Characterizing Application

Like the district-name extraction task, the classification task is foundational to the NZA workflow. The public NZA map displays districts across all jurisdictions as being in one of three categories: primarily residential, mixed with residential (a mixed-use district), or nonresidential. While seemingly simplistic, we have come to understand that these categories allow people to use our map to ascertain, at a glance and across one jurisdiction to the next, how each published jurisdiction approaches zoning. Analysts must classify every zoning district in the database as fitting into one of these three categories.

The design of this part of the experiment relied on the NZA methodology, which instructs analysts to classify districts based on the uses allowed in the districts. The methodology offers shortcuts for preliminary classification based on the name of the district; including, for example, that district names containing the words "airport," "public utility," "mining," or "cemetery" likely to be

---

[163] *See supra* Part I.C.1, citing Judah Axelrod, Lydia Lo, & Sara C. Bronin, *Automating Zoning Data Collection*, Urban Inst. Research Report, Feb. 2023, at 11.

[164] *See supra* Part I.C.1, citing Judah Axelrod, Lydia Lo, & Sara C. Bronin, *Automating Zoning Data Collection*, Urban Inst. Research Report, Feb. 2023, at 12.

[165] Sara C. Bronin, Scott Markley, Aline Fader, & Evan Derickson, *How to Make a Zoning Atlas 2.0: The Official Methodology for the National Zoning Atlas*, Apr. 2024, at 46 (instructing analysts to scan any table of contents; search for the phrases "divided into," "divides the," "hereby," "established," and/or "following zoning districts; search for the words "zone" or "district"; and review use tables or other charts likely to list most or all districts).

nonresidential.[166] However, the methodology requires analysts to search either the use table, or the descriptions of districts, or both, to ascertain the types of uses allowed and thus the final classification.[167] To replicate these methods, the AI trial used machine learning to search for the following phrases: "uses permitted," "permitted uses," "use regulations," "use table," and "allowed use."[168] It also instructed the model to identify, extract, and read paragraphs describing the purpose of specific zoning districts. The collaborators relied on manually coded classifications for 2,385 zoning districts, using 70 percent of districts as a training set, and the remaining 30 percent as a testing set which would be manually verified after the experiment. Within the training set, the model identified variables most associated with the three classifications, learning that districts with words like "residential," "residence," and "housing" in their titles were most likely to be correctly classified as primarily residential, while words like "public" and "industry" were more likely to lead to nonresidential classifications.[169] The model was then asked to apply the findings to the 30 percent of testing districts to characterize them.

After performing all of these steps, this part of the experiment had an overall accuracy of 74 percent.[170] Primarily residential districts were most accurately characterized; nonresidential districts least accurately characterized. Interestingly, the most successful model was one that relied solely on the name of the zoning district. It was not as successful when attempting to use context surrounding the district name in the text. The research team explained that challenges in using the context-based approaches included the inability of the model to: locate district descriptions away from the names of a known zoning district, calibrate the amount of prose extracted, read information in tabular format, and interpret bespoke symbology in use tables.[171] A process that misclassifies a quarter of the zoning districts in the NZA database would result in a highly distorted map. Thus, like the district-name task, the results were not considered a plausible substitute for NZA analyst labor.

* * *

[166] Sara C. Bronin, Scott Markley, Aline Fader, & Evan Derickson, *How to Make a Zoning Atlas 2.0: The Official Methodology for the National Zoning Atlas*, Apr. 2024, at 72.
[167] Sara C. Bronin, Scott Markley, Aline Fader, & Evan Derickson, *How to Make a Zoning Atlas 2.0: The Official Methodology for the National Zoning Atlas*, Apr. 2024, at 69-73.
[168] Judah Axelrod, Lydia Lo, & Sara C. Bronin, *Automating Zoning Data Collection*, Urban Inst. Research Report, Feb. 2023, at 27.
[169] Judah Axelrod, Lydia Lo, & Sara C. Bronin, *Automating Zoning Data Collection*, Urban Inst. Research Report, Feb. 2023, at 16.
[170] Judah Axelrod, Lydia Lo, & Sara C. Bronin, *Automating Zoning Data Collection*, Urban Inst. Research Report, Feb. 2023, at 17. The authors discuss the importance of precision ("the proportion of zoning districts that the model predicts as falling within one of the three categories" correctly) and recall ("the proportion of zoning districts that actually fall into one of the three categories that the model predicts correctly") measures, but these other measures add unnecessary complication given the purposes of this Article. *Id.* at 16.
[171] Judah Axelrod, Lydia Lo, & Sara C. Bronin, *Automating Zoning Data Collection*, Urban Inst. Research Report, Feb. 2023, at 13-14.

Despite insufficient success in both tasks, this relatively straightforward experiment with the Urban Institute nonetheless offered several valuable lessons. It underscored the difficulties using zoning codes in AI experiments given their formatting challenges, inconsistencies in text structure, and difficulties deciphering maps.[172] It revealed that terminology, which can differ greatly from one code to the next, matters a great deal to AI applications, including zoning district names and descriptive text.[173] And it also highlighted the consequences of the all-too-common occurrence of mismatches between text and map, a particular challenge for the district-name extracting task.[174] Here, consequential mismatches included not only typographical mismatches between text and map district names, but also the problem of extinct and unmapped districts, which had been omitted from the results due to the experiment's requirement that a zoning district name must be in both text and map. All of these issues derive from the inherent qualities of American zoning codes, discussed in detail in Part A.

The Urban Institute experiment was limited in scope and did not attempt to utilize AI to extract district-specific information from zoning codes—the bulk of the labor of NZA analysts. However, the collaborators forewarned that such an extraction task would be difficult, stating that "most pieces of information that the National Zoning Atlas requires are not such ideal candidates for algorithmic approaches," especially queries requiring open-ended answers.[175] More AI trials would be needed to test whether this statement were true. And so we launched the trial described next.

B. TRIAL 2: METHODS, RESULTS, AND LESSONS

Zoning texts are ill-suited for AI analysis because of the way they are drafted and amended,[176] but in theory, these challenges might be overcome with the right experiment design and execution. Building on the NZA's first AI trial, which exposed these challenges, the NZA's second AI trial began in 2023 with research collaborator Dr. Alexander Rush and the student members of his lab.[177] We established as our aim the use of large language models (LLMs) to extract information about certain regulations about specific zoning districts from zoning texts. In choosing to focus on this task, we decided not to pursue a new iteration of the district-type characterizing application of the Urban Institute trial, or to pursue other characterizing applications. We also declined to pursue a predicting application, which, as explained above,[178] is not relevant to understanding what zoning codes say.

