

AN ILLUSTRATED GUIDE

Urban Fundamentals



**A friendly guide to cities,
data and urban AI**

Plain English, worked examples and short exercises,
with a glossary and an eight-week study plan.

Twenty-one short chapters · Glossary · Eight-week plan



Welcome

This guide is for anyone who wants to understand how **cities** work and the words people use to talk about them: architects, engineers, data scientists, geographers and students. It gives a plain-English tour of urban form, transport, climate, data, methods and urban AI, one short chapter at a time.

Skills from architecture, engineering, data science or geography all carry over. What a newcomer usually needs is the shared vocabulary of planning, transport and urban data, and a map of who says what. That is what this book provides.

REMEMBER

The goal is fluency, not memorisation: know the main ideas, know the words, and know which questions to ask.

How to read this book

This book is a **fast primer**: about eight hours end to end, so plan on a chapter or two a day. Skim each chapter, look at the pictures, and read the coloured boxes twice. The **30 to 40 minutes a day** belongs to the eight-week plan in Chapter 21, which is where the real reading and hands-on practice happen. Keep the glossary open next to you and tick a term the first time you hear it.

The little boxes you will see

Box	What it means
TIP	A shortcut or a way to remember something.
WATCH OUT	A common trap. Papers fall into these all the time.
REMEMBER	If you keep only one thing from the page, keep this.
SAY IT LIKE THIS	Sentences you can use in a real conversation.
YOU CAN SKIP THIS	Extra detail. Read it on a second pass.
BRIDGE	Where a building-scale background is an advantage.

Five habits for learning a new field

- **State your background.** "I come from the building side." Knowing where you start shows which gaps to fill first.
- **Ask before you claim.** Turn a half-known term into a question.
- **Keep a list of new words.** Check them against the glossary the same night.
- **Claim only what you have read or built yourself.**
- **Notice which words keep recurring** in the papers and talks you meet. Those are the ones to master first.

WATCH OUT

Most numbers and book details here come from general knowledge. A few were checked online in September 2026: Korea's population figures, the latest edition of *Modelling Transport*, the publication details of Bettencourt's book, Howard's book titles, and the names of Lee's seven "sins". Treat everything else as approximate, and verify before you quote it.

1

The City at Every Scale

Zoom in, zoom out: same city, different questions.

A city is read at several scales, and each scale asks different questions with different data. Most confusion in urban conversations comes from two people talking at different scales without saying so. Think of a zoom lens: the city does not change, only how close you stand.

Building / parcel

10–100 m

"How does this building perform?"

Block / street

~100 m

"Is this street good for people?"

Neighbourhood

0.5–2 km

"Can I reach daily needs on foot?"

City

~10 km

"Where do people live, work, travel?"

Metro region

10–100 km

"How do commutes link cities?"

Figure 1. The scale ladder. Ask "which rung are we on?" first.

What is "urban form"?

It is just the physical shape of a city, described with four measurable things:

- **Density.** People or homes per hectare. *Gross* density counts all the land; *net* density counts only the residential land.
- **Land-use mix.** How many different uses (homes, jobs, shops) sit close together.
- **Street network.** Block size and intersection density. A fine grid is walkable; a maze of cul-de-sacs is not.
- **Morphology.** The pattern of streets and plots: grid, organic, radial, superblock.

TIP

Land use is what people *do* on land (housing, industry). **Land cover** is what the surface *is* (concrete, trees, water). Satellites see cover. Use usually comes from plans, or has to be guessed.

Zoning, FAR and coverage: the numbers you already know

Zoning is the legal tool that says what may be built where. Two ratios every architect learns:

- **FAR** (floor area ratio) = total floor area ÷ plot area. In Korean: 용적률 (yongjeokryul).
- **Building coverage ratio** = building footprint ÷ plot area. In Korean: 건폐율 (geonpyeyul).




Figure 2. Same FAR, very different buildings.

BRIDGE

Height limits, setbacks and FAR explain a lot of why a city looks the way it does. AI models trained on street or mobility data often inherit these effects without saying so. A building background is well placed to spot that.

WORKED EXAMPLE (TOY NUMBERS)

Gross vs net density. A district has 10,000 residents on 200 ha of land, of which 80 ha is housing. Gross density = $10,000 \div 200 = 50$ people/ha. Net residential density = $10,000 \div 80 = 125$ people/ha. Same place, two honest numbers. Always ask which one a paper used.

FAR and coverage together. Plot 2,000 m², zoning allows FAR 250% and coverage 60%. Maximum floor area = $2,000 \times 2.5 = 5,000$ m². Maximum footprint = $2,000 \times 0.6 = 1,200$ m². To use all the floor area with the biggest footprint you need $5,000 \div 1,200 \approx 4.2$, so **5 floors**. Two limits, one building shape.

COMMON MYTH

"High density means crowded and bad." Density is a count; crowding is an experience. Some of the world's most liked neighbourhoods are dense, and some low-density suburbs feel cramped. Ask what is being measured and for whom.

"Zoning is just a map." It is a rulebook that shapes prices, building types and who can live where.

TRY IT (10 MINUTES)

Pick the neighbourhood you live in. (1) Name its rung on the scale ladder. (2) Find its population and area and compute gross density. (3) Sketch its street pattern in one line: grid, organic, or cul-de-sac maze. You have just done a small urban-form analysis.

2

Big Ideas Behind Today's Debates

Most "new" arguments are old planning arguments in new clothes.

Five classics

Idea	In one breath	Why it still matters
Garden City Howard, 1898	Planned, self-contained towns ringed by green belts.	Ancestor of new towns. Korea's Sejong is a descendant.
Modernist city Le Corbusier, CIAM	Separate the uses, towers in parkland, the car as organiser.	What most critics push against.
Jane Jacobs <i>Death and Life of Great American Cities</i> , 1961	Eyes on the street, short blocks, mixed uses, old buildings among new.	Modern walkability research mostly tests her claims.
Kevin Lynch <i>The Image of the City</i> , 1960	We read a city through paths, edges, districts, nodes, landmarks.	Ancestor of street-view perception studies.
Christopher Alexander "A City Is Not a Tree", 1965	Real cities are overlapping networks; planned ones are tidy trees.	A doorway into network science.

Five modern concepts

Concept	In one breath	The debate
15-minute city Carlos Moreno	Daily needs within a 15-minute walk or ride.	Easy to measure; harder to show it changes behaviour.
Compact city / TOD	Dense, mixed development around stations. The Seoul and Tokyo model.	Who can afford to live there?
Smart city	Technology-led urban management.	Governance and privacy. Sidewalk Toronto is the usual example.
Resilient city	Designing for heat, flood and energy shocks.	Where climate meets building-scale work.
Digital twin	A live, data-linked model of a city.	Visualisation, prediction, or decision support?

TIP

Short summaries of the classic books are in **Chapter 11**.

SAY IT LIKE THIS

"Which worldview does this model quietly assume? A walkability score, for example, assumes a Jacobs-style idea of what makes a good street."

Which classic explains which modern debate?

Modern debate	Classic voice behind it	The tension
Should a model score a street's "quality" from images?	Lynch, Jacobs	Whose eyes? Legibility and liveliness are not the same thing.
Do 15-minute cities work everywhere?	Howard, Jacobs	Planned self-sufficiency vs organic mixing.
Can a digital twin steer a city?	Lee, Rittel and Webber	Control vs a system that changes as you act.
Why do generic layouts feel lifeless?	Alexander	Tree-like plans erase overlaps.
Are bigger cities "better"?	Bettencourt, Batty	Scaling benefits come with scaling costs.

TRY IT (10 MINUTES)

Take any news headline about smart cities or the 15-minute city. Write down which of the ten ideas above it quietly assumes, and one thing it leaves out.

3

Getting Around: Transport 101

Land use and transport are two sides of one coin.

Where things are decides how people travel. How easily people can travel decides where things get built. Transport people use a shared vocabulary, so here it is.

The four-step model

This is the classic way to forecast travel. Learn the four boxes and you can place almost any mobility paper.



Result: forecast of traffic and transit use on every link of the network

Figure 3. The four-step model.

Its weakness is that it works on zones and averages. **Activity-based models** and **agent-based simulation** (MATSim, SUMO and similar) follow individual synthetic people through a day instead. This is where AI comes in: learned models now try to replace or speed up parts of this pipeline.

