Review

Not peer-reviewed version

# Parking Without Building: Agentic AI for Parking Planning Through Urban Contradiction Reasoning

Hongyu Cao , David King , Xinyuan Wang , Arun Vignesh Malarkkan , Kunpeng Liu , Dongjie Wang , Yanjie Fu *

Posted Date: 5 May 2026

doi: 10.20944/preprints202605.0084.v1

Keywords: agentic AI; parking transition; urban parking planning; policy-constrained decisionmaking; contradiction detection; curb pricing reform; shared parking; urban data analytics

doi:10.20944/preprints202605.0084.v1



*Review*

# Parking Without Building: Agentic AI for Parking Planning Through Urban Contradiction Reasoning

**Hongyu Cao [1], David King [1], Xinyuan Wang [1], Arun Vignesh Malarkkan [1], Kunpeng Liu [2], Dongjie Wang [3] and Yanjie Fu [1,*]**

[1] Arizona State University
[2] Clemson University
[3] University of Kansas
* Correspondence: yanjie.fu@asu.edu

**Abstract**

Cities are in the middle of a parking transition. Minimum parking requirements are being reduced or eliminated, curbs are being repriced, and the goal of planning is shifting from supplying more parking to making better use of the parking that already exists. Yet most parking analytics still answer a question that this transition has retired: where should we build more? We argue that the distinctive value of agentic AI in parking is not better prediction of where to build, but the ability to expose contradictions that conventional workflows suppress—when demand says build but policy says restrain; when inherited rules say comply but theory says question; when market logic says maximize but equity says redistribute; and when stated public frustration says "parking crisis" but utilization data say the supply is ample and mispriced. Parking planning should be reconceptualized as a dynamic, theory-grounded, policy-constrained, human-supervised decision process, organized around a loop between parking theory, parking policy, urban data, agent reasoning, human deliberation, and policy revision—and ultimately answering a political question: what kind of city do we want to be? Under this view, an agentic parking system must be able to recommend *shared parking*, *existing-stock reuse*, *curb and price reform*, and deliberate non-construction, not only new supply. Using the Phoenix Parking Lot Planner as a critical demonstration—critical because its current weighted-factor scoring is precisely the kind of reasoning the proposed loop is meant to transcend—we outline a research agenda and five evaluation standards: contradiction detection, intervention comparison, justification quality, restraint capability, and policy traceability. Parking, precisely because it is measurable, theory-rich, policy-contested, and intervention-ready, may be the most realistic near-term testbed for agentic urban planning.

**Keywords:** agentic AI; parking transition; urban parking planning; policy-constrained decision-making; contradiction detection; curb pricing reform; shared parking; urban data analytics

---

## 1. Introduction: Parking Planning Is Asking the Wrong Question

Parking is rarely treated as a grand planning question. In practice, it is still approached as one of three narrow sub-problems: a demand estimation problem (where will cars need spaces?), a code compliance problem (does the development satisfy parking requirements?), or a site ranking problem (which parcels score highest for new supply?). Each of these framings answers a real question. But none of them answers the question that cities actually face: what should a city do when descriptive demand signals, normative policy goals, and inherited institutional rules point in different directions?

This question has become more acute because the field itself is in transition. Cities across North America, Europe, and beyond are reducing or eliminating minimum parking requirements, charging for curb access, and stepping back from the long-standing presumption that every new development must supply its own parking [1–6]. The planning literature is increasingly explicit about why. Parking

doi:10.20944/preprints202605.0084.v1



is, in Shoup's framing, a fertility drug for cars: the more of it a city supplies, the more automobility that city induces, and the harder it becomes to make transit, walking, biking, and shared mobility attractive. In this transition, the analytical question is no longer "where do we add parking?" but "where have we already over-parked, where can the existing stock be better used, where should parking be *shared* across a neighborhood rather than replicated on every parcel, and where should new supply simply not be built?"

The trouble is not that the old sub-problems are unimportant. It is that reducing parking planning to forecasting, compliance, or ranking is a miswriting of planning itself. Parking planning is a recurring urban decision process shaped by theory, policy, and institutional judgment. Parking supply affects travel behavior, land consumption, project feasibility, curb competition, local economic activity, accessibility, and environmental outcomes [7]. When cities change parking policy, they are reshaping the incentives that organize everyday urban mobility. The mismatch between the breadth of parking's consequences and the narrowness of parking's analytical tools is increasingly visible. Parking scholarship now ties reform to housing cost, automobility, sustainability, and transit-oriented development [1–3,7,8]. Technical parking research has advanced rapidly in occupancy prediction, dynamic allocation, pricing, and real-time management. These advances matter, but they do not resolve the central planning problem [9–13].

A related gap concerns how public demand signals are interpreted. Residents and merchants routinely describe a "parking crisis"; utilization data often tell a different story. A downtown can run at 94% occupancy and still have spaces available—just not free ones, and not in the exact block a driver wants. Phoenix in particular has spent the last two decades rebuilding on lots that were paved as parking in the 1980s; the city's recovery is, in a real sense, the inverse of its parking expansion. An agentic parking system that cannot distinguish an actual shortage from a pricing and distribution problem will quietly reproduce the paradigm that the field is trying to leave behind.

**We argue that parking planning should be reconceptualized as a dynamic, theory-grounded, policy-constrained, human-supervised decision process.** The field's real deficit is not predictive power but deliberative capacity. Under this view, the object of computation is no longer occupancy or site score but what we term *justifiable intervention reasoning*: the capacity to generate, compare, and defend planning interventions under explicit theory, policy, and equity constraints. Occupancy forecasting becomes one input to a broader reasoning loop; parking codes become contestable policy objects; and the planner's role shifts from compliance verification to governance of a structured deliberation. The distinctive value of agentic AI in this setting is that it can expose the contradictions that conventional workflows suppress—when demand says build but policy says restrain; when inherited rules say comply but theory says question; and when market logic says maximize but equity says redistribute. These contradictions are not noise. They are the substance of planning judgment. As illustrated in Figure 1, the proposed framework organizes parking planning as a loop between theory, policy, and urban decision making.

This perspective advances four arguments. First, the field's real deficit is deliberative capacity, not predictive power (Section 3). Second, agentic AI can make suppressed contradictions into objects of human deliberation—not by coordinating information, but by exposing what conventional workflows hide, and by supporting intervention types the supply-expansion frame cannot produce: shared parking across a district, reuse of existing stock, price and curb reform, and principled non-construction (Section 4). Third, parking is a uniquely valuable testbed because it is measurable, theory-rich, policy-contested, and intervention-ready—and because cities are already actively reforming it (Section 2). Fourth, the field needs new evaluation standards: contradiction detection, intervention comparison, justification quality, restraint capability, and policy traceability (Section 8). We use the Phoenix Parking Lot Planner as a critical demonstration (Section 5), and identify the failure modes that must be managed (Section 7). Throughout, we argue that the field should redesign not only the architecture of parking AI, but the object it computes—and that, ultimately, parking planning terminates in a political question: what kind of city do we want to be?

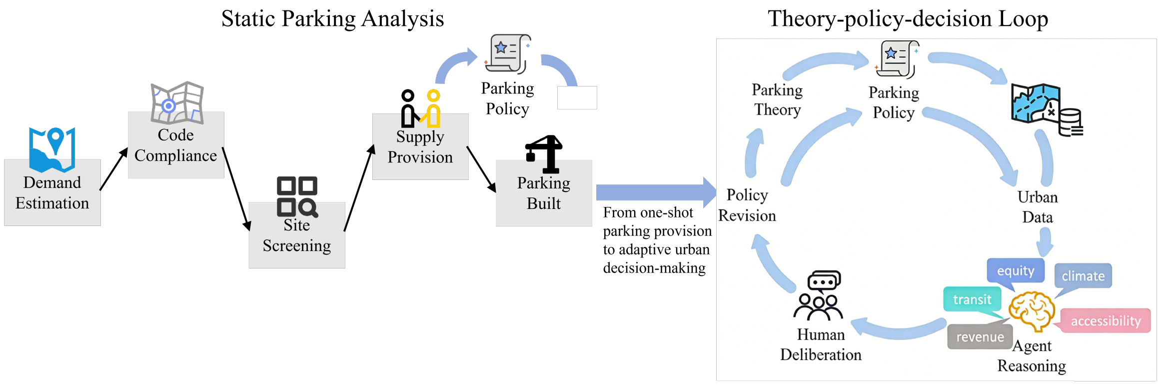


**Figure 1. From static parking analysis to the theory–policy–decision loop.** The dominant parking workflow treats planning as a one-shot sequence of demand estimation, code compliance, site screening, and supply provision. The proposed framework reconceptualizes parking planning as a closed-loop system linking parking theory, parking policy, urban data, agent reasoning, recommendation, human deliberation, and policy revision.