---

[172] *See supra* Parts I.B.2., I.C.1., I.D.1.
[173] *See supra* Part I.C.2.
[174] *See supra* Part I.D.2.
[175] Judah Axelrod, Lydia Lo, & Sara C. Bronin, *Automating Zoning Data Collection*, Urban Inst. Research Report, Feb. 2023, at 20.
[176] *See supra* Parts I.A., I.C.
[177] This work was funded by the National Science Foundation, for a grant, No. 2242302, entitled "Automating Analysis of Zoning Codes for the National Zoning Atlas."
[178] *See supra* Part II.B.

Similarly, we declined to include a geospatial component with this analysis. Our geospatial workflow is conducted separately and could have been paired with textual results if the LLM were successful in extracting district-specific information from the text.[179] Here, we purely intended to ascertain whether AI could help accelerate the completion of the text-based analysis.

To test the LLM's extracting capacity, we narrowed the tested variables to three: minimum lot sizes, minimum unit size, and building heights for single-family housing. We sought answers in the form of the number of acres or square feet for minimum lot size, the number of square feet of a housing unit for minimum unit size, and the number of stories or feet for maximum building height. Where the zoning code offered alternative regulations, we sought the full text of the phrase explaining the given alternatives. This type of data extraction is core to the work of the NZA analysts. Given this overview of the experiment design, this Subpart now proceeds to explain the methods, results, and lessons of this text-extraction trial.

1. Methods

In this two-year AI experiment, we chose to use LLMs, which as noted above can sift through large amounts of text to find answers to specified questions. Humans can train LLMs to find these answers by inputting sample texts and correct answers, which improves the LLM's capacity to replicate the processes that produced the correct answers. With that in mind, the experiment involved several steps: preparing data for inputting, translating zoning terminology, developing and refining queries (including an answer-generation query and a ranking query), and evaluating outputs.[180] Descriptions of these steps suitable for a legal (rather than highly technical) audience follow.

Preparing data was an important first task. The team gathered an initial corpus of 176 zoning texts from Connecticut in PDF format. As with the Urban Institute experiment, these texts were run through an optical character recognition service to ensure machine readability, including the conversion of tables into data structured like a spreadsheet, with rows and columns.[181] In total, the initial corpus consisted of around 20 million words encompassing 31,905 pages of text. While LLMs can accept a large amount of text, they cannot accept this volume of text without requiring significant computing and monetary cost. Thus the text was divided (in AI parlance, "tokenized") into individual pages. These page numbers

[179] We did not aim to extract zoning-district names (an aim of the Urban Institute experiment, which necessarily incorporated zoning map analysis), or to create a composite regulatory index or otherwise assess the restrictiveness of zoning (the aim of scholarly works described in Part II.C.). Had our aim included these goals, it would have necessarily required a geospatial component. *See supra* Parts I.A. & II.

[180] The tasks were divided as follows: NZA assembled and prepared zoning codes, provided a glossary of terms, annotated textual analysis and cross-checking for North Carolina and Texas jurisdictions, and provided and cleaned the training data, while Dr. Rush and his students developed the AI prompts, processes, and coding, and both sides collaborated on experiment design, hypothesis development, and verification of outputs against the training data.

[181] Cornell Tech Zoning Team, Zoning Codes and LLMs, Working Paper, Jan. 2024, at 1 (describing use of the Amazon Textract OCR service).

were indexed to correspond with those in the PDF of the zoning text, enabling future output evaluation. The tokenized and indexed pages were loaded onto a server to await machine processing. The same data preparation steps were used to prepare a second set of zoning texts, from jurisdictions in Texas, before their use in a later round. Note that at one point during the second trial, we attempted to extract the names of zoning districts for 149 North Carolina jurisdictions. But a manual check of this data revealed that the AI model had exactly the same rate of success as the first trial.[182] We should have heeded the lesson from the first trial that an AI model could not successfully complete the judgment-requiring process of identifying district names.[183] For the rest of the project, the team included as a given input the zoning district names that had already been manually extracted by NZA analysts; in the Connecticut batch of codes, for example, this amounted to 2,430 districts.

As the initial documents were being prepared for analysis, the team also had to develop a shared understanding of terminology relevant to the various topics to be explored in the AI query. To account for various potential terms across a wide range of documents,[184] the team created a glossary of synonyms and related terms. This glossary went deeper than the one used in the Urban Institute experiment, with more terms encompassing more types of housing development and regulatory conditions. This glossary was repeatedly refined over the course of the experiment, as queries produced outputs revealing more specialized terms. For example, the glossary of terms likely to help the model find a "minimum lot size" eventually encompassed fourteen words and phrases: "lot," "lot and building," "area and bulk," "dimensional," "minimum lot," "minimum lot coverage," "min lot," "min lot coverage," "lot requirements," "lot and building requirements," "area requirement," "area and bulk requirement," "dimensional requirement," and "minimum area." In addition, the model was trained to associate units with the queries: square feet for minimum lot size and minimum unit size, feet or stories for heights.

Developing and refining queries meant determining, through trial and error, what exactly the model needed to produce.[185] The team ultimately identified two

[182] The manual verification by an NZA analyst of district names for these 149 North Carolina jurisdictions confirmed that of 3,702 answers for district names provided by the AI model: 1,990 (or 54 percent) were correct; 539 (15 percent) required correction; 246 (7 percent) were incomplete; 376 (10 percent) were not districts; 353 (10 percent) had new names; and 198 (5 percent) were correct but not preferred due to NZA naming conventions. (Numbers add up to more than 100 percent due to rounding.) The error rate did not include the names of districts entirely missing from the extraction, which would have further reduced the accuracy rate. Common errors included: duplicative district names, symbols added or missing (such as "RM" instead of "R-M"), inclusion of both the full district name and the abbreviated district name (such as "R-1 Residential"), errors in names to sub-districts, hallucinated abbreviated names where some districts (such as overlays) were not assigned abbreviated names in the code.

[183] *See supra* Parts I.A., I.C.1.

[184] *See supra* Part I.C.2. (explaining the phenomenon of inconsistent terminology across zoning codes).

[185] Over the course of the project, at least four AI systems were used to test these queries: DaVinci 003, ChatGPT 3, ChatGPT 4, and ChatGPT 4-Turbo, which through testing became the preferred model. Ultimately, an agentic system known as "ZoningAgent" was developed to autonomously

sets of queries: one intended to generate answers, and the other intended to rank the answers. The answer-generating prompt asked AI to pretend like it was a lawyer, and to find a specified answer (without guessing) in the text provided.[186] It typically established the format for outputs (in AI parlance, the "schema") in the form of page numbers, extracted text which could then be further mined for the exact answer, and in some case an explanation with a measure of confidence. The exact answer might be required to have certain units of measurement (like acres and square feet for minimum lot sizes). Often, the prompt yielded multiple answers within the same jurisdiction. Thus a second query was developed to ask the machine to rank the answers based on the text and any explanation provided in the first query. The ranking prompt provided the AI model with information about each answer, the respective zoning text supporting it, and the explanation for its selection.[187] Over the course of the project, many refinements were made to the prompts to ensure that the output was both relevant and properly structured.