Words you will hear

Term	Meaning
OD matrix	Trips between every pair of zones (origin to destination).
Modal share	Percent of trips by walking, bus, metro, car and so on.
Induced demand	Build more road capacity and more traffic shows up, so congestion returns.
Level of service (LOS)	A to F grade for how well a road or facility performs.
First / last mile	The leg between a station and the real start or end of a trip.
Micromobility	Shared bikes, scooters and similar small vehicles.
MaaS	Mobility as a Service: one app to plan and pay across modes.
GTFS	The open format transit agencies use to publish schedules.

TIP

Most urban-AI mobility papers (demand prediction, traffic forecasting, trajectory generation) attack **one box** of the four-step pipeline. Work out which box and you can judge the claim.

WORKED EXAMPLE (TOY NUMBERS)

Gravity model (step 2). Zone A produces 1,000 trips. Destination B has 300 jobs and is 2 km away; C has 600 jobs and is 4 km away. Score each as $\text{jobs} \div \text{distance}^2$: $B = 300 \div 4 = 75$, $C = 600 \div 16 = 37.5$. Shares: $B = 75 \div 112.5 = 67\%$, $C = 37.5 \div 112.5 = 33\%$. So about **667 trips go to B and 333 to C**, even though C has twice the jobs. Distance matters that much in this model.

Logit mode choice (step 3). Give each mode a score $U = -0.05 \times \text{minutes} - 0.3 \times \text{cost}$. Car (20 min, cost 4): $U = -2.2$. Bus (30 min, cost 1.5): $U = -1.95$. Probability = $e^U \div \text{sum of } e^U$: car $0.111 \div 0.253 = 44\%$, bus **56%**. Make the bus cheaper and its share rises smoothly, not all at once.

The numbers are invented for illustration. Real models estimate the coefficients from survey data.

COMMON MYTH

"More roads fix congestion." Decades of evidence on induced demand say new capacity often fills up. The better question is what trips the road is enabling.

TRY IT (10 MINUTES)

Look up one paper on traffic or demand prediction. Decide which of the four boxes it works on and what data it uses. Then write one sentence on what it cannot tell a planner.

4

Who Can Reach What?

Accessibility and equity, without the headache.

Accessibility means how many opportunities (jobs, schools, clinics, parks) a person can reach within a given time or cost. It is not the same as mobility, which is only how far or how fast you can move. A city can have fast roads and poor accessibility if everything is far apart.

Ways to measure it

- **Cumulative opportunities.** Count the jobs reachable within, say, 30 minutes by transit.
- **Gravity-based.** Opportunities count less the farther away they are.
- **Isochrone.** A map of everything reachable within a set time. Tools like r5r, OpenTripPlanner and OSMnx make them.
- **Walk Score and similar indices.** Simple and popular. Know their limits.

Equity in three ideas

- **Spatial mismatch.** Jobs far from where the people who need them live.
- **Transport poverty.** When getting around eats a large share of a household's budget or time.
- **Equality vs equity.** Equality is the same for all. Equity gives more where the need is greater.

WATCH OUT

MAUP (modifiable areal unit problem): results change when you change how the areas are drawn. Always say which units you used.

Ecological fallacy: a pattern in an area does not describe each person in it.

Averages hide the very groups that matter. Ask for the breakdown.

WORKED EXAMPLE (TOY NUMBERS)

Why averages hide inequality. A city has three areas. A: 30,000 residents, 120,000 jobs reachable in 30 minutes by transit. B: 30,000 residents, 45,000 jobs. C: 40,000 residents, 8,000 jobs.

Average jobs reachable per resident = $(30 \times 120 + 30 \times 45 + 40 \times 8) \div 100 = 52,700$. That sounds healthy. But **40% of residents (area C) reach only 8,000 jobs**, about one-fifteenth of area A. The average is real and still hides the group that matters.

TRY IT (10 MINUTES)

Pick a service you use (clinic, market, station). How many minutes does it take to reach on foot? What would change for someone who cannot walk far, or who lives in the next district? You have just asked an equity question.

5

Cities, Climate and Energy

Your bridge: the physics you know, repeated across a whole city.

For readers with a building background, this chapter is a bridge. Everything known about a single building (heat, light, wind, energy) scales up. The city is just a very large, very messy building site.




Figure 4. The street canyon. Geometry controls heat and daylight.

- **Urban heat island (UHI).** Cities run warmer than their surroundings: dark surfaces, less vegetation, heat trapped between buildings. Often strongest at night.
- **UTCI / PET and mean radiant temperature.** Outdoor comfort measures that describe what people feel, better than air temperature.
- **Sky view factor and H/W.** The geometry in Figure 4. You already think in these terms.
- **UBEM** (urban building energy modelling). Simulating energy for thousands of buildings from archetypes and geometry.
- **Blue-green infrastructure.** Trees, parks, water and green roofs used for cooling and drainage.
- **Air quality and noise.** Other key layers, measured by sparse stations and increasingly estimated from images and satellites.

BRIDGE

Cities rarely have measured data for every building. Much of the research problem is estimating building attributes (age, height, use, envelope) at scale from images and open data. That is exactly where building expertise meets AI.

WORKED EXAMPLE (TOY NUMBERS)

Street canyon. Buildings 30 m tall on both sides of a 15 m wide street: $H/W = 30 \div 15 = 2$, a deep canyon. Buildings 15 m tall on a 30 m street: $H/W = 0.5$, an open street. Deep canyons shade well by day but can trap heat at night.

Heat island intensity. At 11 pm the city centre reads 24 °C and the surrounding countryside 20 °C. UHI intensity = $24 - 20 = 4$ °C. Studies usually report this figure at night, when it tends to be largest.

Energy use intensity (EUI). A 10,000 m² office using 1,500,000 kWh a year has $EUI = 1,500,000 \div 10,000 = 150$ kWh/m²/year. UBEM does this for thousands of buildings at once, using typical values by building type and age.

COMMON MYTH

"Planting trees always cools a city." Trees usually help through shade and evaporation, but the effect depends on species, water supply and street geometry. In a deep canyon, dense canopy can also trap heat at night. Ask what was measured and when.

TRY IT (10 MINUTES)

Find a street you know on a map. Estimate building height from floors (about 3 m each) and measure street width. Work out H/W . Is it a canyon, a boulevard, or something in between?

6

Urban Data: Where It Comes From

Every dataset tells a story, and hides one.

Urban AI lives or dies on data. For each source, know what it shows and what it hides.

Source	What it shows	Main trap
Census / registers	Who lives where, households, jobs	Slow to update; area units (MAUP)
OpenStreetMap	Roads, buildings, points of interest	Uneven completeness between places
Satellite imagery	Land cover, building footprints, heat	Sees surfaces, not use; clouds, resolution
Street-view images	Streetscape, greenery, facades	Coverage and date vary; camera bias
Phone / GPS traces	Movement and presence	Privacy; who owns a phone; sampling bias
Smart-card / transit data	Actual public-transport use	Only covers transit users
IoT sensors	Air, noise, traffic, energy in near real time	Sparse; needs calibration and upkeep
Social media / reviews	Sentiment, activity, events	Self-selected users; platforms change
Cadastre / registers	Parcels, ownership, building age and use	Errors and gaps; access rules differ

REMEMBER

For any dataset ask four questions: **Who is in it? Who is missing? How old is it? At what spatial unit?** Most bad urban results trace back to one of the four.

WATCH OUT

Mobility and phone data can identify people even when anonymised. Expect talk of aggregation, differential privacy and consent. Ask early what can be shared.

WORKED EXAMPLE (TOY NUMBERS)

Who is in the data? A city has 600,000 adults. Younger adults: 400,000, of whom 30% use a mapping app that shares location, so 120,000 users. Older adults: 200,000, of whom 5% use it, so 10,000 users.

In the population, older adults are $200,000 \div 600,000 = 33\%$. In the data they are $10,000 \div 130,000 = 7.7\%$. They are about **four times under-represented**. Any conclusion about "how people move" is really about younger app users unless you reweight. Bigger datasets do not fix who is missing.

TRY IT (10 MINUTES)

Pick one dataset you have used or read about. Answer the four questions: who is in it, who is missing, how old is it, at what spatial unit. Write one line for each.

7

Urban Science: The Maths of Cities

Cities are messy, yet they obey surprising rules.

A line of research treats cities as complex systems with measurable regularities. It is the theory behind a lot of urban AI.




