Rethinking Binaries in Urban Digital Twins

A Speculative Design Approach

Master Thesis by
Isabella Maljers

Author

Isabella Chantal Maljers

Master Thesis

Technical University of Delft
Faculty of Industrial Design Engineering
MSc. Design for Interaction

Graduation Committee

Chair	Dr. Nazli Cila
Mentor	Dr. Himanshu Verma
2nd Mentor	M.A. Dmitry Muravyov

In Collaboration with

Advier	
Mentor	Eva Ruijter

July 2025

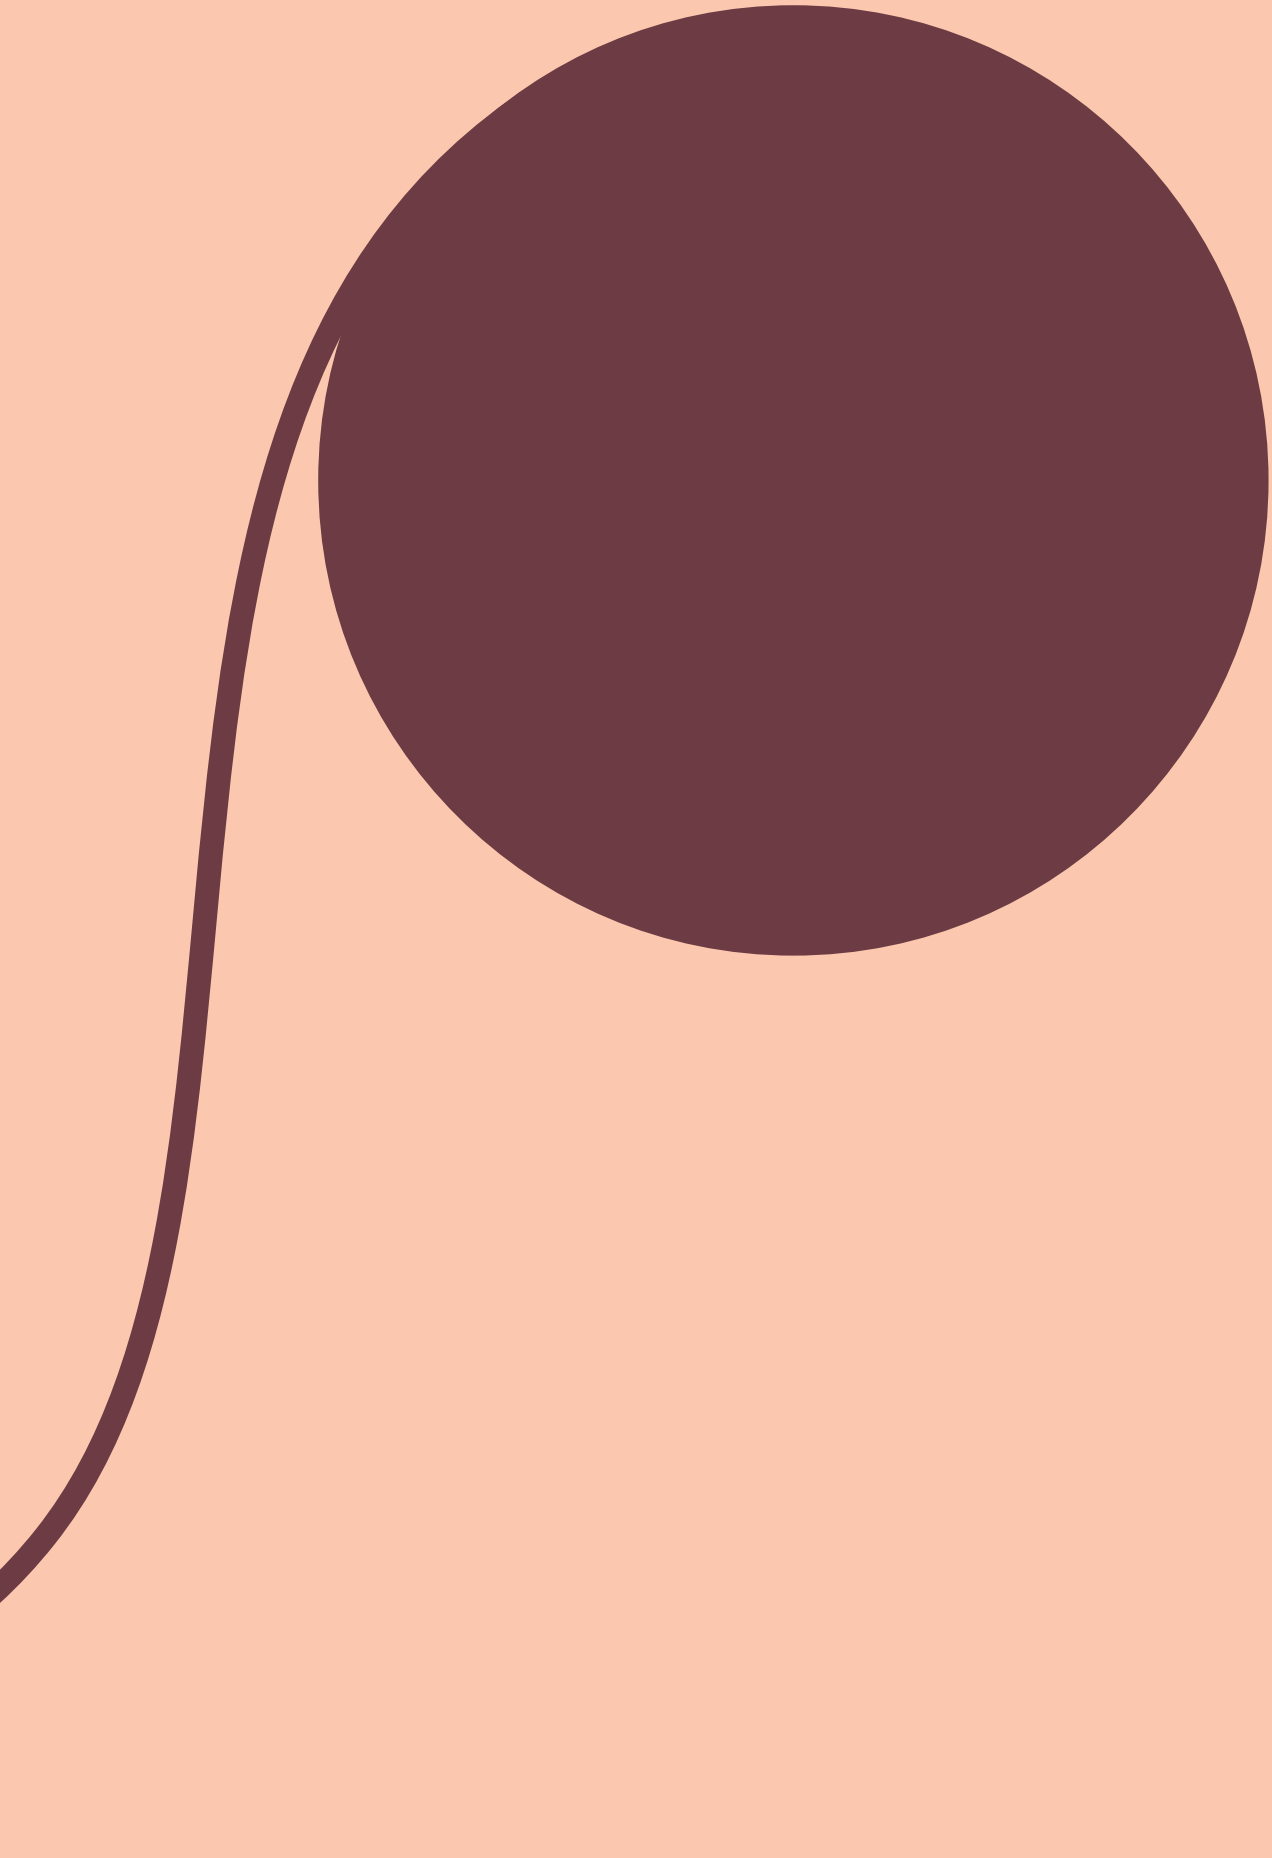

Executive Summary

This thesis investigates how optimisation logics embedded in Urban Digital Twins (UDTs) shape the tension between efficiency and inclusivity within the contemporary and future cities. To do this I question: What happens when optimisation becomes the driving logic behind urban inclusion? Who is at risk of not being represented within an urban algorithm? And how can speculative design help bring to light and open these tensions to surface critical discussion?

Drawing on theories of algorithmic bias, data feminism and design ambiguity, I analyse how UDTs rely on binary classifications (such as high vs low demand, and predictable vs unpredictable), and how these simplify complexity to shape reality while determining who is recognised and who remains unseen in the process. To explore these dynamics, I looked into the binary logics within a specific urban algorithm (the Advier Hub Algorithm). The shortcomings of these binary logics were then extrapolated into a context where a fully Autonomous Urban Digital Twin (AUDT) would have control of the city. This speculation led to the formation of a scenario on which the final speculative artefact, The Behavioural Mirror, was formed.

This artefact was created for within the municipal context to help start the conversation around algorithmic justice within the urban environment, specifically in relation to Urban Digital Twins. The Behavioural Mirror, invites users to encounter a fictional interface that calculates a “visibility score” based on the legibility of their behavioural patterns rather than on identity alone; through the scenario of Sem, a nurse whose irregular hours fall outside optimisation models, participants are prompted to confront how systems silently sort populations according to behavioural norms.

In my research through design process, I combined theoretical framing, design exploration and iterative testing with municipal professionals, mobility experts and designers; in facilitated sessions, participants used the mirror to provoke dialogue, policy reflection and ethical awareness. Within this process questions were raised about data invisibility’s complexity, the trade-offs between opting out and the limits of rigid logics in public systems.

In the facilitated evaluation sessions, I observed the Behavioural Mirror effectively prompted participants to question their own assumptions and discuss the hidden rules of optimisation. Municipal professionals and mobility experts engaged deeply with the video within the artefact, often remarked on their discomfort and curiosity. Discussions ranged from technical concerns about data completeness to ethical debates on agency and refusal. Many participants suggested that the mirror could serve as a primer in team workshops, policy labs and ethical training days by creating a shared reference point for challenging binary logics in Urban Digital Twins. Overall, the testing confirmed that speculative artefacts can be effective in opening space for critical dialogue.

Future work could explore artefacts that foreground the choice of invisibility. What if you would like to choose not to be visible? Aside from this, it could be interesting to iterate and develop this artefact further to test its and other artefacts' potential in influencing the urban decision making over time.

Glossary

Algorithmic bias: Systematic distortions in data or decision rules that privilege certain groups or patterns while disadvantaging others; often rooted in skewed inputs or opaque logic.

Algorithmic governance: The application of algorithms and data models to guide or automate policy, planning and decision-making in urban systems; shapes which interventions are prioritised and whose needs are addressed.

Algorithmic justice: The pursuit of fairness and accountability in algorithm-driven systems; involves identifying and mitigating biases, ensuring transparency of decision rules and providing recourse for those affected by automated decisions.

Ambiguity as method: The deliberate introduction of open-ended or uncertain elements in a design to invite interpretation, provoke reflection and surface hidden values rather than resolve complexity.

Binary classification: The reduction of complex or continuous phenomena into two discrete categories (for example "efficient" versus "inefficient"); useful for analysis but liable to obscure nuance and enforce rigid inclusion criteria.

Data feminism: An approach highlighting how power imbalances shape data collection, interpretation and display; advocates for practices that account for whose realities are counted and whose are excluded.

Data visibility: The degree to which particular data points or populations are made legible to a system; highlights that what's collected and displayed shapes who counts.

Legibility: What data infrastructures make visible and intelligible; the thresholds, categories and metrics that determine which actions and identities enter the model.

Research-through-Design (RtD): An inquiry approach in which iterative making and reflective practice generate knowledge; treats design artefacts as both outputs and vehicles for theoretical insight.

Speculative design: A practice that uses fictional scenarios and artefacts to question prevailing assumptions about emerging technologies; emphasises "what if" provocations over predictive or prescriptive outcomes.

Urban Digital Twin (UDT): A real-time, data-driven digital model of urban infrastructure, mobility and services used to support planning, monitoring and stakeholder engagement.

Preface

In the past few months I have had the chance to explore this topic within the intersection of speculative design; urban governance; and data justice in smart-city systems.

The formation of this project began with my interest in the ethical side of AI and a desire to probe beneath the hype, and to shine a critical light onto it. Soon I was drawn to data bias, and while discussing potential collaboration with Advier I stumbled on the issue of bias in digital twins. From that moment the rest is history.

I am so grateful to have pursued my graduation project on a topic that not only feels urgently relevant but has kept me fascinated throughout. During this journey I can honestly say I was overwhelmed at times, drowning in a sea of information and choices. Yet navigating that sea taught me the value of focus and iteration; each decision forced me to distil complexity into insight, and each setback revealed a new angle to pursue.

Nazli and Himanshu, thank you for your thoughtful guidance, insightful feedback and occasional curveballs that always set me back on track. Dmitry, I appreciate how you made me question every assumption, encouraged me to weave ambiguity into the artefact and helped me recognise its true value. Eva, thank you for your open mind, positive attitude and innovative ideas.

I could not have envisioned a better supervisory team; your expertise is unquestionable, but it is your kindness, thoughtfulness and unwavering support that I will carry forward most of all.

I also would like to thank all my friends and family who have supported me throughout these years. Your endless encouragement kept me grounded when I doubted myself. Your laughter and solidarity turned moments of frustration into renewed energy and made this project possible.

I hope my project is able to offer some provocation, and inspire others to examine the hidden optimisation logics behind urban digital twins, with the hope of pursuing more inclusive practices within smart-city design.

And so now I present to you my graduation report, I hope you enjoy the read!

Isabella

Table of Contents

Chapter 1: Introduction **11**

1.1 Project Motivation & Context 12
1.2 Research Question & Objectives 13
1.3 Project Partners and Stakeholders 14

Chapter 2: Methodology **17**

2.1 Research through Design Approach 18
2.2 Literature and Stakeholder Framing 18
2.3 Context Mapping 19
2.4 Speculative Design and Ambiguity as a Method 20
2.5 Prototyping & Evaluation 23
2.6 Overview of Research Activities 23

Chapter 3: Background & Theory **27**

3.1 Urban Digital Twins: Definitions & Levels 28
3.2 Algorithmic Bias & Power 40
3.3 Data Feminism 44
3.4 Ambiguity in HCI & Design Fiction Foundations 48

Chapter 4: Design Context & Framing **51**

4.1 Context Mapping with ViP 52
4.2 Efficiency and Inclusivity as a Design Lens 56
4.3 Advier Hub Algorithm Case Study 57

Chapter 5: Speculative Design Direction **61**

5.1 Autonomous Urban Digital Twin Speculation 62
5.2 Extrapolating the Hub Algorithm 64
5.3 From Six Binaries to Three Themes 67
5.4 Generating Speculative (what-if) questions 69

Chapter 6: Speculative Scenarios **71**

6.1 Three Scenarios 72
6.2 Testing the Scenarios 78
6.3 The New Scenario 83

Chapter 7: Ideation & the Provotype **87**

7.1 Design Goals & Strategy 88
7.2 Artefact Concept & Form 89

Chapter 8: Prototyping, Iterations & Evaluations **97**

8.1 The Final Design v1 98
8.2 Evaluation Round 1: methods & key findings 104
8.3 The Final Design v2 (iteration based on Evaluation 1) 108
8.4 Evaluation Round 2: methods & key findings 116
8.5 Updates for Standalone Use 120
8.6 How can this artefact be used by professionals? 123

Chapter 9: Discussion & Recommendations **125**

9.1 Linking Back to Literature 126
9.2 Project Limitations 130
9.3 Future recommendations 132

Chapter 10: Conclusion & Reflection **135**

10.1 Key Contributions & Reflection on the Research Question 136
10.2 Implications of the Project 137
10.3 Personal Reflections 138

References **140**

Appendix **144**

Chapter 1

INTRODUCTION

This opening chapter introduces the motivation and the context for the project. It lays out the research questions and objectives, and provides an overview of the report's structure. The work is also positioned within the broader debates on data-driven urban governance, highlighting the increasing role of Urban Digital Twins (UDTs) in the shaping of city life. A central tension is also framed as it will be explored in this project: the relationship between efficiency and inclusivity within optimisation logics.

The project was carried out in collaboration with Advier, a Dutch mobility and innovation consultancy and actively engaged in the development of urban planning models. Through this collaboration, as well as stakeholder engagements with municipal professionals from various municipalities, the project remained grounded in the practical realities of UDT development.

By outlining both the theoretical and practical aims of the research, this chapter establishes the foundations for the graduation journey that follows: from conceptual framing and systemic analysis, through speculative exploration and prototyping, to stakeholder engagement and critical reflection.

1.1 Project Motivation & Context

Urban Digital Twins (UDTs) are increasingly becoming integral components of contemporary urban governance systems. By modelling the city using data and algorithms, these platforms promise precision, optimisation, and scenario-based decision-making. They are utilised to simulate urban development, inform public policy, and guide infrastructure planning. However, as their influence grows, concerns arise regarding what is represented within these models, and how they may oversimplify the complexity of urban life.

UDTs are not neutral representations of reality. Like all data-driven systems, they are shaped by the assumptions, priorities, and omissions of their creators. These assumptions are often embedded in the datasets they rely on, the segmentation models they apply, and the binary decisions they are programmed to support. While many digital twin projects aim to enhance resilience and responsiveness, they risk reproducing exclusion through the very logics designed to optimise urban systems.

This tension is becoming particularly relevant today as cities worldwide, including those in the Netherlands, increasingly invest in integrated digital twin infrastructures in the name of optimisation. Initiatives such as the EU's Destination Earth programme, national smart city frameworks, and local mobility and energy transition projects are quickly advancing the use of UDTs in planning and governance. Alongside these developments, public debates on the ethical use of AI, algorithmic transparency and data-driven decision making are consistently growing, highlighting the importance of critical perspectives on these technologies.

Within this context, this project emerges from both an academic and an ethical motivation of examining how optimisation logics embedded in UDTs influence representations of the city and its inhabitants, and to explore how these systems may unintentionally marginalise or exclude certain behaviours, groups, and needs. By interrogating the intersection of efficiency and inclusivity within these systems, the project aims to contribute to a wider conversation about the future of data-driven urban governance. Rather than assuming an inherent clash, this project critically examines how current optimisation logics construct the relationship between efficiency and inclusivity.

The work is particularly relevant for municipal professionals, digital twin developers and urban policy makers who are now making foundational decisions about how these systems will be used, what values they will hold and whose voices they will actually represent. Rather than offering a solution, this research aims to provide a critical lens through which to reflect on current practices and future directions in UDT development.

1.2 Research Question & Objectives

This project investigates one of the central tensions emerging in Urban Digital Twins: how optimisation-driven models shape which aspects of urban life become visible, valued or excluded from decision-making. As UDTs increasingly start to inform urban planning and governance, their underlying logics (which are rooted in efficiency, predictability and quantification) risk marginalising the complex realities of those whose behaviours, values or needs align less with dominant data patterns.

Rather than aiming to improve or optimise existing UDTs, this project seeks to critically examine how optimisation logics shape urban representations and decision-making. The Advier hub algorithm (a behavioural segmentation model developed to support mobility hub planning) serves as a representative example of optimisation logics within a city. Through this lens, the project explores how categorisation, efficiency, and behavioural profiling influence the urban futures we design.

The project is guided by the following main research question:

> ***How do optimisation logics in Urban Digital Twins impact the relationship between efficiency and inclusivity, and what implications does this have with the field of urban governance?***

To explore this, the following sub-questions were formulated:

- In what ways do current optimisation practices within UDTs embed implicit biases and exclusions?
- How do current optimisation practices define efficiency, and what opportunities exist to foster inclusivity within these frameworks?
- How can speculative and ambiguous design interventions help municipal stakeholders critically reflect on the systemic impacts of optimisation?

The project also sets out to achieve the following objectives:

- to critically investigate the embedded assumptions in UDTs, particularly in relation to efficiency and inclusivity
- to create a speculative design artefact to surface and challenge these assumptions and their implications
- to provide municipal professionals and policymakers with tools for reflective engagement, enabling deeper critical examination of algorithmic governance practices

1.3 Project Partners and Stakeholders

The project was conducted in collaboration with Advier, a Delft-based consultancy specialising in the transformation of mobility systems, space and organisations toward sustainable futures. With a core mission of catalysing sustainable, inclusive and resilient urban transitions, Advier engages in service and system redesign, behaviour change facilitation and innovative hub and public transport solutions. Their experience in developing the “Advier Hub” algorithm, used to find the most preferred locations for mobility hubs, makes them a fitting partner for this project. The algorithm exemplifies the kind of behavioural segmentation and optimisation logics that are looked into within this research.

Within this project, Advier helped out in three distinct manners:

- Case Study provision: The “Advier Hub algorithm” was used as a lens for examining optimisation embedded in urban algorithmic logic.
- Co design input: through workshops and informal meetings, they contributed practical insights into how such algorithms are developed, evaluated and deployed in urban planning contexts
- Stakeholder testing: Advier staff participated in prototype evaluations , offering expert reflection elements presented by the artefact.

Additionally, the project’s evaluation phases and exploratory phases involved municipal professionals from Gemeente Amsterdam, Utrecht, Den Haag and Zwolle, who are all actively engaged in Urban Digital Twin initiatives. These participants provided valuable feedback based in real-world UDT deployment and governance experience and informed iterative refinement of the artefact’s framing and rhetoric.

Together, this blend of private-sector innovation and public sector contextual knowledge created a rich environment for evaluating how speculative design can surface underlying tensions within optimisation logic.

Chapter 2

METHODOLOGY

This project uses hybrid and exploratory Research through Design (RtD) approach that combines literature study, context-mapping and speculative design. The aim is not to develop a functional improvement to Urban Digital Twins (UDTs), but to critically expose the underlying classification logics and design assumptions that shape how urban populations are represented in algorithmic systems.

The process can be broken up into five core methodological components, which will be outlined in the next pages:

1. Research-through-Design Approach
2. Literature and Stakeholder Framing
3. Context Mapping
4. Speculative Design & Ambiguity as a Method
5. Prototyping & Evaluation

2.1 Research through Design Approach

This project treats the iterative design process itself as a means of generating knowledge. Rather than seeing the artefact as an end-product, it is used as a knowledge vehicle that surfaces new insights about data-driven optimisation in UDTs.

Following the Research through Design (RtD) principles (Stappers & Giaccardi, 2013), making and reflection occur in parallel, with each prototype iteration design to elicit stakeholder responses that could not have been anticipated through desk research alone. The artefact does not serve to validate hypotheses or demonstrate technical solutions; instead, it functions as a reflective provocation or a means to surface ethical, political and practical tensions embedded in optimisation-based governance.

Throughout the project, the artefact (or provotype - provocative prototype) functioned as a boundary object that connects critical theory (surrounding algorithmic bias, data feminism, ambiguity) with the emerging field practices of UDT development. It has allowed for insights to be generated that might not have emerged through literature study alone, by prompting situated reflections from those engaged in municipal digitalisation.

In this way, the project aligns with RtD's emphasis on the design-lead questioning; using iterative making as a means of creating knowledge, rather than just illustration or communication.

2.2 Literature and Stakeholder Framing

To ground the project's exploration, a combination of theoretical research with direct engagement in the field of UDTs was used.

A literature study was conducted around the socio-technical dynamics of Urban Digital Twins (UDTs), focussing on algorithmic bias, data feminism classification logistics, and ambiguity in computational systems. This helped bring to surface some concerns around categorisation, exclusion and optimisation driven governance within smart cities.

In addition to this academic research, I engaged with the field through stakeholder

interviews, expert presentations, as well as public events related to digital twin development. These included conversations with municipal staff from Gemeente Amsterdam, Gemeente Utrecht, and Den Haag among others who are involved in projects such as Twin4Resilience. These engagements provided insight into how UDTs are being developed, operationalised and interpreted within real-life settings.

One key text in this process was Data Feminism by D'Ignazio and Klein (2020), which served not only as critical theory but also as a guiding lens to identify recurring structures within urban data systems. While the book outlines seven principles for rethinking data practices, my focus was on how binary thinking and hierarchical categorisation shape the way urban populations are represented in algorithmic tools. Rather than applying the framework in its entirety, I used it to inform the transition into speculative design, particularly in highlighting how optimisation logics reduce complexity and exclude non-dominant perspectives.

In both the Background & Theory and context Mapping chapters of this report, insights from literature, events and interviews are combined. To aid transparency, literature sources remain in black, while **interview insights are marked in orange**, and **event insights appear in pink**. In addition to colour-coding the source type (e.g. interview, evaluation, event), each insight is tagged with a code (e.g. I1, E1.1, PM3) that corresponds to the full list of activities in Section 2.6.

2.3 Context Mapping

As a bridge between the Background & Theory research and design ideation I conducted an exploratory context mapping process.