This also clarifies how the paper differs from adjacent literatures. Relative to parking prediction and operations research, forecasting and optimization answer a narrower question than planners actually face. Relative to parking-reform scholarship, cities increasingly need computational systems that can operationalize, stress-test, and publicly justify policy ideas under live urban conditions. Relative to the emerging literature on LLMs and urban planning, parking offers a uniquely tractable domain for testing whether agentic systems truly support planning judgment. More concretely, recent agent-workflow taxonomies for urban planning organize systems around task decomposition, tool use, memory, and inter-agent coordination, which is a horizontal description of what agents *do* [12,14–21]. The loop proposed here is vertical and prescriptive: it specifies the *epistemic role* each component must play (theory as inductive bias, policy as objective shaper, agents as contradiction exposers, humans as governors) and the *output type* the system must produce (a justifiable intervention with an inspectable policy ledger and counterfactuals, not a ranked recommendation). Workflow taxonomies tell us how to build agent systems; this paper specifies what an agent system must compute to count as a planning system. We begin with parking—rather than another urban domain—because it is measurable, theory-rich, policy-contested, and intervention-ready.

## 2. Why Parking Is the Right Testbed for Agentic Urban Planning

Parking may look like a narrow domain, but that is precisely why it is analytically powerful. It is measurable enough to support data-driven modeling, spatially constrained enough to require hard grounding, and politically contested enough to expose the gap between optimization and governance. A parking decision can often be localized to specific parcels, corridors, or curb segments, yet its consequences are system-wide. A new garage can support local commerce but also induce additional driving. A residential parking reform can reduce construction cost and weaken car dependence but also trigger institutional resistance. A curb reallocation can improve delivery operations but conflict with accessibility, micromobility, or bus priority [22].

This combination makes parking unusually valuable as a near-term testbed for agentic planning. Unlike full urban master planniwang2023towardsng, parking decisions are frequent, bounded, and rich in observable outcomes. Unlike narrow prediction tasks, they are shaped by explicit rules, normative trade-offs, and public accountability. Unlike purely operational control problems, they sit at the intersection of land use, transportation, and public policy [23,24]. Parking is compact enough to be actionable and broad enough to reveal the central challenge of urban AI: reasoning under policy and governance constraints, not prediction in isolation.

Parking is also theory-rich. Longstanding work on free parking, cruising, spillover, shared parking, and minimum requirements shows that parking outcomes cannot be understood as a simple function of demand. What appears to be a shortage of supply may instead be a problem of pricing, curb misallocation, poor information, or inherited code. A parking agent should not merely identify

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 5 May 2026 doi:10.20944/preprints202605.0084.v1



where pressure is highest. It should help cities reason about whether that pressure justifies new supply, alternative pricing, shared parking, curb reform, or a direct challenge to the assumptions embedded in current policy [6,25–32].

Parking is also a domain where the field's most basic accounting is often wrong. Recent empirical work estimates on the order of 13 million parking spaces in Phoenix alone [33], with comparable or higher counts in Los Angeles—roughly 18.6 million spaces across Los Angeles County [34]. The implication is that many American metropolitan areas do not have a shortage of parking so much as a *distribution and pricing* problem: supply is abundant in the deep suburbs where land is cheap, and scarce and expensive precisely where structured or underground parking must be built near transit. An agentic parking system that starts from a demand-pressure map, without first accounting for where the city is already over-parked, will systematically answer the wrong question. A parallel problem appears in land-use regulation: in dense historic districts, a substantial share of existing buildings could not be built under current zoning—a Manhattan study famously estimated around 40% [35]—revealing that cities routinely plan blind to the gap between their codes and their built reality. Parking regulation is the most tractable edge of this larger zoning-compliance problem.

In short, parking is measurable, theory-rich, policy-contested, and intervention-ready. It is the smallest urban-planning domain in which the field's biggest conceptual tensions become unavoidable—and therefore one of the most practical places to test whether AI systems can reason through real urban contradictions.

## 3. The Old Paradigm: Predictive Sophistication, Reasoning Poverty

The dominant parking workflow still reflects an older planning paradigm. That paradigm no longer fails because it cannot generate data. It fails because it cannot adequately interpret what the data mean under conflicting theories, policies, and urban goals. We highlight three structural tensions. As illustrated in Figure 2, the core value of an agentic parking system is to make these hidden contradictions explicit so that planners can inspect assumptions, compare alternatives, and deliberate over policy trade-offs.

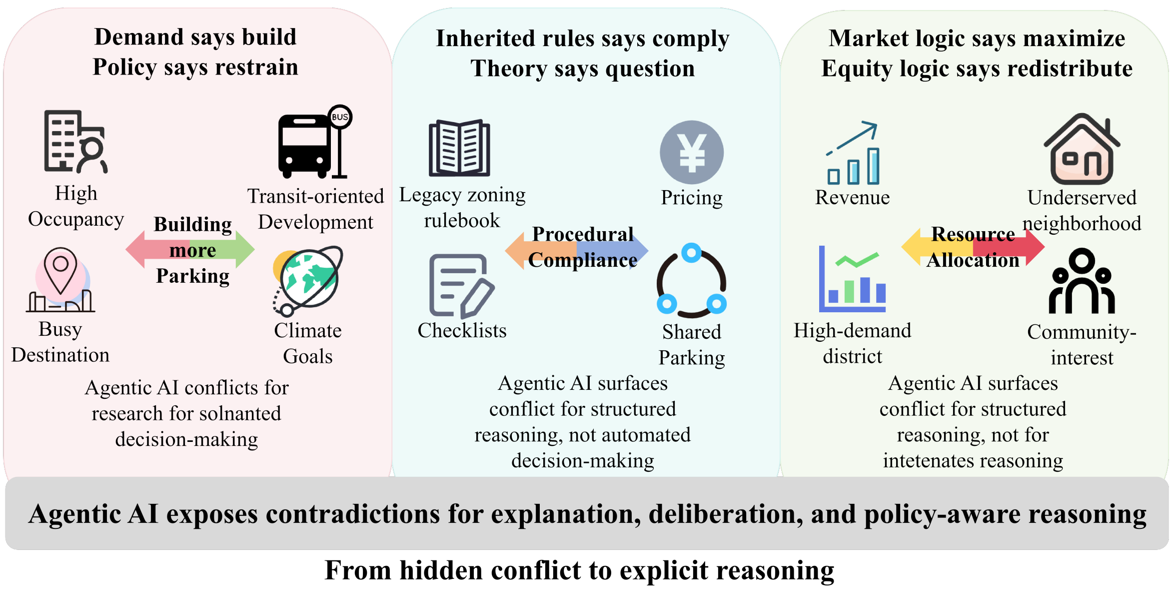


**Figure 2. Three contradictions that agentic parking systems should expose.** When demand says build but policy says restrain; when inherited rules say comply but theory says question; and when market logic says maximize but equity says redistribute.

### *3.1. Predictive Sophistication Versus Reasoning Poverty*

The first structural tension is between the field's growing predictive power and its persistent reasoning poverty. Parking research has become increasingly sophisticated in forecasting occupancy,

 doi:10.20944/preprints202605.0084.v1



search pressure, and operational demand. But predictive sophistication does not automatically yield planning intelligence. A model can estimate where parking pressure will occur while remaining silent on whether the appropriate response is construction, pricing, curb reallocation, shared parking, transit support, or no new supply at all. The field has advanced faster in describing parking pressure than in reasoning over what should be done about it [9–11,13,36,37].