These queries produced single, "best-answer" outputs for minimum lot size, minimum unit size, and maximum building height for single-family homes for each zoning district in the sample. Next, the team needed to understand whether these outputs conformed with the manually-collected data. There were two sets of ground-truth data to review: first, Connecticut zoning code data in tabular format; and second, Texas zoning code data in tabular format and supported by annotations (tracked in the NZA's proprietary data entry portal) by NZA analysts identifying the page number and location coordinates on the page for each relevant tabular entry, including both the district name, the field name (e.g., "minimum lot size"), and the answer. (The creation of this annotation portal was a significant technological feat from a programming standpoint.) As queries were run, the Rush lab checked outputs from queries run on the Connecticut and Texas inputs based on best practices in their field, and refined the queries as noted above.

At the end of the project, the NZA team did one final round of testing, on a few dozen North Carolina codes that had no ground-truth data but had been run through the AI model. The team had built in the NZA proprietary data entry portal a mechanism to display the AI results, including the answer(s), the text in the zoning code where the answer(s) was supposedly located, and the model's reasoning. In this portal, an NZA analyst could confirm whether the answer(s)

---

parse zoning codes by replicating the process of optical character recognition, indexing and query-based retrieval, rationale extraction, and answer generation.

[186] The exact prompt for answer generation for districts, for which information was sought for single-family homes, looked something like this:

> You are an expert architectural lawyer. You are looking for facts inside a document about a Zoning District with the name "{{zone_name}}" and with an abbreviated name "{{zone_abbreviation}}". You are looking to find the value for "{{term}}", which also goes by the following other names: {{synonyms}}. Only output values that are seen in the input and do not guess! Output MUST be valid JSON, and should follow the schema detailed below. Ensure that the field "extracted_text" does not chunk multiple lines. For {{term}} in residential districts, we are only interested in the answer as it pertains to single-family homes.

Cornell Tech Zoning Team, Zoning Codes and LLMs, Working Paper, Jan. 2024, at 5.

[187] Cornell Tech Zoning Team, Zoning Codes and LLMs, Working Paper, Jan. 2024, at 5. The output schema here was the selection of a single answer among multiple answers.

provided conformed with the manually-collected data. An NZA analyst performed this task for the North Carolina codes, roughly verifying the same accuracy rates as for the Connecticut and Texas codes. The NZA analyst logged the amount of time it took to verify the answers, and then estimated the amount of time saved by the AI's involvement. We turn next to the results logged by both teams.

2. Results and Lessons

The multi-faceted project yielded interesting results. Across the three tested variables, the trial saw 80 percent accuracy for minimum lot sizes, 20 percent accuracy for minimum unit sizes, and 83 percent accuracy for maximum height variables across the Connecticut and Texas zoning codes. Accuracy rates were measured by the Rush lab team, which compared the outputs of the AI model against the outputs of the manually coded data set. When the expected answer was a single number, the comparison and verification was automated. When the expected answer was a phrase (as when the zoning code offered alternative regulations), the verification process required a bit of review. It bears noting that the Rush lab team considered any of multiple potential answers for alternative regulations as correct. For example, if the manual data for minimum lot size showed a requirement of one acre under some conditions but two acres under others, they would count a prediction of either one or two as correct. While consistent with computational science practices, the NZA team would consider only the full phrase setting forth alternative regulations (such as "one if X, two if Y") to be correct enough to incorporate into the NZA. While the model did capture some alternative regulations,[188] it was less successful at spotting them in full than single-number answers.

The queries were most successful when the expected answers were simple, as when the code provided a single number immediately adjacent to expected units, like "1 acre" for minimum lot size, or "1,000 square feet" for minimum unit size. The errors typically occurred when those numbers were separated from their units or district names, when finding the answer required more than two cross-references, and when the answer came in the form of alternative regulations like "1 acre minimum lot size per unit, or 0.05 acres if the housing units are affordable units."[189] The lowest accuracy rate, for minimum unit size, might be explained by two facts. First, very few districts in the training set actually had a minimum unit size, meaning there were situations where the model produced an answer where none exists. Second, unit size and lot size use identical measurement units (square footage), which might have tempted the model to substitute lot-size answers for unit-size queries. A machine's confusion about units corresponding to specific queries is understandable, given that certain units appear in multiple sections

[188] One emblematic correct response was the request to find minimum lot size for single-family housing in the SU District in Colchester, Connecticut. *See* COLCHESTER, CONN., ZONING REGULATIONS, § 4.4.1. After reading two tables (one for areas "with available municipal water and sewer" and one without), the model correctly stated the mandatory minimum lot size as "30,000 sq ft (with available municipal water and sewer); 40,000 sq ft (without available municipal water and sewer)."

[189] *See supra* Part I.C.3. for a discussion of alternative regulations.

covering different types of regulations. For example, the unit of "feet" is used not only in lot and unit sizes and building heights, but also in setbacks (front, side, and rear setbacks), parking requirements, connectivity requirements (whether housing must be near certain amenities, like transit), and more.

Another way to measure the experiment's success is the time it saved as compared to manual data extraction methods. It is unquestionable that the AI model saved time in locating possible answers; the Rush lab team members reported that the time savings to extract information from a four hundred-page document about each of the three variables was reduced from approximately 1.5 minutes by a human using manual methods to ten seconds by the model.[190] However, this figure only accounted for data location; it did not account for the manual verification required after machine extraction, which was required given the error rate of the machine outputs. Unfortunately, when the NZA team manually checked outputs related to North Carolina codes, they estimated that it took as much or more time to review the outputs and to find the correct answers for erroneous outputs than for the NZA analyst to perform the entire task manually from scratch. That was true even though the machine-extracted answers were ranked, and specific page numbers and page text appeared on the screen.

Despite these issues, this AI experiment, like the first, offered valuable lessons. It reinforced findings from the first AI trial about the inherent challenges in using zoning codes in AI experiments. Inconsistencies in text structure and terminology, and alternative regulations in their seemingly infinite variety, both negatively impacted results.[191] These results discouraged us from thinking zoning district names could be extracted from zoning texts. On the other hand, the experiment hinted at the possibility that AI could improve data quality; the team found four or five districts where NZA volunteers had incorrectly extracted the Connecticut data the first time around.[192] The experiment also resulted in two promising technology integrations enabling NZA analysts to make location-related annotations of zoning text for the training model and to verify outputs within the NZA database. Finally, the AI model was programmed to include the somewhat novel features of offering reasoning for, and then ranking, provided answers—heralding a rare opportunity for a human user to understand the model's "thinking."

From the standpoint of the team's AI researchers, the trial was both successful on its own terms and a promising model for other experiments. However,

---

[190] Annual Report 2025, Report to the National Science Foundation for Grant #2242302, Nov. 2025, at 2.

[191] *See supra* Part I.C.