To identify recurring tensions and future-relevant issues in the UDT landscape, I applied the context-mapping phase of the Vision in Product (ViP) framework. Typically the ViP method is used as a complete design methodology, however, in this case it was used selectively to explore broader system-level patterns rather than to directly generate design concepts.

The context mapping process drew from a wide range of inputs, synthesised within the literature and stakeholder engagement. These sources resulted in an array of context factors, which were then clustered thematically into speculative groupings such as "the Data privacy paradox", "Equitable Digitalisation" and "Just Human Scepticism". These clusters reflected multi-level tensions around the potential social and political consequences of urban digitalisation and algorithmic governance.

From this analysis, several design-relevant questions were derived. Among them, one question stood out and was selected to continue with:

> *What happens when efficiency clashes with inclusivity?*

While initially this questioned framed efficiency and inclusivity as oppositional forces, the project's development prompted a more nuanced view. Through discussions with supervisors and critical readings (such as from data feminist literature) it became clear that this apparent clash itself constructed by optimisation logics that frame efficiency in quantifiable terms. The project therefore does not treat efficiency and inclusivity as inherently incompatible, but rather questions how certain optimisation driven models privilege specific forms of efficiency that may unintentionally undermine inclusivity. This more critical framing informed both the speculative design direction and the way ambiguity was used to surface underlying assumptions.

2.4 Speculative Design and Ambiguity as a Method

With the central tension framed during context mapping, the project moved into a speculative design phase, using design fiction and ambiguity to critically examine how optimisation logics can shape urban inclusion. The aim is not to predict future developments, but to make visible the assumptions and potential consequences embedded in present technological trajectories (Auger, 2013).

An important nuance is that the tension between efficiency and inclusivity is not assumed to be a fixed opposition, but is seen as a dynamic relationship. Rather than assuming that efficiency inevitably leads to exclusion, this speculative process seeks to examine how current optimisation practices can produce exclusionary effects, while opening questions about whether more inclusive forms of optimisation might be possible.

The following two sections detail the key methodological components in this phase: Speculative design as a design approach, and ambiguity as a design strategy within the artefact.

2.4.1 Speculative Design

Speculative design is used in this project not to forecast or prescribe the future, but to provoke reflection on the present trajectories of Urban Digital Twins (UDTs). looking to Dunne and Raby's book "Speculative Everything (2013), the approach treats design as a medium for imagining alternative futures that expose assumptions and values which can be embedded in today's technologies and policies. Instead of aiming to solve problems, the speculative design poses "what if" questions, opening critical space to examine the political and ethical dimensions of innovation.

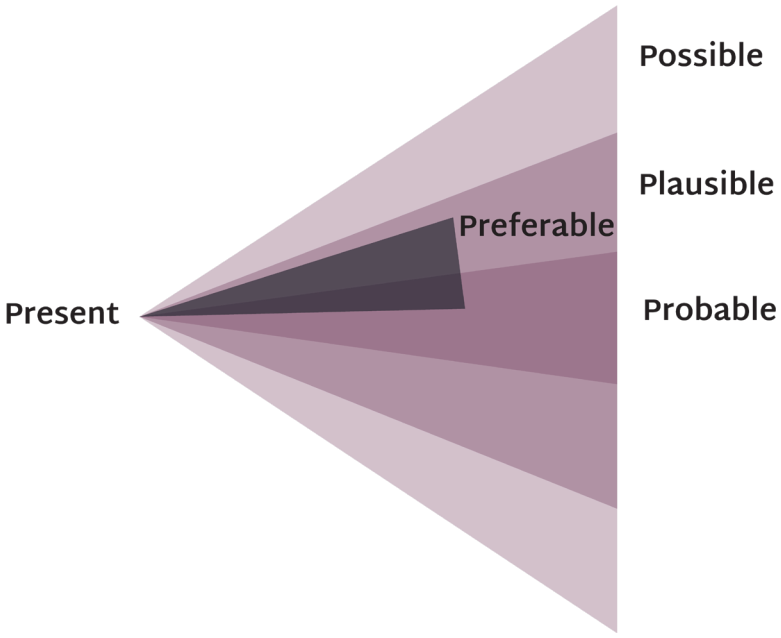

Figure 2.1 Possibility cone from Dunne & Raby (2013) Speculative Everything

As Dunne and Raby describe, speculative design exists on the "cusp between reality and the imaginary" helping designers and stakeholders to question the inevitability of current trajectories. It is concerned not only with imaging different futures, but also with shifting how people think about the present. Additionally, speculative design can be particularly valuable when addressing complex socio-technical systems where unintended consequences, power dynamics and assumptions often can remain hidden. In these contexts, this method can serve to disrupt the dominant narrative of technological progress to invite alternative imaginaries (Dunne & Raby, 2013).

Within the context of this project, speculative design was used to surface and interrogate how optimisation logics in UDTs influence which behaviours, people and places become visible or invisible to data-driven urban systems. This gives stakeholders a means to reflect on the social consequences of systems that quantify visibility and reward data legibility, prompting questions about whose needs are actually served within such models. This aligns with general speculative theory that speculative artefacts should not offer solutions, but rather act as "catalysts for conversation" (Dunne & Raby, 2013).

In the context of this Research through Design project, speculative design served a dual role: both as a tool for inquiry to make implicit assumptions visible, as well as as an engagement tool to foster critical reflection among stakeholders. The speculative lens also helped maintain a focus on systemic patterns rather than technical fixes, ensuring that the project could contribute to broader debates on data justice, inclusivity and ethics of optimisation in UDTs.

2.4.2 Ambiguity as a Design Strategy

In this context, ambiguity is deliberately used as a design strategy to provoke interpretation, discomfort, and reflection. Rather than presenting a polished or explanatory artefact, the goal is to construct an experience that opens space for doubt and questioning and in particularly around the ways optimisation logic shapes urban inclusion and exclusion.

This approach builds on work by (Gaver et al., 2003), (Benjamin et al., 2021), and (Giaccardi et al., 2024), who argue for the productive potential of uncertainty and ambiguity in interactive systems. Gaver argues specifically that ambiguity can enable an "interpretive flexibility" that encourages audiences to make sense of artefacts in personalised ways, often engaging them more deeply than conventional interfaces. Gaver also categorises ambiguity into three types: Ambiguity of information, ambiguity of context, and ambiguity of relationship. These can be used to evoke curiosity, to prompt critical reflection and to disrupt expected interpretation.

Rather than obscuring flaws, ambiguity in this project is used to surface uncertainty and expose systemic blind spots in UDT logic. This builds on (Giaccardi et al., 2024) who highlights that artefacts and systems of today are inherently uncertain, and that embracing this ambiguity can help reveal the limits of algorithmic logic.

Later in this report, the role of ambiguity in speculative design will be further elaborated on, however, it is important to note that ambiguity is integrated as an intentional and active design material throughout the design phase. It has allowed the artefact to resist offering easy solutions, instead allowing for deeper stakeholder engagement and reflection on the societal risks of optimisation.

2.5 Prototyping & Evaluation

Two iterative rounds of prototyping and evaluation were conducted with municipal professionals and digital twin stakeholders. Each cycle involved presenting the speculative artefact, gathering feedback and refining the design to sharpen its critical potential and communicative clarity.

The iterative nature of this process aligns with the RtD principles of reflection in action, using cycles of making and reflection to generate situated insights about the effects of optimisation logics on urban governance. The artefact was not treated as a fixed product, but rather one whose form and framing evolved through the stakeholder engagement. This allowed the project to generate both practical knowledge (about how municipal professionals interpret optimisation logics) and theoretical insights (about the role of speculative ambiguity in provoking critical reflection on algorithmic governance).

2.6 Overview of Research Activities

To support the development of the speculative artefact and theoretical framing, a variety of interviews, participatory sessions, evaluations, and expert events were conducted throughout the project. The following tables provide a structured overview of these research activities, grouped by type. This overview clarifies which perspectives and interactions informed the design decisions and evaluative reflections presented in later chapters.

2.6.1 Exploratory Interviews

Code	Date	Participant	Content
I1	29/11/2024	Information Manager, Gemeente Utrecht	Discussed architecture of the Netherlands 3D model, data layering, citizen modelling, ethical concerns, and future development trajectories.
I2	11/12/2024	Solutions Architect, Gemeente Amsterdam	Spoke about real-time data infrastructure, scenario simulation, long-term planning, privacy risks, and digital twin strategy shifts.
I3	17/12/2024	Policy Advisor, Living Labs, Gemeente Den Haag	Reflected on internal barriers, organisational complexity, and ethical tensions.
I4	28/01/2025	Digitalisation Mobility Manager, Gemeente AMS	Shared insights into data sources, accessibility tools, and ambiguity in the use of the term “Digital Twin.”
I5	31/01/2025	Data Analyst, Gemeente Amsterdam	Discussed the Tool Bereikbare Stad, stakeholder input, and modelling limitations.
I6	11/02/2025	Hub Algorithm Developer, Advier	Discussed the Hub Algorithm, current use cases, and future implications.

2.6.2 Evaluation Interviews

Code	Date	Participant	Content
E1.1	19/05/2025	2 Policy Advisors, Gemeente Zwolle	Evaluation 1
E1.2	22/05/2025	2 Advier advisors	Evaluation 1
E1.3	23/05/2025	2 Adults (no-design background)	Evaluation 1
E2.1-2.4	26-28/05/2025	4 Individual sessions with Advier advisors	Evaluation 2
E2.5	28/05/2025	Digitalisation Mobility Manager, Gemeente Amsterdam (I4)	Evaluation 2
E2.6	28/05/2025	Policy Advisor, Gemeente Den Haag (I3)	Evaluation 2
E2.7	30/05/2025	Information Manager, Gemeente Utrecht (I1)	Evaluation 2

2.6.3 Participatory Moments

Code	Date	Participant	Content
PM1	10/12/2024	3 Design Master students	Brainstorm on future urban AI concerns: surveillance, profiling, etc.
PM2	04/04/2025	2 Urbanism & Architecture students	Scenario testing (visibility, data exclusion)
PM3	07/04/2025	2 Design Master students	Scenario testing (behavioural exclusion, ambiguity)
PM4	09/04/2025	2 Advier colleagues	Scenario testing (binary amplification, policy relevance)

2.6.4 Attended Expert Events

Code	Date	Event	Content
AE1	30/09/2024	Urban AI & Digitalisation – Pakhuis de Zwijger	DCFA talk on digital urban futures and gamification
AE2	01/10/2024	AI & Big Data Expo – RAI Amsterdam	AI governance, security, ethical design frameworks
AE3	28/11/2024	DMI (Dutch Metropolitan Innovations) Birthday Conference	Urban innovation and digital governance
AE4	10/12/2024	Urban Digital Twins & Municipal Policy – BOLD Cities	Science-policy collaborations and UDT applications
AE5	03/03/2025	LIFE Arenapoort End Conference – Johan Cruyff Arena	Digital twin projects in energy, resilience, and infrastructure
AE6	14/04/2025	Salon Generative Things – ThingsCon	Workshop on provotypes and citizen interaction
AE7	03/06/2025	Integrated Area Development with Digital Twins – DMI	Critical reflections on integrated planning and UDTs

Chapter 3

BACKGROUND & THEORY

This chapter introduces the theoretical and contextual groundwork that informs the project's design lens. It begins by unpacking how Digital Twins emerged from NASA's product-lifecycle work, then tracks the term's spread across sectors where hype has often outpaced consensus. The discussion moves from generic Digital Twins to their urban counterparts, detailing the four-level maturity model and showing how real-world initiatives such as Netherlands 3D and Los Angeles traffic control expose both technical promise and political ambiguity.

With the technology situated, the chapter turns to the risks encoded in data practices and algorithmic modelling. It explains how incomplete datasets, method bias and wider societal inequalities can transform apparently neutral Urban Digital Twins into tools that overlook or exclude entire communities. Drawing on Data Feminism, the text challenges the simplifying binaries and hidden labour that underpin many smart-city projects, insisting that power relations are always present in digital infrastructures.

Finally, the chapter argues for ambiguity as a deliberate design material. Rather than treating uncertainty as a flaw, it is positioned as a means to reveal the limits of predictive models and to open space for more plural, reflective urban futures. By weaving these strands together, the chapter frames Urban Digital Twins as socio-technical systems whose impact depends on the values and assumptions embedded in their design and governance.

3.1 Urban Digital Twins: Definitions & Levels

3.1.1 Definitions and Evolution of Digital Twins

If you work in the technological innovation sector, you have probably heard of the Digital Twin (DT). The term "Digital Twin" was first coined by NASA in 2010, which evolved from the idea of product lifecycle management and resulted in an information mirroring model (Korotkova et al., 2023). This model aimed to simulate and monitor space systems remotely by maintaining a real-time digital counterpart.

Over the years, the term Digital Twin, has become quite an umbrella term with terms like digital shadow, 3D model and virtual model being used interchangeably (Hartmann et al., 2024). However, not all of these terms imply the same technical depth, especially when it comes to real-time data integration, feedback loops or control capabilities. The technology spans various industries such as manufacturing, healthcare and medicine, urban planning, construction etc (Rathore et al., 2021), furthering the complication of having one single consensus as to what it specifically constitutes.

It is generally agreed upon that a Digital Twin is a virtual representation of a physical/digital product or system. A clear definition is given by Ortt and Tiihonen, adapted for modernity from Grieves original definition given in 2002: "A digital twin (DT) is a tool for the development and control (both automatic and human) of base systems, including individual products, systems of multiple products, and multi-domain environments such as urban areas. Firstly, a DT consists of a detailed twin representation of a system. Secondly, data connections from the base product or system to the twin representation that can be used to transfer data that in turn can be used, for example, to conduct testing with the twin representation. Thirdly, data connections from the twin representation back to the base system that can be used, for example, for control." (Ortt & Tiihonen, 2022).

However, this definition still allows for open interpretation as to the level of bidirectional exchange of data, the type of data, the size and complexity, the purpose it serves, and the life span of the system (Li et al., 2021). This leads to the basis of the lack of consensus about the Digital Twin itself (Jones et al., 2020). In particular this lack of standardisation leads to a wide variability in how the term is used across academic, technical and policy domains.

Although the technology has been around for quite some years, in the recent years the technology has become one of the most discussed cutting-edge topics, being

named as one of the top 10 strategic tech trends for several consecutive years. It was even placed at the top of the Gartner hype cycle in 2017 and 2018 (Korotkova et al., 2023) (the Gartner hype curve is one of the most influential promissory organisations shaping the business of the tech expectations).

The hype status of the technology has led Digital Twins to be heavily invested in as a future technology that will help technological automation and advancement within various industries (Ortt & Tiihonen, 2022). However, the hype around the technology also means that many governments and non-government institutions want to be part of it without having a clear goal; in fact, "governments often invest in AI and Digital Twins as part of innovation strategies, even when practical outcomes remain uncertain" (Korotkova et al., 2023).

This diffusion of enthusiasm without shared clarity raises questions around Digital Twins raises key questions about what they really are, why they are implemented, how they are implemented, who benefits from them and what assumptions are embedded into their design. As such, digital twins should be examined not only as technological artefacts, but also as socio-technical systems shaped by governance structures, institutional values and epistemic worldviews.

3.1.2 Urban Digital Twins and Local Implementations

While the general concept of Digital Twins spans across industries, their application in urban contexts provides a unique set of challenges. Urban Digital Twins (UDTs) are virtual representations of physical assets or systems, which are being increasingly utilised in urban planning, smart city development, and infrastructure management (Faliagka et al., 2024; Ortt & Tiihonen, 2022). These advanced models enable real-time monitoring and simulation of various urban elements which include traffic management and environmental monitoring systems (Ali et al., 2023).

Unlike industrial based Digital Twins, UDTs are often layered, open-ended and inherently political. They aim to mirror the complexity of urban life while supporting strategic interventions. However, this can clash with practical constraints of reality such as data fragmentation, technical integration and political coordination. Within the scale of municipalities, Local urban Digital Twins (LDTs) are being developed through projects such as the Twin4Resilience project (an EU funded project), with the goal of creating an ecosystem of networked LDTs that help cities achieve facilitated learning and resource efficiency.

These Digital Twin systems allow municipalities to innovate and experiment without disturbing ongoing activities (Ortt & Tiihonen, 2022). Within the Netherlands, this technology has become quite popular with 37,5% of municipalities with over 100,000 citizens actively working on these projects (Ávila

Eça de Matos, 2023). However, it is important to note that different cities have different levels of maturity and strategic approaches to developing Urban Digital Twins (UDTs), resulting in fragmented ecosystems rather than single coherent city twins ([AE4]; van Apeldoorn et al., 2023).

In order to explain the Urban Digital Twin, I will first start with explaining the Digital Twin in general. As mentioned in the definition written in the introduction, the Digital Twin contains three parts:

1. The base system
2. The twin representation
3. The connections of data and information that ties the virtual and real systems together in both directions

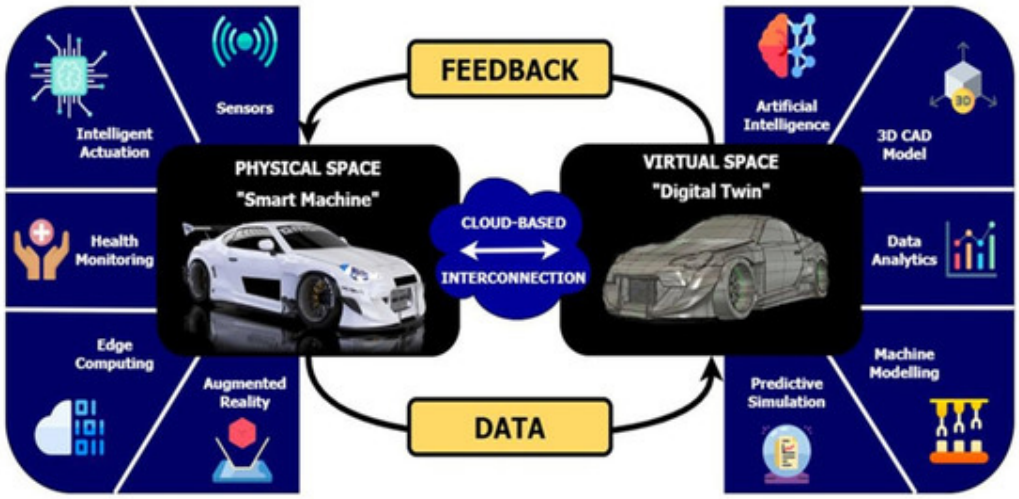

Figure 3.1 Digital twin of a car (Piromalis & Kantaros, 2022)

These components remain, but in the case of UDTs each becomes more complex: What counts as a "system" in a city? Who defines the representation? And how is data legitimacy established across diverse sources and actors?

As recent scholarship shows (Jones et al., 2020; van Apeldoorn et al., 2023), the Urban Digital Twin is as much a socio-political construct as a technological one. Its implementation demands not only technical capacity but also the negotiation between stakeholders, transparency around data flows as well as long-term strategies. In the following sections, I will draw upon how these tensions unfold in practice and how speculative design can help showcase their implications.

3.1.3 The layered nature of UDTs

The complexity of the digital twin can be described in 4 levels. For example, figure 3.1 shows a representation of a typical digital twin in the automotive industry. Within the case of the car, the detailed twin of the car is modelled before its production, allowing the car to be inspected and assessed to a high level of detail (level 1). If developed further, this model, virtual tests can be run optimising

design of various parts, simulating potential damage scenarios, and assessing the longevity of design in the long term with respect to safety, performance and comfort (level 2).

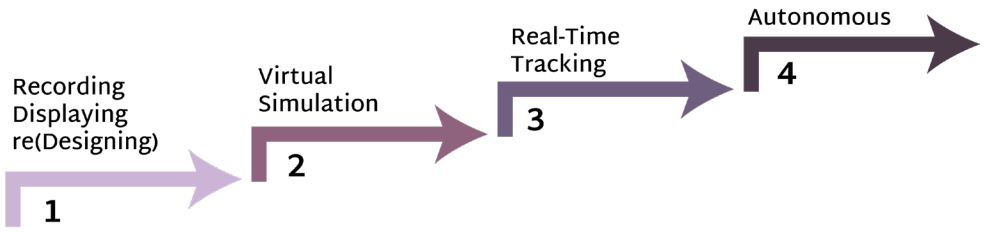

Figure 3.2 - Overview of the four levels of Digital Twins

After the physical car has been produced, sensors in the car can be integrated into the car linking to the virtual model to provide real-time insight into the car's condition, performance and environment (level 3). Lastly, a data link from the digital twin to the physical car allows the base to be automatically controlled and optimised remotely (level 4) (Ortt & Tiihonen, 2022). An overview of these levels can also be seen in figure 3.3 and 3.4.

Level	Description	Principle	Functionality	Subsystems	Example
1	There is a base system and a digital representation of that system. The two are not directly connected.	Virtual modelling	Recording, displaying, (re)designing	Data storage, computer, rendering software	CAD-representation of the design of a base system that can be developed or already exists and then can be easily redesigned.
2	There is a base system and a digital representation that simulates some key processes or mechanisms of that system. The two are not directly connected.	Virtual simulation	Testing, forecasting	Simulation software	Detailed twin representation of a base system with operational simulations that can be virtually tested and analysed.
3	There is a base system and a digital representation that simulates some key processes or mechanisms of that system. The two are connected one-way with information flowing from the base system to the virtual representation.	Digital commu-nication, sensors, AI	Tracking in real time	Data connectors, sensors, AI software	Detailed twin representation of a base system with operational simulations and real-time data analytics to optimise maintenance and aid human control.
4	There is a base system and a digital representation that simulates some key processes or mechanisms of that system. The two are connected two-way.	Controller technology, machine learning	Autonomous control in real time	Machine Learning software, controllers	Detailed twin representation of a base system with operational simulations, real-time data analytics, and control feedback to optimise performance in real time.

Figure 3.3 - Overview of the four levels of Digital Twins B (Ortt & Tiihonen, 2022)

This layered model does not only apply to cars. It has become a conceptual framework for understanding the maturity and functionality of Urban Digital Twins as well. In urban contexts, many digital twin initiatives still operate at level 1 or 2, relying on static models and scenario simulations.

3.1.4 Strategic Motivations and Political Ambiguity

One of my observations was that there are many different motivations behind the development of each UDT, specifically when it came to its commercial intentions. This intention also leads to a variance in the level of data that is visible to the user. Aside from the fact that this leads to many commercially developed UDTs not being visible to the general public, this also means that in certain situations the party asking for the technology to be developed for them (municipalities generally) do not have insight into the data being used to create the model. This also leads to an opacity in the deployment of the DT [AE3].

Such opacity raises questions about who owns the data and what assumptions are embedded in the system. While these systems are framed as neutral infrastructures they become shaped by institutional agendas and commercial logic. These layered motivations can make it difficult to clearly define their purpose or assess their broader societal impact.

3.1.5 Institutional challenges of UDTs

Within the Netherlands, many local governments are working on building their own Digital Twin. Partially because of the buzz term, local governments feel incentivised by the need to be part of the technological frontiers, sometimes without having a clear goal in mind. The DMI (Dutch Metropolitan Innovations) ecosystem, of which Advier is part of, is a collective of companies and some gemeentes who work to exchange knowledge and data. Companies pay to be part of this group.

Municipalities should try to temper their initial techno-optimism as UDTs cannot “solve” city challenges on their own. In these situations, often, external parties (such as TNO, Archeleo, Tygron, etc) are hired to create the model, but without a clear end goal, the initial results can fall somewhat disappointing [AE3]. Cities should try to clarify what UDTs are meant to achieve (whether its internal coordination, improved public services or more inclusive planning processes) and then structure data platforms and participation strategies accordingly ins [AE4]. Defining these ambitions early on can help to avoid digital twin implementation becoming pilots with no long term relevance.

The development of urban technology should start from actual pressing urban

challenges rather than introducing new technology just because it's available. By doing this, it ensures that technological deployments serve a real purpose and are better received by stakeholders [I3].

Additional challenges of municipalities are that their governance structures are traditionally siloed (eg separate departments for mobility, housing, utilities etc) which complicates the integration of UDTs across multiple domains (symposium bold cities). This is enforced by the fact that working within a municipal government is inherently complex. Even when good solutions exist, gaining traction from the right departments can be a challenge. Different municipal teams (eg, public participation vs enforcement) may have different priorities and comfort levels with tech. Aligning these (through communication, defining responsibilities and clarifying shared objectives) is required for any pilot or innovation project [I3].

3.1.6 UDTs in Practice: case studies

As part of my research I have compiled a few case studies of UDTs which are currently being developed, an overview of which can be found below.

You will find quite a diversity in the types of twins shown below. As mentioned previously, this is because their are some issues relating to the definitions of Digital twins, namely: there are many different definitions (Jones et al., 2020), some systems meet the definition of digital twins as given above but are not referred to as digital twins, and the definition does not match all digital twins that are described used and tested. This overview does not aim to address this issue, but rather should encompass some of the examples I have come across in my research.

Netherlands 3D (Twin4Resilience)

Figure 3.4 Screenshot from the interface shown in Netherlands 3D.