Figure 5. Scaling laws (Bettencourt, West and colleagues). Exponents are approximate and debated.

- **Scaling laws.** Double a city's population and infrastructure grows by less than double (bigger cities are more efficient), while wages, patents and also crime grow by more than double.
- **Rank-size rule (Zipf).** City sizes follow a regular pattern: the second-largest is roughly half the first. Real countries deviate, and Seoul's dominance in Korea is a famous case.
- **Network science.** Streets, transit and social ties as graphs. Words to know: centrality, betweenness, small-world. **Space syntax** is the design-side cousin.
- **Complexity and emergence.** Patterns that come from many small decisions rather than one plan. This links back to Jacobs and Alexander.
- **Agent-based models.** Simulate many simple actors to see city-level results, such as Schelling's segregation model.

WATCH OUT

A warning from 1973. Rittel and Webber called planning problems "**wicked**": no single correct formulation, no clear end point, every solution changes the problem. That same year Douglass Lee published *Requiem for Large-Scale Models*, arguing that big city models failed because they were opaque, data-hungry and oversold.

This is the standard caution for anyone building urban foundation models or digital twins today. Citing it marks you as someone who knows the history.

WORKED EXAMPLE (TOY NUMBERS)

Scaling by hand. Suppose a city doubles from 1 million to 2 million people. With an exponent of 0.85, infrastructure grows by $2^{0.85} \approx 1.80$, so **80% more** pipes and roads for **100% more** people: about 10% less per person. With an exponent of 1.15, wages and output grow by $2^{1.15} \approx 2.22$, which is about **11% more per person**.

Rank-size rule. If the largest city has 10 million people, the rule predicts the second at about 5 million, the third at about 3.3 million, and the tenth at about 1 million (largest $\div$ rank).

COMMON MYTH

"Scaling laws let you predict any single city." They describe averages across many cities. Individual cities sit above or below the line, and that gap is often the interesting part.

TRY IT (10 MINUTES)

Take ten cities in one country. Plot population against rank on log-log axes, or against total road length. Is the line straight? Which city sits far from it, and why?

8

Urban AI: The Vocabulary

The AI ideas, in urban language.

If you come from AI, this chapter lines up the urban words. If you come from the built environment, it lines up the AI words. Either way, it lets you place any paper on the map.

Term	What it means	Ask this
Urban computing	Umbrella field: collecting and analysing city data (sensors, traces, images) to solve urban problems.	What problem, for whom?
GeoAI	Machine learning on spatial data.	Did you handle spatial autocorrelation?
Spatial cross-validation	Test a model on <i>regions</i> it has not seen, not random points.	How was the split done?
Urban foundation model	Pre-trained model over urban data (images, trajectories, text, maps).	Which cities? Does it transfer to a new one?
Agentic digital twin	A twin that mirrors the city and also has AI agents acting in it.	What decision does it support, and who checks it?
Synthetic population	Artificial people matching census statistics, to drive simulation.	Matched to what, from which year?
Urban perception	Predicting how people feel about a place from images.	Whose perception was collected?
Causal inference	Telling whether a policy or design <i>caused</i> an outcome.	Correlation or cause?

TIP

Spatial autocorrelation just means nearby places tend to look alike. It breaks the usual "independent samples" assumption, so a random train/test split leaks. This is the single best question to ask about a GeoAI paper.

TIP

A handy way to sort any paper: **white box** (physics-based simulation such as CFD, daylight, energy), **black box** (pure machine learning), or **grey box** (a hybrid of both). Most digital-twin work lives in the grey zone.

BRIDGE

Building attributes (height, age, use, energy) are the missing layer in most city models. Reconstructing them, then feeding energy, heat or mobility models, is a real gap and a natural entry point for anyone with a building background.

A checklist for reading any urban AI paper

Question	Why it matters
What is the task? Prediction, generation, simulation, decision support?	Different tasks need different evidence.
What data, from where and when?	Coverage and date decide what the model can know.
At what scale? Building, street, neighbourhood, city, region.	Results rarely transfer across scales.
How was it tested? Random split or by region and by city?	Random splits inflate accuracy (spatial autocorrelation).
Against what baseline?	A simple model may do nearly as well.
Who benefits, and who is missed?	Errors are rarely spread evenly.

WORKED EXAMPLE (TOY NUMBERS)

Why random splits mislead. A model predicts street greenery from images in 10 neighbourhoods. With a random split, it scores $R^2 = 0.90$ on test points, because near-identical neighbours of each test point are in training. With a split that holds out whole neighbourhoods, the same model scores $R^2 = 0.50$. Both numbers are honest; only the second says what happens in a new place. These figures are invented to show the pattern. Real gaps vary, but the direction is typical.

COMMON MYTH

"A foundation model works for every city." Ask which cities and which data it was trained on, and test transfer to a city it has not seen.

TRY IT (10 MINUTES)

Read the abstract of any recent urban AI paper and fill the checklist above in six short lines. If you cannot answer one, that is your question for the authors.

9

Korea: A Quick Guide

A few facts and a few words go a long way.

Big picture (approximate; verify before quoting)

- Very high urban share and density. The Seoul metropolitan region holds about half the population.
- Rapid postwar growth made the **high-rise apartment complex** (아파트 단지) the standard housing form. Rebuilding ageing complexes is a major policy topic.
- **New towns**: the 1990s first-generation towns around Seoul, and later Sejong as the administrative capital.
- Strong public transit, integrated fare cards, and heavy investment in smart-city pilots.
- National planning laws set FAR and coverage limits by zone, so Chapter 1 applies directly.

Mini Korean glossary

Korean	Romanised	Meaning
도시계획	dosi gyehoek	urban planning
용적률	yongjeokryul	floor area ratio (FAR)
건폐율	geonpyeyul	building coverage ratio
아파트 단지	apateu danji	apartment complex
신도시	sindosi	new town
대중교통	daejung gyotong	public transport
도시재생	dosi jaesaeng	urban regeneration
동 / 구 / 시	dong / gu / si	neighbourhood / district / city

TIP

Korean statistics are published by 동, 구 **and** 시. These are the spatial units you will work with, so learn them first.

How Korean planning is organised (in brief; verify details)

- Land planning is governed by national law, with plans at national, regional and city level. Cities produce a long-term **basic plan** (도시기본계획) and a detailed **management plan** (도시관리계획) that sets zoning and infrastructure.
- Land is divided into broad zones, such as urban areas, management areas, agricultural and forest areas, and conservation areas. Urban zones are split further into residential, commercial,

industrial and green zones, each with its own FAR and coverage limits.

- Big, planned housing supply has long been a national tool: new towns, apartment complexes, and later regeneration and reconstruction programmes.

Where Korean urban data tends to live

Source (Korean name)	What you get
KOSIS (국가통계포털)	National statistics portal: population, housing, economy, by administrative area.
Census output areas (집계구)	The smallest statistical units, used for fine-grained population maps.
Building register (건축물대장)	Official record of each building: use, floors, area, approval year.
City open-data portals	Seoul and other cities publish transit, sensor and land data online.

TIP

A quick way to learn the system: find your own address in the building register and read every field. You will meet FAR, coverage, use, floors and age in one place.

10

Asking Good Questions

Good questions beat borrowed answers.

Questions that show you understand the field

- "At what spatial scale is this model meant to work?"
- "How did you split train and test: randomly, or by region?"
- "Which cities is it trained on, and would it transfer to one that is not?"
- "Who is missing from this dataset?"
- "What decision would a planner make differently because of this?"
- "Is this correlation, or can you say the design caused the effect?"
- "How do you get building attributes at this scale?"

SAY IT LIKE THIS

"I come from the building side. Can you explain how planners define that here?"

"I have read about it but not used it. What should I look at first?"

"Let me check that and come back to you tomorrow."

"That is outside what I know yet. Who has worked on it?"

WATCH OUT

Avoid quoting a number you have not verified, defending a half-known term, or pretending to have followed a debate you only skimmed. A colleague who says "I do not know yet" and returns next day with the answer earns more trust than one who bluffs once.

Habits for learning in a new group

- **Learn names and projects.** For each person, note one question about what they are working on.
- **Ask for one paper each.** "Which single paper should a newcomer read first?" Read those before anything else.
- **Keep a "words I heard" list.** Check it against the glossary that night.
- **Say one useful thing early.** A small concrete skill (a script, a dataset, a figure) earns more respect than a long speech.
- **Write things down.** Notes on what you were told, and what you promised to check.