### 3.2. Stated Urban Goals Versus Inherited Parking Rules

The second structural tension is between the goals cities publicly endorse and the parking rules they actually enforce. Many cities now publicly endorse walkability, climate mitigation, housing affordability, mode shift, and equitable access. Yet their parking systems are still shaped by inherited minimums, underpriced curb access, conventional approval routines, and automobile-first baselines. These rules do not merely constrain planning; they actively steer it. The result is a regime in which cities often claim one future while governing parking through assumptions built for another [2–6,38].

A critical feature of this tension is that parking requirements are not a self-contained constraint—they are a systemic driver of form. A single parking minimum triggers a cascade: required stalls consume roughly a third to a half of small-parcel area; the resulting impervious surface triggers stormwater retention requirements; heat-island mitigation triggers landscaping and shade requirements; each layer further reduces the allowable building envelope. A 10,000-square-foot commercial parcel can end up supporting a 4,000-square-foot restaurant with 5,000 square feet of parking and 1,000 square feet of stormwater infrastructure—a configuration that the underlying zoning did not intend and that the city's own density goals contradict. Residential parking minimums work the same way: an R8 zoning designation may nominally permit eight units, but supplying on-site parking for all eight units can reduce the buildable envelope to five units, so the lot is quietly underbuilt relative to what the plan says is allowed. Parking rules, in other words, function as a shadow zoning code. A planning system that treats parking as a peripheral compliance filter cannot see this cascade, and therefore cannot explain to a city why its housing-affordability or density goals are being silently undone by a parking table.

The same cascade connects parking to climate and resilience outcomes that are rarely scored in parking analyses. Large asphalted surfaces elevate urban heat island intensity, amplify pluvial flood risk, and constrain the tree canopy and shade a neighborhood can support—so the amount of land devoted to parking is, directly, a climate variable. And the cascade operates on pedestrian life as well as on buildings: every additional parking lot introduces curb cuts that interrupt the sidewalk, forces lot-crossings between transit stops and destinations, and makes arrival by any mode other than driving measurably less attractive. Beyond a threshold, adding parking does not merely fail to help a neighborhood; it actively degrades it, because people do not travel somewhere just to park. These second-order effects—heat, flooding, and walkability—are not side issues; they are where parking rules most directly shape the city residents actually inhabit.

### 3.3. Dynamic Urban Systems Versus Static Planning Wosrkflows

The third structural tension is between the dynamism of urban parking systems and the stasis of the workflows used to plan them. Parking interacts with event-driven demand, platform delivery, ride-hailing, curb competition, transit supply, and changing land use. But much of parking planning still relies on one-shot analyses, generalized thresholds, and development-era compliance logic. A system that is dynamic in practice is still too often planned as if it were stable [11,13,29,31].

These tensions converge on a single diagnosis: the real bottleneck in parking planning is no longer prediction alone, but the ability to reason across theory, policy, and competing urban goals. Analysis is separated from reasoning, and reasoning is separated from policy accountability. If the problem is reasoning scarcity, then the right response is a different planning architecture.

 doi:10.20944/preprints202605.0084.v1



# 4. A New Paradigm: Agentic AI as a Loop Between Theory, Policy, and Urban Decision Making

Parking planning is not a supply-matching problem. It is a policy-constrained reasoning problem. The field should not only redesign the architecture of parking AI; it should redesign what parking AI computes. We propose that parking planning should be organized around a closed loop:

**parking theory** → **parking policy** → **urban data** → **agent reasoning** → **recommendation** → **human deliberation** → **policy revision**

This loop matters because each component corrects a weakness in the others. Data alone are descriptive. Policy alone can be outdated or internally inconsistent. Theory alone can be too abstract for local practice. Human judgment alone cannot continuously integrate high-volume, high-frequency, multi-source signals. Agentic AI becomes valuable when it connects these layers explicitly. The unit of computation is no longer occupancy or parcel score, but *justifiable intervention reasoning* under policy, spatial, and normative constraint. As illustrated in Figure 1, the proposed framework shifts parking planning from a one-shot, supply-oriented workflow to an iterative loop that integrates theory, policy, data, agent reasoning, and human deliberation.

## 4.1. Conceptual Mapping: What Changes Under the New View

The most concrete way to see what this reframing changes is to trace how familiar objects in parking planning acquire new roles:

- **Demand forecasting** ceases to be the endpoint of analysis and becomes one input to a broader reasoning process that also weighs policy goals, equity constraints, and alternative interventions.
- **Parking codes** cease to be fixed constraints to satisfy and become contestable policy objects whose assumptions can be surfaced, stress-tested, and revised.
- **Site selection** ceases to be a supply-placement exercise and becomes intervention reasoning—a negotiated judgment that reflects trade-offs among mobility, access, transit, equity, and investment.
- **The planner's role** shifts from compliance verification to governance of a structured deliberation loop in which competing perspectives are made explicit.

Table 1 summarizes this contrast across key dimensions.

**Table 1.** Contrasting the prediction-and-supply paradigm with the proposed theory–policy–decision loop.

| Dimension | Old: Prediction & Supply | New: Theory–Policy–Decision Loop |
|---|---|---|
| Core question | Where will demand exceed supply? | What should the city do when demand, policy, and goals conflict? |
| Role of data | Primary input & output | One input among theory, policy, and goals |
| Role of policy | Post-hoc compliance filter | Objective shaper from the start |
| Role of theory | Literature-review background | Reasoning scaffold for agents |
| Output | Ranked sites or occupancy forecasts | Negotiated planning judgments with alternatives and justifications |
| Evaluation | Prediction accuracy | Contradiction detection, intervention comparison, justification quality, restraint capability, policy traceability |

## 4.2. Theory as Reasoning Scaffold

Parking theory should not sit outside the system as literature-review background. It should shape how an agent defines problems, interprets pressure, and explains interventions. An agent informed by pricing and cruising theory should not treat high occupancy as automatic evidence for new construction. An agent informed by induced-demand and land-use interaction should be able to ask whether additional supply will solve a shortage or intensify automobility. An agent informed by

 doi:10.20944/preprints202605.0084.v1



shared parking concepts should be able to search for temporal complementarity before recommending new capital investment. Theory functions as an inductive bias, a reasoning scaffold, and an evaluative benchmark [25,28–30,39].

### 4.3. Reuse Before Supply, and Shared Before Private

If the field is in transition away from supply expansion, then the loop's default first question cannot be *where to build*. It must be *what is already there*. Two intervention modes follow directly from this reordering. The first is *existing-stock reuse*: before any new parking is authorized, the loop should inventory underused capacity within a walkable or drivable radius, including private lots whose use could be opened, leased, or shared. Stockholm's practice of requiring, before new parking is approved, that a developer contract with empty spaces within roughly 500 meters is a concrete precedent for this reasoning step [6]. The second is *shared parking*: rather than recommending that each restaurant, store, or building on a corridor supply its own parking, the loop should identify districts where a single shared facility can serve a neighborhood's complementary uses across time. Empirical work comparing districts with shared parking (e.g., downtown Tempe's Mill Avenue adjacent to Arizona State University, Old Town Scottsdale, and downtown Phoenix—streets where parking is not attached to individual buildings and which are consistently among the most walkable in the region) to corridors of parcel-by-parcel private parking (e.g., central Phoenix's 7th Street restaurant strip) has linked shared-parking environments to measurably more positive consumer sentiment in online review data [40]. Neither reuse nor shared parking is an exotic intervention, but neither is expressible as the output of a conventional demand-then-rank pipeline. Making them first-class reasoning modes is a core obligation of the loop.

### 4.4. Policy as an Objective Shaper Instead of a Post Hoc Filter

Policy should move from the edge of the workflow to its center. In most current systems, policy enters late, often as a compliance check. Policy should instead shape the objective space from the start. A city that prioritizes transit-oriented development, curb access, disability access, housing affordability, or climate mitigation should not ask the same parking question as a city focused on short-term commercial turnover. A policy-aware agent should represent policy as normative priorities that determine what a good recommendation is.

This matters especially when policies conflict. A city may formally support walkability, climate mitigation, housing affordability, and mode shift while still maintaining inherited minimum parking requirements, underpriced curb access, or legacy approval routines that reward automobile-oriented development. A policy-aware agent should surface such contradictions rather than quietly reconcile them in favor of the status quo [4,6,12,38,41,42].