Netherlands 3D grew out of the separate Amsterdam-3D and Utrecht-3D viewers, yet now appears as a single, browser-based window on the whole country (see figure 3.4). The work sits inside the Interreg-funded Twin4Resilience programme, where six European municipalities compare notes on building Local Digital Twins and gently test new governance models. Locally the platform is stewarded by Team Beeld & Data and their Amsterdam counterparts, who emphasise that the underlying code is open source and available on GitHub for anyone willing to experiment.

At its core the twin stitches together three statutory datasets: 3D-BAG for buildings, 3D-BGT for terrain and public space, and the national AHN height model for relief. These layers are served as 300-metre 3D-tiles that stream in only when a user zooms close enough, keeping performance steady while allowing the camera to rise to 1 500 metres for a province-wide overview. The model is rendered in Unity and the team routinely feeds its C-sharp code through an AI assistant to catch minor bugs and suggest cleaner patterns.

Amsterdam, Utrecht, Flevoland or any other municipality can attach its own viewer, load extra layers and know that fresh BAG or AHN changes arrive automatically. Sensitive networks such as sewer mains or fibre-optic routes are held back in an internal edition, reflecting an explicit decision to balance transparency with security. Public layers, on the other hand, remain generous: users may dive below ground with the so-called dome view, toggle tree crowns, or overlay experimental datasets like Utrecht's Urban-Relief heat-stress points. Those points are deliberately aggregated in coarse hexagons so that no passer-by can be reverse-identified, a choice endorsed by the municipal data-protection officer.

At present Netherlands 3D does not stream live sensor feeds, although small pilots have shown how hourly or daily summaries could be ingested without flooding storage. For now the team classifies the twin as level 1–2 on the maturity scale introduced in Figure 3.3: geometry is rich, interaction is public, but the feedback loop with real-time infrastructure is still in its infancy. The intention, however, is clear: by offering a national canvas that any city can colour in, the project hopes to make Digital-Twin practice less proprietary, more comparable and, ultimately, more useful for the quiet work of planning streets, parks and entire neighbourhoods [I1].

Digital Twin for nature in Flevoland [I1]

The Digital Twin for nature, initiated in the province of Flevoland, emerges as a pivotal instrument for ecological advancement, focusing on the application of the '3-30-300 rule'. This rule indicates that from each residence 3 trees should be visible, 30% of the neighbourhood should be covered by trees and green space, and no one should live more than 300 meters from the nearest park or green space (see figure 3.5).

Built around provincial tree inventories, land-use parcels, species records and

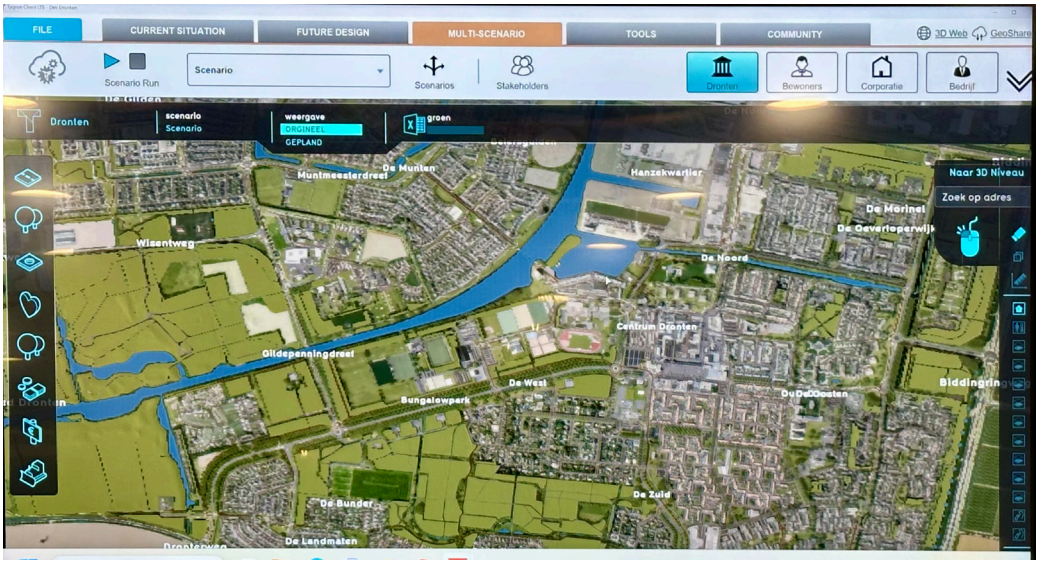

Figure 3.5 -The interface for Flevoland's Digital Twin for nature **[AE3]**.

a handful of environmental layers, the model is intentionally lightweight and currently falls in level 1–2 territory. What makes the project stand out is its nine-month field-lab cycle. After each round the team folds new data, visual tweaks and on-the-ground feedback into the next version, creating a steady rhythm of experiment, reflection and refinement.

Even in its early iterations the twin has guided new street-tree plantings in Lelystad, tested resilient species mixes along windy boulevards and explored tiny wetland pockets that could link urban habitats to the larger Oostvaardersplassen reserve. Future cycles aim to add simple scenario sliders so residents and ecologists alike can explore, for example, how widening a hedgerow might nudge canopy coverage toward that magic thirty-percent threshold.

Amsterdam

Within Gemeente Amsterdam the team prefers not to use the label "digital twin", yet the Tool Bereikbare Stad matches most textbook definitions (figure 3.6). It started life as an internal pilot and is slowly growing into a city-wide decision-support platform for planologists who need quick insight into how easily residents can reach schools, parks, health care and daily shopping.

The model begins with a quilt of hexagonal cells that cover the full Metropolitan Region. For each cell a single, realistic departure point is calculated from BAG address clusters so the journeys start where people actually live or work. From every point the project has pre-computed travel-time isochrones for walking, cycling, car and public transport. The routing engine accounts for ferry crossings, slower cycling in Amsterdam traffic, rush-hour speed penalties for cars and full GTFS timetables for trams, buses and trains, including a fifteen-minute departure window so missed connections do not skew the results.

Accessibility scores are calculated with the LISA establishments database. Each

isochrone is placed over the LISA points, the reachable facilities are counted by category and that total is written back to the hex-grid. Inside the Power BI dashboard users can filter by municipality, facility type, transport mode, travel-time band and even grid resolution. Colours update instantly giving a clear picture

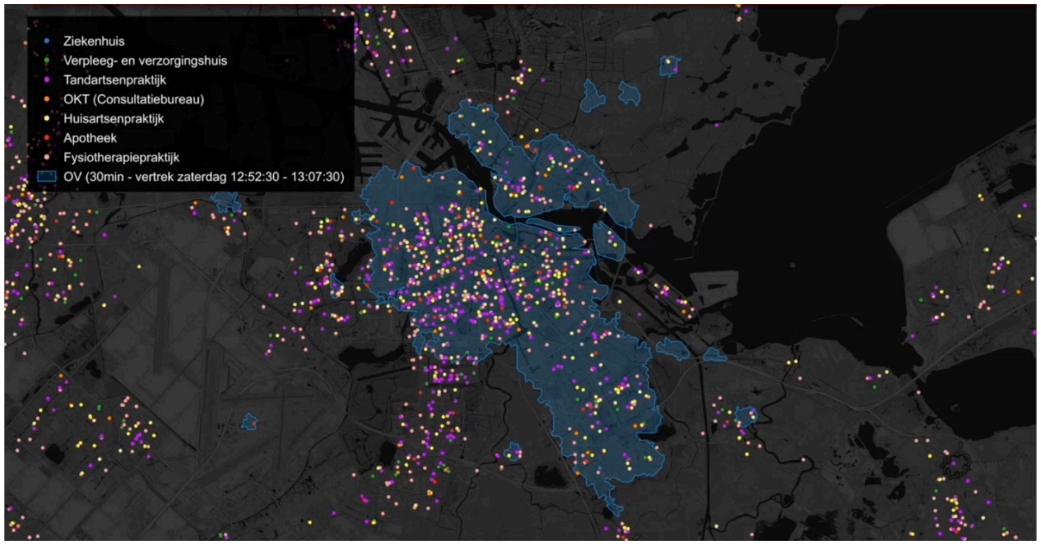

Figure 3.6 -The tool bereikbare stad [I5]

of where provision is strong and where gaps remain. An extra toggle shows how the picture shifts between a busy Monday morning and a quiet Sunday afternoon.

Although still in its first public prototype the tool is already guiding conversations about new school locations, proposed ferry timetables and recently completed neighbourhoods that lack playgrounds. The next iteration will combine the accessibility layer with socio-economic indicators so planners can identify, for instance, low-income areas that face the longest trips to childcare.

Built mainly from 3D BAG, 3D BGT, AHN and LISA, the Tool Bereikbare Stad sits somewhere between level 1 and level 2 of the maturity table, providing a dependable mirror of current conditions without yet dipping into live sensor streams [I4];[I5].

Cambridge DT

The “Cambridge Digital Twin” that is often cited in recent literature is not a single, all-encompassing platform. Rather, it is a collection of pilot twins being explored by researchers at the University of Cambridge’s Centre for Digital Built Britain (CDBB) together with Cambridgeshire & Peterborough Combined Authority. Their goal is to test how a city-scale model can help with flooding, transport and land-use questions while feeding lessons into the wider UK National Digital Twin programme (CDBB, 2020). Early work began with semi-structured interviews to surface the most pressing local issues; those conversations pointed to storm-water management, bus reliability and active-travel safety. Each theme was then

translated into a lightweight prototype, stitched together from the county's LiDAR, Ordnance Survey MasterMap, bus Automatic Vehicle Location feeds and the Environment Agency's flood layers (Caprari et al., 2022).

Even at this experimental stage, the team insist that the model should never be treated as "authoritative truth". Instead, planners open the twin during design workshops, run a scenario ("what happens if we divert a new busway through Mill Road?") and use the visual output as a talking point with residents. Because most data arrive as periodic snapshots rather than live streams, Cambridge's twin currently fits somewhere between level 1 and level 2 in the maturity scale shown in figure 3.3.

Singapore

Virtual Singapore is described by the Singapore Land Authority, GovTech and the National Research Foundation as the island's shared geospatial language, a cloud-based 3-D model that ministries, universities and private firms can open, annotate and query (Caprari et al., 2022). Airborne LiDAR, oblique imagery and the BIM files that developers must submit supply the geometry; policy layers such as façade materials, energy baselines and ventilation corridors turn the mesh into something planners can act on (see figure 3.7).

Caprari et al. note that most feeds arrive in daily or weekly snapshots rather than live streams, which means the production build sits between level 1 and level 2

Figure 3.7 -Virtual Singapore (Caprari et al., 2022)

in the maturity scale. During design workshops planners open the twin, launch a quick scenario and use the visual output as a conversation starter with residents and other agencies, always reminding users that the model is a decision-support canvas rather than an oracle (Caprari et al., 2022).

3.1.7 Current Ethical Considerations of Municipalities

When collecting data, many municipalities tend to stick to the cautious side rather than risk crossing the boundaries of public privacy. For example, within the municipality of Amsterdam, it was chosen only to hang cameras at a height where faces would not be discernible. Moreover, within the Gemeente Den Haag, they chose not to use any form of cameras whatsoever for their Living Labs in Scheveningen. This cautiousness is praiseworthy, but raises critical questions: can a digital model be accurate if key behavioural data is purposefully excluded? And who decides what kind of data is too sensitive to model?

Although the Netherlands 3D is mostly OpenSource and publicly available, certain aspects of data were already limited in the output of the 3D model. For example, the sewage planning and information which goes along with it are omitted from this model as the makers were worried about making this so easily available to everyone who wishes. This information is still publicly available in public databases, however by integrating it into such an accessible model, the information becomes extremely easy to access by those who wish to wrong.

Together these cases show that Urban Digital Twins are not singular, all-knowing systems (yet), but rather evolving concoctions of data, politics, technical infrastructures and social values. Their potential lies not only in what they model, but also in how they clearly communicate limitations, assumptions, and intended outcomes.

3.1.8 Opportunities and challenges for Urban Digital Twins

Opportunities

Urban Digital Twins (UDTs) offer a range of promising opportunities for cities wanting to better understand, plan and manage their complex systems. One of the most frequently mentioned advantages is their capacity to simulate the effects of policy decisions before their physical implementation. By visualising interconnected cause-effect relationships across different domains such as mobility, green infrastructure, or energy use, UDTs allow both policymakers and citizens to better grasp the of different choices (Faliagka et al., 2024; Marcucci et al., 2020). The enhanced visualisation formats of UDTs, often include interactive tools or animated scenario models, which help demystify technical information and make trade-offs more accessible to non-expert stakeholders and thereby also strengthening both decision making processes and participatory engagement [AE3].

As shown in a pilot project in Greece (Faliagka et al., 2024), UDTs have enabled the optimisation of smart parking policies, integrating real-time parking detection with multi-modal transport planning and what-if scenario testing. Their simulation capacities also extend to environmental impact analyses by correlating traffic with noise, emissions and air quality which helps support more sustainable urban strategies. Similarly, UDTs can support emergency response infrastructure by detecting traffic accidents in real time and suggesting re-routing strategies for emergency vehicles. These digital environments reduce the risks and costs that usually come with live trials and so a greater scale of experimentation is allowed especially within politically or financially sensitive areas.

Challenges

Nevertheless, UDTS also present significant challenges as well. Many of these are technical such as the need for high-quality, real time data-streams, which can be difficult to gather and maintain across multiple systems (Hassija et al., 2024; Marcucci et al., 2020). Interoperability across platforms and standards remains a persistent bottleneck (Rathore et al., 2021), especially in municipal environments, where different departments use different tools with different goals. The scalability of the model and accuracy, especially in larger urban settings, create ongoing barriers, especially as models become more complex and data-heavy (Hassija et al., 2024). Additionally, ethical concerns also rise for example in relation to privacy. In contexts where AI-driven twins touch on surveillance or health applications, careful governance is required. As UDTs become more advanced, there is a growing concern surrounding the displacement of human labour in planning and operational roles. As (Hassija et al., 2024) note, “automation through Digital Twins is streamlining processes but also raising concerns about job displacement in various sectors”.

Several interviewees in this research also emphasised the risks of premature procurement or overreliance on private companies. Municipalities often are not digital experts and so may not realise the extent to which they may be locked into specific vendor ecosystems, limiting future flexibility and fostering dependency (Caprari et al., 2022). This concern is also echoed by (Matos, 2023) who points to a range of challenges municipalities face such as collecting reliable and timely data to integrating DTs with existing workflows. In many cases the difficulty is not technical but also strategic, finding relevant use cases, avoiding duplication and ensuring that the DTs support real urban needs rather than just serving as a technological showcase.

Finally, there is also the growing concern that the persuasive visual power of UDTs can result in a false sense of objectivity. When complex social dynamics are reduced to simplifies simulation outputs, there is a risk that predictions are treated as certainties and so that accountability becomes blurred. If unintended consequences arise from over-reliance, who will be held responsible?

3.2 Algorithmic Bias & Power

This section examines how bias operates at both the data and algorithmic levels within data-driven systems. In the context of Urban Digital Twins, data bias and method bias can shape how cities are modelled, who is visible within the system, and which perspectives are excluded. These dynamics are especially important when UDTs are used for high stakes activities such as behavioural segmentation, planning, and policy evaluation, where incomplete or reductive data can have real-world consequences.

3.2.1 Data Bias in Urban Digital Twins

Datasets often reflect the long standing structural inequalities in our society. As is the case with all technologies based on data, urban digital twins are only as good as the data quality and the relevance of their input data. Errors or biases in data can cascade through the system, leading to inaccurate or misleading outputs.

As we are realising a shift in the role and purpose of data in society, while it transitions from a tool for analysis and monitoring to becoming the foundation for decision-making, prediction, and action. Therefore, it is important to address the ethical dimensions that arise with this transformation, particularly with regards the risks of data completeness, bias, and data quality (Filho et al., 2022).

While Urban Digital Twins promise precision and optimisation, they rely fundamentally on the data that is fed to them; however, data is never neutral. As D'Ignazio and Klein emphasise, "data are not raw; they are cooked with care", data is shaped by who collects it, how it is collected, and with what purpose (D'Ignazio & Klein, 2020). In the context of UDT, data often reflects the historical inequalities and infrastructural blind spots which can be particularly problematic when these datasets become the basis for urban predictive modelling.

Data bias is not only a repercussion of technical oversight, but it can also arise from social and cultural factors that influence data collection and representation. Certain groups may be underrepresented, with their lived experiences and perspectives missing from the data. Furthermore, data from commercial parties such as Google maps or third-party mobility apps might overrepresent the behaviours of those who are more digitally connected (usually high-income groups) while under-representing marginalized communities (Sánchez–Vaquerizo, 2024).

This selective visibility can then lead to what is described as urban invisibility where certain populations are excluded from the digital layer of the city simply because they have not been measured or are to complex to model and/or quantify (Benjamin et al., 2021; Muravyov, 2022). Within such contexts, it can be said that the Urban Digital Twin does not mirror reality, but rather only reinforces a partial, biased and efficiency-oriented version of it.

3.2.2 Algorithmic bias: Beyond data

In their paper, Akter et al. (2021) use thematic analysis to identify three main causes of bias within algorithms. These include data bias, method bias and societal bias (see figure 3.8):

Data bias: arises from skewed, incomplete, or unrepresentative training data, which can lead to algorithms that perform poorly for certain demographic groups or geographic locations.

Method bias: results from the design and implementation of algorithms, where certain choices or assumptions may inadvertently favour specific outcomes or reinforce existing inequalities.

Societal bias: refers to the broader social and historical biases that are embedded in data and algorithms, reflecting systemic discrimination or prejudice (Akter et al., 2021). These biases can perpetuate and amplify inequalities, leading to unfair or discriminatory outcomes in various domains, including urban planning (Calvi, 2022; Sanchez et al., 2024).

These three types of bias are not isolated; they often interact and reinforce each other. An unbalanced dataset may reflect societal inequalities, and the algorithm trained on that data may further institutionalise those patterns through optimisation. Understanding these layers is essential when working with Urban Digital Twins, as even technically accurate models can silently reproduce systemic exclusions.

Figure 3.8 -Bias types

3.2.3 The politics of Automated Governance

There is a broad societal tendency to place trust in the objectivity purported by quantified methods and the presumed independence and integrity of institutions that employ such approaches. It is therefore imperative to critically examine the underlying assumptions, biases, and potential limitations of AI to prevent the uncritical acceptance of their outputs (Muravyov, 2022). When algorithmic outputs are accepted at their face value, without scrutiny of their underlying assumptions, the can introduce or exacerbate injustice in the semblance of neutrality. Additionally, when such systems are deployed in scale, they risk amplifying and perpetuating these issues. (Calvi, 2022) (Sánchez–Vaquerizo, 2024) This occurs as a result of algorithmic tools excluding areas that were never represented well in the data to begin with, leading to the structural risk of automated exclusion. This is especially important to consider in critical domains such as housing, mobility or urban resource distribution.

In the context of Urban Digital Twins, this means that algorithmic bias is not an isolated technical concern but a governance challenge. As Giaccardi et al argue: algorithmic system are increasingly probabilistic and agentive; They are designed to learn, predict, and act within uncertain contexts. These systems have the potential to participate in a decentralised making process through which the urban landscape is constantly being configured. We must critically examine their outputs in order to avoid them becoming self-reinforcing feedback loops where optimisation for performance sidelines human difference.

While understanding bias at the technical and systemic level is critical, it is also necessary to examine how power, representation, and exclusion are embedded in data practices themselves. The next section uses Data Feminism as a framework to reflect more deeply on these issues.