Quick self-check (answers at the bottom)

1. A plot is 1,000 m², the footprint is 400 m², and the building has 5 floors of 400 m². What are the coverage and the FAR?
2. Name the four steps of the four-step model.
3. What does it mean to split data by region rather than randomly, and why do it?
4. Which is larger for a big city as it doubles in size: infrastructure or wages?
5. What is the difference between land use and land cover?
6. What does MAUP warn you about?
7. Which of Lee's seven sins is "hungriness"?
8. What does a street's H/W ratio measure, and what does a high value mean?
9. Give one reason a smart-city claim deserves scrutiny.
10. Why is an average accessibility number not enough for equity?

Answers: 1) coverage 40%, FAR 200%. 2) generation, distribution, mode choice, assignment. 3) It tests on places the model has not seen, because nearby places are similar (spatial autocorrelation) and random splits leak. 4) Wages grow faster (exponent about 1.15 vs about 0.85 for infrastructure). 5) Use is what people do on the land; cover is what the surface physically is. 6) Results can change when the area boundaries change. 7) Needing vast amounts of data. 8) Building height divided by street width; a high value is a deep canyon with less sky, less daylight and trapped heat. 9) Governance, privacy and who benefits (any of these). 10) An average can hide a large group with very poor access.

11

The Classics in Ten Minutes Each

No time to read the books? Here is what is in them.

Each summary gives you the big idea, the concepts worth borrowing, the main criticisms, and a sentence you can use. They are short retellings in plain words, written from general knowledge rather than checked page by page against the books. They are enough to follow a conversation. To quote a book, open the book.

The Death and Life of Great American Cities

Jane Jacobs · 1961

A city is a living, messy system, and planners who tidy it up kill it.

The big idea

Jacobs looked at real sidewalks in New York and other US cities and asked why some streets are safe, lively and loved while planned projects feel dead. Her answer: cities are problems of **organised complexity**, with countless small interacting parts. You cannot design them from above like a machine; you have to nurture the conditions that let life happen.

Ideas to borrow

- **Eyes on the street.** Safety comes from many ordinary people watching and passing by, not from police or empty green buffers. This needs a mix of shops, homes and people at different hours.
- **Four conditions for diversity.** (1) Mixed primary uses so people are around all day. (2) Short blocks, so many routes exist. (3) Buildings of different ages and conditions, so rents vary and small businesses survive. (4) Enough density of people.
- **Sidewalk life and trust.** Small daily contacts (a shopkeeper, a neighbour) build public trust between strangers.
- **Against "urban renewal".** Clearing old neighbourhoods and rebuilding as towers in parkland destroys the economy and social networks that made the place work.
- **Gradual vs cataclysmic money.** Slow, small investment lets areas improve; sudden huge funding tends to wreck them.

Criticism and limits

Based mostly on a few US neighbourhoods. Critics say she underplayed race, class and how the very qualities she praised can lead to gentrification. Her four conditions are hypotheses; modern walkability and street-view studies are essentially testing them.

SAY IT LIKE THIS

"That walkability score seems to assume a Jacobs-style idea of a good street: mixed use, short blocks. Has it been tested outside that setting?"

The Image of the City

Kevin Lynch · 1960

People navigate by a mental picture of the city, and good design makes that picture clear.

The big idea

Lynch interviewed residents of Boston, Jersey City and Los Angeles and had them sketch maps and describe routes. He found people build a simplified mental image. Places that are easy to picture (**imageable** or **legible**) help people orient, feel comfortable, and connect emotionally.

The five elements

Element	What it is	Example
Paths	Channels people move along	Streets, transit lines, walkways
Edges	Boundaries that are not paths	Rivers, railways, walls, shorelines
Districts	Areas with a recognisable character	A market quarter, a campus
Nodes	Junctions or focus points	Squares, station forecourts, major crossings
Landmarks	Distinct reference objects	A tower, a clock, a unique building

Criticism and limits

Small samples from three US cities, mostly educated residents. It describes how places are read, not what people value. Still, mental maps and legibility are the ancestors of today's work on urban perception and wayfinding.

SAY IT LIKE THIS

"Which of Lynch's elements is the model picking up when it says a street feels memorable: a landmark, a path, or a district?"

A City Is Not a Tree

Christopher Alexander · 1965

Real cities are overlapping networks; planned cities are tidy trees, and that is why they feel lifeless.

The big idea

Alexander compares two ways of organising the parts of a city. In a **tree**, every part belongs to exactly one branch, and groups never overlap. In a **semilattice**, groups can overlap: a newsstand belongs to the street's life *and* to the pavement and to a corner crowd. Natural, historic cities are semilattices. Modern planned cities (his examples include Brasília and Chandigarh) are trees, which forces a rigid separation and loses the rich overlaps that make city life.

Ideas to borrow

- **Overlap is a feature.** Uses and people that share space create the connections that make a city work.
- **Structure matters.** How parts connect matters as much as what the parts are. This is the seed of network thinking about cities.
- **Design is hard because the real structure is complex.** Simple diagrams are usually the wrong ones.

Criticism and limits

Short and abstract, and he gave few tools for designing a semilattice. He later developed the idea in *A Pattern Language* (1977), a catalogue of building and town design patterns. Modern network science supplies the maths he lacked.

Garden Cities of To-morrow

Ebenezer Howard · 1898 (retitled 1902)

Combine the best of town and country in planned, self-contained new towns.

The big idea

Howard imagined "three magnets" pulling people: town (jobs, culture, but crowding), country (nature, but isolation), and a third, **town-country**, that has both. His answer was a new, planned town of modest size (about 32,000 people) surrounded by a permanent agricultural belt, with land held in common so its rising value benefits the community. When a town filled up, start another nearby; a cluster of towns linked by rail forms a "social city".

Why it still matters

- Gave us the **new town** idea: Letchworth and Welwyn in England, later Korea's new towns and Sejong.
- The **green belt** and neighbourhood planning ideas trace back to it.
- Shows the tension planners still face: control and design vs organic growth (Jacobs' critique).

Requiem for Large-Scale Models

Douglass Lee · 1973

Why the first wave of giant city models failed, and how not to repeat it.

The big idea

In the 1960s, planners built huge computer models of whole cities, hoping to predict and steer them. By the early 1970s they had mostly failed. Lee listed seven "sins". Read them today as a checklist for urban foundation models and digital twins.

Sin (Lee's word)	What it means	Modern echo
Hypercomprehensiveness	Tries to model everything at once	"One twin for the whole city"
Grossness	Too coarse to answer real questions	Averages that hide local effects
Hungriness	Needs more data than exists	Models that need data no city has
Wrongheadedness	Theory does not match how people really behave	Correlations treated as causes
Complicatedness	Outputs hard to interpret, and tweaked to look realistic	Black-box deep models
Mechanicalness	The machinery itself adds errors (Lee cites computer rounding)	Numerical and pipeline bugs nobody checks
Expensiveness	Costs too much for what it delivers	Twins nobody maintains

SAY IT LIKE THIS

"Lee's 1973 critique still applies: what decision does this twin help make, who can check it, and who will keep the data updated?"

Dilemmas in a General Theory of Planning

Horst Rittel and Melvin Webber · 1973

Planning problems are "wicked": no clean formulation, no clear solution.

The big idea

Simple ("tame") problems have a clear definition and a clear finish, like solving an equation. Social and urban problems are **wicked**. The authors list about ten features. The ones most worth remembering:

- **No definitive formulation.** Defining the problem is already part of choosing a solution.
- **No stopping rule.** You stop when you run out of time or money, not when it is solved.
- **Solutions are better or worse, not right or wrong.** Different stakeholders judge differently.
- **No immediate test.** Effects play out over years and ripple out.
- **Every attempt counts.** You cannot try and undo; the city changes as you act.
- **Every problem is unique** and a symptom of another problem.

TIP

When someone claims an AI model will "solve" congestion or housing, ask which of these wicked features it handles. It is a polite way to test a big claim.

Modelling Transport

Juan de Dios Ortúzar and Luis Willumsen · 5th edition, 2024

The standard textbook for how travel demand is modelled.