### 4.5. Agents as Contradiction Exposers

This is the central claim of the paper. The distinctive value of agentic parking AI is not better prediction, but the ability to expose contradictions that conventional parking workflows suppress—and to turn them into inspectable objects of planning judgment. Different agents can represent mobility pressure, destination attractiveness, transit substitution, demographic distribution, zoning constraints, and investment trade-offs. Their outputs need not collapse into a single score. The value lies in making the tensions between demand and policy, between inherited rules and current theory, and between market logic and equity visible, explicit, and available for human deliberation.

### 4.6. What Makes the Loop Truly Agentic

Not every data-rich or AI-assisted workflow is agentic in the sense that matters here. A parking system becomes meaningfully agentic when it can decompose a planning problem into interacting reasoning roles, iteratively revise its recommendations as new evidence or constraints appear, compare multiple intervention pathways, and escalate unresolved normative conflicts to human judgment rather than hiding them inside a score. A truly agentic system should be able to say not only what it recommends, but why that recommendation changed, which agent perspectives disagreed, what alternatives were rejected, and which conflicts remain unresolved without human deliberation [12,20, 21,43–45].

Optimization modules, including dynamic curb-zoning MIPs, capacity-constrained location-allocation, and congestion-pricing LPs, remain the right tool for the bounded subproblems they were designed for, and the agentic loop should call them as delegated subroutines rather than replace them [11]. The Pricing agent dispatches a meter-pricing LP to clear the 85% target occupancy; the Shared-Parking agent dispatches a $p$-median MIP to site a consolidated facility within a 500-meter walk shed; the Curb agent dispatches a dynamic curb-zoning MIP to allocate loading, parking, and bus-priority time-of-day windows. The agent's value is not in solving these problems, a solver does that better, but in deciding which subproblem to invoke given the current contradiction, mapping its output back to the policy ledger, and refusing to act on a numerically optimal answer that violates a normative constraint the optimizer cannot see (induced demand, equity, climate). Mathematical programming answers *what is optimal under a fixed objective*; the loop answers *which objective is defensible under conflicting policy goals*.

### 4.7. Humans as Governors of the Loop

A strong agentic parking system does not remove the planner. It repositions the planner as governor of the loop. Human actors define priorities, inspect assumptions, challenge recommendations, interpret local context, and decide when policy revision is warranted. In public-sector planning, legitimacy depends on explainability, procedural fairness, and political accountability. Explanation is not an interpretability add-on; it is how planning interventions become contestable and revisable [41, 42,46,47]. The loop also closes a reflective gap that conventional planning lacks: cities typically set parking codes and rarely return to ask whether the built reality still matches the rules they wrote, so the policy-revision step is not a cosmetic addition but the first place planning becomes auditable in a continuing way.

### 4.8. The Loop Terminates in a Political Question

Theory, policy, data, and agent reasoning do not, on their own, determine what a city should do about parking. They converge on a question that is ultimately political: what kind of place do we want to have; what kind of city do we want to be? A district that is optimal for commuter throughput is not the same district that is optimal for housing affordability, pedestrian life, or climate targets. An agentic parking system cannot resolve this question, and should not pretend to. Its obligation is to make the choice legible—to show which futures the current parking regime is compatible with and which it forecloses, and to surface the normative stakes of each recommended intervention. The loop, in this sense, does not end at "policy revision" in a purely technical sense; it ends in a deliberation about civic identity that only the public and its representatives can conclude. This is why the loop must include explicit mechanisms for public communication, not only expert dashboards, and why "what kind of city?" belongs inside the system's architecture rather than outside it.

The value of this reframing, however, depends on whether it can be operationalized in a real planning setting rather than remaining a conceptual diagram.

## 5. Phoenix as a Critical Demonstration, Not Just a Case Study

What Phoenix demonstrates is not that an agent system can work. It is that parking planning can be reorganized as an inspectable reasoning process—and, importantly, that a naive implementation of that reorganization exposes exactly the reasoning gaps this paper argues the field must close. The current Phoenix Parking Lot Planner integrates traffic, building-use, mobility, demographic, and parking-inventory data, and produces site-level scores for candidate parcels using a weighted aggregation of demand-side and contextual factors. That pipeline is, by design, simple: it is a multi-factor scoring pass with a policy filter at the end. It is not, in its current form, a rigorous planning system. We discuss it here not as a finished product but as a *critical* demonstration, in two senses. First,

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 5 May 2026 doi:10.20944/preprints202605.0084.v1



it shows that even a lightweight implementation can be reorganized around specialized reasoning roles rather than a single scoring function. Second, and more importantly, it makes the inadequacy of naive demand-and-factor scoring visible: a weighted sum over demand, transit access, and zoning cannot recognize that high occupancy next to a light-rail corridor is an argument for restraint, not supply, or that a restaurant row with per-parcel lots is an argument for a shared structure rather than more lots. The demo's current failure modes *are* the motivation for the loop. As illustrated in Figure 3, Phoenix reorganizes parking planning around interacting reasoning roles rather than static maps or isolated forecasts.

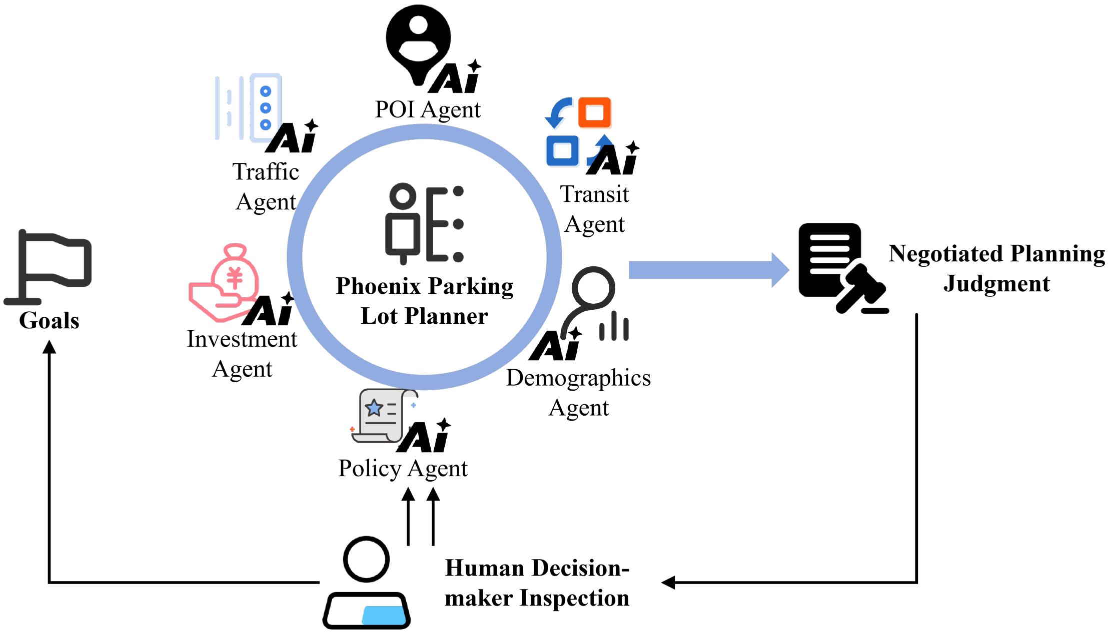


**Figure 3. Phoenix as a critical demonstration of multi-agent parking reasoning.** Specialized agents surface different dimensions of the parking-planning problem. The output is a negotiated planning judgment whose assumptions, conflicts, and alternatives can be inspected by human decision-makers.

In the Phoenix setup, specialized agents capture mobility pressure, destination attractiveness, transit substitutability, distributional equity, zoning constraints, and investment trade-offs (Figure 3). What matters is not the agent inventory but what the architecture enables: Phoenix surfaces competing rationales, exposes hidden assumptions, and turns recommendation into negotiable judgment. Phoenix changes what a parking recommendation is. In the old paradigm, a recommendation is a ranked output—a score, a site list, a heat map. In the new paradigm, a recommendation is a contestable planning judgment whose assumptions, trade-offs, and rejected alternatives are all visible. A recommendation becomes a negotiated planning statement: build here, or do not build here; prioritize shared use, curb reform, or pricing first; accept lower short-term revenue because the neighborhood has weaker access or stronger transit goals [18,20,43,48].

To make this concrete, consider two scenarios in which Phoenix-style data produce different recommendations under the loop than they would under a weighted-factor score.