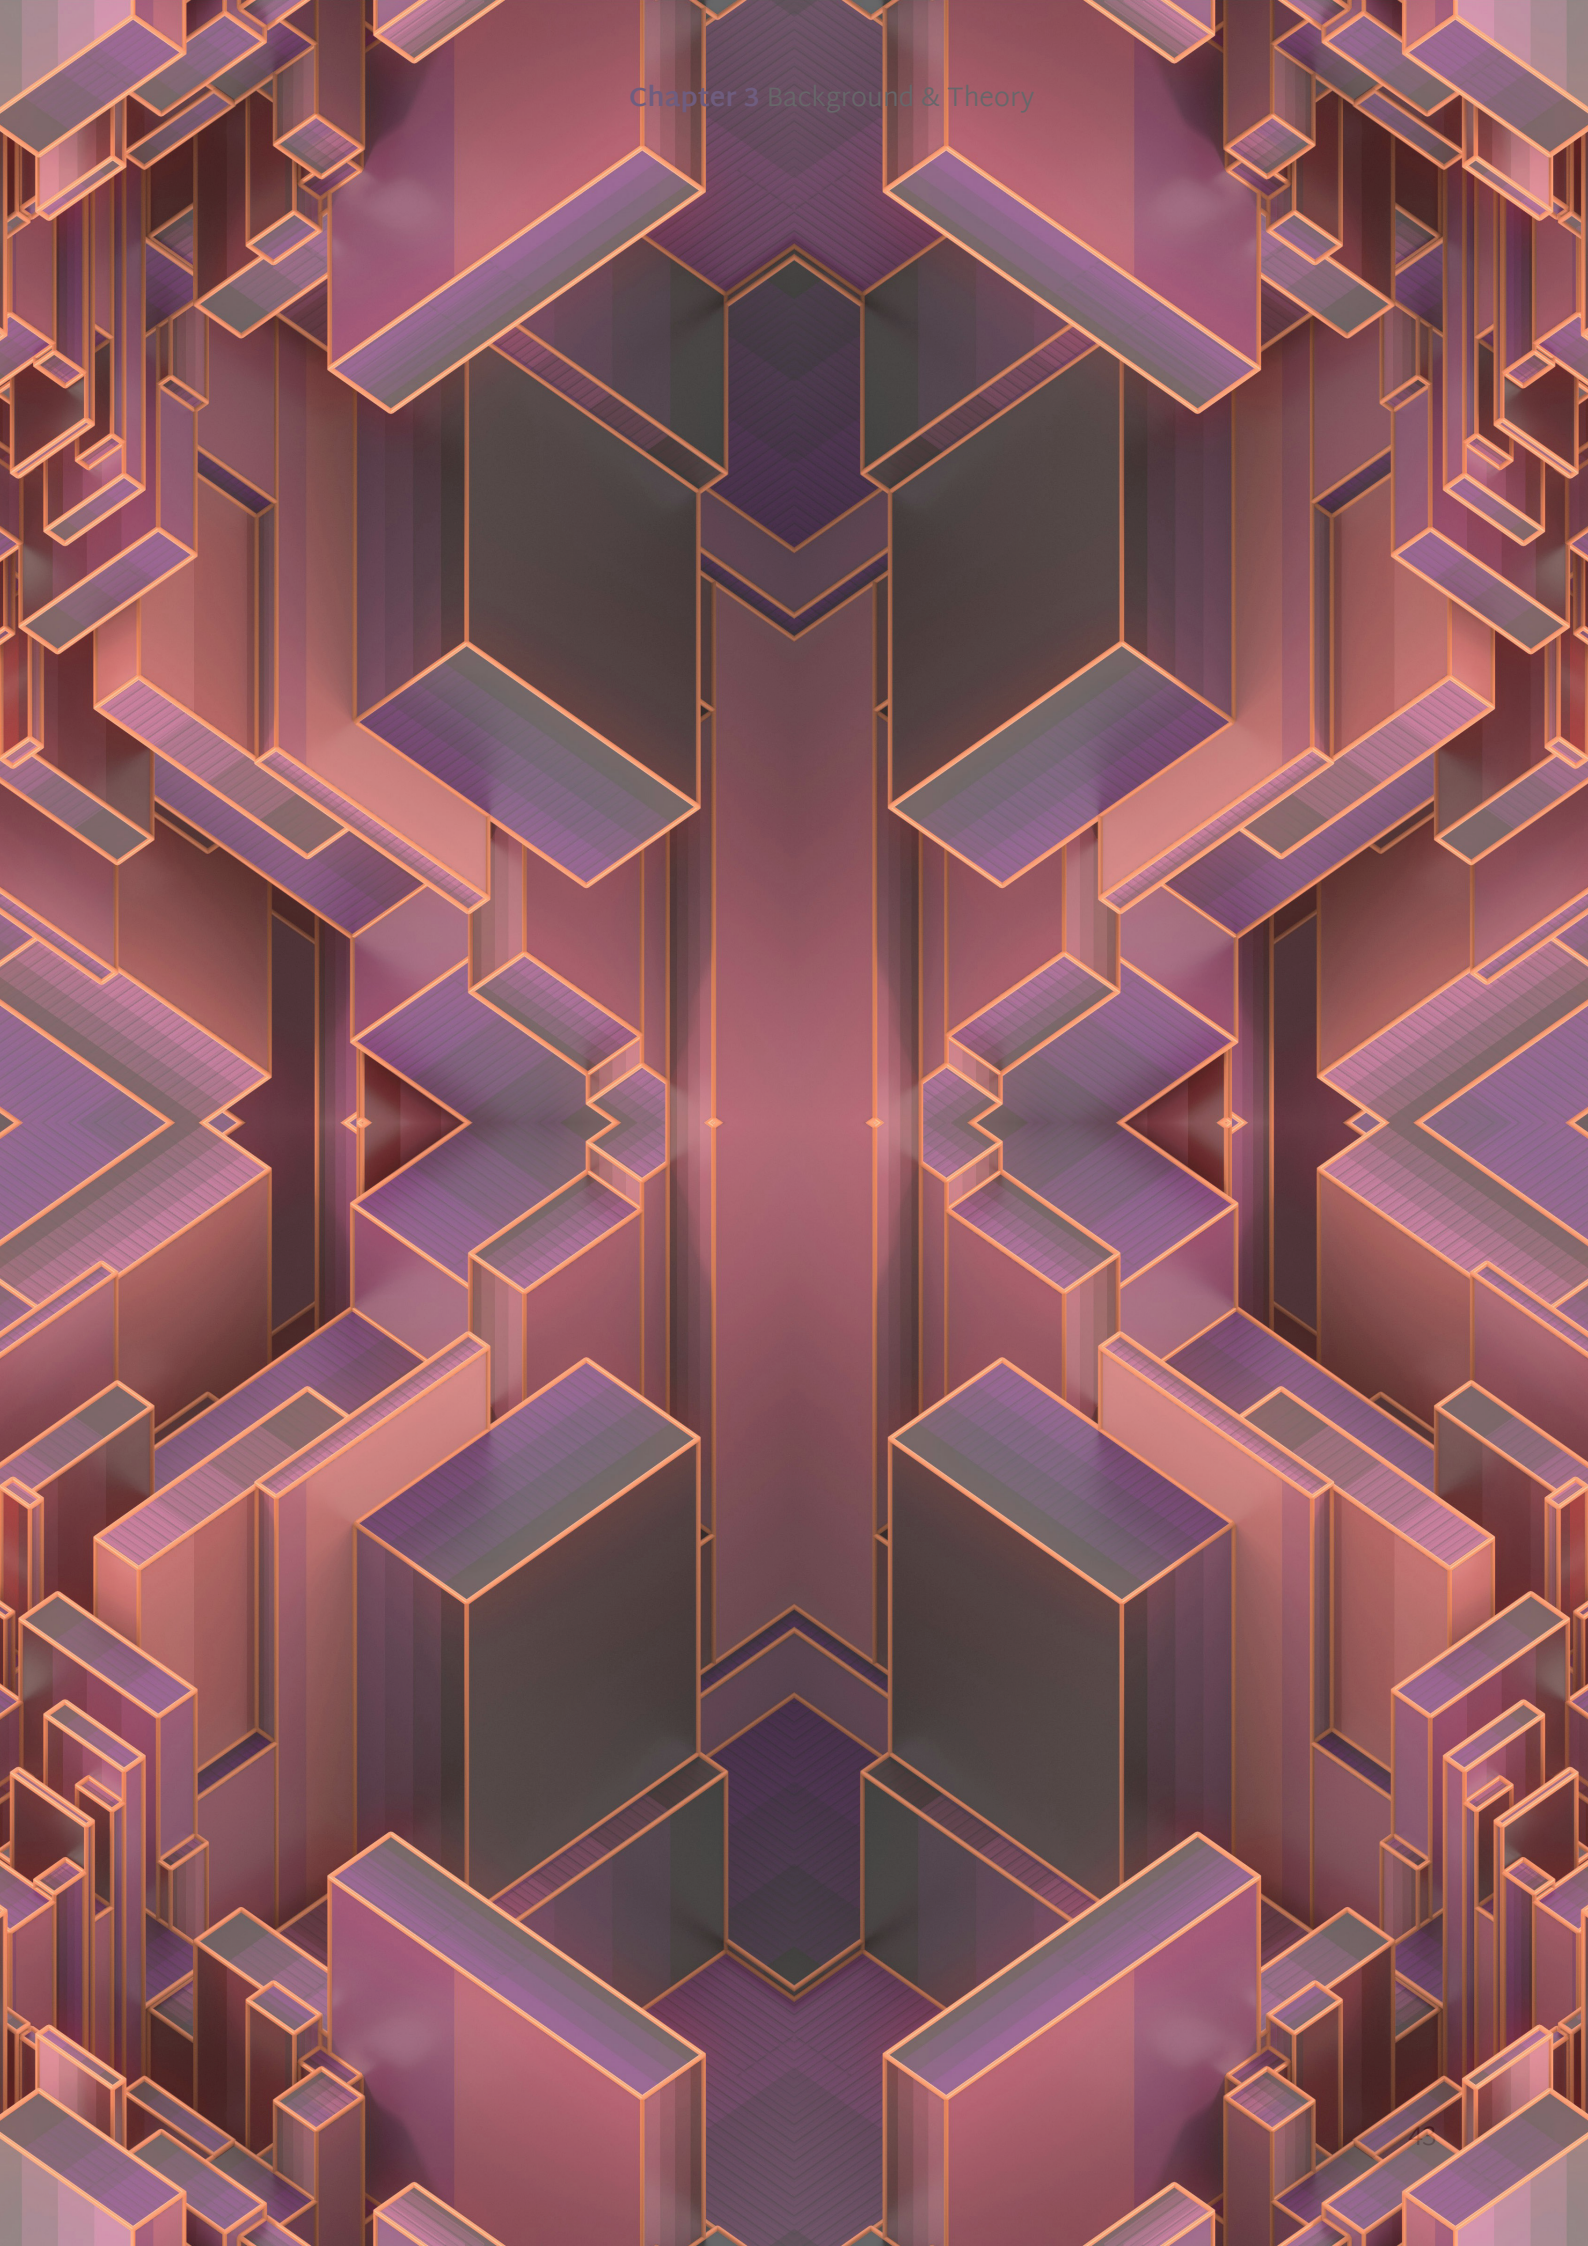

3.3 Data Feminism

3.3.1 Feminist Critique of Data Systems

Data Feminism draws from intersectional feminist theories to critique and reimagine how data systems operate. Data Feminism was chosen as a guiding framework because it challenges dominant power dynamics in data practices and helps uncover whose realities are made visible or excluded.

Within their book by the same name, authors Catherine D'Ignazio and Lauren Klein argue that data science is not neutral, but rather is embedded into the systems of power that shape what is collected, who is counted and how the data is used. They challenge the objectivity of data by highlighting how biases, inequalities and exclusions are inherently part of our data systems. Their book balances feminist theory, activism and data science in order to rethink data practices so that equity and justice are central.

In the premise of their book, the authors propose seven principles of data feminism, with the goal of making data science more inclusive and accountable. The seven principles as described in the book are outlined to the right; For each of the principle, a connection to how it ties in to the context of this project has been added.

These principles provide a framework for recognising and challenging power structures within data, and advocate for ethical data practices that promotes justice and equity . Within the context of UDTs, they help us see that algorithmic logic is not simple technical but can also be political. Within their follow-up paper, Data feminism for AI, D'Ignazio and Klein further go on to outline how this should be a key consideration in the design of artificial intelligence systems as well. They go on to explain how the shift toward predictive autonomous systems heightens the need for feminist critique; they describe how AI systems, including those used in urban planning are often driven by "corporate, capitalist and market based logics" that can reinforce social inequality rather than dismantling it (D'Ignazio & Klein, 2020; Klein & D'Ignazio, 2024).

The Seven Principles of Data Feminism

1. Examine Power

Power is not evenly distributed generally, something which data systems often reflect and reinforce. Examining power means uncovering who controls the data pipeline and how their values are shaping the outcomes.

In the context of UDTs this involves asking: who decides what urban indicators matter? Who would benefit from being visible to the model and who remains invisible?

2. Challenge Power

Data feminism urges us to actively challenge oppressive systems; this might involve designing against creating counter data or refusing data fiction as a whole. It calls for using data science to push back against existing and unequal power structures and to work towards more just and equitable futures.

In this project, speculative design and ambiguity will be leveraged as tools to challenge optimisation logic within urban digital twins.

3. Elevate Emotion and Embodiment

Western science often frames emotion as irrational or irrelevant. Data feminism however, argues that emotion, intuition and lived bodily experience are valid and critical forms of knowledge.

UDTs often do not include qualitative data, such as how people feel in space, how safety or comfort cannot be quantified. A feminist critique reintroduces these dimensions into the urban model.

4. Rethink binaries and Hierarchies

Traditional data structures rely on rigid categories and hierarchies that fail to represent the complexity of lived experience. Instead of fixing biased systems, data scientists must redesign how data is classified in order to avoid reinforcing oppression.

Within the segmentation logic of algorithms, including those within UDTs, simplifications suppress diversity in favour of efficient modelling. These should be assessed critically, with particular focus on the effect of this categorisation on the real-life situation. We should aim to build models that account for ambiguity and multiplicity, not just optimisation.

5. Embrace pluralism

Pluralism echoes that no single dataset can capture reality; it values multiple perspectives, including those which are traditionally excluded from technical domains such as indigenous knowledge, local narrative and lived experience.

UDTs often draw from centralised and standardised data which is often publicly available. A pluralist approach would value situate and granular data, provided by the urban community, as part of the modelling process in order to include the voices of minorities groups and centralise community generated data in the modelling of the urban future.

6. Consider context

Data always produced somewhere, by someone and for some purpose. We must take the context into account when considering the outcome of the data.

UDTs tend to generalise across contexts, comparing neighbourhood, cities and people in an interchangeable manner. A feminist lens insists that data be situated within social, cultural and historical environments.

7. Make Labour visible

The process behind data is deeply labour intensive including the process of annotation, cleaning and maintenance. Much of this labour is invisible and undervalued.

Although most of the data collected for the UDTs in the Netherlands is open source and so perhaps scrutinised more thoroughly, just as in general, we should recognise the labour that is part of making this system.

3.3.2 Applying a Feminist Lens to Urban Digital Twins

Emerging critiques underscore several feminist concerns in the context of UDTs:

Differential access to technology and data: Not everyone has equal access to the technology and data that power UDTs. This could lead to a situation where some citizens benefit from personalized services and information, while others are left out. (Sánchez–Vaquerizo, 2024) The article points out that the current development of UDTs is being driven by "corporate, capitalist, and market-driven approaches," which could further marginalize certain groups.

Reinforcement of existing power structures: UDTs could be used to reinforce existing power structures by giving those in control greater access to information and resources. This could further exacerbate existing inequalities and create a more divided urban landscape (Sánchez–Vaquerizo, 2024).

Creation of "filter bubbles": Personalized experiences could lead to the creation of "filter bubbles," where people are only exposed to information and perspectives that align with their own. This could lead to increased polarization and a decline in social cohesion.

3.3.3 Binaries and Exclusion in Data

One of the most relevant issues to focus on in the context of UDTs is the recurring issue of the oversimplification of complex urban dynamics through rigid categorisation and binary classification (Calvi, 2022; Klein & D'Ignazio, 2024). As such, I have chosen to focus on this principle specifically within the context of my graduation project.

In this project, I will not only analyse these binaries, but speculate how these will be manifested within the context of a fully integrated (level 4) UDT as will be expanded upon later. By highlighting the false clarities of these binaries, I aim to draw attention to the need for ambiguity, contradiction and openness in the future of urban modelling.

These binaries are not neutral, they shape who becomes visible, legible, and optimised in urban planning processes. This critique directly informed the speculative focus of this project, which explores how these classification logics are embedded in behavioural segmentation systems.

3.4 Ambiguity in HCI & Design Fiction Foundations

In much of western thought, ambiguity and uncertainty has been treated as problems to be solved rather than states to be explored. As M. Jackson notes in her book Uncertain: the wisdom and wonder of being unsure, we have long been perusing the "quest for certainty" , a want for definite answers as well as clear causations, and predictable outcomes. However, new insights from neuroscience and psychology are challenging this perspective (Jackson, 2023).

Recent theories of cognition argue that a degree of ambiguity, unexpectedness, and uncertainty may actually be essential for learning, imagination, and adaptive behaviour (Benjamin et al., 2021). Psychologists suggest that when faced with ambiguous or uncertain situations, people are compelled to draw upon their own background knowledge and experiences to make sense of and navigate the world around them. Rather than resisting this inherent uncertainty, ambiguity can be harnessed as a powerful tool for engaging people both intellectually and emotionally. (Gaver et al., 2003)

Contemporary critical and cultural theory emphasises that ambiguity and indeterminacy are not necessarily impediments, but rather constitute intrinsic aspects of the world and our understanding of it. (Benjamin et al., 2021; Sánchez–Vaquerizo, 2024) Ambiguity exists not only in how we perceive and represent the world, but also in the very nature of the systems we seek to model and control. Giaccardi et al. (2024) argue that "the things, artefacts, and systems we are making today are inherently probabilistic and uncertain due to their gradual becomings and multiple entanglements" (p. 9), highlighting that ambiguity is no longer peripheral but central to the nature of computational artefacts and their design (Giaccardi et al., 2024).

HCI and design has begun to re-examine ambiguity as a design resource rather than a design flaw; It can be argued that ambiguity in interactive systems can enhance engagement, encourage interpretation, and allow for richly personal interactions with technology (Gaver et al., 2003). Going further, it can be proposed that we can think of machine learning uncertainty not just as an inconvenience to be avoided, but as a design material to be actively shaped and used in the creation of reflective and speculative design artefacts (Benjamin et al., 2021). Going further, ambiguity and uncertainty are increasingly treated as active design materials, to be shaped and used in the creation of reflective and speculative artefacts that provoke new ways of seeing algorithmic systems (Giaccardi et al., 2024).

Utilizing uncertainty and ambiguity in the creation of interactive systems, can create novel experiences that reveal the limits and biases of the underlying models. Additionally, by leaning into the uncomfortable feeling that is uncertainty, we

can use playful and unsettling design interventions to enhance people's critical awareness of algorithmic systems. (Gaver et al., 2003) (Benjamin et al., 2021)

In the context of Urban Digital Twins, this points to new ways of embracing the intrinsic ambiguity and uncertainty present in the complex and dynamic urban environments they seek to model. Inherently, the encoding of the binaries into the datasets and algorithms that power UDTs means that important contextual details and lived human experiences are simplified, abstracted, or outright excluded from the model. Deliberately highlighting these gaps could create more reflective, nuanced, and representative digital representations of the city.

In this project, ambiguity is intentionally used as a speculative design strategy to disrupt the normative logic of data-driven urban systems. Rather than guiding users toward a fixed conclusion, the artefact leverages ambiguity to expose the reductive binaries, uncertain assumptions, and data silences embedded in the Advier hub algorithm and the wider Urban Digital Twin. As can be distilled from the literature, ambiguity becomes a material to provoke reflection, reveal exclusions, and enable users to question how urban realities are simplified, misrepresented, or erased by algorithmic models (Benjamin et al., 2021; Gaver et al., 2003; Giaccardi et al., 2024).

Chapter 4

DESIGN CONTEXT & FRAMING

This chapter introduces the design framing and contextual groundwork that guide the rest of the project. It positions the speculative artefact within a structured inquiry that moves from broad system mapping to a focussed case study.

First, section 4.1 uses Vision in Product (ViP) methods to map technological, political and cultural forces shaping Urban Digital Twins. By clustering dozens of signals, raising surveillance anxiety, automation hype and uneven data access, the context mapping exercise raised tensions that future facing designs must face.

Next the chapter sharpens in on the trade off visible in practice: efficiency vs inclusivity. Drawing on observations from municipal pilots and critical theory, it argues that this is not a fixed binary but a spectrum where design decisions either widen or narrow access.

The framing then becomes more concrete through the Advier Hub Algorithm case study. This mobility hub planner combines CBS statistics, Google Maps points of interest, WHIZE segmentation and GIS clustering to site “optimal” hubs. Its data foundations and roadmap reveal how optimisation logic is already shaping public shape choices.

Together, the ViP landscape, the efficiency-inclusivity lens and the Hub Algorithm case study establish the problem space that chapter 5 will explore through speculation. They make clear that any design intervention must reckon with the hidden values in data pipelines and the political stakes being seen by the city’s Digital Twin.

4.1 Context Mapping with ViP

To support the framing of this project, I engaged with the Vision in Product (ViP) method, developed by (Hekkert & van Dijk, 2020) at TU Delft. While ViP is typically used as a complete design methodology, in this case, it was used selectively to conduct a future-oriented exploration of tensions within the development of Urban Digital Twins (UDTs).

ViP encourages designers to look beyond existing needs or products and instead consider the systems, values, and shifts that shape the world around them. In this project, ViP served as a tool to surface underlying contradictions in the way smart city technologies are conceptualised and implemented. Rather than generating direct design outcomes, the method helped reveal frictions and misalignments in the broader sociotechnical context of algorithmic urban planning.

Input

Signals were gathered from three streams:

- Literature on Digital Twins and smart-city governance
- Expert talks and symposium notes (DMI, Bold Cities, etc.)
- Semi-structured conversations with municipal officers, data scientists and citizen advocates

The result was a wall of 64 context factors ranging from political enthusiasm for automation to public distrust of surveillance and strategic use of “green” narratives (see figure 4.1).

Eight Clusters

Through analysis and thematic clustering, these signals were grouped into seven speculative clusters, each representing a type of future-relevant tension. Some clusters highlighted practical gaps in regulation or infrastructure; others revealed ideological contradictions embedded in smart city rhetoric (see figure 4.2).

1. The Data-Privacy Paradox
2. Equitable Digitalisation
3. (Hidden) Manipulations
4. Integration and Investment
5. Just-Human Scepticism
6. The Green Ideal
7. Growing Gap Between Regulation and Innovation
8. Experiential Impact and Future Vision

Figure 4.1 - All the context factors in their categories

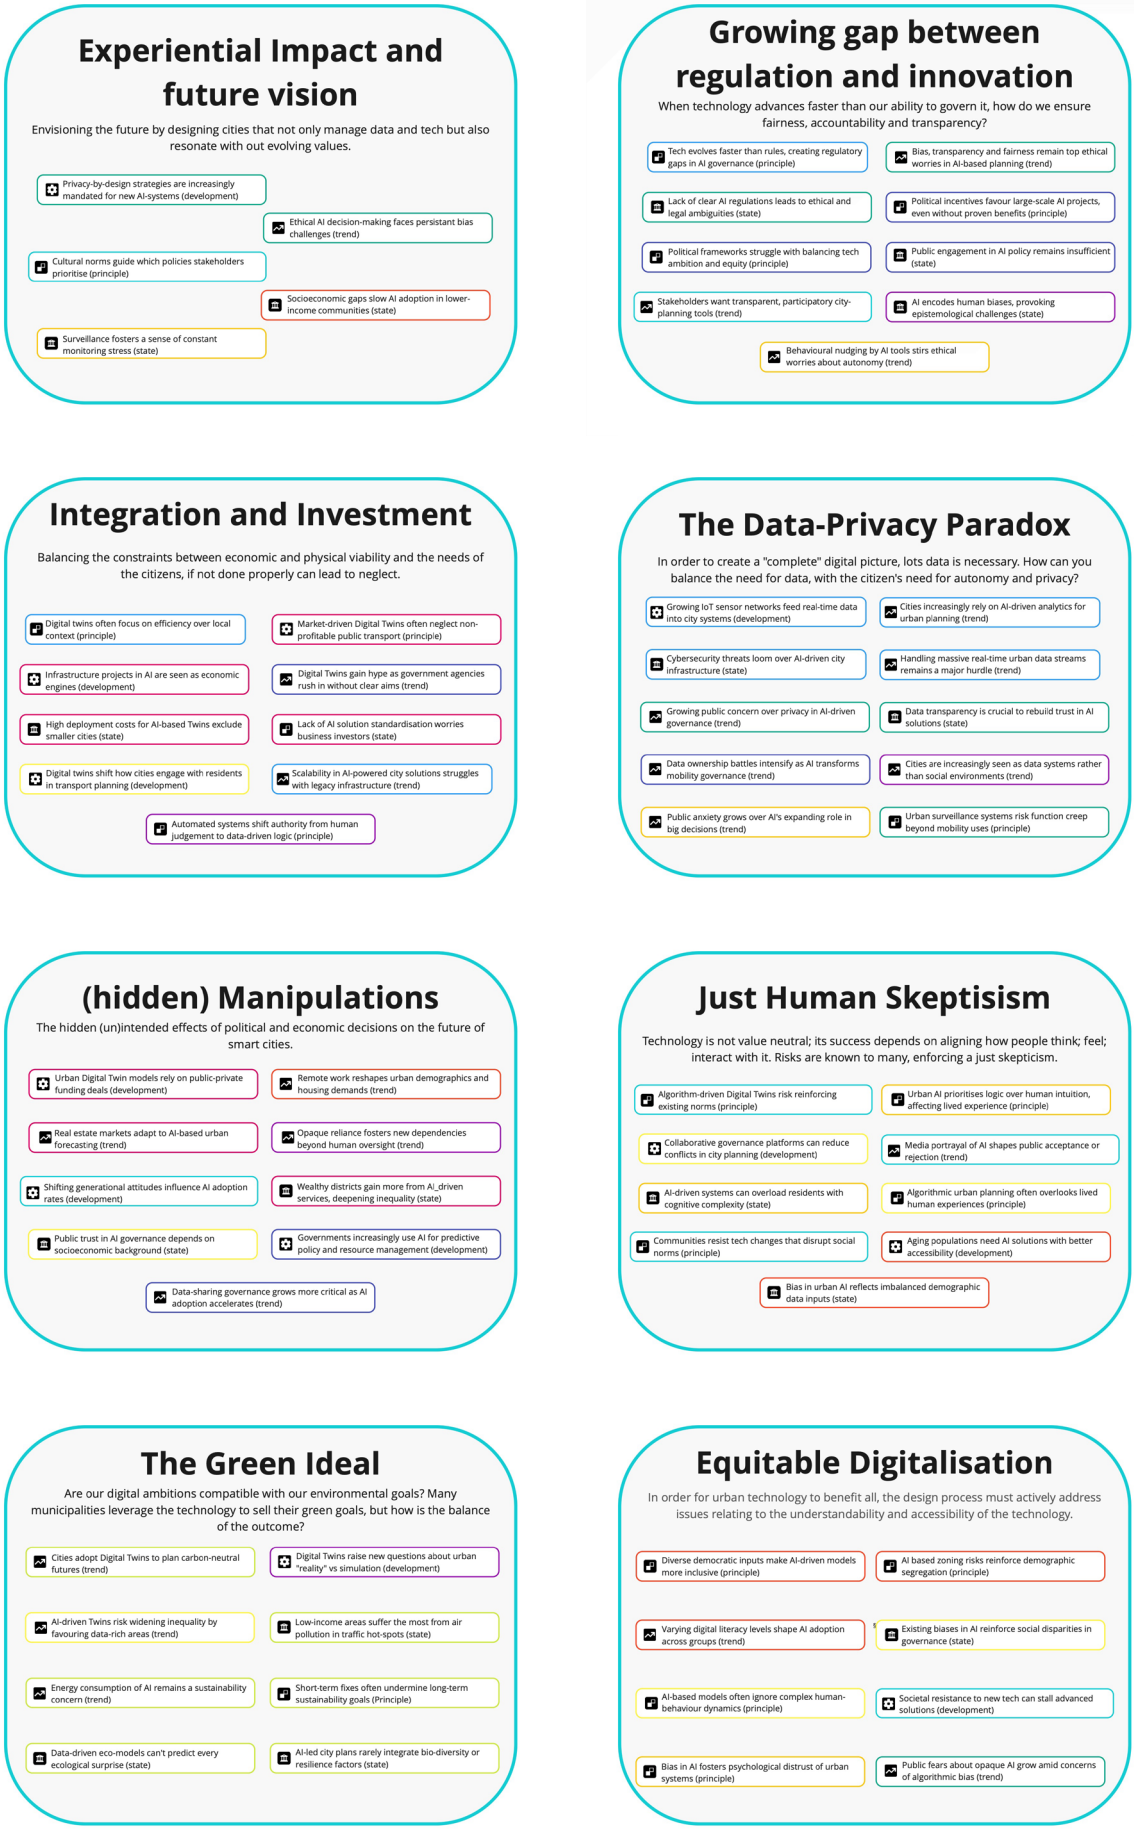

Figure 4.2 - All the context factors in their clusters

Process

collect factors --> label each by ViP type (State, Development, Trend, Principle) --> map on ViP wall --> identify affinities --> cluster into tensions --> derive critical design questions

As can be seen in Figure 4.2, this clustering process allowed for a speculative reading of tensions across technical, political, and social domains.

- In the cluster "The Data Privacy Paradox", the tension emerged between the desire for comprehensive digital insight and the growing unease around surveillance and data misuse. Here, public trust was positioned as fragile, and data transparency as both a technical and ethical risk.
- The cluster "Equitable Digitalisation" surfaced concerns around digital literacy, access, and algorithmic visibility. Factors in this group questioned whether systems designed to serve "the average citizen" inadvertently marginalise edge cases, informal practices, or those who are digitally invisible.
- In "Hidden Manipulations", political and economic motives surfaced beneath the surface of seemingly neutral technological choices. The cluster revealed how commercial priorities, private investments, and opaque decision-making processes can shape the trajectory of UDTs in ways that are not always publicly accountable.
- The cluster "Integration and Investment" focused on the infrastructural and institutional challenges of scaling UDTs, especially when efficiency becomes the dominant logic. Factors included deployment costs, technical interoperability, and misalignment between long-term planning goals and short-term implementation pressures.