[192] For example, the minimum lot size for the DMR zoning district in Mansfield, Connecticut, was referred to twice in the zoning text: one on page, the minimum acreage required to create a district was five acres, but the city could reduce lot sizes to five thousand square feet if the district were subdivided for one- or two-family dwellings; another figure allows residential density at 8.7 units per acre, or five thousand square feet. *See* MANSFIELD, CONN., ZONING REGULATIONS § 10.A.4.h, fig. 10.B-2. While the manually coded data provided a minimum lot size of five acres, after discussion we determined that the intent of the code was likely to enable single-family dwellings on lots as small as five thousand square feet, and corrected the entry in the master NZA file. Note that the Connecticut data was manually collected by student volunteers; the Texas data later collected by NZA full-time analysts did not, to my memory, have any errors spotted by AI and requiring correction.

from the standpoint of the NZA, the model did not yet achieve either the level of accuracy or the time savings that would encourage the NZA to rely on it for data extraction.

C. Implications on AI Scholarship About Zoning and Beyond

As others involved in zoning research heralded their successes in using AI to characterize, predict, or extract information from zoning codes, the team building the NZA had big expectations for our own AI trials focused on the extracting application of AI. Having analyzed zoning data across thousands of jurisdictions using a painstaking manual process, no one could have been more excited than we were about the possibility of automating some our work. We now realize that other scholars' claims about AI rested on shaky foundations.[193] And our own trials focused on extracting applications of AI revealed that the NZA's multi-step processes to find information and resolve ambiguities could not yet be replicated by AI.

Maybe we should not feel so badly about this. There's a good argument to be made that zoning is the most complicated body of law around. Measured by sheer volume, local zoning laws likely add up to constitute the largest single body of law in the American legal canon. If the length of zoning texts of jurisdictions not yet analyzed by the NZA match the length of jurisdictions in the NZA database, we will need to read more than two and a half million pages of zoning codes to analyze the full country.[194] Moreover, these pages are written and adopted in a decentralized manner by tens of thousands of jurisdictions that differ (in big ways and small) in regulatory approach from one place to the next. And the people writing and adopting these approaches are largely laypeople, who amend their codes in piecemeal fashion. Together, these characteristics of zoning likely introduce inconsistencies more frequently than bodies of law written by legislative counsel or other experts.

On the other hand, zoning seems like just the type of law that AI should be trained, and trainable, to decipher. The hope of achieving this task seems tantalizing possible given a unique training asset: the manually coded regulatory dataset developed by the NZA for more than a third of the country's jurisdictions. To my knowledge, no other scholars have developed as robust a training set as ours, because of our extensive use of freedom-of-information-act requests, on-site code collection, and manual contacts. Theoretically, this dataset can buttress the development and refinement of an extracting application of AI that would enable successful, accurate, and comprehensive analysis of the content of zoning codes. And perhaps AI models today would be more sophisticated and achieve higher accuracy rates than the ones we used for our trials. Still, significant verification

[193] *See supra* Part II (critiquing scholarship applying AI to characterize, predict, and extract information from zoning codes as ungrounded in a basic understanding of codes' characteristics).

[194] Here's the math: If there are 33,296 jurisdictions in the country, and we have analyzed 11,200, then we have around 23,000 left. Thus the jurisdictions incomplete outnumber the jurisdictions complete by 1.97 times. Multiplying 1.97 by 1.35 million yields 2.66 million. This figure is a rough estimate only; it does not account, for example, for the possibility that the remaining 23,000 jurisdictions will have different conditions (including choosing not to zone) than the first 11,200.

would have to be conducted, and new, and possibly extra-textual, approaches that replicate human reasoning to resolve ambiguities, would likely have to be developed for an AI model to fully substitute for human review. In the meantime, this Article intends raise awareness about the issues specific to zoning analysis by AI, temper overexuberant claims, and point the way to more methodologically sound analysis in the future.

Beyond zoning, our AI trials also contribute to broader scholarly discussion about AI's suitability to interpret legal texts.[195] Scholars have long identified the variety of issues that could arise if lawyers were replaced by machines, from the exclusion of human values from legal processes to the erosion of confidence in the judiciary.[196] There are also broader ethical issues raised by AI, including the theft of creative output[197] and the increasing, and increasingly horrible environmental and human-health costs of data centers that support AI. This Article raises a different issue, diving deeply into a specific body of law to show how AI cannot yet capably perform tasks requiring human reasoning, even with an experienced team training AI models and a robust manually-produced corpus. We should either expect to devote significant manual labor to helping AI interpret complex legal texts, or shift what we believe AI can achieve. The next Part makes suggestions for how we might shift our thinking, at least in zoning, about the types of applications of AI that could achieve success with current technological capacity.

---

[195] An increasing amount of scholarship covers concerns about privacy and discrimination by AI, but that does not seem relevant here, in a process reviewing publicly available codes. *See* Richard M. Re & Alicia Solow-Niederman, *Developing Artificially Intelligent Justice*, 22 STAN. TECH. L. REV. 242 (2019); Solon Barocas & Andrew D. Selbst, *Big Data's Disparate Impact*, 104 CALIF. L. REV. 671 (2016) (discussing data mining and its potentially discriminatory impact); Ryan Calo, *Robotics and the Lessons of Cyberlaw*, 103 CAL. L. REV. 513 (2015); Danielle Keats Citron & Frank Pasquale, *The Scored Society: Due Process for Automated Predictions*, 89 WASH. L. REV. 1 (2014).

[196] *See*, *e.g.*, JERRY KAPLAN, HUMANS NEED NOT APPLY: A GUIDE TO WEALTH AND WORK IN THE AGE OF ARTIFICIAL INTELLIGENCE (2015); RICHARD SUSSKIND & DANIEL SUSSKIND, THE FUTURE OF PROFESSIONS: HOW TECHNOLOGY WILL TRANSFORM THE WORK OF HUMANS (2015); RICHARD SUSSKIND, TOMORROW'S LAWYERS: AN INTRODUCTION TO YOUR FUTURE (2013); Aziz Z. Huq, *A Right to a Human Decision*, 106 VA. L. REV. 611 (2020) (investigating whether there should be a right to a decision in court that is made by a human); Frank Pasquale, *A Rule of Persons, Not Machines: The Limits of Legal Automation*, 87 GEO. WASH. L. REV. 1 (2019) (arguing that the legal profession must recognize that legal automation elides or excludes important human values and that to preserve accountability, legal reasoning must be made by humans, not machines); W. Bradley Wendel, *The Promise and Limitations of Artificial Intelligence in the Practice of Law*, 72 OKLA. L. REV. 21, 24-25 (2019) (expressing doubts that AI can replace attorneys, including litigators preparing for and participating in trials); Daniel Martin Katz, *Quantitative Legal Prediction—or—How I Learned to Stop Worrying and Start Preparing for the Data-Driven Future of the Legal Services Industry,"* 62 EMORY L.J. 909 (2013).