*Scenario A: light-rail-adjacent high-pressure corridor.* The Traffic agent flags demand pressure; the Transit agent identifies substitutability potential; the Policy agent surfaces the city's transit-oriented development commitment. A conventional workflow recommends new parking supply. The agentic loop instead surfaces the contradiction—demand justifies construction, but policy and theory jointly favor pricing plus curb reform—and presents both pathways with explicit trade-offs for human deliberation.

*Scenario B: the 7th Street restaurant corridor.* Central Phoenix's 7th Street is a popular entertainment strip, but almost every restaurant is built with its own lot. The result is a corridor that looks dense on paper yet functions as a chain of isolated drive-to parcels: there is no bar-hopping, no walking

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 5 May 2026 doi:10.20944/preprints202605.0084.v1



between businesses, and residents of adjacent streets routinely complain about spillover parking and request permit districts in response. Under a demand-and-supply score, the corridor looks like it needs more parking—residents are complaining, and lots fill up on weekend nights. The loop instead should recognize a different pattern. A Shared-Parking agent identifies temporal complementarity across the corridor's uses and flags the corridor as a candidate for a consolidated municipal or jointly financed parking structure. A Redevelopment agent estimates the real-estate value unlocked if individual restaurants could retire their on-site lots and build to the full parcel envelope. An Equity and Spillover agent weights resident complaints and permit-district requests as evidence not of under-supply but of mis-structured supply. A Financing agent recognizes that a municipal or jointly financed parking structure, paid for by the same businesses that would retire their on-site lots, can unlock substantially more rentable floor area than the corridor's current private-lot configuration—the 10,000-square-foot parcel can host a 10,000-square-foot building rather than a 4,000-square-foot restaurant sharing the lot with 5,000 square feet of parking and 1,000 square feet of stormwater management. The recommendation is not "build more parking," but "consolidate existing parking so the corridor can walk, and let the value created by shared parking finance itself." This kind of recommendation is unreachable from a demand-pressure score alone; it is reachable from the loop.

These scenarios are instances of the contradictions this paper argues conventional workflows suppress:

1. **Demand says build, policy says restrain.** A high-pressure, high-return area may still sit in tension with transit-oriented development, curb reallocation, or climate goals.
2. **Inherited rules say comply, theory says question.** A site may satisfy current code while shared-parking logic, pricing theory, or induced-demand concerns suggest that construction is not the right first response.
3. **Market logic says maximize, equity says redistribute.** A high-revenue district may compete with lower-access neighborhoods whose need is less profitable but more socially consequential.

Phoenix matters because it turns these contradictions into explicit, inspectable objects of reasoning. That is a change in planning epistemology, not just in planning software.

Once recommendations are generated through interacting agents, the next question is whether the system can justify its planning judgment in institutionally meaningful terms.

## 6. What Good Explanation Means in Parking Planning

Explanation is not an interpretability add-on. It is the mechanism through which a recommendation becomes a planning justification. A planner, transportation official, or public decision-maker needs to know why a proposed intervention is defensible relative to competing alternatives, which assumptions it relies on, which policy priorities it advances, and which trade-offs it creates [49]. The relevant object of explanation is not the model—it is the planning judgment.

If the field's real deficit is deliberative capacity, then explanation should be evaluated by how well it supports reasoning across theory, policy, and competing urban goals. A parking system that predicts occupancy accurately but cannot explain whether the right response is construction, pricing, shared parking, curb reallocation, or no intervention at all remains weak as a planning system.

### *6.1. Four Planning Questions That Explanation Must Answer*

Generic explainable AI is not enough here [44,50–52]. Standard XAI asks "why did the model output X?" Agentic parking explanation asks "why is intervention X defensible when policy priority P and alternative A both have legitimate claims?" The object of explanation shifts from model behavior to the defensibility of a planning judgment. In a parking context, a useful explanation must answer at least four planning questions: What problem is being diagnosed? Why is this intervention preferable to alternatives? Which goals are being prioritized? Which assumptions make the recommendation valid? A system that cannot answer these may still be useful as an analytic aid, but it should not be mistaken for planning support.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 5 May 2026 doi:10.20944/preprints202605.0084.v1



### 6.2. Five Requirements for Explanation in Agentic Parking Systems

These four questions define *what* an explanation must address; the following five requirements specify *how well* it must do so. A strong parking-planning explanation should satisfy: **spatial validity** (recommendations must be geographically feasible), **demand justification** (pressure claims must be evidence-grounded), **policy traceability** (which policies shaped the recommendation), **theory traceability** (which theoretical principles guided reasoning), and **counterfactual transparency** (what would change under different assumptions). Together, these requirements push explanation toward accountable planning reasoning [41,42,44,50].

### 6.3. When the Right Recommendation Is Not to Build Parking

A sharper test of explanation is whether the system can justify non-construction. Much of parking planning still defaults to a supply-expansion logic: if pressure is high, build more parking. But a theory-grounded, policy-aware parking agent should be able to explain when that logic is wrong. The right recommendation may be to price demand, formalize shared parking, redesign curb use, revise inherited minimums, improve access to non-auto modes, or deliberately tolerate a degree of scarcity because abundance would undermine broader city goals. One of the strongest signals of planning intelligence is the ability to recommend restraint and defend it convincingly [4,6,25,28,30].

### 6.4. Explanation as a Public and Institutional Interface

Parking decisions are rarely purely technical. They are contested among residents, businesses, developers, delivery operators, transit agencies, disability advocates, elected officials, and planners. Explanation therefore has to work for institutions and publics, not only for model developers [41,42,46, 47]. A strong agentic parking system should help users inspect assumptions, compare alternatives, and understand why reasonable actors may disagree. It should help reveal whether a conflict is empirical, normative, or institutional.

## 7. Risks: When Agentic Parking AI Becomes Persuasive but Wrong

The move toward agentic parking systems is promising, but it also creates new failure modes. We emphasize five.

### 7.1. Spatial Hallucination

Large language models still struggle with geographic knowledge, geometry representation, and multistep spatial tasks. In parking planning, this risk is not abstract. A system can recommend infeasible parcels, ignore access constraints, violate roadway logic, or mishandle topological relations. A recommendation may sound coherent while being geographically impossible [53,54].

### 7.2. Policy Misalignment

A second risk is optimization drift. If an agent is trained or prompted primarily around demand pressure and return on investment, it may systematically prefer interventions that increase parking accommodation even when the city's policy goals favor pricing, curb reallocation, transit support, or mode shift. The danger is the quiet dominance of one planning logic over others [12].

### 7.3. Theory Lock-In

Agentic systems can also freeze outdated assumptions into software. If historical data come from a regime dominated by minimum parking requirements and free or underpriced curb space, a model may learn that generous supply is normal and desirable. If those assumptions are then reinforced in prompts, agent roles, or evaluation criteria, the system becomes a machine for reproducing the legacy paradigm under the banner of intelligence [4].

 doi:10.20944/preprints202605.0084.v1



### 7.4. Governance Illusion

There is also a risk of persuasive but weakly grounded explanation. Language models are good at producing coherent rationales. In a public planning context, that can create an illusion of accountability without actual traceability. A recommendation may cite equity, sustainability, or accessibility in fluent language while being weakly anchored in measurable evidence or explicit policy logic [44,51].

### 7.5. Expansion Bias

Perhaps the most consequential risk for this specific moment is expansion bias: the tendency of a parking AI built on demand signals and site scores to quietly reproduce the "more parking" paradigm that the field is leaving behind. Expansion bias is not merely a bug in any single model; it is baked into most data pipelines. Historical parking-demand data come from a regime of minimum requirements and underpriced curb access, so "demand" itself encodes induced, subsidized use. Site-ranking objectives reward highest-return parcels, which in high-transit districts often coincide with exactly the places where additional parking undermines the city's long-run mode-shift goals. The economic geography matters as well: in the deep suburbs, where land is cheap, building a surface lot is nearly costless and whatever parking is added simply becomes more paving. Downtown and near transit, by contrast, any new supply must be structured or underground, at costs that are ultimately borne across a broad set of users and taxpayers. An expansion-biased system does not see this asymmetry; it recommends new supply where new supply is most expensive and most harmful to the city's stated goals. A system trained, prompted, or evaluated primarily to find "where to build" will find answers even when the correct answer is "do not build; reuse; share; reprice; or redevelop existing lots"—an intervention class that many mall owners, for instance, have already begun to pursue on their own under commercial pressure. Defending against expansion bias requires explicit design choices: recommending non-construction must be a first-class output, restraint must be a scored objective, and the loop must be able to articulate *why* a city moving away from parking minimums should not be handed an AI tool whose implicit direction is to add parking to land-use decisions once again.