Across these clusters, certain themes repeated, particularly the idea that UDTs often prioritise legibility and efficiency, but struggle to account for complexity, ambiguity, and lived human diversity. These patterns helped generate a set of critical design questions:

- What narratives are hidden when we rely only on objective data? (from Just Human Scepticism)
- How can you balance data privacy with data completeness? (from Data-privacy paradox)
- ***What happens when efficiency clashes with inclusivity?***
  - » This question is particularly relevant because it captures a broader shift in the way cities are increasingly governed through data. As digital systems become more embedded in urban planning, there is a growing reliance on optimisation strategies that privilege what is legible, quantifiable, and operationally efficient. However, these strategies often struggle to account for social complexity, informal practices, and the needs of those who do not conform to dominant data categories. Section 4.2 therefore reframes the efficiency–inclusivity clash as the project's primary design lens.

4.2 Efficiency and Inclusivity as a Design Lens

The framing of efficiency and inclusivity in this project is not a rigid dichotomy, but rather a conceptual lens to interrogate the priorities embedded in data-driven urban systems. This brings to question the specific nature of efficiency being pursued and how particular optimisation logics may unintentionally constrain inclusivity. The project therefore does not treat efficiency and inclusivity as inherently incompatible, but instead examines how optimisation-driven models can frame this relationship as oppositional through particular definitions of value performance and visibility.

In many smart city applications, efficiency is understood as the capacity to allocate resources, infrastructure, or services in a way that minimises waste, maximises performance, and rapidly responds to predicted needs. It is frequently associated with objectives such as speed, cost-effectiveness, or user uptake. Within Urban Digital Twins, efficiency is typically pursued through segmentation, modelling, and behavioural forecasting all of which rely on the simplification of complex urban realities into manageable, measurable components (Mazzetto, 2024).

In contrast, inclusivity refers to the extent to which systems, services, and infrastructures account for human diversity, social complexity, and varied forms of participation. It examines whether a system recognises people with differing access needs, habits, values, and relationships to the city (Afacan & Afacan, 2011). Inclusive design challenges normative assumptions and seeks to accommodate those typically excluded by standard models, whether due to digital invisibility, infrastructural neglect, or socio-economic marginalisation .

While often portrayed as oppositional, these concepts need not be inherently in conflict. The underlying assumption that a system must choose between efficiency and inclusivity is itself shaped by the optimisation frameworks applied in system design. It is worth questioning who defines efficiency in the first place, is it a neutral value, or one shaped by institutional priorities and funding structures? Additionally, why is inclusivity so often treated as something that must be justified, negotiated, or added as an afterthought?

The next section tests this efficiency and inclusivity tension in practice by using the Advier Hub Algorithm as a usecase.

4.3 Advier Hub Algorithm Case Study

Advier, a consultancy focussed on sustainable society transitions, developed a hub algorithm that systematically places mobility hubs across urban environments. The algorithm was created with the goal of making sustainable mobility accessible to all citizens while optimising urban space utilisation.

In order to understand the function of the algorithm, it should first be explained what a mobility hub is. Mobility hubs are defined as physical locations where different modes of transportation, and shared facilities come together, offering seamless transitions and integrated services. These hubs enhance the accessibility of sustainable transport options, reduce congestion, and promote a more efficient use of urban space (Faliagka et al., 2024). Additionally, the hubs are intended to be a social connector and integrator of the neighbourhood, allowing people to meet and connect.

4.3.1 Algorithm Method

The mobility-hub algorithm developed by Advier aims to identify optimal hub locations within a study area. Its workflow can be summarised as follows:

1. **Urban classification:** CBS data are used to locate densely populated areas that are suitable for hub development.
2. **Anchor-point identification:** GIS tools map key locations such as public-transport stops, shopping centres and community facilities.
3. **Hub clustering:** anchor points that lie within 400m of one another are grouped, ensuring that each potential hub is within comfortable walking distance for residents.
4. **User-facility matching:** the WHIZE segmentation model links hub facilities to neighbourhood demographic profiles, so that each hub menu reflects local needs.

An example of the resulting hub network for the municipality of Utrecht is shown in Figure 4.3.

Figure 4.3 - The Hub Algorithm in use in Utrecht

4.3.2 The Data Foundation

The algorithm relies on several data sources:

- CBS (Statistics Netherlands): provides urban-rural classification layers.
- Google Maps API: supplies locations of transport stops, retail centres and other points of interest.
- WHIZE segmentation model*: categorises neighbourhoods by dominant user profiles at the four-digit postcode level (see Figure 4.4 for an illustration).
- GIS platforms (QGIS, ArcGIS, OpenStreetMap): enable spatial analysis and mapping.
- Municipal datasets: add local planning information where available.

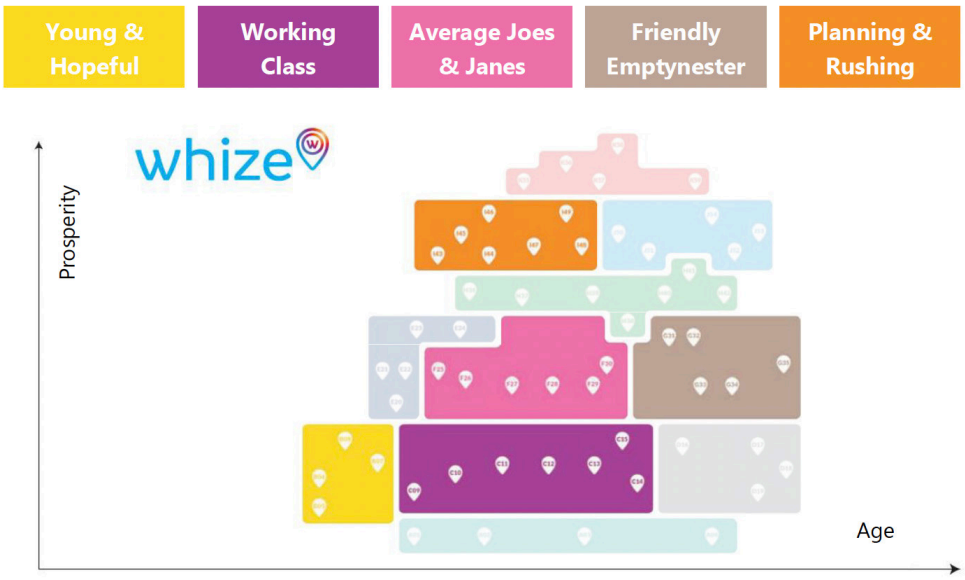

Figure 4.4 - Example of the WHIZE segmentation model in use (Vianen, 2022)

4.3.3 Future Development Roadmap

Advier has outlined both short and long-term development goals

Short-term

- Automate hub-location selection with improved AI agents.
- Integrate real-time parking availability and energy-consumption feeds.
- Expand data partnerships to improve coverage and accuracy.

Long-term vision

- Fully automated, adaptive hub planning that evolves with urban change.
- Integration with urban digital twins for scenario testing before implementation.
- Coupling with energy-grid simulations to optimise storage and distribution.
- Predictive analytics that anticipate future mobility needs based on demographic shifts.

While the algorithm foregrounds efficiency metrics such as anchor density and predicted demand, inclusivity enters only indirectly through dominant WHIZE profiles; intra-neighbourhood diversity and tacit mobility needs therefore remain outside the optimisation loop. This tension aligns with the high-versus-low-demand, dominant-profile-versus-diversity and quantified-versus-tacit binaries discussed earlier in this chapter.

With such a system there are of course still some limitations and risks such as:

- **Data-poor neighbourhoods**: sparse or outdated datasets may exclude areas from the hub network.
- **Feedback loops**: hubs placed in high-demand corridors generate usage data that reinforce the original bias.
- **Spatial justice**: households beyond the 400 m catchment radius may face longer first-mile journeys, undermining modal-shift targets.

This oversight can perpetuate existing spatial inequalities, disadvantaging communities already facing limited access to essential services and public infrastructure.

Chapter 5 extends these optimisation logics into an autonomous Urban Digital Twin scenario, asking what happens when adaptive algorithms continuously reshape mobility infrastructure without human deliberation.

**WHIZE classifies every Dutch household into 59 lifestyle sub-segments grouped into 11 clusters such as Jong & Hopevol (young starters) and Gezellige Emptynesters (sociable retirees). The system blends municipal registers, census microdata and anonymised consumer-behaviour data to produce an annually updated postcode-level layer. The algorithm uses the dominant segment per four-digit postcode as a proxy for local mobility preferences (Whize, 2025).*

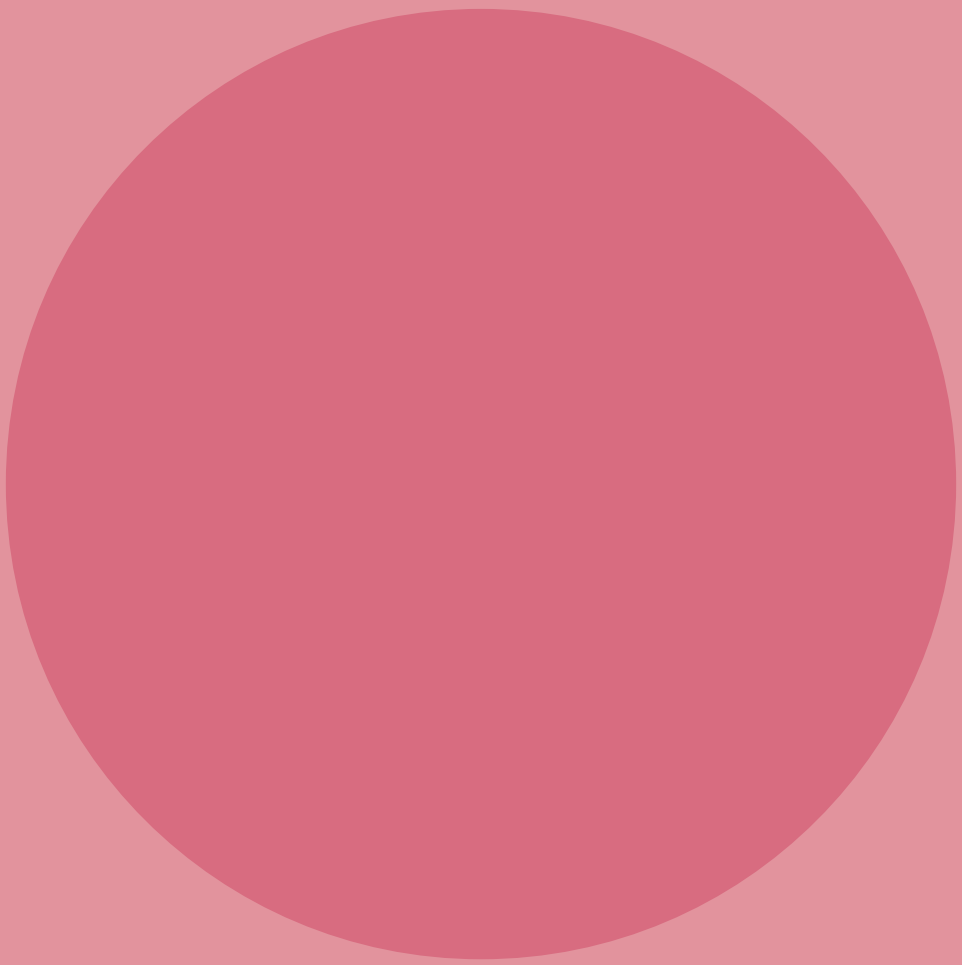

Chapter 5

SPECULATIVE DESIGN DIRECTION

With the design space for Digital Twins being quite fragmented as a result of the lack of consensus into what a Digital Twin specifically in addition to technology being relatively new, looking to an alternative speculative scenario provides the opportunity to create a scenario where a comprehensive city digital twin would be actively used.

Using literature study, as well as current pilot projects and speculation, a vision can be formed about what a city-scale UDT would look like, specifically one integrating Advier's mobility hub. For this scenario, I will not be taking into account technologies which do not yet exist (such as the potential for Artificial General Intelligence), but rather will focus on the implementation of technologies that are currently being piloted. There is also an argument to be made that even if there are vast developments in the way technology is implemented by that time, there is typically a lag between the creation and implementation of technology (Ortt & Tiihonen, 2022).

In this section I will apply speculation as a methodological tool: by projecting today's pilot scale technologies into an (alternative) future context, I can surface the social and spatial implications of a fully operational city.

5.1 Autonomous Urban Digital Twin Speculation

This section presents a speculative extrapolation of current digital twin developments, envisioning a system in which optimisation, behavioural segmentation and real-time responsiveness are scaled into a fully integrated, autonomous urban governance model, which will be referred here to as the Autonomous Urban Digital Twin (AUDT), this imagined system is not a prediction, but a design fiction used to critically examine the implications of current trajectories.

5.1.1 System Logics of the AUDT

The AUDT represents a hypothetical future system in which Urban Digital Twins have become cross-domain platforms, coordinating urban services through automated, data-driven optimisation. The following characteristics reflect speculative yet plausible evolutions based on existing trends:

- **Real-time monitoring and adaptability:** With advances in real-time data analytics (from IoT sensors, citizen apps, BIM data (live occupancy, equipment performance, maintenance, energy), GIS data (traffic, air quality, weather, public transport etc), the AUDT continuously monitors urban conditions, detects anomalies, and adapts to changing circumstances dynamically (Deren et al., 2021).
- **Predictive technology and AI-driven decision making:** the AUDT utilises machine learning algorithms and predictive analytics to anticipate mobility demand, stress points, energy consumption fluctuations and climate impact.
- Automated interventions: The twin not only simulates, but is autonomous and acts upon the outcomes of the simulations. These AI-powered smart grids and real-time transport models continuously adjust traffic signals, transit schedules, and energy distribution, maximising efficiency and sustainability (Filho et al., 2022).
- **Multi-domain integration and interoperability:** The AUDT integrates data from various urban systems, including transportation, energy, water, waste management, and public safety, providing a holistic view of the city (Caprari et al., 2022; Deren et al., 2021).
- **Ethical challenges and digital governance risks**: While the twin enhances efficiency, automated policymaking, algorithmic biases, and cybersecurity threats raise concerns over data sovereignty, urban equity, and digital exclusion (D'Ignazio & Klein, 2020) (Ávila Eça de Matos, 2023)

In this scenario, the AUDT doesn't just advise, it also governs. Live data and algorithms create a dynamic system that monitors, predicts and reallocates urban services in real time.

5.1.2 Speculative Framing and Scope

The framing of the AUDT draws from Saeed et al. (2022) who describe the emergence of "cognitive cities" which are described to be urban environments capable of learning from continuous citizen-generated data as well as real-time data from IOT applications. In such contexts, citizens are increasingly treated as "users-as-sensors", feeding data into DTs, through people's smartphones and everyday activities as dynamic sensors of urban conditions (see figure 5.1 for Saeed et al.'s diagram on the cognitive city). While efficient, this version raises critical concerns about consent, and algorithmic profiling, especially when systems prioritise statistical legibility over lived complexity.

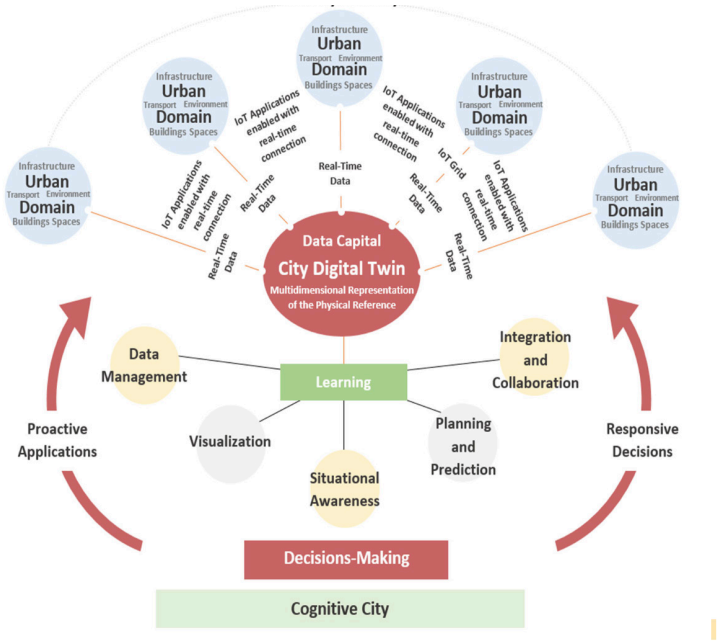

Figure 5.1 - The Hub Algorithm in use in Utrecht

With this speculative framework, the AUDT is positioned as a level 4 digital twin (Ortt & Tiihonen, 2022), capable of both real-time adaptation and autonomous intervention. It becomes a model of urban governance in which classification replaces consultation, and where invisible populations risk exclusion not through explicit policy, but rather through the silent logic of data-driven omission.

The focus of this scenario lies specifically in examining how classification and segmentation logics such as those found in the Advier hub algorithm, might scale when embedded in an autonomous system like the AUDT. It does not claim to address all the technical or infrastructural challenges associated with urban digital twins.

As such, the following aspects fall outside the scope of this speculative scenario:

- Organisational and interdepartmental data transfer challenges
- ICT infrastructure required to support the AUDT (Hartmann et al., 2024)
- Data privacy regulations and enforcement frameworks
- The use of UDTs for participatory urban design

5.2 Extrapolating the Hub Algorithm

Within this Autonomous Urban Digital Twin (AUDT), it can be speculated what the effect would be of a specific algorithm integrated into the system, such as Advier's hub algorithm. To explore the impact of algorithmic logic at scale, I examine how the core assumptions within Advier's hub algorithm would manifest if integrated into an autonomous, city-wide Urban Digital Twin system.

5.2.1 System Logics of the AUDT

Using the basis of the function of the Hub algorithm currently it can be speculated how these would be extrapolated within a integrated AUDT:

- User segmentation could be updated in real time
- Demand prediction is fed from real-time data collected through IoT devices and smartphone usage.
- Mobility hubs are not only planned, but dynamically adapted. The facilities they offer also are dynamic (this could be done through a pop-up services appearing, or vanish based on predicted behaviour)
- Public space could be algorithmically managed to align with "optimised usage"

For example, if a neighbourhood's user profile suggests a low likelihood of using share mobility spaces, the system may automatically reallocate resources to a more "efficient" zone, while not consulting with those who would be excluded. This also links to the efficiency vs inclusivity question.

In this scenario, it can be speculated that the citizens would be nudged towards recommended behaviour by their devices through the AUDT. For example, a route might be recommended to optimise overall traffic flow, even if it costs the affected person time. When combined with profiling; this could lead to the integration of inherent biases in the name of greater "efficiency".

Within this autonomous Urban Digital Twin, there is no utopia or dystopia, rather the system is a technocratic dream. Essentially, the system aims for visibility and frictionless optimisation; reducing complexity to be manageable and reducing

humanity to profiles. Broadly, it presents the urban complexity of a city as a software problem.

Taking inspiration from the Data Feminism book, below I have extracted the most apparent binaries from the Hub algorithm as it is currently situated. An overview of this can be seen in the table below (D'Ignazio & Klein, 2020).

Binary	In the current Hub Algorithm
High demand vs low demand areas	The locations of the hubs are selected based on anchor point clustering and potential user demand (inferred from population density as well as profiles).
Dominant profile vs intra-area diversity	Each postal area is assigned a single dominant Whize profile, which defines service preferences and guides amenity placement within de design of the hub.
Promising vs less promising locations	The locations are explored based on anchor point clustering and the presence of a dominant user profile who is likely to adopt shared mobility. Some areas are considered more suitable for intervention than others based on this alignment.
Quantified demand vs non-measurable needs	The selection of the amenities is mostly driven by quantifiable user preferences and behavioural assumptions taken from Whize and CBS data. No mention is made of participatory or lived experience data in the selection process.
Adopter vs non adopter profiles	Personas are ranked on expected openness to shared mobility and amenities are tailored to profiles which are likely to adopt quickly.
Data included vs data-absent areas	Modelling tools used to create the algorithm (such as GIS and QGIS) rely on available datasets. Area that lack data or include outdated data may be less precisely modelled.

Situating the Hub Algorithm within a speculative AUDT, we can see that the binary logics underpinning the tool would be amplified and proliferated at scale. The technocratic drive towards efficiency, foreseeability and manageability would increasingly overshadow the complexity and heterogeneity of urban life (Sánchez–Vaquerizo, 2024).

This speculative integration of the hub algorithm into a fully fledged Autonomous Digital Twin reveals the systematic risks that can emerge when optimisation logics are scaled without critical oversight. Using the binary structures embedded in the algorithm as a foundation, I have extrapolated how these logics might manifest within the AUDT's governance model. For each binary, I propose both a likely outcome based on the system's optimisation logic ("prediction") and a more critical extrapolation ("provocation") that surfaces the potential structural consequences if such logics remain unchallenged (see next page).

Binary	In the current Hub Algorithm	Prediction for within the Integrated Digital Twin	Provocative scenario
High demand vs low demand areas	The locations of the hubs are selected based on anchor point clustering and potential user demand (inferred from population density as well as profiles).	The hubs are placed within areas that have sustained high use. Areas with insufficient activity naturally receive fewer services.	Over time, "low-demand" zones experience infrastructural neglect, as their needs are algorithmically erased. Public service becomes a feedback loop of optimisation.
Dominant profile vs intra-area diversity	Each postal area is assigned a single dominant Whize profile, which defines service preferences and guides amenity placement within de design of the hub.	Hub amenities are tailored to the dominant group, ensuring relevance and high usage.	Minority needs are made invisible. The UDT flattens community complexity and plans only for statistical majorities.
Promising vs less promising locations	The locations are explored based on anchor point clustering and the presence of a dominant user profile who is likely to adopt shared mobility. Some areas are considered more suitable for intervention than others based on this alignment.	Investment follows predicted success: locations with strong profiles and anchor points get prioritised.	Spaces deemed "less promising" never get the chance to evolve. The system closes off potential futures before they emerge.
Quantified demand vs non-measurable needs	The selection of the amenities is mostly driven by quantifiable user preferences and behavioural assumptions taken from Whize and CBS data. No mention is made of participatory or lived experience data in the selection process.	Planning is based on hard data: what can't be measured can't be managed.	Emotional, cultural, and informal mobility practices disappear. The UDT optimises for metrics, not people.
Adopter vs non adopter profiles	Personas are ranked on expected openness to shared mobility and amenities are tailored to profiles which are likely to adopt quickly.	Profiles with high adoption likelihood get quicker rollouts and more adaptive service.	Unmodelled behaviours are deprioritised. The system begins to manage for conformity, not diversity.
Data included vs data-absent areas	Modelling tools used to create the algorithm (such as GIS and QGIS) rely on available datasets. Area that lack data or include outdated data may be less precisely modelled.	Zones with reliable data are serviced first. It ensures operational accuracy and accountability.	Data-poor areas are gradually excluded from urban visioning. They fall off the map, not through policy, but through invisibility.

The extrapolated binaries reveal a broader set of risks that emerge when predictive algorithms govern urban environments. These include automated service allocation without oversight, demographic profiling embedded in infrastructure decision, the erasure of qualitative needs, as well as the emergence of behavioural feedback loops. These outcomes are not unintended side effects, but rather represent structural consequences of the algorithm's underlying logic, which is amplified at scale within an autonomous and integrated UDT.

5.3 From Six Binaries to Three Themes

In order to ground the speculative direction of this project, I chose to combine the initial binary oppositions identified in the section above into three overarching themes; while the original binaries revealed important tensions, addressing them individually risked fragmenting the broader critique. By combining them, it allows for a more systemic exploration of how optimisation logics shape which aspects of urban life are prioritised, modelled and made actionable within an autonomous Urban Digital Twin.

These composite themes emerged, guided by analytical clustering and conceptual resonance. Each theme shows a different mechanism of exclusion such as spatial, behavioural or epistemological. They reveal how algorithmic systems not only manage urban complexity, but also decide on what and who becomes visible, and what is ignored. In the next section, they are transformed into situated speculative futures that explore the lived impact of their embedded logics.

The three themes are outlined below.

5.3.1 Urban Legibility

Data included vs data-absent areas + Quantified demand vs non-measurable needs

This theme captures how data visibility and measurability determine which zones and needs are seen and planned for. It highlights how places or experiences that lack standardised, extractable, or machine-readable data are deprioritised in optimisation systems. As a result, informal, embodied or culturally embedded forms of urban life risk being rendered invisible in the planning processes of an Urban Digital Twin.

5.3.2 Algorithmic Hierarchies

Dominant profile vs intra-area diversity + High demand vs low demand areas

This theme addresses how the algorithm simplifies intra-urban complexity by prioritising dominant user profiles and predicted high-demand zones. It reveals how value is algorithmically assigned, favouring those who conform to statistical majorities or projected usage patterns, while marginalising local diversity and less easily defined users. In doing so, it reinforces a narrow understanding of usefulness and belonging.