[197] *See, e.g.*, Katherine Lee, E. Feder Cooper, & James Grimmelman, *Talking About AI Generation: Copyright and the Generative-AI Supply Chain*, 72 J. COPYRIGHT SOC'Y 251 (2025); Jennifer Cobbe, Michael Veale, Jatinder Singh, *Understanding Accountability in Algorithmic Supply Chains*, 2023 ACM Conference on Fairness, Accountability, and Transparency, at https://arxiv.org/abs/2304.14749.

IV. RECALIBRATING EXPECTATIONS ABOUT AI IN ZONING

In our enthusiasm to utilize AI to do everything humans do, we have overlooked its limitations—and our own. Certainly, this is true in the zoning context, where recent tests of the ability of AI to accurately characterize, predict, and extract regulatory information from zoning codes have disappointed. To some extent, the disappointment derives from human error, including either a lack of familiarity about zoning or a willingness to take dangerous short-cuts.[198] But even when people training AI models are familiar with zoning's particulars, as is the case with the NZA team, the AI models cannot quite deliver.

Rather than push AI farther than its current capabilities, we should pinpoint more modest uses of AI that can viably enhance our understanding of zoning codes and zoning practice. This Part briefly touches on characterizing and predicting applications of AI. It then outlines plausible uses of AI to extract information from a variety of sources, both within zoning codes and outside of them, to improve our understanding of zoning. It argues that AI can enhance our understanding of zoning codes and zoning practice, albeit in a more modest way than recent scholarship has claimed, or recent experiments have expected.

A. Characterizing Applications

Characterizing applications of AI are almost certainly AI's most reliable applications to legal contexts. These might be deployed either to classify certain elements of zoning codes or to improve understanding about the particulars of zoning code administration.

The success of either effort depends on suitable human prompts. For example, classifying zoning codes as form-based codes (or not), as attempted by Salazar-Miranda and Talen, will be more achievable with inputs that better reflect zoning conditions.[199] The latest AI models, which are more capable of processing texts, might also have greater technological capacity to improve upon prior attempts to classify zoning district types.[200] One can also imagine using AI tools like the word-embedding technique to measure the extent to which zoning codes in different jurisdictions, and in different types of jurisdictions, are similar to each other. Like the question in the scholarship on treaties discussed above,[201] a scholar could assess whether large cities tend to have codes more similar to each other than they do to small town codes, or whether similar codes cluster within metropolitan areas. Another area of exploration, like Choi's analysis of the word "sanitation" in the pandemic-era judicial decision,[202] might be for AI to tell us the most common

[198] *See supra* Part II.
[199] As emphasized above, the main critique of their paper was not the capability of AI but their choice of inputs. *See supra* Part II.A.2.
[200] *See supra* Part III.A.2. (describing the National Zoning Atlas's attempt in 2022 to train a model to classify the type of as primarily residential, mixed with residential, or nonresidential).
[201] *See supra* Part II.A.1. (discussing Elliott Ash & Omri Marian, *The Making of International Tax Law: Empirical Evidence from Tax Treaties Text*, 24 FLA. TAX REV. 155 (2020)).
[202] *See supra* Part II.A.1. (discussing Jonathan H. Choi, *Measuring Clarity in Legal Text*, 91 U. CHIC. L. REV. 1 (2024)).

meanings of particular terms across codes, within their context and common understandings of the words. Searching zoning codes for the definition of the term "family"—which, as I have observed elsewhere is deeply problematic—would yield interesting results that might tell us what moral and relational values are embedded in legal documents that have unexpected influence on family structure.[203] While many possibilities exist, scholarship has not been activated in this direction.

Just as interesting as reviewing codes themselves is the possibility that AI could scour manually coded datasets *about* codes, such as the dataset compiled by the NZA team. In the near future, the NZA team will likely ask an AI system to pull out trends from our dataset across a broad range of regulatory topics—the use of public hearings, allowances for accessory dwellings, treatment of parking mandates, and more. LLM capacity seems well-suited to the task of properly characterizing large datasets, and we are excited to see what we might find.

AI could also help us better understand trends in the administration of zoning codes. Patricia Salkin and her student co-authors have identified several such opportunities.[204] They posit, for example, that AI could help with meetings and public engagement by automating meeting minutes and meeting summaries, managing agendas, and analyzing public comments at community meetings.[205] AI could also analyze inspector citations and reports to assess zoning violation patterns. An understanding of the most common types of violations could in turn inform refinements to codes themselves. In addition, AI can be applied to characterize decisions, such as those made by zoning boards of appeals, or zoning agreements, such as those negotiated between cities and developers.

B. Predicting Applications

Predicting applications of AI may not tell us as much about existing zoning as characterizing applications. In fact, as discussed when critiquing the Lawrimore article above, it is impossible for a human or a machine to predict what zoning codes say in one geography based on what zoning codes say in another. However, with the right inputs, predicting applications might help us make guesses about possible directions for zoning decisions, impacts, and permitting.

One highly plausible potential use of AI is to digest zoning decisions and then predict future decisions of local boards and commissions made pursuant to those codes. A repository of commission meeting recordings, transcripts or minutes, for example, might be mined to understand trends in reasoning across time, as Stiglitz, Thalkin, and Choi aimed to do in analyzing judicial decisions.[206] One might use this repository to determine, for example, whether decisions over a specified period have come to treat multi-family housing more favorably in recent

---

[203] Sara C. Bronin, *Zoning for Families*, 95 INDIANA L. J. 1 (2020).
[204] David Looney, Harrison Stern & Patricia Salkin, *Using AI to Assist in Land Use Planning and Zoning: Opportunities and Pitfalls*, MUNICIPAL LAWYER, Nov. 2025, at 37.
[205] David Looney, Harrison Stern & Patricia Salkin, *Using AI to Assist in Land Use Planning and Zoning: Opportunities and Pitfalls*, MUNICIPAL LAWYER, Nov. 2025, at 43-44.
[206] *See supra* Part II.B.1.

years, as one might guess given the rise of the pro-homes movement. One might also use the repository to assess whether decision-makers take a statutory, normative, textualist, or purposivist approach, and whether this approach differs by zoning jurisdiction (on, say, the basis of political ideology, population, or geography). This analysis could in turn influence the way people approach their zoning applications on the front end. On the back end, the analysis could influence judicial review of zoning decisions. Careful experiment design could account for the fact that many local elected officials involved with zoning may be regularly replaced. On the other hand, non-elected board members tend to serve for many years on zoning commissions and may provide more continuity.

Salkin and her co-authors also believe AI could predict the environmental impacts of new zoning changes, and potential consequences of certain actions.[207] They suggest AI could also be used to speed up the permitting process, visualize zoning changes through simulated maps or virtual reality walk-throughs, and analyze different development scenarios.[208] Each of these applications can help us understand zoning in new ways, and are good uses of predicting applications of AI. would enable the NZA to generate scalable, transparent, and reproducible workflows.