What is inside

- The **four-step model** in detail (see Chapter 3), with the maths of each step.
- **Discrete choice and random utility.** People pick the option with the highest "utility" (time, cost, comfort), with some randomness. The multinomial **logit** model is the workhorse.
- **Generalised cost.** Convert time, money and comfort into one number, using the **value of time**.
- **Revealed vs stated preference.** What people actually do versus what they say they would do in a survey.
- **Calibration and validation.** Fit parameters on one dataset, test on another. The same discipline AI papers need.

TIP

Read only the chapters on the four-step model and on discrete choice. That is enough to understand almost any mobility paper.

Introduction to Urban Science

Luis Bettencourt · 2021

A framework for treating cities as complex systems that can be studied like a science.

The big idea

The book (MIT Press, 2021) builds on social heterogeneity, complex networks, scaling, circular causality and information. Cities are networks of people interacting in space, supported by infrastructure. Because people interact, cities generate more of what they make (ideas, jobs, wages) as they get bigger, and use infrastructure more efficiently. This gives measurable regularities across cities and countries: the **scaling laws** in Chapter 7.

Ideas to borrow

- **Compare cities fairly.** Divide by population, and ask whether a city is above or below what its size predicts.
- **Cities as social networks in space.** The maths of interaction explains growth and innovation.
- **Theory plus data.** Data alone will not do; you need models of how cities work.

Criticism and limits

The scaling exponents vary with how a city is defined, and critics argue that some fits are weaker than the headline numbers suggest. Treat the exponents as approximate.

The New Science of Cities

Michael Batty · 2013

Cities are made of flows and networks, so measure interactions, not just places.

The big idea

Batty argues that a city is best seen as a system of **interactions**: flows of people, goods and information along networks. Older planning looked at places and land use; the new science looks at how they connect and change, using complexity theory, network analysis and big data.

Ideas to borrow

- **Bottom-up.** City patterns emerge from many small decisions, so use agent-based and cellular models.
- **Networks and flows.** Study connections, not just locations.
- **Real-time data and the "smart" city.** New sensors allow new measurement, and a need for theory to make sense of it.

REMEMBER

If you keep one thread from all these books: cities are **complex, interacting systems** (Jacobs, Alexander, Batty, Bettencourt), planners can rarely control them (Rittel and Webber), and big models fail when they forget this (Lee).

12

Housing, Land and Urban Economics

Why prices, rents and locations look the way they do.

A city is also a market. Land is scarce and fixed in place, so **location** is what people pay for. Even a technical urban AI project quickly runs into prices, rents and who can afford to live where. Here are the ideas you need.

Five ideas from urban economics

- **Bid-rent theory.** Different uses (shops, homes, warehouses) bid different prices for land depending on distance to the centre. Uses that gain most from a central location outbid the others there, which produces rings of land use.
- **Agglomeration economies.** Firms and workers gain from being close: sharing suppliers and infrastructure, matching jobs to workers, and learning from each other. This is a main reason cities exist and grow.
- **Externalities.** Costs or benefits that land on others, such as noise, congestion, or a park that lifts nearby values. Much of planning tries to correct them.
- **Housing supply and demand.** Prices rise when demand grows faster than the supply of homes. Rules that limit building (zoning, height limits, long approvals) make supply respond more slowly. How much this matters is debated.
- **Land value capture.** Public investment, such as a new metro station, raises nearby land values. Value capture recovers part of that gain through taxes, fees or development bonuses to pay for the investment.

Words about housing

Term	Meaning
Affordability	Often measured as price-to-income ratio, or rent burden (a common rule of thumb flags households spending more than about 30% of income on housing).
Tenure	Whether people own, rent privately, or live in social/public housing.
Gentrification	Rising costs and changing residents in a formerly lower-income area. Displacement is the main worry.
Hedonic pricing	A regression of prices on features (size, distance to station, noise, greenery) to estimate what each feature is worth.
Jeonse (전세)	A Korean lease type: the tenant pays a large refundable deposit instead of monthly rent. Very different from most countries.
Wolse (월세)	Monthly rent, often with a smaller deposit.

WORKED EXAMPLE (TOY NUMBERS)

Bid-rent by hand. Suppose retail can pay 100 per m² at the centre, falling 20 per km. Housing pays 60, falling 4 per km. Warehouses pay 30, falling 1 per km. Retail beats housing until $100 - 20d = 60 - 4d$, so $d = 2.5$ km. Housing beats warehouses until $60 - 4d = 30 - d$, so $d = 10$ km.

Result: **retail 0-2.5 km, housing 2.5-10 km, warehouses beyond.** Change any slope and the rings move. That is the whole idea.

Hedonic pricing. A model estimates that 10 more points of street-level greenery raises apartment prices by 2%, all else equal. On a 500 million won apartment that is about 10 million won. This is how street-view AI connects to housing economics.

WATCH OUT

Correlation is not cause here either. Expensive areas often have more trees because rich areas plant trees, not only the other way round. Hedonic studies must control for that.

TRY IT (10 MINUTES)

Choose two neighbourhoods you know: one with high rents, one with low. List five differences that could explain the gap (transport, jobs, schools, greenery, age of buildings). Which could a dataset measure?

13

Public Space and Urban Design

How streets and squares work for people, and how to observe them.

Urban design is the craft of shaping streets, squares and buildings so that public life works. It gives urban AI its most human questions: what makes a place feel safe, lively or pleasant? A few names and ideas cover most of the field.

Voice	Main idea
Jan Gehl <i>Life Between Buildings</i> (1971), <i>Cities for People</i> (2010)	Design for people at walking speed. Study how people actually use spaces, by counting who passes, who stays, and who is missing (children, older people).
William H. Whyte <i>The Social Life of Small Urban Spaces</i> (1980)	Filmed plazas with time-lapse cameras. Found that people prefer movable seating, sun, food, water and trees, and that people attract people.
Gordon Cullen <i>Townscape</i> (1961)	A city is experienced as a sequence of views. He called this "serial vision".
Ian McHarg <i>Design with Nature</i> (1969)	Overlay maps of soils, water, slope and habitat to find the best land use. The ancestor of GIS layer analysis.
Bill Hillier Space syntax	Measure how connected each street is in the network, and relate that to movement and activity.

Design ideas that recur

- **Human scale.** Buildings, doors and windows that relate to a person, not only to a car.
- **Active frontage.** Ground floors with shops, doors and windows facing the street, rather than blank walls.
- **Permeability.** How easily one can move through an area; many small links beat few large barriers.
- **Edge effect.** People like to sit and stand at the edges of spaces, not in the empty middle.
- **Public, semi-public, private.** Real spaces blend these; privately owned public space (POPS) raises questions of rights.
- **Placemaking and tactical urbanism.** Quick, cheap trials (paint, planters, temporary closures) to test an idea before building it.

WORKED EXAMPLE (TOY NUMBERS)

Doors per 100 m. A common street-life measure. Street A has 24 entrances along a 100 m stretch; Street B has 3. A is more likely to feel lively and offer many reasons to stop, and B is more likely to feel blank. The count is simple; interpreting it needs local knowledge.

Why it matters for AI. Perception models trained on street-view images (research on this includes MIT's Place Pulse project) try to predict how people rate places from images, and studies measure green view, sky view and enclosure. Ask: which design idea from this chapter is the model measuring?

TRY IT (10 MINUTES)

Sit in a public place for 10 minutes. Count people who pass and people who stay. Note where they sit and why. You have just run a small Gehl-style public life study.

14

Governance and the Planning Process

Who decides what gets built, and how.

Data and models do not decide what a city does. People and institutions do. Understanding the process tells you where an urban AI tool could actually be used, and where it will be ignored.

The cast

- **National and regional government:** laws, funding, big infrastructure.
- **City government and planning agencies:** plans, zoning, permits, services.
- **Developers and landowners:** build and finance most of the city.
- **Residents and community groups:** live with the results and often push back.
- **Utilities and transport operators:** deliver networks that shape growth.
- **Researchers and consultants:** supply evidence, models and tools.

Tools planners use

Instrument	Example
Regulation	Zoning, building codes, height limits.
Incentives	Extra floor area for providing public space; tax breaks.
Investment	Building a metro line, a park or a school.
Pricing	Congestion charges, parking prices, development fees.
Information	Maps, dashboards, campaigns, open data.
Land readjustment	Owners pool land, the public sector builds infrastructure, and owners get back smaller but more valuable plots. Used in Korea, Japan and Germany.