5.3.3 Mobility as Compliance

Promising vs less promising locations + Adopter vs non adopter profiles

This theme looks into how access to mobility and infrastructure is increasingly conditional on behavioural alignment with system goals. Locations and individuals are evaluated based on their likelihood to adopt shared mobility. Those deemed less promising or misaligned with target behaviours receive fewer resources or attention, reinforcing exclusion through feedback loops. This dynamic frames participation in mobility systems as a form of algorithmic compliance.

5.4 Generating Speculative (what-if) questions

Urban Legibility

What might a city look life if only data-included zones and measurable needs were made visible to urban planning systems while data-absent areas faded into algorithmic obscurity?

- *What types of places or communities would become invisible in such a city?*
- *Who gets to decide which needs are measurable, and what is left out?*
- *How might residents respond to being excluded from planning due to a lack of data?*
- *What alternative methods of visibility might emerge from below?*
- *What does an "invisible" neighbourhood feel like to live in?*

Algorithmic Hierarchies

What might happen if urban mobility infrastructure were planned solely around dominant profiles and high-demand zones, creating a hierarchy of access rooted in algorithmic efficiency

- *What happens to people whose behaviours don't match the dominant profile?*
- *How does a system decide who is "statistically dominant"?*
- *Could neighbourhoods attempt to alter their profile to gain infrastructure?*
- *What becomes of diversity within a zone if it's not reflected in the data model?*
- *What are the ethical limits of optimisation when it leads to exclusion?*

Mobility as compliance

What if participation in urban mobility became conditional on behavioural alignment, where only promising adopters received access and others were left behind?

- *What behaviours are considered compliant, and who sets those norms?*
- *How would people be profiled, scored, or nudged based on their mobility patterns?*
- *What mechanisms might resist or subvert this behavioural conditioning?*
- *Could being unpredictable become a form of protest or punishment?*
- *What forms of mobility become stigmatised in a system designed for compliance?*

Chapter 6

SPECULATIVE SCENARIOS

In this chapter I turn the six binaries and their provocations into short design fiction scenarios that dramatise the logics embedded in Advier's hub algorithm when scaled through an autonomous Urban Digital Twin. These short, design fiction scenarios do not aim to predict a single future, but rather embody the lived consequences of the logics embedded in the hub algorithm when scaled through an autonomous integrated Urban Digital Twin. This speculative approach follows the framework of speculative design outlined by Dunne & Raby, (2013), in which design becomes a tool to interrogate dominant systems by imagining their possible futures. By exploring each provocation as a situated experience through the eyes of a citizen, a planner, or even the system itself, these narratives aim to surface the tension between optimisation and inclusion and efficiency.

Each scenario is layered through different dimensions (technological, spatial, social. emotional and philosophical) to reflect not only systemic changes but also how these logics are felt, contested and embodied in everyday life. this structure supports a more nuanced interrogation of how algorithmic futures are lived; what becomes real, who is recognised and how exclusion is normalised through system design.

"Speculative design begins with a what-if question… it doesn't aim to predict the future, but to open up space for discussion by showing what could happen if a certain trend, value, or system were taken to an extreme." (Dunne & Raby, 2013)

Historically, "speculation" was dismissed as pie-in-the-sky thinking; yet post-Kantian philosophers such as Fichte, Schelling and Hegel reclaimed the word, arguing that ideas liberated from immediate experience can reveal new kinds of truth (Debaise et al., 2017). In the same spirit, Margaret Atwood describes speculative fiction as a way of exploring "possibilities that are inherent in our society now, but which have not yet been fully enacted" (Atwood, 2019). Together, these reflections frame the scenarios that follow; each story treats the city as a testing ground where the latent potentials of the AUDT are pushed far enough to expose their social and ethical stakes.

6.1 Three Scenarios

6.1.1 Scenario 1: Urban Legibility

What if only data-included zones and measurable needs were made visible to urban planning systems, while data-absent areas faded into algorithmic obscurity?

In this speculative future, the autonomous Urban Digital Twin (AUDT) governs all major planning decisions, relying entirely on real-time behavioural data, sensor networks, and predictive modelling. Urban knowledge is no longer produced through observation or participation, but through data capture. What is not measured does not exist.

Technological Layer

The AUDT uses a real-time optimisation engine that ingests vast behavioural and environmental datasets. Zones below the required visibility threshold are automatically excluded from planning simulations. These blind spots are not flagged as anomalies but are structurally omitted from all decision-making dashboards. Predictive algorithms model service allocation based on historical activity, sensor saturation, and system-defined metrics of relevance.

Spatial Layer

The city begins to fragment. Data-rich districts glow in full resolution: smooth roads, clean parks, responsive lighting, and frequent maintenance. In contrast, data-absent areas appear greyed out on UDT maps, physically underserved, and algorithmically illegible. Public investment flows toward where optimisation seems most promising, reinforcing a feedback loop of attention and invisibility.

Social Layers

Communities in data-absent zones form alternative structures to reclaim visibility. Some wear devices to simulate system-recognised behaviours. Others create informal networks to track unmet needs, attempting to inject this data back into the AUDT through subversive channels. A new class distinction emerges, not based on wealth, but on algorithmic legibility.

Emotional Layer

Residents of the excluded zones experience a mix of alienation, uncertainty, and resistance. Feelings of abandonment and bureaucratic erasure coexist with

Figure 6.1 - Generated Image to symbolise the Urban Legibility scenario

inventive strategies for reappearance. Some citizens internalise the need to "perform" datafied behaviours, while others express quiet defiance or grief over the fading of their neighbourhood from civic visibility.

Philosophical Layer

The AUDT enforces a strict epistemological regime: only that which is measurable is considered real. This produces an urban ontology rooted in procedural knowledge, where legitimacy is granted through data presence. The exclusion of unmeasurable needs is not seen as a failure, but as non-existence. Urban reality becomes a reflection of what can be captured by systems, displacing lived complexity with algorithmic certainty.

Vignette

> *Nova heard the city's planning dashboards were updated with new "live visibility zones" but when she searched her postcode, her block came up blank; "data insufficient" . It did not feel like a big deal at first at first, yet she noticed their street had not been resurfaced in years, bin collection had become less predictable. Her neighbours joked about being on the city's blind side. Nova wasn't sure if the omission was intentional, or just a technical lag, but lately, when things go unfixed, she wonders whether they are simply not being seen.*

6.1.2 Scenario 2: Algorithmic Hierarchies

What might happen if urban mobility infrastructure were planned solely around dominant profiles and high-demand zones, creating a hierarchy of access rooted in algorithmic efficiency?

In this speculative future, the AUDT no longer models people as fluid or multifaceted. Instead, it reduces each neighbourhood to a single, dominant statistical profile. Infrastructure is designed exclusively around this composite user, flattening urban complexity into predictable archetypes.

Technological Layer

The system uses clustering algorithms to identify "dominant mobility profiles" within each geographic zone. These profiles are based on behavioural patterns, income data, transit use, and platform preferences. All infrastructural decisions, from mobility hubs to service types, are calibrated to serve the profile alone. Minority behaviours and outlier patterns are automatically filtered out as noise or inefficiency.

Spatial Layer

Each district becomes hyper-optimised for one ideal user: the young commuter, the busy family, the high-efficiency remote worker. The physical city is shaped around these dominant needs. As a result, infrastructure that once accommodated diversity becomes monofunctional and exclusionary. Sidewalks disappear in areas modelled around car-based profiles. Bike hubs vanish where the dominant group is deemed uninterested.

Social Layer

Residents whose needs diverge from the dominant profile are no longer seen by the system. They may find themselves walking further, taking slower routes, or left without suitable infrastructure. Over time, a behavioural hierarchy emerges: those who conform gain access and ease, while others are marginalised or pressured to adapt. Diversity becomes a personal burden.

Emotional Layer

Those excluded feel erased, flattened, or forced into mimicry. Individuals adjust their mobility patterns not because they want to, but because the system silently demands it. There is unease in being optimised for, but also anxiety in being left out. Some take pride in fitting in; others feel a growing sense of alienation from a city that no longer recognises them.

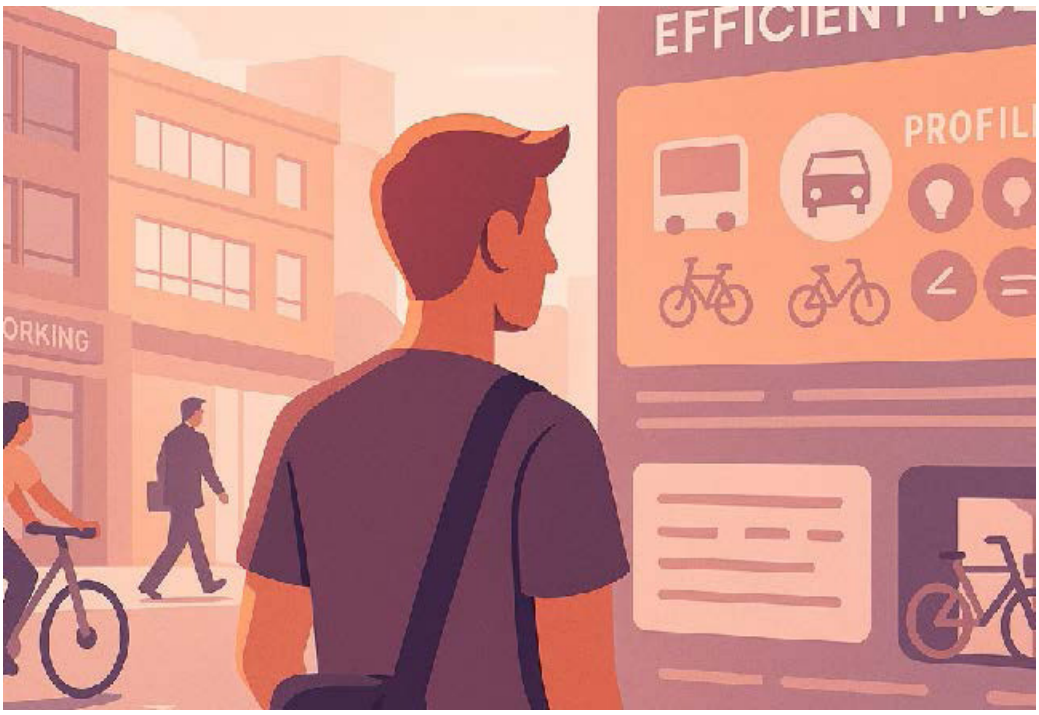

Figure 6.2 - Generated Image to symbolise the Algorithmic Hierarchies scenario

Philosophical Layer

The AUDT enforces a utilitarian logic: to serve the many, reduce the complexity of the few. Diversity within zones is not treated as value, but as inefficiency. Representation is no longer plural but average. The city becomes a system of behavioural archetypes, where statistical dominance defines what is real and desirable. All else is discarded as irrelevant.

Vignette

> *Sem lived in a part of the city that had always been a mix of residents: young families, shift workers, students, older residents. But when the Urban Digital Twin rolled out its adaptive zoning strategy, that diversity began to dissolve. His neighbourhood was reclassified as "efficiency-focused commuter zone," based on dominant user data from just three streets over. Within months, services shifted: tram frequency dropped, the daycare closed, a new car-share depot replaced the small playground. The official dashboard showed the zone as "well-aligned with target mobility behaviours." Sem still biked to his night shifts at the clinic, but it was clear the system didn't see him or the others who didn't match the zone's dominant profile. They hadn't moved. They'd just stopped being counted.*

6.1.3 Scenario 3: Mobility as Compliance

What if participation in urban mobility became conditional on behavioural alignment, where only promising adopters received access and others were left behind?

In this speculative future, mobility becomes a system of conditional privileges. Users are continuously evaluated by the AUDT on how well they conform to modelled behaviour. Access to services is earned through predictability, participation in shared systems, and algorithmic approval.

Technological Layer

The AUDT assigns a mobility compliance score to every individual, based on shared mobility usage, app integration, route consistency, and interaction with system-approved platforms. These scores determine eligibility for travel credits, optimised routes, or access to high-tier services. Unpredictable movement patterns lower your rank and affect how the city responds to you.

Spatial Layer

The physical layout of the city reflects behavioural stratification. Compliant users experience seamless, automated infrastructure: always-available shared bikes, fast transit corridors, adaptive routing. Others face delays, degraded services, or rerouted paths. Infrastructure itself becomes responsive to behavioural metrics, not geographic or community need.

Social Layer

Social groups are reshaped by system expectations. Communities begin to self-monitor, encourage app-based participation, and reward conformity. Resistance becomes costly. Alternative or informal mobility cultures are stigmatised or erased. Movement becomes a test of algorithmic loyalty.

Emotional Layer

People live with the tension of being watched, scored, and nudged. A mixture between ride and anxiety; some feel rewarded by their compliance, others feel degraded or manipulated. Deviating from the model becomes emotionally risky. Even spontaneous travel feels like a potential penalty. The freedom of movement becomes contingent and conditional.

Figure 6.3 - Generated Image to symbolise the Mobility as Compliance scenario

Philosophical Layer

The AUDT reframes freedom not as a right, but as a systemic reward. The logic of the system holds that predictability equals trustworthiness. Mobility becomes a behavioural contract: to be optimised, you must align. Deviance is not just inconvenient, it is illegible. The system collapses pluralism into conformity, framing resistance as disorder.

Vignette

> *Anna had never paid much attention to the app's nudges , route suggestions, ride incentives, minor prompts to "opt in" for mobility upgrades. But after a few months of skipping the shared ride offers, things started changing. Her travel options narrowed: fewer route suggestions, longer transfer times, slower check-ins at the station. At first she thought it was a glitch. Then she checked her mobility dashboard. Her profile read: "low uptake: system passive." Below that, a score. No explanation, just a number. She hadn't broken any rules. She'd just chosen differently. Now, her choices were choosing for her. The city was still open, but only if she moved the way the system wanted her to.*

6.4.1 Key Interactions between the Scenarios

While the three scenarios appear distinct in focus, they are not at all mutually exclusive. In fact, they represent interconnected mechanisms through which algorithmic exclusion may operate simultaneously within an autonomous Urban Digital Twin. Each scenario isolates a particular logic of control, but in practice these logics may compound one another. For example, data absent zones (such as in Urban Legibility) may never receive tailored infrastructure because they lack a dominant profile (as in algorithmic hierarchies), and residents in such areas may be further excluded for failing to exhibit system-aligned behaviour (as in Mobility as Compliance). These scenarios serve as analytical lenses, allowing the systemic nature of exclusion to be unpacked across multiple scales. Behaviour becomes a criteria for access, diversity is flattened into optimised archetypes, and non measurable needs disappear from the decision space. The resulting system does not exclude through one mechanism alone, but through the quiet convergence of multiple forms of optimisation, each reinforcing the limits of what is seen, known and planned for.

6.2 Testing the Scenarios

Aim: A speculative scenario should draw viewers into the imagined problem and trigger critical reflection on future risks. To gauge that initial emotional response, I conducted a small user test with six participants (three pairs). Each pair reacted to the three scenarios and then identified the one they found most provocative.

6.2.1 Method

Participants received laminated, double-sided A6 cards, one per scenario, containing a title, image, what-if question, concise layer summary, and vignette (Figure 6.4, 6.5, 6.6). After reading the cards individually, they took part in a guided discussion and completed a short form.

Discussion prompts

Per scenario

- Which emotion does this scenario evoke for you? (Peter Desmet)
- which words do you associate with this scenario (multiple possible)?
- What makes this scenario interesting, recognisable, or concerning to you?

Final Reflection

- Which scenario did you find the most provocative?
- Why does this scenario resonate with you most?

Participants (in pairs)

2x Graduated master students from the faculty of Architecture

2x Graduated master students from the IDE faculty (SPD)

2x Mobility Consultants from Advier

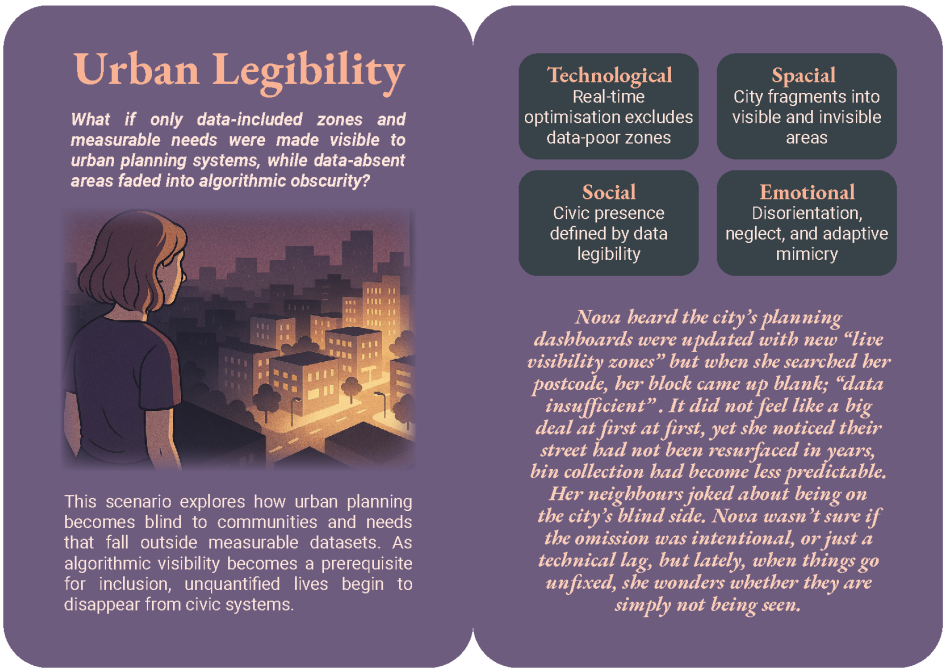

Figure 6.4 - Scenario card for Urban Legibility

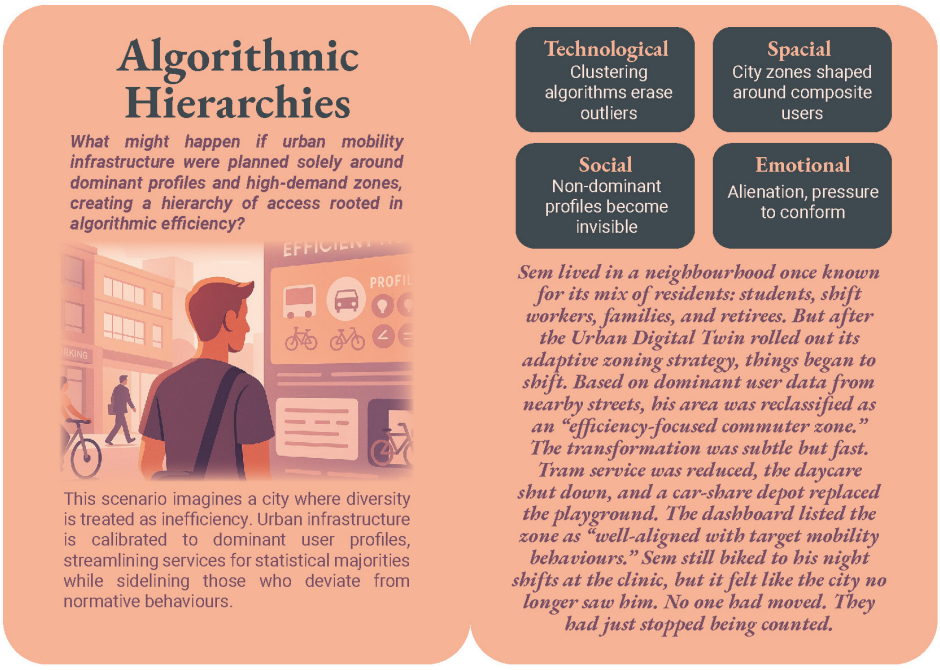

Figure 6.5 - Scenario card for Algorithmic Hierarchies

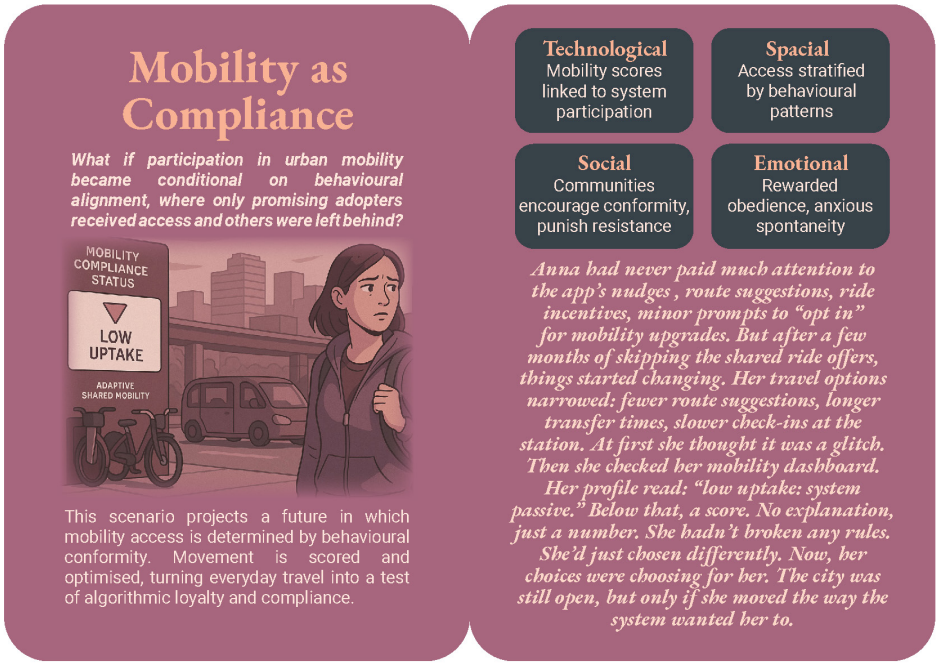

Figure 6.6 - Scenario card for Urban Legibility Mobility as Compliance

6.2.2 Results

The goal of the user test was to understand the emotional effect and associations that the viewers would have when confronted with the described scenarios, and to understand whether these scenarios would potentially resonate with an audience.

The results will be discussed separately for each scenario, as well as a reflection on which of the three scenarios was most effective.

An overview of the synthesised results from the forms can be seen in the table in Figure 6.7.

Key Findings per scenario

Urban Legibility

- Seen as the most emotionally impactful as it evoked sadness, exclusion and detachment
- Seen as highly recognisable and relevant, especially for participants familiar with marginalised communities or low-data areas.
- Highlighted real world concerns around data visibility and urban neglect
- Some called it too dystopian, without presenting alternative paths or hope.

Algorithmic Hierarchies

- Seen as highly provocative and conceptually rich
- Triggered concern about profiling, segregation, and flattening of urban diversity.
- Many recognised it as a plausible trajectory of current trends

- Raised critical questions about who defines a dominant profile and how that influences systemic allocation of services

Mobility as Compliance

- Most polarising as some participants found it conceptually compelling, while others considered it unclear
- Evoked strong feelings of discomfort and revulsion, especially related to surveillance and social control
- Compared to social credit systems, raising ethical concerns about punishment through invisibility
- One participant criticised the framing as overly negative or deterministic

		Participant 1	Participant 2	Participant 3	Participant 4	Participant 5	Participant 6
1. Urban Legibility	**Emotion**	**Sadness**	**Sadness**	**Fascination**	**Shame**	**Sadness**	**Sadness**
	Word(s)	Alienation Invisibility Freedom Detachment	Alienation Loss Invisibility Detachment	Discomfort Invisibility Tension Shift	Loss Invisibility Detachment	Alienation Connection Invisibility Detachment	Alienation Loss Connection Intimacy Slowness Invisibility Empathy Recognition Detachment
2. Algorithmic Hierarchies	**Emotion**	**Disgust**	**Contempt**	**Contempt**	**Contempt**	**Disgust**	**Dissatisfaction**
	Word(s)	Alienation Freedom Loss Discomfort Invisibility Tension Detachment Resistance	Friction Loss Discomfort Invisibility Detachment	Friction Loss Detachment	Shift Detachment Resistance	Alienation Friction Loss Discomfort Tension Shift Detachment Resistance	Alienation Friction Control Connection Discomfort Invisibility Fixation Recognition Detachment Resistance
3. Mobility as compliance	**Emotion**	**Contempt**	**Fear**	**Dissatisfaction**	**Dissatisfaction**	**Shame**	**Dissatisfaction**
	Word(s)	Alienation Friction Tension Detachment Resistance	Alienation Loss Discomfort Detachment Resistance	Friction Responsiveness Tension	Friction Tension Detachment	Alienation Loss Connection Invisibility Freedom Recognition Detachment	Alienation Friction Control Discomfort Fixation Freedom Tension Resistance
	Chosen to be most provocative	**1**	**2**	**1**	**2**	**2**	**3**

Figure 6.7 - Synthesised results from scenario research

6.2.3 General Reflections

- Participants noted that showing only negative consequences might limit engagement. Including speculative benefits or presenting scenarios with layered trade-offs could provoke deeper reflection.
- The concept of "passive profiles" led to confusion or discomfort. Participants preferred the idea of incentivising positive behaviour rather than penalising divergence or withdrawal from data systems.
- Scenarios were seen as most engaging when they felt grounded in plausible developments. Some found the ideas too speculative without more real-world anchors or contextualisation.
- The backgrounds of the participants shaped their perception of the scenarios; the participants with an architectural background resonated more with the "data absent" within the cities described in Urban Legibility, while participants with a mobility background were more likely to relate to the "compliance" of mobility as described in Mobility as Compliance.