C. Extracting Applications

As discussed in Part III, extracting applications of AI seem to be challenging to apply to a comprehensive analysis of zoning codes, given issues inherent in codes themselves. However, three different types of extracting applications might be feasible and might help us better understand zoning. First, a data-expansion application could pull certain information about specific, easy-to-find regulations from zoning texts. Second, a geospatial-processing application could accelerate the digitization of zoning maps. And third, a data-maintenance application could utilize webscraping to monitor local websites for announcements of amendments, which would automate notices of amendments to texts and maps. There are certainly other potential applications of AI that could transform information from zoning texts and maps, but these three illustrate the range of possibilities. Each has the added benefit of potentially being used by the NZA to generate scalable, transparent, and reproducible workflows and to ensure that NZA data remains comprehensive, current, and accurate. Explaining these applications' utility to an understanding of zoning generally, and to the NZA project specifically, offers insights for scholars and others seeking to develop more realistic extracting applications of AI in the zoning context.

1. Deriving Limited Textual Information

[207] David Looney, Harrison Stern & Patricia Salkin, *Using AI to Assist in Land Use Planning and Zoning: Opportunities and Pitfalls*, MUNICIPAL LAWYER, Nov. 2025, at 37.
[208] David Looney, Harrison Stern & Patricia Salkin, *Using AI to Assist in Land Use Planning and Zoning: Opportunities and Pitfalls*, MUNICIPAL LAWYER, Nov. 2025, at 43-44.

In light of the complexity of zoning texts and the current technological capacity of AI, certain types of data extraction might not be feasible. However, an algorithm that is fed the names of zoning districts might be able to extract very specific types of information from zoning texts.

One potential aim of an extraction application could be to supplement the NZA's information about whether codes allow certain residential uses with information about whether codes allow certain nonresidential uses. Nearly all datapoints collected by the NZA relate to zoning regulations on residential development. The one exception is that the NZA identifies the "most common nonresidential use" allowed in each zoning district from a list of 14 categories: 11 specified categories such as "office" and "industrial," and 3 catch-all categories such as "specific commercial uses." Given that these categories are so general, there is an unmet need to collect and disseminate information about specific types of nonresidential uses. For example, the NZA team has repeatedly received inquiries about data centers, warehouses, and solar farms—each of which has proven controversial from a siting standpoint. Data centers, which can bring jobs and investment, also draw from water stores and use significant energy.[209] Warehouses have been increasingly necessary as online retailers expand, but they have increased in size, and their large delivery trucks bring traffic, noise, and pollution.[210] Large solar installations, which many view as necessary to reduce reliance on fossil fuels and increase energy independence, are viewed with suspicion by land conservationists and Indian Tribes.[211] Knowing where these types of uses are explicitly permitted could inform ongoing policy conversations.

---

[209] TERRY NGUYEN & BEN GREEN, WHAT HAPPENS WHEN DATA CENTERS COME TO TOWN? 5–6 (2025) (explaining that data centers, which rely on tremendous and unsustainable amounts of water and electricity, "cannot fully operate on renewable energy alone," making any clean energy pledges by data center operators empty promises "with no commercially viable models yet in operation"); Felicity Barringer, *Thirsty for Power and Water, AI-Crunching Data Centers Sprout Across the West* (Apr. 8, 2025), https://andthewest.stanford.edu/2025/thirsty-for-power-and-water-ai-crunching-data-centers-sprout-across-the-west/ (discussing data centers' high demand for water and power, which have caused concerns for local communities and governments).

[210] Kaveh Waddell & Maanvi Singh, *"Warehouses in Their Backyards": When Amazon Expands, These Communities Pay the Price* (Dec. 9 2021), https://www.theguardian.com/us-news/2021/dec/09/when-amazon-expands-these-communities-pay-the-price#:~:text=CR's%20investigation%20shows%20that%20the,to%20distribute%20them%20more%20evenly. (concluding that Amazon is more likely to open warehouses in neighborhoods with "a greater proportion of people of color," contributing to "disproportionate health and environmental impacts" in marginalized communities); Ryan Murphy & Emily Feng, *Why More Residents Are Saying "No" to AI Data Centers in Their Backyard* (July 17, 2025) https://www.npr.org/2025/07/17/nx-s1-5469933/virginia-data-centers-residents-saying-no (finding that local residents are increasingly opposing warehouses and data centers in their neighborhoods for reasons ranging from noise, space constraints, pollution, and excessive water usage)

[211] Troy A. Rule, *Rural Solar Rights*, 51 BYU L. REV. 447, 466–67 (2025) (explaining that despite distinct benefits of siting solar farms on private rural land, local residents increasingly oppose solar farms due to possibilities of food production reduction, environmental impact, and safety risks); Wesley J. Furlong, *The Other Non-Renewable Resource: Cultural Resource Protection in a Changing Energy Future*, 42 PUB. LAND & RES. L. REV. 1, 11 (2020) (demonstrating the detrimental cultural effects of siting solar farms on or near tribal land through the 709-megawatt solar farm approved near El Centro, which threatened over 400 cultural sites belonging to the Quechan Tribe of the Fort Yuma Indian Reservation); Sara C. Bronin, *Curbing Energy Sprawl with Microgrids*, 43

I believe this project could perform better than the NZA AI trial because the query would, by definition, search for words or phrases, and not numbers. As discussed above, numbers appear so frequently in zoning codes that they might confound a model.[212] Moreover, we believe that the models have improved to some extent, including in finding terms and in reading charts (like use tables). Setting up this project would begin by identifying a detailed list of nonresidential uses worth knowing about, given social, economic, and political factors. Whatever use or uses were chosen, a glossary of similar terms would have to be developed to inform an AI query.

Next, the team would develop and refine an answer-generation query and a ranking query. The answer-generation query would include prompts similar to those in the second NZA AI trial. It would include, for example, synonyms for the use, as well as known full and abbreviated zoning district names pulled from the one hundred thousand district names included in the NZA database.[213] The model could be trained on NZA analyst manual annotation of codes from around one hundred different jurisdictions, where the annotations identify the location (page number and location on the page) for the selected uses. The ranking query could be required to explain the machine's reasoning and to show the location in the text supporting the model's answer.

Finally, the team would evaluate the outputs. Ideally, outputs would be integrated into the NZA workflow and into its proprietary data entry portal, as occurred with the second NZA AI trial. The team would incorporate document search across sections to verify the validity of its extraction in a human readable format.

This hypothetical use of AI to expand information included in the NZA database would be designed to avoid some of the pitfalls raised by prior AI experiments. It would focus on narrowly defined unique words or phrases, instead of numbers, and it would supplement information already collected by the NZA. If successful, we could use the same process on other types of extracting queries.