Ideas about the process

- **Survey, analysis, plan.** The classic cycle: understand the place, work out the options, choose and act, then monitor and revise.
- **Arnstein's ladder of participation** (1969). Eight rungs from manipulation up to citizen control. Asking "which rung is this?" reveals whether public engagement is real or for show.
- **Cost-benefit and multi-criteria analysis.** Two ways to compare options. Cost-benefit puts everything in money; multi-criteria weighs different goals openly.
- **Scenario planning.** Explore several possible futures instead of predicting one.
- **NIMBY and YIMBY.** "Not in my back yard" opposition to nearby development, and the opposite movement that supports building more homes.

WORKED EXAMPLE (TOY NUMBERS)

A cost-benefit check. A metro line costs 1,000 units. It brings 120 units of benefit a year for 20 years. At a 5% discount rate the benefits are worth $120 \times 12.46 \approx \mathbf{1,495}$, a benefit-cost ratio of about 1.5. At 8% they are worth about 1,178 (ratio 1.18). At 10%, about 1,022 (ratio 1.02).

The project looks good at low discount rates and marginal at 10%. Always ask which assumptions the answer is sensitive to. Numbers are invented for illustration.

SAY IT LIKE THIS

"Who would use this tool in the planning process, at which step, and what would they do differently because of it?"

TRY IT (10 MINUTES)

Find one recent planning decision in a city you know (a new station, a rezoning, a park). Write down who decided, who was consulted, and what evidence was cited.

15

Infrastructure, Sustainability and Climate

The systems that keep a city running, and the pressures on them.

Cities depend on networks that are mostly invisible until they fail. They also produce much of the world's emissions and face climate hazards. This chapter gives you the vocabulary.

Infrastructure systems

- **Water supply, wastewater and stormwater.** Drinking water, sewage, and drainage of rain.
- **Energy.** Electricity grids, gas and heat networks, and growing rooftop solar and batteries.
- **Waste and materials.** Collection, recycling and disposal.
- **Transport and telecoms.** Roads, transit, cables and mobile networks.
- **Interdependence.** A power cut stops pumps; a flood knocks out power. Failures cascade, so resilience thinking studies the links between systems.

Sustainability and climate words

Term	Meaning
SDG 11	The UN Sustainable Development Goal on sustainable cities and communities.
Mitigation vs adaptation	Mitigation cuts emissions. Adaptation prepares for the climate change that is already coming.
Territorial vs consumption-based emissions	Territorial counts what is emitted inside the city. Consumption-based counts what residents' purchases cause anywhere.
Embodied vs operational carbon	Embodied is from making and building materials. Operational is from running the building.
Risk = hazard, exposure, vulnerability	A flood is a hazard; how many people and assets are in its path is exposure; how badly they are hurt is vulnerability.
Sponge city, blue-green infrastructure	Using soil, plants, ponds and permeable surfaces to absorb rainwater.
Urban metabolism	Tracking the flows of energy, water and materials in and out of a city, as a body would.

WORKED EXAMPLE (TOY NUMBERS)

Grid decarbonisation. A district uses 200 GWh of electricity a year. With a grid emission factor of 0.45 kg CO₂ per kWh: $200,000,000 \text{ kWh} \times 0.45 = 90,000,000 \text{ kg} = \mathbf{90,000 \text{ t CO}_2}$. If the grid improves to 0.30, the same use gives **60,000 t**, a third less with no change in behaviour.

Per-person figure. A city that emits 5 million tonnes a year with 2.5 million residents is at 2 t per person. Check whether a paper means territorial or consumption-based before comparing.

BRIDGE

Building energy, heat and daylight are the point where infrastructure, climate and building expertise meet. Urban building energy models, heat maps and retrofit targeting are all practical areas for AI work.

TRY IT (10 MINUTES)

Find one city's published climate action plan. Which sector is largest (buildings, transport, industry)? What single measure does the plan lean on most?

16

GIS and Spatial Data: The Practical Basics

Enough to open a map file without breaking anything.

Almost every urban AI project handles geographic data. Most bugs come from a small set of basics: data types, coordinate systems and a handful of standard operations.

Two kinds of data

- **Vector.** Points (a bus stop), lines (a road) and polygons (a building footprint or district), each with a table of attributes.
- **Raster.** A grid of cells with values, such as a satellite image, elevation or temperature. Its **resolution** is the size of one cell.

Coordinate reference systems (CRS)

- A CRS says how numbers map to places on Earth. **Geographic** systems use latitude and longitude in degrees (WGS 84, code EPSG:4326, is the common one). **Projected** systems use metres on a flat map.
- Measure areas and distances in a projected CRS suited to the region. In Korea, national systems such as EPSG:5179 and the belt-based systems are widely used (check the right code for your data).
- The popular web-map system (EPSG:3857) distorts area, so avoid it for measuring.

Common formats and tools

Item	Notes
Shapefile	Old and widespread, but limited (short field names, several files).
GeoJSON, GeoPackage, GeoParquet	Modern vector formats. GeoParquet is compact and good for large data.
GeoTIFF / COG	Standard raster formats; Cloud Optimised GeoTIFF can be read in parts over the web.
QGIS	Free desktop GIS. The best place to look at data before writing code.
GeoPandas, Shapely, rasterio, PostGIS	Python and database tools for vector, raster and spatial queries.
OSMnx, momepy, PySAL	Street networks, urban morphology measures, and spatial statistics.

The operations you will use constantly

- **Buffer:** everything within a distance of a feature (all homes within 500 m of a station).

- **Clip and intersect:** keep only the part of one layer that falls in another.
- **Spatial join:** attach attributes from one layer to another by location.
- **Dissolve and overlay:** merge polygons, or combine layers to find where conditions coincide.
- **Nearest neighbour and routing:** distance to the closest feature, and travel time along a network.
- **Zonal statistics:** summarise a raster inside polygons (mean temperature per district).

WORKED EXAMPLE (TOY NUMBERS)

The degrees trap. At a latitude of about 36°N (central Korea), one degree of latitude is about 111 km, but one degree of longitude is only about $111 \times \cos 36^\circ \approx 90 \text{ km}$. A step of 0.001° is therefore about 111 m north-south and 90 m east-west. Computing distance in degrees gives distorted answers; project first.

Buffer and join. Three homes have 4, 1 and 0 transit stops within 500 m. Two out of three (67%) have at least one stop. This is a first accessibility indicator in three lines of code.

WATCH OUT

Mixed CRS is the classic silent bug: two layers that look aligned but are in different systems, or an empty result from a join. Always print the CRS of each layer.

Axis order: some tools want (latitude, longitude), others (longitude, latitude). Check with one known point.

Invalid geometry: self-intersecting polygons break operations. Repair before analysing.

TRY IT (10 MINUTES)

Open QGIS and load building footprints (from OpenStreetMap) for your neighbourhood. Draw a 500 m circle, project the layer, and compute total footprint area divided by circle area. That is a coverage ratio for the neighbourhood.

17

The Methods Toolkit

What people actually do with the data.

If you know machine learning, this chapter lines it up with the statistical and spatial methods that urban researchers use every day, and adds the ones that are specific to places.

Family	Methods to know	Typical use
Statistics	Descriptive statistics, regression (linear, logistic), hypothesis tests	Relate outcomes to features; test differences
Spatial statistics	Moran's I, local indicators (LISA), hot-spot analysis, spatial lag and error models, geographically weighted regression (GWR)	Is a pattern clustered? Do relationships change across space?
Clustering	k-means, DBSCAN / HDBSCAN, hierarchical	Group neighbourhoods or trips into types
Network analysis	Shortest path, centrality, betweenness, community detection	Street and transit structure
Simulation	Agent-based, microsimulation, system dynamics	Explore "what if" without experiments
Causal inference	Difference-in-differences, regression discontinuity, instrumental variables, synthetic control	Did a policy or design change cause the outcome?
Machine learning	Gradient boosting, CNNs, transformers, graph neural networks (GNNs), spatio-temporal forecasting	Prediction, perception, generation

Habits that separate good from weak work

- **Always a baseline.** For forecasting, "predict the same as last week" is a strong baseline. A complex model must beat it.
- **Spatial (or temporal) splits.** Split by region or by time, not at random, so the test resembles real use.
- **Watch for leakage.** Features that contain the answer, or neighbours that appear in both training and test.
- **Report uncertainty,** not only a single accuracy number.
- **Explain with care.** Feature-importance tools (such as SHAP) describe the model, not the world.
- **Make it reproducible.** Share code and data where allowed, and fix random seeds.