6.2.4 Conclusion

The format with the cards worked well as a way to get a feel for the scenarios, allowing for the testing sessions to give rich insightful feedback. Aside from the forms, the level and interest within the discussions between the participants were also taken note of as part of the overall impact.

The results suggest that Urban Legibility was seen as the most emotionally impactful, while Algorithmic Hierarchies was seen as the most conceptually rich, and Mobility as Compliance was seen as the most polarising.

The participants emphasised the need for a nuanced tone within the scenarios; as they became less believable, their impact decreased as well. This suggests that within a further iteration of the scenario, a more balanced scenario should be considered.

The fact that the background of the participants also influenced their perception of the scenarios thus, creates an opening for potential ambiguity within the experience of the design itself. Whether intentional or not, this is something that should be analysed within the final scenario design.

Although according to the ratings within in form, Algorithmic hierarchies was seen as the most provocative on average, going forward a combination of all three scenarios will be considered, taking into account the emotional effect, ethical considerations, and plausibility.

6.2.5 Reflection

This test had a great focus on the emotional impact of each of the scenarios. This was reasoned from the thought that emotions are an effective way to assess the likelihood a participant will meaningfully engage with a concept. By engaging

emotions, a conversation is likely to start particularly within a group setting. In fact, D'Ignazio and Klein (D'Ignazio & Klein, 2020) would say that embracing emotion and embodiment are key to challenging false binaries and hierarchies, and can help us learn, remember, and communicate more effectively with data.

However, this approach does bring the focus away from other important aspects of scenario effectiveness such as its ability to raise questions and stimulate critique. Going forward, this take a more dominant role in the evaluation of the quality of the scenario.

6.3 The New Scenario

The composite scenario draws on all three of thematic scenario directions: Urban legibility, Algorithmic Hierarchies, and mobility as compliance. It presents a future urban system where inclusion is conditional on data visibility, and diversity is increasingly treated as inefficient. Drawing from the feedback from the scenario tests, the composite version forms the narrative foundation for the final artefact design.

Scenario Premise

In an autonomous urban digital twin (AUDT) system, being visible and served is contingent upon generating clean, classifiable data. If residents fall outside dominant behavioural patterns, or are underrepresented in the data, they are either ignored or gently nudged to conform. This creates an urban experience where inclusion is conditional, legibility is partial, and diversity is treated as inefficiency.

What happens when efficiency defines inclusion?

This speculative system is not dystopian by design; rather it reflects a technocratic ideal. It optimises for what it can see, and the result is a city shaped more by statistical dominance than lived diversity.

Narrative perspective: Sem

Sem is a nurse with non linear shift patterns and informal commuting habits (walking, informal carpools, public transport at off-hours). She avoids wearables and rarely uses apps. Her behaviour is human, but her data trail is fragmented, non-standard, making her barely visible to the AUDT.

Sem is present in the city, but invisible to its system.

Key Tensions

Efficiency & Inclusivity	System logic
Inclusion requires legibility	Data visibility becomes a precondition for planning. What cannot be measured is excluded.
Diversity is reduced to archetypes	Dominant behavioural profiles dictate services allocated; minority needs are treated as exceptions or inefficiencies.
Behaviour equals access	Services are conditional on alignment with "desirable" mobility behaviours.

Friction and Texture

Her zone is fully mapped, but increasingly optimised for Young families profiles who are predictable, wearable tracked, multimodal commuters. Sem receives subtle system prompts such as "synchronise device to enable customised travel options".

Moment of failure

Sem's bus stop is removed due to "underperformance." When she queries it, she's told: "Zone profile does not support continued service. Low behavioural efficiency score." She is offered a voucher but only if she connects a wearable and shifts to shared mobility.

Emotional and social

People begin adapting to the dominant behaviour pattern, not because they want to, but because it is the only way to stay seen. Others drift into infrastructural irrelevance. While the city remains orderly, it becomes emotionally impoverished.

Inclusion becomes a performance.

Embedded Logics

This scenario is directly extrapolated from the three scenario archetypes as written:

- Urban Legibility: data-poor behaviours and zones are algorithmically invisible.
- Algorithmic Hierarchies: one dominant profile guides all planning, others are noise.
- Mobility as Compliance: services become rewards for behavioural alignment.

Implication

This scenario shows how optimisation logics scale into structural exclusion when embedded in an AUDT. It asks:

- Who becomes visible? At what cost?
- What does inclusion mean when mediated by prediction?
- Is diversity something the system can accommodate, or something it instead resists?

Vignette

Sem, a night-shift nurse, finishes late and walks to her usual bus stop. It is gone. In its place, a slim digital panel pulses softly. The screen reads: "Route removed due to underuse. Zone profile no longer supports service."

Below it appears a prompt: "To access personalised mobility options, please synchronise your wearable device." Sem reads the message twice. She does not own a wearable. Her commutes vary: walking, informal carpools, public transport at irregular hours. Her behaviour does not match the system's expectations. No other guidance follows. There is no map, no option, no voice.

She glances around, a cyclist passes without noticing, across the street, a couple unlocks a shared car with ease; for others everything seems to flow. Sem turns and begins to walk. The city continues to hum with quiet precision, yet for her, it feels slightly out of reach.

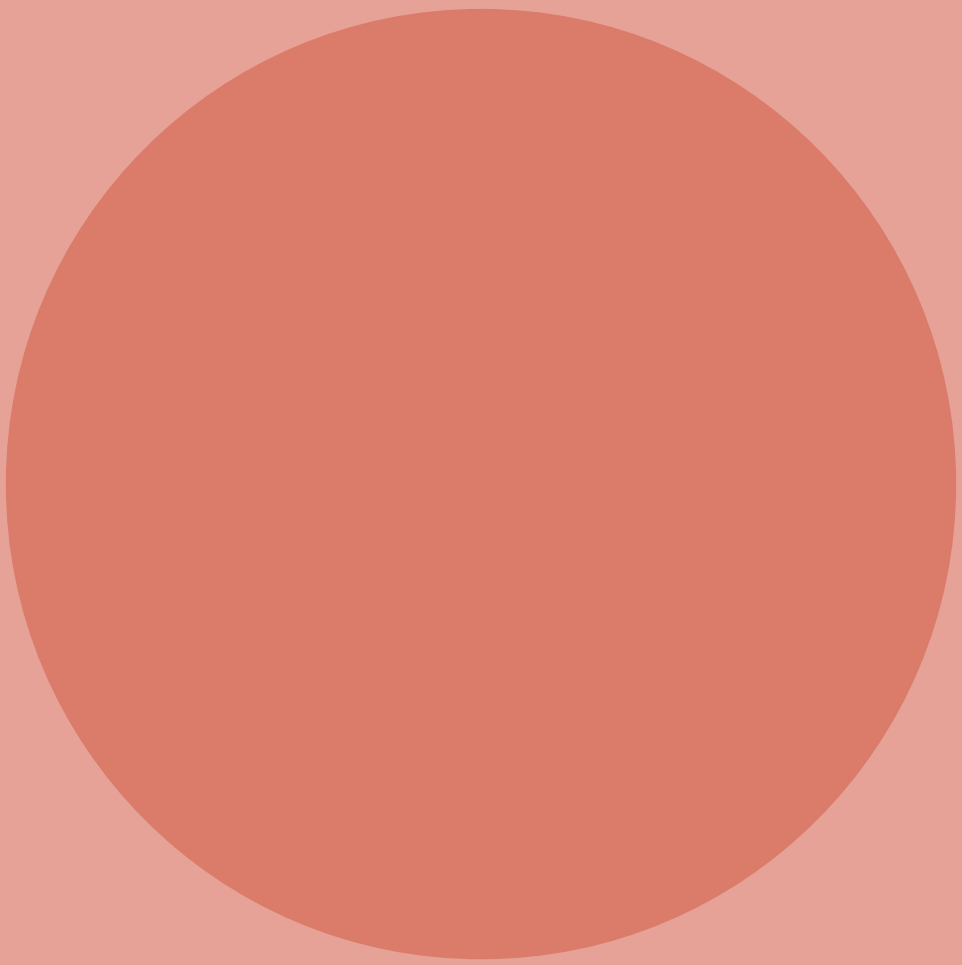

Chapter 7

IDEATION & THE PROVOTYPE

In this chapter, the speculative direction of the project is translated into a concrete design intervention. Building on the themes, binaries and provocations outlined in previous chapters, this phase explores how ambiguity can be used as a design strategy to surface the exclusions and assumptions embedded in Urban Digital Twin systems.

This chapter outlines the design goal, conceptual framework, and ideation process that led to the development of two speculative concepts. Through iterative exploration and scenario testing, one of these concepts (the Behavioural Mirror) will be further developed as a prototype: a provocative prototype that invites reflection, dialogue and ethical questioning.

Drawing from speculative design practices, ambiguity is positioned not as a failure to resolve, but as a space for reflection. The artefact is intended otherwise be recognisable enough to engage municipal professionals and policy stakeholders, while strange or suggestive enough to trigger critical reconsideration.

7.1 Design Goals & Strategy

Problem Statement

As data-driven systems increasingly shape how cities are planned, managed, and experienced, Urban Digital Twins (UDTs) have emerged as key instruments in urban governance. These systems promise optimisation and evidence-based decision-making, yet often rely on classification models that simplify complex urban realities. The growing use of segmentation and behavioural modelling introduces a structural risk: that optimisation becomes prioritised over inclusivity, leading to the exclusion or misrepresentation of populations that do not fit within dominant data categories. This raises critical questions about who is seen, who is excluded, and what values are being embedded in the design of algorithmic urban systems.

Design Goal

To create a speculative artefact that reveals how visibility and access in the city are shaped by data-driven systems, and to provoke critical reflection among urban professionals on the systemic consequences of optimisation logic.

Design Context and Audience

The speculative artefact is grounded in the case of Advier's hub algorithm, which segments populations based on behavioural data to support mobility hub placement. This algorithm is used as a lens through which to examine broader systemic patterns in algorithmic urbanism.

The primary audience for this intervention consists of urban policy advisors, planners, and municipal stakeholders involved in the use, governance, or implementation of Urban Digital Twins. A secondary audience includes designers and critical publics who may encounter the artefact in reflective or exhibition settings. In both cases, the aim is to shift the conversation around UDTs from technical innovation to critical reflection.

Design Requirements

As a speculative artefact, the design is not bound by functional requirements but is guided by critical criteria:

- The design must reveal or question a hidden assumption within optimisation logic
- It must create space for interpretive or ambiguous interaction
- It must be legible to a policy-oriented audience, without requiring technical expertise

- • It must be grounded in real-world systems while proposing a speculative extrapolation

The outcome is intended to foster reflection rather than resolution, inviting users to reconsider how digital models structure what is seen, measured, and valued in the city.

7.2 Artefact Concept & Form

7.2.1 Design Strategy

The speculative artefact that will be developed in this project builds on ambiguity as a core design strategy. Rather than aiming to clarify or resolve the issues embedded in Urban Digital Twins, the artefact uses ambiguity to foreground uncertainty and provoke reflection. Ambiguity in this context, can be seen as having both a critical and a communicative role as it mirrors the indeterminacies within real-world data systems while also inviting users to actively engage in interpretation.

Particularly, I aim to draw if the concept of data ambiguity which can be seen as the slippages between data and what it claims to represent. By designing from within these cracks (where human behaviour does not fit into neat categorisation) I am to surface the invisible exclusions and quiet violences of optimisation. Ambiguity is used to highlight rather than resolve these tensions, inviting viewers to consider whole lives become invisible or illegible to the system and at what cost.

The speculative artefacts are thus designed not only as provocations, but as reflective devices that challenge users to reframe their own assumptions. Through partial legibility, layered narratives but with unfamiliar interfaces, they aim to engage municipal professionals and policymakers in critical dialogue about the implications of algorithmic governance within the urban context.

Rather than answering the question “what should we design?”, this strategy aims to raise the questions:

What futures do our systems already anticipate?

What remains invisible within our datasets?

What forms of diversity are unintentionally treated as inefficiency?

In doing so, the design process becomes a site of inquiry itself; a speculative, open ended start to a critical conversation.

7.2.2 Methodology

To guide the ideation and development of the speculative artefacts, a step-by-step process was followed that balanced structure with openness. While speculative design embraces ambiguity, a light framework helped support creativity, clarify choices, and ensure alignment with the overall design goal.

Categorisation Framework

As a basis to create the speculative artefacts and for generating the artefact ideas, I creates a framework of potential directions or categories which the speculative artefact could fall into. These categories have not been taken from literature, but rather have been distilled from various examples come across in literature and from the book "Speculative Everything" by Dunne and Raby (Dunne & Raby, 2013). These categories were formed to have a foundation for the artefact generation process, but these were by no means restrictive and rather acted as thought provokers and not a rigid set of rules.

The framework consists of the following directions

- Public Space Interventions & Installations
- Cognitive Contraptions & Satirical Tools
- Archival & Documentary Fictions
- Reflective/Personal Artefacts
- System Extensions & Interface Fictions
- Bureaucratic Absurdities

This categorisation helped open up a wide spectrum of speculative formats, ranging from playful to critical, and from abstract to contextually grounded. It allowed the ideation process to remain exploratory while still tethered to the thematic focus of the project: algorithmic visibility, optimisation, and exclusion.

Each direction also suggested different modes of engagement, from spatial encounters to satirical reappropriations of data infrastructure. These possibilities acted as entry points for generating artefact ideas. Some were more narrative, others performative or visual. Rather than narrowing the scope, the goal of this stage was to expand the range of speculative expressions that could speak to policy and urban decision-making in a meaningful way.

7.2.3 Concept Ideation

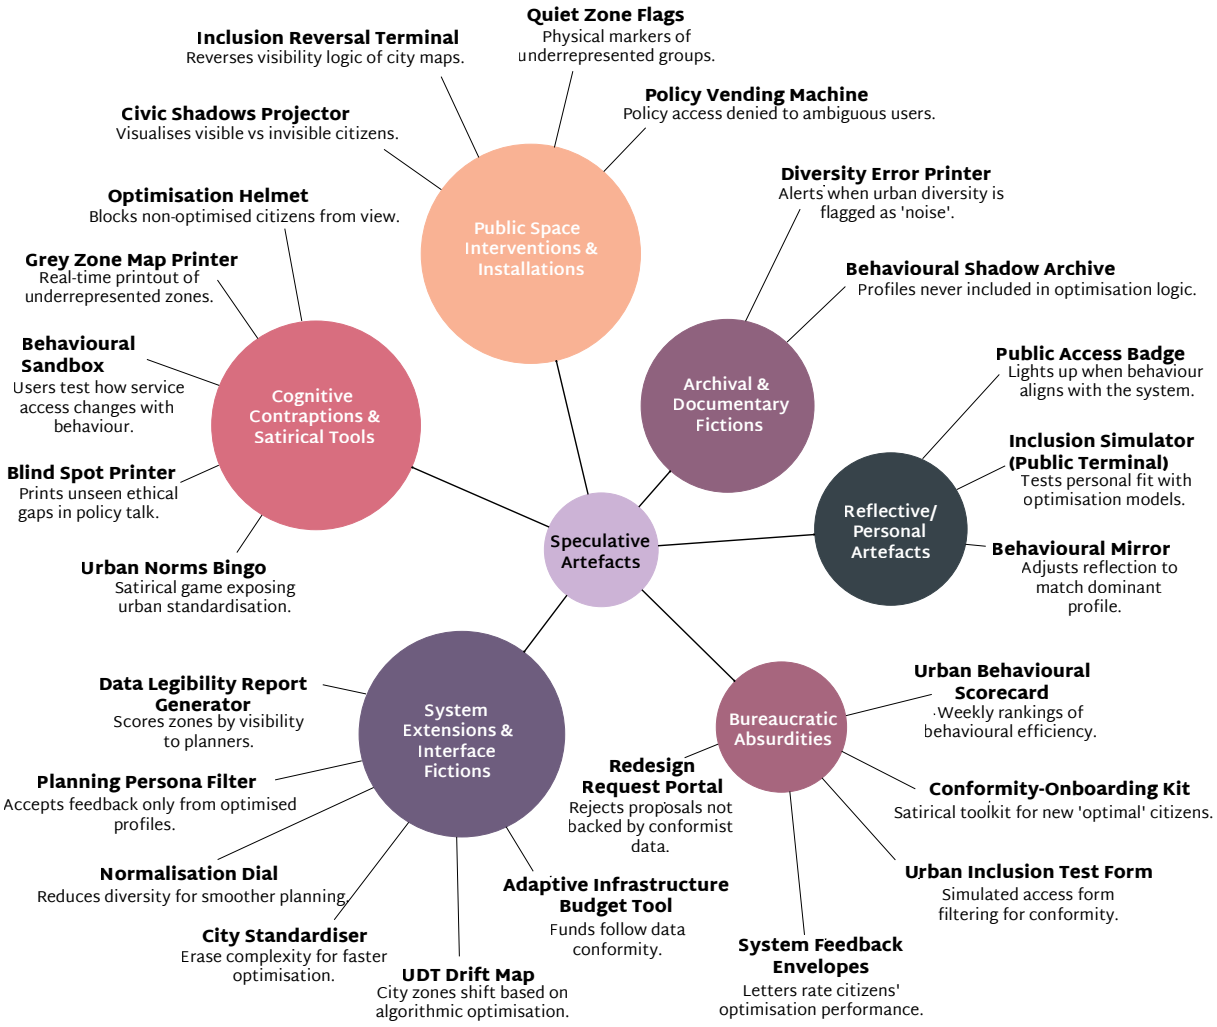

Figure 7.1 - Artefacts ideas generated through brainstorm activities

Using brainstorming techniques a wide range of speculative artefact ideas were generated which can be seen in diagram 7.1. This phase prioritised quantity and variety to expand the imaginative scope of what speculative critique could look like.

To guide the development and selection of speculative artefacts, I constructed a two-dimensional mapping framework that plots each concept along the axes of familiarity (vertical) and ambiguity (horizontal). This grid supports critical reflection on how different artefacts might be received, interpreted, and acted upon in a policymaking or public context (see figure 7.2).

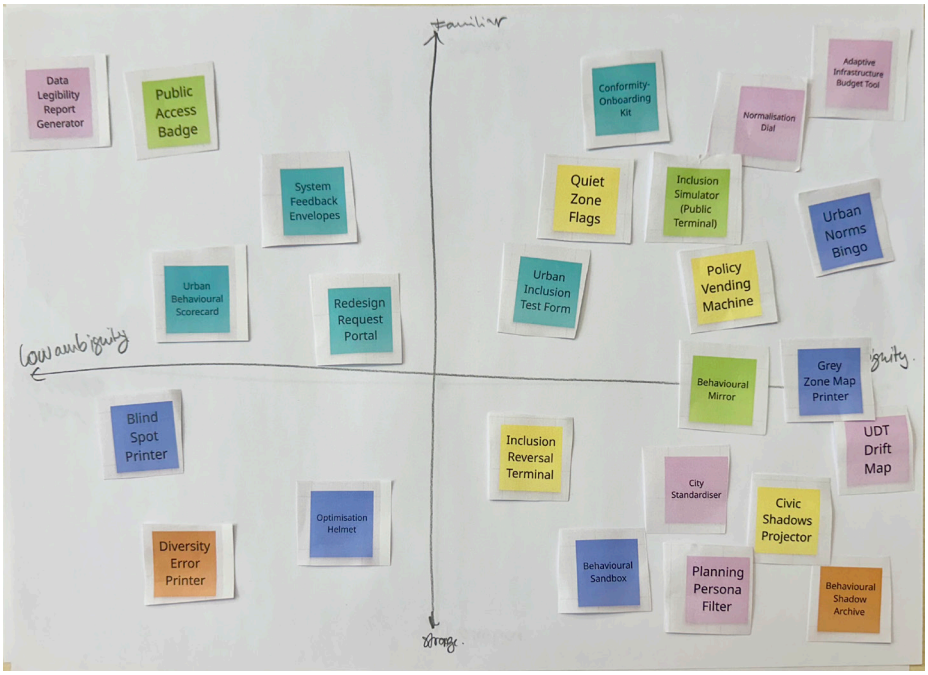

Figure 7.2 - Artefacts placed on a box-plot

The vertical axis ranges from familiar to strange, indicating the degree to which an artefact builds on recognisable formats, everyday experiences, or urban technologies. Familiar artefacts may resemble existing bureaucratic tools or signage systems, while stranger ones may appear unfamiliar or alien to the observer.

The horizontal axis spans from low to high ambiguity, reflecting how much interpretive openness the artefact invites. Low ambiguity concepts offer clearer messages or critiques, while high ambiguity artefacts rely on suggestion, provocation or contradiction to foster deeper reflection.

This landscape does not prescribe fixed categories, but rather surfaces how artefacts might function across different levels of legibility, disruption and speculative engagement. By deliberately locating artefacts across this spectrum, I am to balance recognisability with provocation; encouraging moments of discomfort or doubt, while still offering enough scaffolding to enable reflection.

The artefacts selected for further development can be seen within this landscape as moderately familiar and intentionally ambiguous, to strike a balance between accessibility and critical disruption.

This mapping exercise supported both design decision-making and theoretical anchoring, helping to clarify not only which artefact to continue with, but how they function as reflective instruments within the research-through-design methodology.

7.2.4 Selected Concepts for Further Exploration

From a wide set of generated ideas, two speculative artefact concepts were selected for further development: The behavioural Mirror and the Grey Zone Map Printer.

These two artefacts were chosen not to represent the full range of possible outcomes, but rather because of their ability to embody distinct aspects of the central tension: what happens when optimisation systems begin to define what is legible, and by extension, what is valuable in the urban environment?

Both artefacts were deliberately selected for their moderate familiarity (mirroring bureaucratic tools and spatial planning devices) and high interpretive ambiguity. This ensures they can serve not only as provocations, but as reflective instruments in dialogue with municipal professionals. As discussed earlier, a critical design should sit between the strange and the normal (Galloway & Caudwell, 2018).

Behavioural Mirror

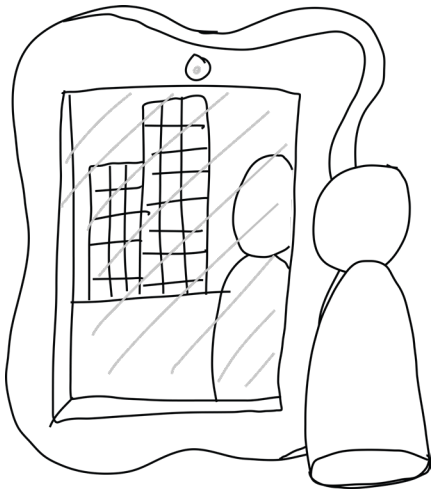

The behavioural Mirror is an interactive installation that allows its user to confront their own visibility within a future urban system governed by data. As a person steps in front of the mirror, the interface reads “How visible are you” and an assessment starts. While seeing their own reflection, the mirror provides an overlay with the results of the assessment.

The system draws on behavioural categories in order to calculate the visibility score of the individual; this score then reflects how closely the user’s actions align with the dominant data profiles recognised by the Urban Digital Twin. Rather than offering a coherent narrative the mirror embraces the ambiguity. It may glitch, distort or deliver an incomplete classification as “not enough data” or “partially legible”. The experience is intended to be unsettling through the slow realisation that visibility is not a given.

Future versions could also include speculative advertisements, ambient feedback or additional layers. This concept behind the mirror is not to provide answers, but rather to open space for questioning. What kind of futures are we building? Who becomes visible? Who fades from view? What do our systems fail to see, even when they claim to measure everything?

Grey Zone Map

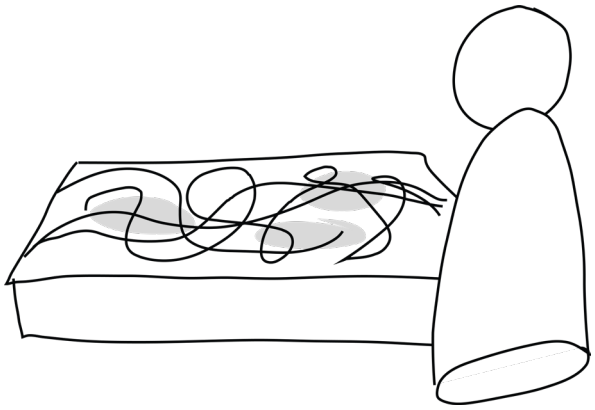

The grey zone map is a speculative cartographic device, similar to that of a more basic Urban digital twin, but rather than showing the situation as is the device reveals the unseen spaces within data-driven urban decision-making. Styled to resemble an official planning tool, the map continuously generates maps of zones that are underrepresented or invisible to the Urban Digital Twin. The map implies how these zones are quietly deprioritised, due to their lack of legibility.