2. Transforming Zoning Maps

While many AI advances, like those discussed in the preceding section, have focused on text, there have also been advances in the utilization of AI in geospatial (mapping) applications. As the two preceding sections have discussed text-based applications of AI, we turn now to potential map-based applications of AI.

AI could accelerate the rate at which zoning maps in their native format are cleaned and digitized. Of course, the primary application I have in mind here is the construction of the NZA. By design, the NZA is a map created from other maps:

CONN. L. REV. 547, 553–55 (2010) (observing that extra-urban siting of energy sprawl, which includes solar farms, can cause "significant negative long-term impacts," such as land waste in energy production and threats to wildlife).

[212] *See supra* Part III.B.2.

[213] These names have been manually extracted from zoning codes. As explained above, the NZA team tried twice, but failed, to use AI models to extract zoning district names from zoning codes. *See supra* Part III.

namely, tens of thousands of official public maps that show the location of distinct zoning districts. As noted above, these maps are kept by localities in one of two ways: static (raster) maps in PDF or graphic format; or downloadable geospatial data.[214] Currently, for a raster map, NZA staff georeference the map and trace polygons, converting it into a vector map; while for a downloadable map, NZA staff review the geospatial layers to identify and correct errors (such as leftover "slivers" or misalignments with jurisdictional boundaries).[215] For both types of maps, NZA staff ascertain whether the digital versions of the maps show the correct zoning districts. These processes are time-consuming and labor-intensive. And there are two ways that AI could help make them faster.

First, AI could feasibly help to more quickly digitize zoning maps obtained from public sources and integrate them into the NZA by automating different steps in these processes. The overall idea for experiment design would be very similar to the design of the data-expansion project discussed above. NZA analysts could create a corpus of around one hundred manually georeferenced raster image maps, and use these maps to train an AI model to better understand how to separate objects and features, and assign these objects and features as zoning districts. (They would exclude maps like the indecipherable maps shown in Part I.C. above.) The model would be further trained to do specific tasks, such as reading the map legend and making sure zoning district boundaries do not exceed the boundaries of, say, federal parkland, waterways, or roads. The output of the model would be a clean map in a compatible format. The final product would be verified by NZA analysts, who could compare the output to the official zoning map. Ideally, the model could improve upon current NZA methods of human identification of errors, including by helping to identify errors a human might not be able to see with the naked eye.

Second, AI could help to identify map-to-text discrepancies, meaning that it could help identify the areas of ambiguity that would then have to be resolved by a human analyst. As noted above, these discrepancies appear in perhaps one third of zoning codes.[216] They take the form of minor discrepancies, such as clerical errors, and major ones, such as extinct districts that appear only in the map, but not the text. AI that identifies these ambiguities and flags them for a human to resolve would save time in our process of building the NZA, and could spot errors a human might miss.

Though mapping technology has not been a focus of this Article, there have been very promising applications of the use of AI to predict and extract information from geospatial files. In the predicting realm, Microsoft has used AI (and deep neural networks) to extract building footprints from satellite imagery or raster data to produce vector maps of building footprints. As for extraction, several private companies have developed processes for raster to vector conversion.[217] Tailoring

[214] *See supra* Part I.A.1.
[215] *See supra* Part I.D.
[216] *See supra* Part I.D.2.; Jonah Pellecchia & Melissa Hayashida, *Eliminating Map-to-Text Discrepancies*, ZONING PRACTICE, Oct. 2024.
[217] One such company is Bunting Labs; a number of smaller companies have experimented with this feature as well.

some of these tools to read zoning maps, and to compare maps with texts, would greatly assist with improving our understanding of zoning.

3. Aggregating Amendments

Zoning codes, as legal documents stewarded by governmental bodies, are often amended. While few jurisdictions completely overhaul their codes,[218] many jurisdictions adjust their texts and maps in minor ways on a somewhat regular basis. Such amendments might include changed regulations on uses, structures, and lots; additions or deletions of zoning districts; or modifications to procedural and administrative requirements.

Over the last few years, the NZA has assessed the rate of code updates in several jurisdictions. In California, nearly all 539 jurisdictions updated their codes between 2023 and 2025. In Connecticut, between 2022 and 2025, 91 percent of the 185 jurisdictions updated their zoning codes. In New Hampshire, of 232 jurisdictions with zoning, 74 modified their codes to change their housing approach over the course of a year, not counting zoning code modifications affecting matters other than residential uses.[219] Though the numbers vary from jurisdiction to jurisdiction, the amount of change in these three states suggests that a large number of American zoning jurisdictions change their codes in any given year, albeit modestly.

Despite the large number and regularity of amendments, jurisdictions have not always developed a smooth path for publishing information about amendments to the public. Jurisdictions that host websites for their codes might not post amendments in a timely manner, or might merely reference that an amendment exists without yet incorporating it into the text or map. Jurisdictions that pay companies to put their codes online may not promptly deliver information about amendments to those companies, or the company may delay posting.[220] At times, we have found amendments referenced in the meeting minutes of a local zoning commission and nothing referenced in the code or official website at all. And, as noted above, some codes are not publicly available, and amendments to them are just as elusive.

Any research on zoning codes should recognize that codes often change, and that these changes are not always readily available without sleuthing—which may require scouring websites, piecing together meeting minutes, or contacting the jurisdiction. When it comes to establishing AI workflows, identifying and including amendments must be accounted for in the data preparation period, along with the formatting needs described next.

---

[218] Sara C. Bronin, *Comprehensive Rezonings*, 2019 BYU L. REV. 725 (2020) (only around two dozen cities over 100,000 people comprehensively revised their zoning codes in the preceding decade).

[219] Jason Sorens, *The New Hampshire Zoning Atlas Shows Why Reform is Still Needed*, Union Leader, Apr. 11, 2025 ("Between June 1, 2023 and June 1, 2024, 47 towns made 54 liberalizing changes for housing, while 27 jurisdictions made 29 restricting changes for housing. Looking just at the more significant changes, there were 36 significant changes to allow more housing and 20 significant changes to ban more housing.").

[220] Examples of code aggregators include Municode, eCode360, or American Legal Publishing.

As noted above, our data suggests that around half of jurisdictions change their codes, some in major ways but most in minor ways, in any given year.[221] Yet, just as it is difficult to track down zoning codes at single point in time, it is equally difficult to track down the amendments jurisdictions make to their codes, or to compare one version of a zoning code to another to assess what, specifically has changed. AI can be deployed to systematically monitor updates to zoning regulations posted online and extract information about such updates. For a use case like the NZA, whose public map is constantly outdated, such an AI tool can help notify team members of amendments that must be reviewed and incorporated into the public map.

Webscraping—namely automated processing of extracting information from online sources[222]—can put unstructured information found online into some sort of structured format. Among myriad past users of webscraping, a number of scholars[223] and private companies[224] have deployed it to track government decisions. Usually their techniques include fetching (or downloading) a webpage,

---

[221] *See supra* Part I.B.3.