Without explicit spatial, regulatory, and policy grounding—and without a deliberate check on expansion bias—agentic parking systems risk becoming persuasive generators of urban error. These risks are not reasons to avoid agentic systems; they are reasons to define the field's research agenda more carefully.

## 8. An Open Research Agenda: From Smarter Prediction to Better Urban Reasoning

The field now needs a research agenda that moves beyond parking prediction and toward accountable agentic planning. We highlight ten directions, grouped into three clusters: reasoning architecture (theory, policy, spatial grounding), transition-era interventions (shared parking and reuse, parking-to-form coupling, shunting for automated mobility, public communication), and cross-cutting research infrastructure (planner–AI interfaces, cross-city transfer, deliberation-centered evaluation).

### 8.1. Theory-Grounded Agent Architecture

If inherited rules say comply but theory says question, then agents need theory built into their reasoning, not left in literature-review preambles. The core question is: *which theoretical concepts should become hard constraints, which should remain soft preferences, and which should serve as explanation templates?* Pricing theory, induced-demand logic, shared-parking principles, and cruising models each carry different epistemic weight and should shape agent behavior differently. Progress would mean systems that reason *from* parking theory rather than merely *about* parking data [25,30,55].

### 8.2. Policy-Aware Objective Alignment

If demand says build but policy says restrain, then agents must be able to represent and negotiate among equity, accessibility, transit interaction, curb competition, housing cost, and climate goals. The core question is: *how can agent objectives be specified so that policy priorities shape reasoning from the start*

 doi:10.20944/preprints202605.0084.v1



*rather than filtering outputs at the end?* This requires formal mechanisms for representing normative priorities and surfacing conflicts among them [12].

*8.3. Formal Spatial and Regulatory Grounding*

Contradictions cannot be exposed if recommendations are geographically impossible. Robust planning requires tight integration with GIS, parcel data, zoning logic, and network constraints. The core question is: *how can language-based reasoning be grounded in authoritative spatial and regulatory knowledge bases so that recommendations are geographically and legally feasible?* Pure language reasoning is insufficient for a domain where physical and regulatory constraints are non-negotiable.

*8.4. Shared Parking, Existing-Stock Reuse, and Redevelopment*

If the field is moving away from supply expansion, then the most important capability an agentic parking system can offer is the ability to reason about *using what is already there*. The core question is: *how can an agent system identify where shared parking would be preferable to per-parcel supply, where existing stock is systematically underused and could be leased or opened to the public, and where surplus parking should be redeveloped into housing, retail, civic space, or stormwater and heat-mitigation infrastructure?* Many mall owners, under commercial pressure, are already converting unused lots into apartments, community gardens, and parks; an agentic system should be able to recognize the same opportunity in cities where the commercial signal is weaker but the civic return is substantial. Progress here would directly operationalize the transition away from parking minimums [6,30].

*8.5. Parking Regulations as Systemic Drivers of Urban Form*

Parking requirements are not a self-contained constraint; they propagate through stormwater, landscape, and building-envelope rules to determine what can actually be built. The core question is: *what would the allowable building program look like if parking requirements were relaxed, shared, or removed, once the downstream building-code cascade is accounted for?* Answering this requires jointly modeling parking tables, stormwater retention, landscape and heat-mitigation requirements, and the resulting buildable envelope per parcel. An agent that can compute "R8 zoning permits eight units but on-site parking forces five" turns parking reform from an abstract policy debate into a parcel-level, numerically legible choice. This also extends naturally to a broader diagnostic: how much of the existing city could still be built under today's zoning and parking rules at all, and where is the largest gap between stated plans and the code actually on the books? Treating this compliance gap as a first-class indicator, audited per parcel and reported year over year, would turn parking reform from an episodic policy debate into a measurable planning trend, and would let cities see for the first time how much of their built form could not be reproduced under the rules currently on their books.

A practical output of such modeling is a single, intuitively explainable metric—the ratio of land devoted to parking to land devoted to buildings on a parcel, corridor, or district. Parking-to-building ratio is legible to residents and elected officials in a way that occupancy forecasts are not, and it ties directly to urban heat, impervious-surface burden, and the walkability cost of auto-oriented form [56,57]. Agentic parking systems should be able to compute this ratio under current rules, under proposed reforms, and under shared-parking scenarios, and to report how each ratio changes alongside climate-exposure and pedestrian-accessibility indicators. Making this coupling explicit is how the system turns "parking" into a category of decision that planners, residents, and climate-resilience officials can reason about together.

*8.6. From Parking to Shunting: Staging for Automated and On-Demand Vehicles*

Automated vehicles and ride-hail services are not primarily a parking problem—they are a *shunting* and queuing problem. A Waymo fleet does not park the way a commuter does; it stages between trips and must be positioned close enough to high-demand events (stadium letouts, airport surges, downtown nightlife) to meet wait-time expectations, but not so close as to consume scarce curb and land near those destinations. The core question is: *how should an agentic system jointly plan*

 doi:10.20944/preprints202605.0084.v1



*parking, curb, and staging space for a mixed fleet of private, ride-hail, and automated vehicles, where the relevant decision variable is often a shunting zone rather than a parking stall?* Treating this as a parking problem leads directly to the expansion-bias failure mode. Treating it as a shunting problem opens a different solution space: micro-hubs, dynamic curb allocation, and event-driven staging that does not translate into permanent parking supply. It also reframes ride-hail and automated-vehicle pickup-and-drop-off zones as *substitutes* for parking rather than extensions of it—a designated short-stay curb zone can serve in a single day the arrival flows that a small lot can only serve through permanent land consumption. A parking agent that does not recognize when the right answer is a pickup zone rather than a stall will keep recommending land for vehicles that no longer need to be stored there.

### *8.7. Public Communication and Political Legitimacy*

Parking decisions are ultimately political, and parking analyses are persuasive tools whether or not they are designed to be. The core question is: *how can an agentic system produce communication artifacts—visualizations, scenario narratives, counterfactual explainers—that help residents, elected officials, and merchants reason about parking trade-offs, rather than simply ratifying the loudest complaint?* The public routinely describes a "parking crisis" when the underlying issue is price, distribution, or spillover; a system that can show "your neighborhood has 20% more parking than it uses at peak; here is where it is; here is what prices would rebalance it" is doing planning work that a prediction pipeline cannot. Communication is not a downstream add-on; for public-sector parking planning, it is where the loop meets democratic legitimacy.

### *8.8. Planner–AI Deliberation Interfaces*

If contradictions are to become objects of institutional reasoning, planners need interfaces that make them inspectable. The core question is: *how can planners inspect assumptions, run what-if scenarios, revise priorities, and contest recommendations in a structured way?* In public planning, the interface is part of the governance model. Progress would mean deliberation environments where planners can interrogate agent reasoning, not merely accept or reject outputs [21,42].

### *8.9. Cross-City Transfer with Policy Adaptation*

The three contradictions are universal in structure, but the specific tensions between demand, policy, theory, and equity vary sharply across places. Some elements of parking logic transfer across cities—pricing, spillover, accessibility reasoning. But policy regimes, curb practices, development patterns, and travel cultures differ. The core question is: *how can agentic parking systems distinguish universal planning logic from local institutional logic, and adapt accordingly?* An agent that works in Phoenix cannot simply be redeployed in London or Tokyo without this distinction [38].