WORKED EXAMPLE (TOY NUMBERS)

Reading Moran's I. The value runs roughly from -1 to $+1$. Near **+1** means similar values cluster together (rich areas next to rich areas). Near **0** means no spatial pattern. Near **-1** means a checkerboard, with high next to low. A significant positive value is a warning that random splits will leak.

Difference-in-differences. Rents in an area that got a new station rise from 100 to 112. In a comparable area without one, rents rise from 100 to 105. The estimated effect of the station is $(112 - 100) - (105 - 100) = 7$. It assumes both areas would have followed the same trend without the station (the "parallel trends" assumption), so check that first.

SAY IT LIKE THIS

"What is the baseline, and how is the test set separated from the training set in space and time?"

TRY IT (10 MINUTES)

Take any city dataset with a value per district (income, temperature, rent). Compute Moran's I with PySAL, or just map the values and judge by eye whether high values cluster. Then say what that means for a random train/test split.

18

Images, Satellites and Digital Twins

Seeing the city from above, from the street, and in 3D.

Remote sensing in five lines

- **Optical and multispectral** sensors record reflected light in several bands. **Thermal** bands estimate surface temperature. **SAR** (radar) sees through cloud. **LiDAR** measures distance with laser pulses and gives heights.
- **Resolution has three faces:** spatial (pixel size), temporal (how often it revisits), spectral (how many bands).
- Common free sources: **Sentinel-2** (about 10 m pixels, revisit of days) and **Landsat** (about 30 m, a record going back decades).
- **NDVI** (vegetation index) = $(\text{NIR} - \text{Red}) \div (\text{NIR} + \text{Red})$. High values mean healthy vegetation.
- Remote sensing sees surfaces, not uses. A roof that looks the same can be a home or a factory.

Computer vision tasks in cities

Task	Meaning	City example
Classification	One label per image	Is this a commercial street?
Detection	Boxes around objects	Counting cars, bikes, pedestrians
Semantic segmentation	A label for every pixel	Share of sky, trees, road, building in a street-view image
Regression / scoring	A number from an image	Perceived safety score; building height
Vision-language models	Answer or describe in words	Describe a streetscape; query images with text

Segmentation models are commonly trained on street scene datasets such as Cityscapes and ADE20K. Street-level images come from services such as Google Street View and Mapillary; check the licence and terms of use before collecting or redistributing them.

3D city models and digital twins

- **Levels of detail (LOD).** A 3D city model can be block-like (low LOD) or include roofs, windows and interiors (high LOD).
- **CityGML** (an OGC standard) describes 3D city objects and their meaning. **CityJSON** is a compact JSON version. **IFC** (from buildingSMART) is the standard for BIM data at building scale. **3D Tiles** is a format for streaming large 3D scenes to a browser.

- **BIM and GIS** differ in scale, coordinates and meaning of objects. Joining them is a hard, real problem.
- A **digital twin** is more than a 3D model. It needs data that stays in sync with reality, a model that can predict or simulate, an interface people can use, and a decision loop.
- A common way to talk about maturity: **descriptive** (what is it?), **diagnostic** (why?), **predictive** (what will happen?), **prescriptive** (what should we do?), then **autonomous**. It is a way of thinking, not a formal standard.

WORKED EXAMPLE (TOY NUMBERS)

NDVI by hand. A pixel of grass reflects 0.50 in near-infrared and 0.10 in red: $\text{NDVI} = (0.50 - 0.10) \div (0.50 + 0.10) = \mathbf{0.67}$. A pixel of concrete reflects 0.25 and 0.20: $\text{NDVI} = 0.05 \div 0.45 = \mathbf{0.11}$.

Green view index. A segmentation model labels 3,000,000 of 10,000,000 pixels in a street-view image as vegetation: green view = 30%. Averaged over many images along a street it becomes a measure of street greenery. Ask which season and camera height the images had.

WATCH OUT

A pretty 3D model is not a digital twin. Ask what data updates it, what it predicts, and who acts on it.

TRY IT (10 MINUTES)

Open a free web map with satellite view and street-level images for one block. Decide which of the tasks in the table above you could run on it, and what each would tell you that the other could not.

19

The Field Map: People, Journals, Data and Tools

Where the conversation happens, and what to use.

Any young field has a small set of places where its ideas circulate. Knowing them saves months. The lists below are starting points, not rankings, and they are from general knowledge. Names exist, but check current details before relying on them.

Names to know (examples)

- **Urban computing:** the term popularised by Yu Zheng and colleagues (2014 survey in ACM Transactions on Intelligent Systems and Technology).
- **Labs with visible urban analytics or city science work:** MIT Senseable City Lab, MIT City Science, UCL CASA (founded by Michael Batty), the Urban Analytics Lab at NUS, and the Future Cities Lab in Singapore.
- **Open tools by researchers:** OSMnx (Geoff Boeing), momapy and PySAL (spatial and morphology tools).

Journals

- **Urban analytics and computing:** Computers, Environment and Urban Systems; Environment and Planning B; Urban Informatics; Nature Cities.
- **Planning and cities:** Cities; Urban Studies; Journal of the American Planning Association; Landscape and Urban Planning.
- **Transport:** the Transportation Research family (Parts A-F).
- **Buildings and environment:** Building and Environment; Energy and Buildings; Sustainable Cities and Society.

Conferences and workshops

- **Spatial and data:** ACM SIGSPATIAL; KDD and its urban computing workshops; GIScience and AGILE.
- **AI venues:** NeurIPS, ICML and ICLR, where urban and climate workshops appear.
- **Transport:** the TRB Annual Meeting; IEEE ITSC.
- **Planning and modelling:** CUPUM (computational urban planning and management).
- **Building simulation and design computing:** IBPSA Building Simulation, SimAUD, eCAADe, CAADRIA, ACADIA, BuildSys.

Datasets and tools

Type	Examples
Maps and buildings	OpenStreetMap; Overture Maps; Microsoft and Google building footprints; Global Human Settlement Layer (GHSL)
Imagery	Sentinel-2, Landsat; Mapillary and Google Street View (check terms)
Mobility	GTFS feeds; NYC taxi trip records; traffic benchmarks such as METR-LA and PEMS-BAY
Perception and vision	Place Pulse; Cityscapes; ADE20K
Analysis	QGIS, GeoPandas, PySAL, OSMnx, momepy
Transport and accessibility	r5r, OpenTripPlanner, SUMO, MATSim
Buildings and climate	EnergyPlus, Ladybug Tools, CityBES, UMI
Visualisation and twins	Kepler.gl, deck.gl, Cesium

Three habits for staying current

- **Follow three journals** or labs with alerts, not thirty.
- **Keep a paper log** with columns for task, data, scale, split, main finding, and what you doubt.
- **Reproduce one small result** per month. Nothing teaches a field faster.

TRY IT (10 MINUTES)

Choose one dataset and one tool from the table. In an hour, load the data with the tool and make one map or chart. Note every obstacle. That log is your first real lesson.

20

Ethics, Privacy and Law

Doing urban AI without harming the people in the data.

Urban data is data about people, even when it looks like it is about places. Trust and public support depend on handling it well.

Privacy

- **Location data is hard to anonymise.** A well-known study of mobile phone data found that a handful of location points was enough to single out most individuals. Removing names is not enough.
- **Ways to reduce risk:** aggregate to coarse areas or times, add noise (**differential privacy**), require that each record looks like several others (**k-anonymity**), and limit who can access raw data.
- **Laws.** The EU has GDPR; Korea has its Personal Information Protection Act (PIPA). Each sets rules on consent, use and transfer. Ask your institution about what applies before you collect or share personal data.

Fairness and power

- **Whose data, whose benefit.** Missing groups get worse service from models trained without them (see Chapter 6).
- **Surveillance creep.** Sensors and cameras installed for one purpose can be reused for another. Ask who can see the data and for how long.
- **Contested projects.** The Sidewalk Toronto smart-city project became a well-known case of concern about data governance and who controls public space.
- **Human in the loop.** Keep a person accountable for decisions that affect people, and make model limits clear to them.

Licences and rules

- **OpenStreetMap** data is under the Open Database Licence (ODbL), which requires attribution and has share-alike terms for derived databases.
- **Commercial imagery and map services** often restrict bulk downloading and redistribution. Read the terms before building a dataset.
- **Open science.** Where possible, share code and documentation, and state limits honestly.