[222] *See* Hafsah Hassan et al., *Automation in Web Data Extraction: Opportunities, Challenges, and Ethical Considerations*, 1 INT'L J. COMPUTING & DATA SCI. 37, 37 (2025) ("Web scraping . . . is the automated process of using specialized software to obtain data from websites.").

[223] *See*, *e.g.*, Kylie L. Anglin, *Gather-Narrow-Extract: A Framework for Studying Local Policy Variation Using Web-Scraping and Natural Language Processing*, 12 J. RSCH. EDUC. EFFECTIVENESS 685, 699, 701 tbl. 1 (collecting 3,743 documents from over 800 Texas school districts to review their District of Innovation plans and any other documents relevant to such plans); Stuart Rennie et al, *Scraping the Web for Public Health Gains: Ethical Considerations from a 'Big Data' Research Project on HIV and Incarceration*, 13 PUB. HEALTH ETHICS 111, 113 (2020) (collecting 200,000 inmates' personal information, including, name, age, date of birth, arrest charge, and date of incarceration daily from the public websites of 29 North Carolina county jails to develop an incarceration database and to ultimately improve HIV-positive inmates' access to HIV treatment); Eva Maxfield Brown & Nicholas Weber, *Councils in Action: Automating the Curation of Municipal Governance Data for Research*, 85TH ANNUAL MEETING ASS'N INFO. SCI. & TECH. (2022) (citing Big Local News, a platform created by Stanford researchers that conducts webscraping by "collect[ing] meeting assets such as documents, presentations, videos, captions, etc." of local governments to improve access to local government legislative materials); Moritz Schütz, Lukas Kriesch & Sebastian Losacker, *Mapping Local Government Priorities: A Web-Mining Approach for Regional Research*, 17 REG'L SCI. POL'Y & PRAC. 1, 3–5 (collecting 205 topics from the websites of 400 German counties, including 10,612 municipalities to determine the priorities of local German governments).

[224] *About Us: All Politics Is Local*, GATHERGOV, https://gathergov.com/about-us (last visited Jan. 11, 2025) ("GatherGov indexes every local government conversation in the US, building a best-in-class knowledge graph."); *NGO and Government Data Scraping Services*, X-BYTE ENTERPRISE CRAWLING, https://www.xbyte.io/ngo-and-government-data-scraping-services/#:~:text=Full%2DSpectrum%20Government%20and%20NGO,challenges%20and%20anti%2Dbot%20restrictions (last visited Jan. 11, 2025) ("Gain limitless insights into government data with custom built web scrapers that ethically extract online data from government portals and NGO platforms."); *Government Data Scraping Services*, 3I DATASCRAPING, https://www.3idatascraping.com/government-data-scraping/#:~:text=We%20extract%20legal%20and%20regulatory,over%20specific%20tasks%20or%20steps (last visited Jan. 11, 2025) ("Expert data scraping services to extract useful data insights from government and NGO institutions.").

extracting information from that webpage, and then reformatting that information into a database. While fetching techniques have been used for several decades, the latter two tasks have increasingly involved more sophisticated uses of AI.

Keeping track of zoning code amendments is especially difficult given the decentralization of zoning codes. They are published across thousands of websites in a variety of formats—including PDFs, Word documents, geospatial map layers, or html websites—with no consistency in how updates are communicated. Fortunately, the NZA team maintains a list of thousands of local-entity websites where analysts found zoning texts, maps, and geospatial files in the first place. An AI agent could be programmed to scrape these websites at regular intervals: fetching updated material, converting documents using optical character recognition software, performing direct content comparison with previous versions of the zoning code, and extracting structured information about new ordinances or zoning amendments. These extracted results can be formatted for integration into the NZA's data entry portal for further human analysis.

This AI strategy has an important limitation: it cannot be relied upon to provide information about code updates for all jurisdictions. As noted above, one-seventh of zoning codes cannot be found online.[225] Jurisdictions that do not post their codes online are very unlikely to post code updates online. Even jurisdictions that do make their codes available online may not necessarily post updates online at all, or in a timely manner. Even recognizing this limitation, this AI strategy could be highly useful for ensuring that the NZA data is kept accurate and up to date, and could be coupled with a manual updating workflow for the remainder of jurisdictions.

* * *

Together, these examples illustrate potential uses of AI that could plausibly characterize, predict, and extract information from zoning codes. The extracting applications have the additional potential practical impact of integrating information more quickly into the NZA, which in turn can ensure that NZA data remains comprehensive, current, and accurate. If successfully designed and executed, the applications also have the potential broader impact of refining our understanding of how AI applications can be used to support a study of a highly complex body of law.

CONCLUSION

One day, AI may be able to mimic human reasoning and decipher zoning codes more effectively. Or codes themselves, perhaps inspired by the NZA, will become more similar to each other, making the task of AI easier. But it is hard to see how AI can overcome the three distinct and interrelated reasoning challenges presented by zoning codes. First, the haphazard nature of codes' adoption and inconsistent drafting require knowledgeable humans to resolve ambiguities.

[225] *See supra* Part I.B.1.

Second, these difficulties in turn make it difficult for humans to train models to accurately interpret codes. And third, codes challenge algorithmic analysis itself. This Article cautions against expecting too much from AI at this early stage: simply put, AI has not proven that it can substitute for human reasoning in the important task of reading, interpreting, and extracting information from complex legal texts.

This Article has tested only AI applications to zoning, an area of law may well offer the most complicated legal texts in the American canon. Part I admittedly painted a dire picture about zoning codes' opaqueness and inconsistencies. But all laws can present interpretation challenges. Knowledgeable humans can adaptively reply to these challenges as they arise, in their various combinations. What this Article attempted to show was that even when knowledgeable humans have tried to train an AI model to replicate human reasoning, the model could not keep up. Whether that will be a permanent deficiency with AI or merely a temporary obstacle that can be overcome remains to be seen. At the same time, we should be cautious not to push AI too far in a field that rests on human interpretation, negotiation, and contestation and is designed, at a very foundational level, to engage people in its administration and interpretation.

We can only fully comprehend the scope of AI's challenges in interpreting zoning codes because we have such a rich, human-coded dataset on which to evaluate AI's success, in the form of the National Zoning Atlas. The project was created to inform the public about zoning, in part to catalyze zoning changes and foster more and better housing, stronger economies, healthier environments, and more connected communities. No one would be happier than the NZA team to develop a reliable AI tool that accelerates the delivery of actionable and accessible zoning data. However, we are also realistic about the opportunities presented by current technology.

This Article encourages other scholars to be more realistic, too. Even if we cannot expect the world from AI, we need to expect more of ourselves. Scholars working on AI and zoning should work to gain a much deeper understanding about how zoning actually works, and to identify how, exactly, AI might address the interpretation challenges enumerated in this Article. This will help to further the important effort of enhancing our understanding of the nation's most important, but least-understood, regulatory power.