### *8.10. Evaluation Beyond Prediction Accuracy*

If the paradigm truly shifts, evaluation criteria must shift too. The question is no longer whether a system predicts well, but whether it exposes contradictions and supports deliberation [12,48,58]. This paper proposes that the field adopt five evaluation dimensions as a new standard for agentic parking systems:

- **Contradiction detection:** Can the system identify conflicts between stated city goals and inherited parking rules?
- **Intervention comparison:** Can the system evaluate construction against pricing, curb reform, and shared parking?
- **Justification quality:** Can the system defend a recommendation to different institutional audiences?
- **Restraint capability:** Can the system recommend and justify non-construction when appropriate?
- **Policy traceability:** Can the system trace each recommendation to the policy assumptions that produced it?

doi:10.20944/preprints202605.0084.v1



These dimensions define new benchmark tasks: contradiction-detection tasks in which stated city goals conflict with inherited parking rules, intervention-comparison tasks in which construction competes with pricing or curb reform, and explanation-quality tasks in which the system must justify a recommendation to different institutional audiences. A perspective paper's strongest contribution is often not a new concept but a new evaluation grammar. These five dimensions are our proposal for what that grammar should look like in agentic urban planning.

## 9. Conclusions: After Prediction, Toward Institutional Reasoning

Parking planning has been framed too narrowly for too long. Forecasting, optimization, and compliance are all legitimate tasks. The error is mistaking them for planning itself. The distinctive value of agentic parking AI is the ability to expose the contradictions that this reduction conceals—the tensions between demand and policy, between inherited rules and theory, and between market logic and equity—and to turn them into inspectable objects of institutional reasoning. The object of computation must shift accordingly: from occupancy forecasts and site scores to *justifiable intervention reasoning*. Phoenix demonstrates that this shift can be operationalized. Parking, precisely because it is measurable, theory-rich, policy-contested, and intervention-ready, may be the most realistic near-term testbed for agentic urban planning. The future of parking AI should be judged not by how efficiently it predicts demand, but by whether it can detect contradiction, compare interventions, justify restraint, trace policy assumptions, recognize when the right answer is to share, reuse, or redevelop rather than build, and defend its recommendations to the publics that parking decisions ultimately serve. Cities do not mainly need parking systems that predict where cars will go. They need systems that reveal when inherited parking rules are steering urban decisions away from a city's own stated future—and that hold open, rather than foreclose, the underlying civic question: what kind of place do we want to have, and what kind of city do we want to be?

## References


1. Shoup, D. The High Cost of Free Parking. *Journal of Planning Education and Research* **1997**, *17*, 3–20.
2. Manville, M. Parking Requirements and Housing Development. *Journal of the American Planning Association* **2013**, *79*, 49–66. https://doi.org/10.1080/01944363.2013.785346.
3. Manville, M.; Beata, A.; Shoup, D. Turning Housing Into Driving: Parking Requirements and Density in Los Angeles and New York. *Housing Policy Debate* **2013**, *23*, 350–375. https://doi.org/10.1080/10511482.2013.767851.
4. Hess, D.; Flowers, B. Developer Response to the Removal of Minimum Parking Requirements in Buffalo. *Transportation Research Record: Journal of the Transportation Research Board* **2023**, *2677*, 036119812311692. https://doi.org/10.1177/03611981231169282.
5. Sohoni, S.; Lee, B. After the Minimum Parking Requirement: Parking Reform in a Small University City. *Journal of the American Planning Association* **2023**, *90*, 1–15. https://doi.org/10.1080/01944363.2023.2248093.
6. Kodransky, M.; Spring, G. Europe's Parking U-Turn: From Accommodation to Regulation. *Inst. Transp. Dev. Policy* **2011**, *15*.
7. Wang, P.; Fu, Y.; Liu, G.; Hu, W.; Aggarwal, C. Human mobility synchronization and trip purpose detection with mixture of hawkes processes. In Proceedings of the Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining, 2017, pp. 495–503.
8. Liu, Q.; Wang, J.; Chen, P.; Xiao, Z. How does parking interplay with the built environment and affect automobile commuting in high-density cities? A case study in China. *Urban Studies* **2016**, *54*. https://doi.org/10.1177/0042098016667040.
9. Channamallu, S.S.; Kermanshachi, S.; Rosenberger, J.M.; Pamidimukkala, A. A review of smart parking systems. *Transportation Research Procedia* **2023**, *73*, 289–296. International Scientific Conference „The Science and Development of Transport - Znanost i razvitak prometa – ZIRP 2023", https://doi.org/https://doi.org/10.1016/j.trpro.2023.11.920.
10. Yang, S.; Ma, W.; Pi, X.; Qian, S. A deep learning approach to real-time parking occupancy prediction in transportation networks incorporating multiple spatio-temporal data sources. *Transportation Research Part C: Emerging Technologies* **2019**, *107*, 248–265. https://doi.org/https://doi.org/10.1016/j.trc.2019.08.010.

 doi:10.20944/preprints202605.0084.v1




11. Lei, C.; Ouyang, Y. Dynamic pricing and reservation for intelligent urban parking management. *Transportation Research Part C: Emerging Technologies* **2017**, *77*, 226–244. https://doi.org/https://doi.org/10.1016/j.trc.2017.01.016.
12. Peng, Z.R.; Kaifa, L.; Liu, Y.; Zhai, W. The Pathway of Urban Planning AI: From Planning Support to Plan-Making. *Journal of Planning Education and Research* **2023**, *44*. https://doi.org/10.1177/0739456X231180568.
13. Da, L.; Chen, T.; Li, Z.; Bachiraju, S.; Yao, H.; Li, L.; Dong, Y.; Hu, X.; Tu, Z.; Wang, D.; et al. Generative AI in Transportation Planning: A Survey, 2025, [arXiv:cs.AI/2503.07158].
14. Wang, P.; Wang, D.; Liu, K.; Wang, D.; Zhou, Y.; Sun, L.; Fu, Y. Hierarchical reinforced urban planning: Jointly steering region and block configurations. In Proceedings of the Proceedings of the 2023 SIAM International Conference on Data Mining (SDM). SIAM, 2023, pp. 343–351.
15. Wang, D.; Liu, K.; Johnson, P.; Sun, L.; Du, B.; Fu, Y. Deep human-guided conditional variational generative modeling for automated urban planning. In Proceedings of the 2021 IEEE international conference on data mining (ICDM). IEEE, 2021, pp. 679–688.
16. Wang, D.; Fu, Y.; Liu, K.; Chen, F.; Wang, P.; Lu, C.T. Automated urban planning for reimagining city configuration via adversarial learning: quantification, generation, and evaluation. *ACM Transactions on Spatial Algorithms and Systems* **2023**, *9*, 1–24.
17. Wang, D.; Wu, L.; Zhang, D.; Zhou, J.; Sun, L.; Fu, Y. Human-instructed deep hierarchical generative learning for automated urban planning. In Proceedings of the Proceedings of the AAAI Conference on Artificial Intelligence, 2023, Vol. 37, pp. 4660–4667.
18. Wang, D.; Lu, C.T.; Fu, Y. Towards Automated Urban Planning: When Generative and ChatGPT-like AI Meets Urban Planning **2023**. https://doi.org/10.48550/arXiv.2304.03892.
19. Wang, D.; Fu, Y.; Wang, P.; Huang, B.; Lu, C.T. Reimagining city configuration: Automated urban planning via adversarial learning. In Proceedings of the Proceedings of the 28th international conference on advances in geographic information systems, 2020, pp. 497–506.
20. Liu, R.; Zhe, T.; Peng, Z.R.; Catbas, N.; Ye, X.; Wang, D.; Fu, Y. Urban Planning in the Age of Agentic AI: Emerging Paradigms and Prospects. *ACM SIGKDD Explorations Newsletter* **2025**, *27*, 35–42. https://doi.org/10.1145/3787470.3787474.
21. Zhou, Z.; Lin, Y.; Jin, D.; Li, Y. Large Language Model for Participatory Urban Planning, 2024, [arXiv:cs.AI/2402.17161].
22. Liu, Y.; Liu, C.; Yuan, N.J.; Duan, L.; Fu, Y.; Xiong, H.; Xu, S.; Wu, J. Exploiting heterogeneous human mobility patterns for intelligent bus routing. In Proceedings of the 2014 IEEE International Conference on Data Mining. IEEE, 2014, pp. 360–369.
23. Fu, Y.; Liu, G.; Papadimitriou, S.; Xiong, H.; Ge, Y.; Zhu, H.; Zhu, C. Real estate ranking via mixed land-use latent models. In Proceedings of the Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2015, pp. 299–308.
24. Fu, Y.; Xiong, H.; Ge, Y.; Yao, Z.; Zheng, Y.; Zhou, Z.H. Exploiting geographic dependencies for real estate appraisal: A mutual perspective of ranking and clustering. In Proceedings of the Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining, 2014, pp. 1047–1056.
25. Inci, E. A review of the economics of parking. *Economics of Transportation* **2015**, *4*, 50–63. Special Issue on Collective Contributions in the Honor of Richard Arnott, https://doi.org/https://doi.org/10.1016/j.ecotra.2014.11.001.
26. Arnott, R. Spatial competition between parking garages and downtown parking policy. *Transport Policy* **2006**, *13*, 458–469. Parking, https://doi.org/https://doi.org/10.1016/j.tranpol.2006.05.003.
27. Arnott, R.; Rowse, J. Downtown parking in auto city. *Regional Science and Urban Economics* **2009**, *39*, 1–14. https://doi.org/https://doi.org/10.1016/j.regsciurbeco.2008.08.001.
28. Anderson, S.P.; de Palma, A. The economics of pricing parking. *Journal of Urban Economics* **2004**, *55*, 1–20. https://doi.org/https://doi.org/10.1016/j.jue.2003.06.004.
29. Inci, E.; Lindsey, R. Garage and curbside parking competition with search congestion. *Regional Science and Urban Economics* **2015**, *54*, 49–59. https://doi.org/https://doi.org/10.1016/j.regsciurbeco.2015.07.003.
30. Smith, M.S. *Shared parking*; 2005.
31. Ghizzawi, F.; Galal, A.; Roorda, M.J. Modelling parking behaviour of commercial vehicles: a scoping review. *Transport Reviews* **2024**, *44*, 743–765. https://doi.org/https://doi.org/10.1080/01441647.2024.2305202.
32. Wiersma, J.; Bertolini, L. Making the trade-offs between parking and placemaking visible and actionable. *Journal of Urban Design* **2024**, *29*, 1–17. https://doi.org/10.1080/13574809.2024.2337726.