WORKED EXAMPLE (TOY NUMBERS)

k-anonymity in one table. A "de-identified" trip dataset has three fields: age band, home district, work district. If the combination (65+, District A, District Z) appears only once, that person is easy to find. Requiring $k = 5$ means every combination must appear at least five times. Rare combinations must be merged into wider groups or removed.

SAY IT LIKE THIS

"What would happen to the people in this dataset if it leaked, and can we get the same result with less detailed data?"

TRY IT (10 MINUTES)

Take a dataset you know. List its fields. Which combination could identify a person? What is the coarsest version of the data that would still answer your question?

21

An Eight-Week Study Plan

Small steps. No burnout.

Weeks 1-2

Ideas, movement

Ch. 1-4, 11

Weeks 3-4

Climate, data, economy

Ch. 5-7, 12-13

Weeks 5-6

Governance, GIS, methods

Ch. 14-17

Weeks 7-8

Frontier, field, ethics

Ch. 8, 18-20

Week	What to do
1. Ideas	Chapters 1-2, then the Jacobs, Lynch, Alexander and Howard summaries in Chapter 11.
2. Moving and reaching	Chapters 3-4 and the Ortúzar summary. Build one isochrone with r5r or OSMnx.
3. Climate, data, science	Chapters 5-7 and the Lee, Rittel and Webber, Bettencourt and Batty summaries.
4. Money and design	Chapters 12-13. Do the ten-minute public-life study from Chapter 13.
5. Governance and climate	Chapters 14-15. Read one planning decision and one climate action plan.
6. Practical GIS and methods	Chapters 16-17. Do the QGIS exercise, then compute Moran's I on one dataset.
7. Frontier	Chapter 18 and Chapter 8. Read two recent surveys of urban computing or urban foundation models.
8. Field map	Chapters 19-20 and Korea in Chapter 9. Read a research group's recent papers, and reproduce one figure.

TIP

Best use of your time: read the recent papers of the group or field you are targeting **before** any of the above. They show which parts of this book matter most for you.

Optional: the original books (by title; find current editions yourself)

- Jane Jacobs, *The Death and Life of Great American Cities* (1961)
- Kevin Lynch, *The Image of the City* (1960)
- Christopher Alexander, "A City Is Not a Tree" (1965)
- Ortúzar and Willumsen, *Modelling Transport*
- Luis Bettencourt, *Introduction to Urban Science* (2021)
- Michael Batty, *The New Science of Cities* (2013)
- Douglass Lee, "Requiem for Large-Scale Models" (1973)
- Rittel and Webber, "Dilemmas in a General Theory of Planning" (1973)

Glossary A to Z

Tick a term the first time you hear it.

✓	Term	One line
☐	3D Tiles	Format for streaming large 3D scenes.
☐	Accessibility	How many opportunities you can reach in a given time or cost.
☐	Active frontage	Ground floors with doors and windows that face the street.
☐	Activity-based model	Travel model simulating each person's daily schedule.
☐	Adaptation	Preparing for climate impacts already coming.
☐	Agent-based model	Simulation of many simple actors to see system-level results.
☐	Agglomeration	Benefits firms and workers get from being close together.
☐	Arnstein's ladder	Eight rungs of public participation, from token to real control.
☐	Bid-rent theory	Uses bid for land by distance to the centre, forming rings.
☐	Buffer	Everything within a set distance of a feature.
☐	Building coverage ratio	Footprint ÷ plot area (건폐율).
☐	Cadastre	Official register of land parcels and ownership.
☐	Centrality	Network measure of how important a node or street is.
☐	CityGML / CityJSON	Standards for describing 3D city objects and their meaning.
☐	Compact city	Dense, mixed-use, transit-served urban form.
☐	Cost-benefit analysis	Comparing options by putting costs and benefits in money.
☐	CRS	Coordinate reference system: how numbers map to places on Earth.
☐	Difference-in-difference s	Compares change in a treated area with change in a control area.
☐	Differential privacy	Adding calibrated noise so individuals cannot be picked out.
☐	Digital twin	Data-linked model of a city or system that updates with reality.
☐	Discount rate	Rate that converts future benefits into today's value.
☐	Ecological fallacy	Assuming an area-level pattern holds for every individual in it.
☐	Embodied carbon	Emissions from making and building materials.
☐	Equity	Fairness in who benefits; more where need is greater.
☐	FAR	Total floor area ÷ plot area (용적률).
☐	First / last mile	Short leg between a station and the true trip end.
☐	Foundation model	Large pre-trained model adapted to many tasks.
☐	Garden City	Howard's planned self-contained town with a green belt.
☐	Gentrification	Rising costs and changing residents in a lower-income area.
☐	GeoAI	Machine learning applied to spatial data.
☐	Gravity model	Flow rises with size, falls with distance.
☐	GTFS	Open format for public-transport schedules.

✓	Term	One line
☐	GWR	Geographically weighted regression: relationships that vary by place.
☐	Hedonic pricing	Regression that values features of a home or place.
☐	IFC	Standard file format for building information models (BIM).
☐	Induced demand	New road capacity attracts extra traffic.
☐	Isochrone	Map of everything reachable within a set time.
☐	Jeonse	Korean lease with a large refundable deposit instead of rent.
☐	k-anonymity	Each record looks like at least k-1 others.
☐	Land cover	What the surface physically is.
☐	Land readjustment	Owners pool land; public sector builds infrastructure; owners get plots back.
☐	Land use	What people do on the land.
☐	Land value capture	Recovering part of the land value that public investment creates.
☐	Level of service	A–F grade of how well a road or facility works.
☐	LiDAR	Laser-based measurement of distance and height.
☐	LOD	Level of detail of a 3D model.
☐	MaaS	Mobility as a Service: one app across modes.
☐	MAUP	Results change when the area units change.
☐	Micromobility	Shared bikes, scooters and similar.
☐	Mitigation	Cutting emissions to limit climate change.
☐	Modal share	Share of trips by each mode.
☐	Mode choice	Step that picks walk, bus, car and so on.
☐	Moran's I	Measure of spatial autocorrelation, from about -1 to +1.
☐	Morphology	Pattern of streets, plots and buildings.
☐	NDVI	Vegetation index from near-infrared and red reflectance.
☐	Network science	Study of systems as graphs of nodes and links.
☐	NIMBY / YIMBY	Opposition to nearby building / support for building more homes.
☐	OD matrix	Trips between every pair of zones.
☐	OpenStreetMap	Open, volunteer-built map database.
☐	Placemaking	Improving public places, often through quick, low-cost trials.
☐	Rank-size rule (Zipf)	Regular pattern in city sizes.
☐	Raster	Grid of cells with values, such as a satellite image.
☐	Resilience	Ability to absorb and recover from shocks.
☐	Scaling law	Power-law link between city size and a quantity.
☐	SDG 11	UN goal for sustainable cities and communities.
☐	Sky view factor	Share of sky visible from a point.
☐	Smart city	Technology-led urban management.

✓	Term	One line
☐	Space syntax	Design-side analysis of street network structure.
☐	Spatial autocorrelation	Nearby places tend to be alike.
☐	Spatial cross-validation	Testing a model on regions it has not seen.
☐	Spatial join	Attach attributes from one layer to another by location.
☐	Spatial mismatch	Jobs far from those who need them.
☐	Street canyon	Street between tall buildings; described by H/W.
☐	Synthetic population	Artificial people matching census statistics.
☐	TOD	Transit-oriented development around stations.
☐	Trip assignment	Step that chooses routes.
☐	Trip distribution	Step that decides where trips go.
☐	Trip generation	Step that decides how many trips zones make.
☐	UBEM	Urban building energy modelling.
☐	Urban computing	Analysing city data to solve urban problems.
☐	Urban form	Physical shape of a city and its measurable properties.
☐	Urban heat island	Cities run warmer than their surroundings.
☐	Urban metabolism	Tracking flows of energy, water and materials through a city.
☐	Urban perception	Predicting how people feel about places from images.
☐	UTCI / PET	Outdoor thermal comfort indices.
☐	Vector	Points, lines and polygons with attributes.
☐	Walkability	How supportive an area is of walking.
☐	Wicked problem	No single correct formulation or end point.
☐	Zoning	Legal rules for permitted uses and limits.