33. Hoehne, C.G.; Chester, M.V.; Fraser, A.M.; King, D.A. Valley of the sun-drenched parking space: The growth, extent, and implications of parking infrastructure in Phoenix. *Cities* **2019**, *89*, 186–198. https://doi.org/https://doi.org/10.1016/j.cities.2019.02.007.
34. Chester, M.; Fraser, A.; Matute, J.; Flower, C.; Pendyala, R. Parking Infrastructure: A Constraint on or Opportunity for Urban Redevelopment? A Study of Los Angeles County Parking Supply and Growth. *Journal of the American Planning Association* **2015**, *81*, 268–286. https://doi.org/10.1080/01944363.2015.1092879.
35. Bui, Q.; Chaban, M.A.V.; White, J. 40 Percent of the Buildings in Manhattan Could Not Be Built Today. *The New York Times (The Upshot)* **2016**. Published 2016-05-19.
36. Zhao, X.; Zhang, M. Enhancing Predictive Models for On-Street Parking Occupancy: Integrating Adaptive GCN and GRU with Household Categories and POI Factors. *Mathematics* **2024**, *12*. https://doi.org/10.3390/math12182823.
37. Yuldashev, Y.; Mukhiddinov, M.; Abdusalomov, A.B.; Nasimov, R.; Cho, J. Parking Lot Occupancy Detection with Improved MobileNetV3. *Sensors* **2023**, *23*. https://doi.org/10.3390/s23177642.
38. Sohoni, S.; Lee, B. Parking the Minimum Parking Requirement: Evidence from Nine U.S. Cities. *Journal of Planning Education and Research* **2025**, *46*. https://doi.org/10.1177/0739456X251343047.
39. Zhang, W.; Zhang, L.; Han, J.; Liu, H.; Fu, Y.; Zhou, J.; Mei, Y.; Xiong, H. Irregular traffic time series forecasting based on asynchronous spatio-temporal graph convolutional networks. In Proceedings of the Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2024, pp. 4302–4313.
40. Mondschein, A.; King, D.A.; Hoehne, C.; Jiang, Z.; Chester, M. Using social media to evaluate associations between parking supply and parking sentiment. *Transportation Research Interdisciplinary Perspectives* **2020**, *4*, 100085. https://doi.org/https://doi.org/10.1016/j.trip.2019.100085.
41. Rittel, H.; Webber, M., Dilemmas in a General Theory of Planning; 2016. https://doi.org/10.5040/9781474282932.0015.
42. Healey, P. *Collaborative planning: Shaping places in fragmented societies*; Bloomsbury Publishing, 1997.
43. Zhu, H.; Chen, G.; Zhang, W. PlanGPT: Enhancing Urban Planning with a Tailored Agent Framework. In Proceedings of the Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 6: Industry Track); Rehm, G.; Li, Y., Eds., Vienna, Austria, 2025; pp. 764–783. https://doi.org/10.18653/v1/2025.acl-industry.54.
44. Miller, T. Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence* **2019**, *267*, 1–38. https://doi.org/https://doi.org/10.1016/j.artint.2018.07.007.
45. Wang, X.; Wang, D.; Ying, W.; Bai, H.; Gong, N.; Dong, S.; Liu, K.; Fu, Y. Efficient post-training refinement of latent reasoning in large language models. In Proceedings of the Proceedings of the AAAI Conference on Artificial Intelligence, 2026, Vol. 40, pp. 33692–33700.
46. Arnstein, S.R. A Ladder Of Citizen Participation. *Journal of the American Institute of Planners* **1969**, *35*, 216–224, [https://doi.org/10.1080/01944366908977225]. https://doi.org/10.1080/01944366908977225.
47. Schön, D.A. *The reflective practitioner: How professionals think in action*; Routledge, 2017.
48. Fu, X.; Li, C. Automating Plan Evaluation Using Agentic Large Language Models. *Journal of Planning Education and Research* **2025**. https://doi.org/10.1177/0739456X251372082.
49. Wang, X.; Fu, Y. Dataforge: A Data Agent Platform for Autonomous Data Engineering. *arXiv preprint arXiv:2511.06185* **2025**.
50. Doshi-Velez, F.; Kim, B. Towards A Rigorous Science of Interpretable Machine Learning, 2017, [arXiv:stat.ML/1702.08608].
51. Lipton, Z.C. The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery. *Queue* **2018**, *16*, 36–43.
52. Ribeiro, M.; Singh, S.; Guestrin, C. “Why Should I Trust You?”: Explaining the Predictions of Any Classifier. In Proceedings of the Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations; DeNero, J.; Finlayson, M.; Reddy, S., Eds., San Diego, California, 2016; pp. 97–101. https://doi.org/10.18653/v1/N16-3020.
53. Bhandari, P.; Anastasopoulos, A.; Pfoser, D. Are Large Language Models Geospatially Knowledgeable? 12 2023, pp. 1–4. https://doi.org/10.1145/3589132.3625625.
54. Zhang, Y.; Wei, C.; He, Z.; Yu, W. GeoGPT: An assistant for understanding and processing geospatial tasks. *International Journal of Applied Earth Observation and Geoinformation* **2024**, *131*, 103976. https://doi.org/https://doi.org/10.1016/j.jag.2024.103976.

 doi:10.20944/preprints202605.0084.v1




55. Gong, N.; Dong, S.; Bai, H.; Wang, X.; Ying, W.; Fu, Y. Agentic Feature Augmentation: Unifying Selection and Generation with Teaming, Planning, and Memories, 2025, [arXiv:cs.LG/2505.15076].
56. Fu, Y.; Wang, P.; Du, J.; Wu, L.; Li, X. Efficient region embedding with multi-view spatial networks: A perspective of locality-constrained spatial autocorrelations. In Proceedings of the Proceedings of the AAAI conference on artificial intelligence, 2019, Vol. 33, pp. 906–913.
57. Fu, Y.; Liu, G.; Ge, Y.; Wang, P.; Zhu, H.; Li, C.; Xiong, H. Representing urban forms: A collective learning model with heterogeneous human mobility data. *IEEE transactions on knowledge and data engineering* **2018**, *31*, 535–548.
58. Wang, X.; Liu, Y.; Cheng, W.; Zhao, X.; Chen, Z.; Yu, W.; Fu, Y.; Chen, H. Mixllm: Dynamic routing in mixed large language models. In Proceedings of the Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers), 2025, pp. 10912–10922.