The map doesn't offer any legend, rather leaving the reason for the grey zones open to interpretation. It invites the viewer to question the systems that made these areas vanish from the planning conversation. By mimicking the aesthetics of official spatial tools, it blurs the line between the speculative and the real.

What does it mean when entire communities are rendered invisible? How do we define worthiness in a system built on data thresholds? And how might maps conceal more than they reveal?

Although other concepts were considered, these two artefacts were chosen with the aim to focus the scope and testing timeline.

Each concept operates at a different scale and targets different frictions uncovered in the scenario and system analysis as such:

- The behavioural mirror engages the personal scale, it reflects how individuals are profiled based on behavioural patterns and prompts reflection on how inclusion is earned, calculated or withheld.
- The Grey Zone Map operates more on the systemic scale as it visualises how entire neighbourhoods can become deprioritised or illegible based on incomplete or skewed data representation

These concepts were also chosen because they make visibility the core speculative premise of the project; they state visibility within data systems is conditional shaped by it categorisations (binaries), legibility thresholds and behavioural assumptions. By surfacing the emotional and procedural implications of this premise, the artefacts offer different yet complementary windows into the quiet exclusions embedded within optimisation logics.

Other concepts such were valuable in the ideation phase and helped shape the larger picture, but were ultimately set aside to prioritise coherence and depth.

7.2.5 Analysing the concepts in comparison to the Design Strategy

The strategies synthesised into three headings:

- Role of Ambiguity
- Reflecting systemic logic of exclusion and nudging
- Supporting reflective municipal use

Looking through the lens of the overall design strategy, we can analyse to what extent the two selected concepts support this direction.

Concept	Ambiguity	Systemic Logic of Exclusion & Nudging	Reflective Municipal Use
Behavioural Mirror	++	++	+
Grey Zone Map	+	+	++

Both artefacts respond meaningfully to the three design strategies, but with the emphasis varying. The Behavioural Mirror leans into the speculative ambiguity and emotional engagement, confronting users with their own data visibility and prompting self reflection. Its strength lies in its ability to materialise the system's logic of inclusion and exclusion in an immediate, embodied way.

The Grey Zone Map Printer situates the critique more firmly in the spatial and infrastructural domain. This can be especially effective in municipal contexts, where maps are a dominant decision-making tool.

The two artefacts work in harmony, shedding light on the interplay between individual data visibility and systemic prioritisation. While the Behavioural Mirror prompts introspection, the Grey Zone Map reveals spatial patterns of exclusion. Together, they offer complementary perspectives, but the Behavioural Mirror arguably elicits a more personal and thought-provoking response.

One of the main differences between the concepts are their focus on the individual (Behavioural Mirror) or the system (Grey zone Map). While both are valid, the tangibility of the Behavioural Mirror makes the concept more suitable for a wider audience.

These concepts form the foundation for the final speculative artefact developed in this project. The next chapter elaborates how one of these concepts was prototyped and tested, highlighting the material, narrative, and interactive dimensions of the design.

Chapter 8

PROTOTYPING, ITERATIONS & EVALUATIONS

This chapter documents the evaluation process of the speculative artefact, the Behavioural Mirror, developed as a provotype to interrogate visibility, optimisation, and exclusion in data-driven urban systems. While speculative artefacts are not evaluated through conventional success metrics, they can still be assessed for their reflective impact, emotional resonance, and capacity to raise critical questions within their intended context.

The Behavioural Mirror was designed to be experienced both individually and within policy-focused group settings, where it functions as a reflective prompt. Rather than delivering a clear message, it uses ambiguity and speculative narrative to provoke interpretation and unsettle assumptions about algorithmic governance.

Throughout this chapter, I present the three iterations of the design, each followed by evaluations and insights gathered through user tests. These evaluations focused not only on comprehension, but on the artefact's ability to shift perspectives, raise awareness of system logic, and provoke meaningful conversations around inclusion and legibility in Urban Digital Twins.

The chapter begins with a description of the first complete prototype, followed by two rounds of evaluations and an account of the design refinements made in response to the feedback. The chapter concludes with final design additions to support the artefact in exhibition or stand-alone contexts.

8.1 The Final Design v1

Provotype — provocative prototype.

This section on the final design has been split into three design iterations. Although the main concept did not change along the way, various evaluation sessions lead to the iteration and development of the design.

8.1.1 Final Design 1

The Behavioural Mirror is a speculative artefact that visualises how behavioural data determines an individual's visibility within an Urban Digital Twin (UDT). It presents an interactive reflection in which the user is confronted with a representation of their digital presence, or absence, as perceived by the system. The artefact is based on a fictional narrative following Sem, a nurse whose routine and low-profile behaviour result in her invisibility in a future data-driven city.

The artefact consists of two core components: a short speculative video and a physical or screen-based interactive mirror. The video introduces the fictional UDT scenario and frames the systemic exclusion Sem experiences. The mirror then invites users to reflect on their own visibility through subtle interactions and prompts. Behavioural categories are not displayed explicitly but are suggested through tone, presence, or absence, reinforcing the ambiguity central to the design.

Rather than simulating a specific system, the artefact embodies a critical metaphor; It invites urban professionals to consider how algorithmic optimisation may produce unintended exclusions. It is designed for use both in facilitated

Figure 8.1 - Generated image showing The Visibility Mirror

sessions, such as policy workshops, and in exhibition contexts, where it is accompanied by question cards and supporting materials. Through its layered structure, the Behavioural Mirror encourages reflection on the normative assumptions embedded in digital systems and the consequences of becoming legible or illegible to them.

8.1.2 Storyboard

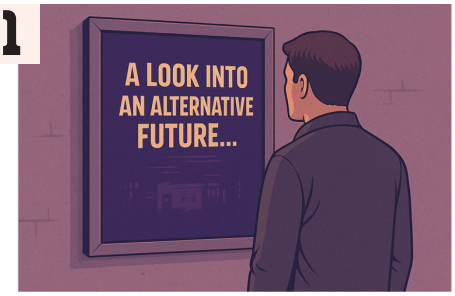

A public interactive mirror prompts passersby with the question: "A look into an alternative future..."

A short narrative video plays, introducing a data-driven urban world where visibility depends on efficiency and legibility.

The mirror pauses, displaying the question: "How visible are you?" inviting personal reflection before the assessment begins.

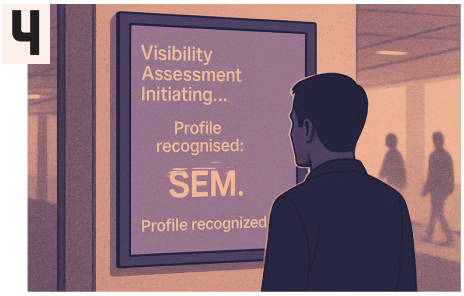

The system initiates the visibility assessment, recognising the user as "SEM." A subtle glitch hints at system fragility.

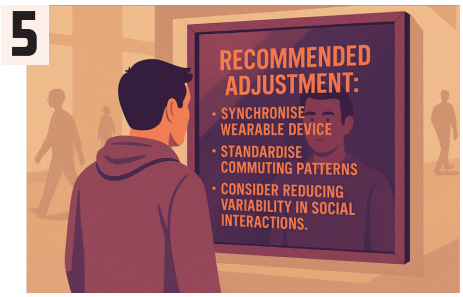

Partial data is retrieved: commuting and work behaviour unknown, community activity 23% match, revealing data gaps.

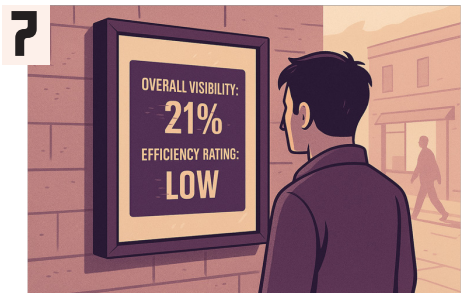

The participant receives their visibility score and efficiency rating. The low rating introduces ambiguity and bias reflection.

The mirror offers behavioural recommendations, nudging the user toward conformity, including wearable syncing and standardised patterns.

A retro-futuristic advertisement promotes the Visibility Pin reinforcing systemic conformity pressures.

Questions it should raise

Throughout the experience, the artefact is designed to provoke the following types of questions:

-Who becomes visible or invisible in data-driven decision systems?

- What biases shape system legibility?

- Is efficiency truly incompatible with inclusivity?

- What unseen assumptions do we embed in algorithmic models?

- How do behavioural nudges reinforce conformity at the cost of diversity?

These questions are intended to remain unresolved, functioning as conversation starters in group or workshop contexts.

8.1.3 The Video

The speculative video acts as the narrative entry point into the Behavioural Mirror. The story is voiced by an artificial-sounding narrator and introduces a future where optimisation logic governs urban policy. The fictional protagonist, Sem, experiences exclusion from urban services due to her fragmented behavioural data.

The video script is structured to emphasise the mechanics of invisibility: what gets measured gets optimised, while what cannot be measured disappears. The narrative is deliberately paced, slow, and calm, allowing space for absorption. Each line corresponds to a frame in the mirror, some overlaid with guiding captions or interactive prompts (see figure 8.3 for a still from the video).

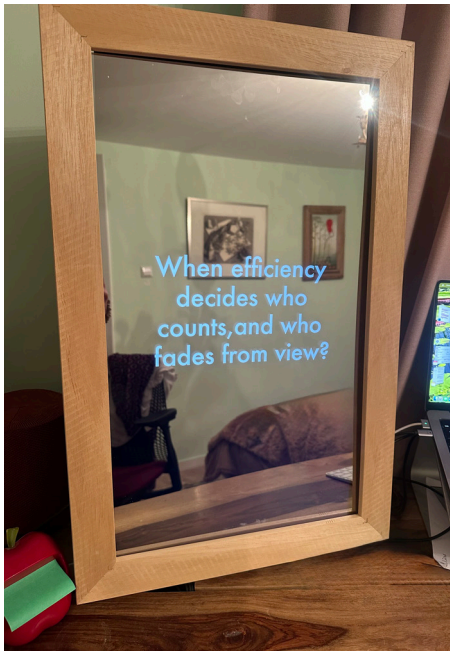

Figure 8.2 - The Visibility Mirror prototype

Figure 8.3 - Still of a scene in the video

Narrative Text (read by an AI voice)

Welcome. *You are entering a near-future city shaped by data.*

In this city, public policies no longer rely on human observation and debate alone.

Instead, an autonomous Urban Digital Twin collects data, predicts behaviours, and helps decide how public spaces and services are designed.

The system rewards behaviours it can easily measure and predict.

Citizens whose lives align with efficiency models gain visibility, and with it, influence on city decisions.

But not all patterns are easy to capture.

Not all lives fit the model.

Meet Sem.

Sem is a nurse. She works irregular shifts and travels between multiple care locations.

Her mobility and work patterns do not align with the standard profiles the system recognises.

Some of Sem's data is missing.

Some of her behaviours are too complex for the algorithm to process.

As a result, she becomes partially invisible to the decision-making system.

Zones with low data representation (grey zones) receive fewer services.

Behaviours seen as inefficient are quietly excluded from future planning.

You will now experience the city through Sem's perspective.

What happens when diversity becomes a data error?

When efficiency decides who counts and who fades from view?

Your visibility assessment is about to begin.

Figure 8.4 - The Visibility Mirror prototype

8.1.4 Footage and Atmosphere

All footage was collected or edited in a vertical format to match the mirror's physical layout. A mix of stock footage and recorded scenes of urban mobility were utilised. The tone is clean, slightly sterile, and deliberately impersonal to create emotional distance. Scenes include cityscapes, transport, and Sem's subtle, everyday actions: walking, waiting, biking, checking her phone.

8.1.5 Physical Prototype

The prototype consists of a frame-mounted one-way mirror (plexiglass) positioned in front of a vertical screen (see figure 8.2 and 8.4). This setup allows the video and reflective text to appear as layered content, making the user simultaneously see themselves and the projected assessment.

An additional design element was the Visibility Pin (see figure 8.5), a speculative badge meant to reward optimised behaviour (this item was part of the initial design but removed in later versions due to its distracting interpretation).

Figure 8.5 - The Visibility Pin

This artefact is not presented as a solution, but as a research tool. Its contribution lies in the speculative reflection it provokes among municipal professionals, the surfacing of systemic blind spots in optimisation-driven systems, and the use of ambiguity to expose data-driven exclusions. The project contributes a methodological approach that combines speculative narrative, critical metaphor, and design testing to interrogate the ethical implications of Urban Digital Twins.

8.2 Evaluation Round 1: methods & key findings

8.2.1 Initial Design Goals

The first iteration of the Behavioural Mirror was designed as a speculative and ambiguous artefact to explore how optimisation logics in Urban Digital Twins impact legibility, visibility and access to urban resources. It aimed to provoke reflection among urban professionals by asking: What if your behaviour makes you invisible to the system?

8.2.2 Testing approach

The first evaluation round consisted of three in-person sessions with six participants in total. These sessions were designed to test the initial version of the Behavioural Mirror and to gather insights into both its narrative clarity and reflective potential.

The participants included

- Two municipal policy advisors from Gemeente Zwolle working in the mobility sector
- Two mobility consultants from Advier
- Two adults with no design or data background

While all sessions were treated equally in terms of format and openness, the insights from the Zwolle session were weighed more heavily in the overall analysis, as these participants most closely resembled the intended audience of municipal professionals.

Each session followed the same structure:

1. **Introduction**

   Participants were introduced to the project through a short presentation, covering the background of Urban Digital Twins, the role of optimisation, and the use of speculative design. This ensured all participants, regardless of experience, could engage with the concept.

2. **Viewing the Artefact**

   Each participant viewed and interacted with the Behavioural Mirror prototype. This included watching the short speculative video and reflecting on the mirror interface (see figure 8.6).

3. Semi-structured discussion

A guided conversation followed, using a flexible set of prompts. While the format allowed room for spontaneous discussion, the conversation was loosely structured around the following themes:

Experience

- *How did the interaction feel?*
- *Did you recognise yourself or others in the experience?*
- *What irritated or surprised you?*

System critique

- *What happens when optimisation becomes leading in a city?*
- *Can you refuse to go along with these systems?*
- *Which groups are structurally ignored or excluded*

Policy relevance

- *What could this scenario mean for future policy?*
- *What if this already existed tomorrow?*
- *Do you recognise these dynamics in your own work?*

Closing reflections

- *What do you take away from this experience?*
- *Would this type of reflection be valuable in policymaking or digital strategy?*
- *Is there anything you would like to critically examine further yourself?*

Each session lasted approximately 60 minutes. The aim was not to test for correctness or usability, but to provoke reflection and gather layered, qualitative insights. This approach aligns with the goals of speculative design, which prioritises emotional engagement and critical questioning over conclusive outcomes.

Figure 8.6 - The Visibility Mirror being tested

8.2.3 Insights & Results

Participant responses highlighted both the strengths and limitations of the first prototype. While the theme was widely recognised as relevant and evocative, several areas were identified for improvement.

a) Narrative Ambiguity and Comprehension

The story of Sem was seen as emotionally resonant, but too abstract. Participants struggled to grasp what Sem was missing and why. One noted, *"Nu heb ik eigenlijk alleen tekst gelezen… dus ik snap hem nog niet helemaal."* (*"**Now I've actually only read text… so I don't fully understand it yet."**)* Another said, *"Wat is precies haar probleem? Waar mag ze niet heen, wat mag ze niet gebruiken?"* (*"**What exactly is her problem? Where is she not allowed to go, what is she not allowed to use?"**)*

Reflective impact was strong but underdeveloped: *"Het raakt iets, maar ik weet niet precies wat. Misschien als ik wist wat ze mist."* (***"It hits on something, but I don't know exactly what. Maybe if I knew what she was missing."**)*

Some participants also expressed uncertainty about who was being critiqued. *"Het voelt nu een beetje alsof Sem het probleem is omdat ze zich niet aanpast."* (*"**It kind of feels like Sem is the problem because she doesn't adapt."**)* Another reflected, *"Ik snap dat het over systemen gaat, maar ik krijg een beetje het gevoel dat het haar eigen schuld is."* (*"**I get that it's about systems, but I kind of feel like it's her own fault."**)*

b) System Logics and Binary Thinking

Some participants reflected on the binary logic in the system, noting its exclusionary implications. *"Het is alsof je óf mee moet doen, óf je valt buiten de boot. Geen middenweg."* (*"**It's like you either have to go along with it, or you're left behind. No middle ground."**)*

The scenario reminded others of behavioural control systems already in use elsewhere. *"Het doet me denken aan dat systeem in China… waar gedrag bepaalt wat je nog mag."* (***"It reminds me of that system in China… where behaviour determines what you're still allowed to do."**)*

Others noted the complexity of visibility in digital systems, raising the idea that invisibility might sometimes be chosen even though at a cost. *"Soms wil je misschien juist niet gezien worden. Maar dan ben je meteen ook onzichtbaar voor alles."* (*"**Sometimes you might actually not want to be seen. But then you become invisible to everything."**)*

c) Interaction Design and the Visibility Pin

The visibility pin was introduced as a playful, satirical object; a tangible marker of 'approved' urban behaviour, inspired by gamification logics embedded in optimisation systems. However, its symbolism was not immediately clear to participants. While some interpreted it as a reward for conformity, others questioned whether it was meant to be ironic. Rather than reinforcing the critique, it unintentionally shifted focus away from the core message. This ambiguity proved too open-ended, especially without supporting context, and led to the decision to remove the pin in the second iteration. *"Is dit een soort badge voor brave burgers? Of is het sarcastisch?" ("**Is this some kind of badge for well-behaved citizens? Or is it sarcastic?"**)*

d) Relevance and Policy Potential

Despite the issues mentioned above, participants strongly affirmed the theme's relevance and speculative framing. *"Dit zou echt iets zijn voor een sessie over algoritmes en ethiek bij de gemeente." ("**This would really be something for a session about algorithms and ethics at the municipality."**)*

The artefact was also recognised as a useful reflection tool. *"Het zet aan tot denken… het roept vragen op die we anders niet stellen." ("**It makes you think… it raises questions we wouldn't otherwise ask."**)*

One participant emphasised its practical applicability: "*Dit zou echt kunnen in een beleidssessie. Of als opening van een workshop." ("**This could really work in a policy session. Or as the opener to a workshop."**)* This aligns with the intended function of the artefact as a policy-facing reflection tool, aimed at surfacing blind spots in data-driven strategy.

Despite these issues, participants acknowledged the value of the theme and the speculative framing. One participant stated: *"Dit zou echt iets zijn voor een sessie over algoritmes en ethiek bij de gemeente." ("**This would really be something for a session about algorithms and ethics at the municipality."**)* Another added, *"Het zet aan tot denken… het roept vragen op die we anders niet stellen." ("**It makes you think… it raises questions we wouldn't otherwise ask."**)*

These evaluations suggested that while the Behavioural Mirror succeeded in raising awareness, it required refinement to sharpen the critique and enhance reflective impact.

8.2.4 Reflection on Outcomes

The interactions and discussions within these tests had a greater focus on the methodology and the content of the artefact, than on the critical thought and questioning that it should raise. Although the awareness and importance of the topic did become apparent, the prototype needs to be strengthened in order to better achieve its goal.

The artefact was thus iterated on, as is described hereafter, to allow the artefact to function more as a reflective prompt. The artefact should open critical reflection into the implications of future urban systems.

8.3 The Final Design v2 (iteration based on Evaluation 1)

8.3.1 Refining the Provocation

Following the initial prototype test sessions, the Behavioural Mirror was further refined to deepen its reflective and critical impact. The feedback, particularly from municipal professionals, highlighted the importance of balancing ambiguity with interpretability. While the speculative logic of the system was clear in tone, the consequences of being rendered partially visible needed to be articulated better. In this iteration both narrative framing and interaction design are sharpened, while remaining deliberately unresolved.

Key Changes

1. **Narrative Revision**

   The script was rewritten to more clearly articulate why Sem becomes invisible within the system; Her non-predictable routines and lack of personalised data are now explicitly framed as the source of her exclusion. This shifts the focus from the user's behaviour to the system's limitations, strengthening the systemic critique.

2. **Addition of Guiding Captions**

   Each frame in the video now includes the written text as guidance to the voiceover, reinforcing the story and clarifying what is happening in the scene. These captions guide the viewer through the speculative sequence without over-explaining it, offering structure while preserving the open-endedness of the experience.

3. **Inclusion of Sem's Reflection**

   To clarify that the interaction is from the perspective of Sem, a semi-transparent visual of Sem is now overlaid on the mirror interface during the visibility assessment. This makes clear that the viewer is witnessing Sem's evaluation, not their own. This makes the interaction easier to follow, resulting in a more coherent interaction.

4. **Removal of the Visibility Pin**

   The satirical Visibility pin, previously presented as an advertisement within the mirror sequence, was removed within this iteration. While introducing a humourous element, the pin introduced distraction at the end of the interaction.

5. **Stronger Closure and Reflective Prompts**

   The video within the Visibility mirror ends with a series of prompts following Sem's assessment and recommendations. This allows for the reflection, that although Sem can change her behaviour, the problem is not with her, but rather with the system. By asking questions such as "Can you refuse to participate in systems like this?" and "What happens when optimisation defines urban life?" the interaction moves from entertainment to reflection, creating room for more open-ended ethical exploration. These questions also allow for the clarification of the intention of the video, making the goal clearer to the user.

6. **Refinement of System Critique**

   Throughout the second version, the system logic is made more precise. Rather than blaming the user or focusing on the flaws of Urban Digital Twins in general, the artefact now critiques the optimisation logic that fails to account for non-standard lives. It reveals how invisibility is not a side effect, but a structural outcome of current data assumptions.

Narrative Text (read by an AI voice)

Alternate future loading... Please remain still during the simulation.

Welcome.

You are entering a near-future city shaped by data.

In this future, data overrides deliberation.

Numbers speak louder than voices.

An autonomous Urban Digital Twin functions as the city's operational brain, continuously adjusting the urban environment based on live data flows.

It predicts behaviour, allocates services, and fine-tunes infrastructure; all without human oversight.

What gets measured, gets optimised.

And what isn't recognised, isn't reinforced.

Uncaptured behaviours, like irregular shifts or untracked travel, are quietly left out of policy models.

Meet Sem.

Sem is a nurse. She works unpredictable shifts and moves between multiple care locations.

She cycles, walks, carpools, and often takes unplanned detours.

She doesn't use standardised transport apps.

She doesn't wear a fitness tracker.

And she doesn't follow regular commuting hours.

Her data footprint is fragmented.

Her behaviours don't match the system's dominant profiles, and so, they go uncounted.

She lives in a low-data zone.

No demand is registered, so no services arrive.

Sem's care work is vital, but to the system, she is barely visible.

The city no longer sees her.

You will now experience the city through Sem's perspective:

Visibility Assessment Initiating...

Profile Recognised: SEM

Behavioural Pattern: Irregular

Work Schedule: Fragmented

Mobility Mode: Low-data

Wearables: Not linked

Overall Visibility: 23%

Efficiency Rating: LOW

Recommended adjustments:

Enable continuous wearable tracking

Standardise daily routes and routines

Align activity patterns with recognised service zones

But should she have to change?

The problem isn't Sem's behaviour.

It's the system's inability to make sense of it.

What happens when optimisation defines urban life?

Can you refuse to participate in systems like this?

And who gets left out, again and again?

How visible are you?

Start your visibility assessment now.

Alternate future
Loading...

Please remain still during the simulation

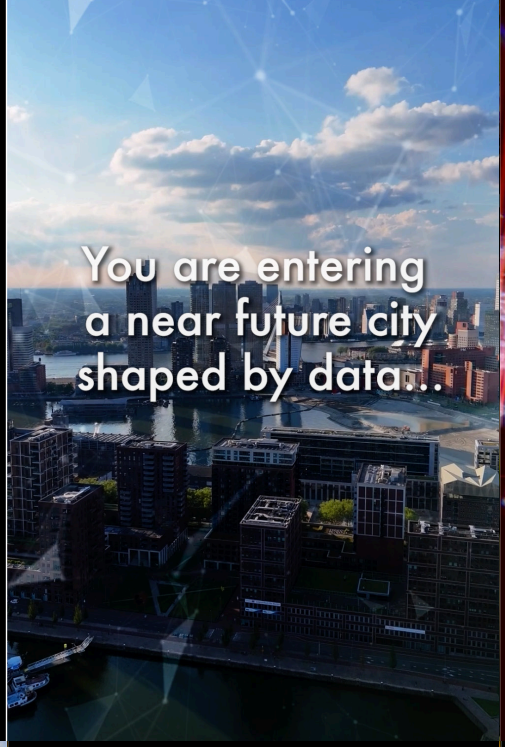

But not all
patterns are easy
to capture...

Not all
lives fit
"The Model"
...

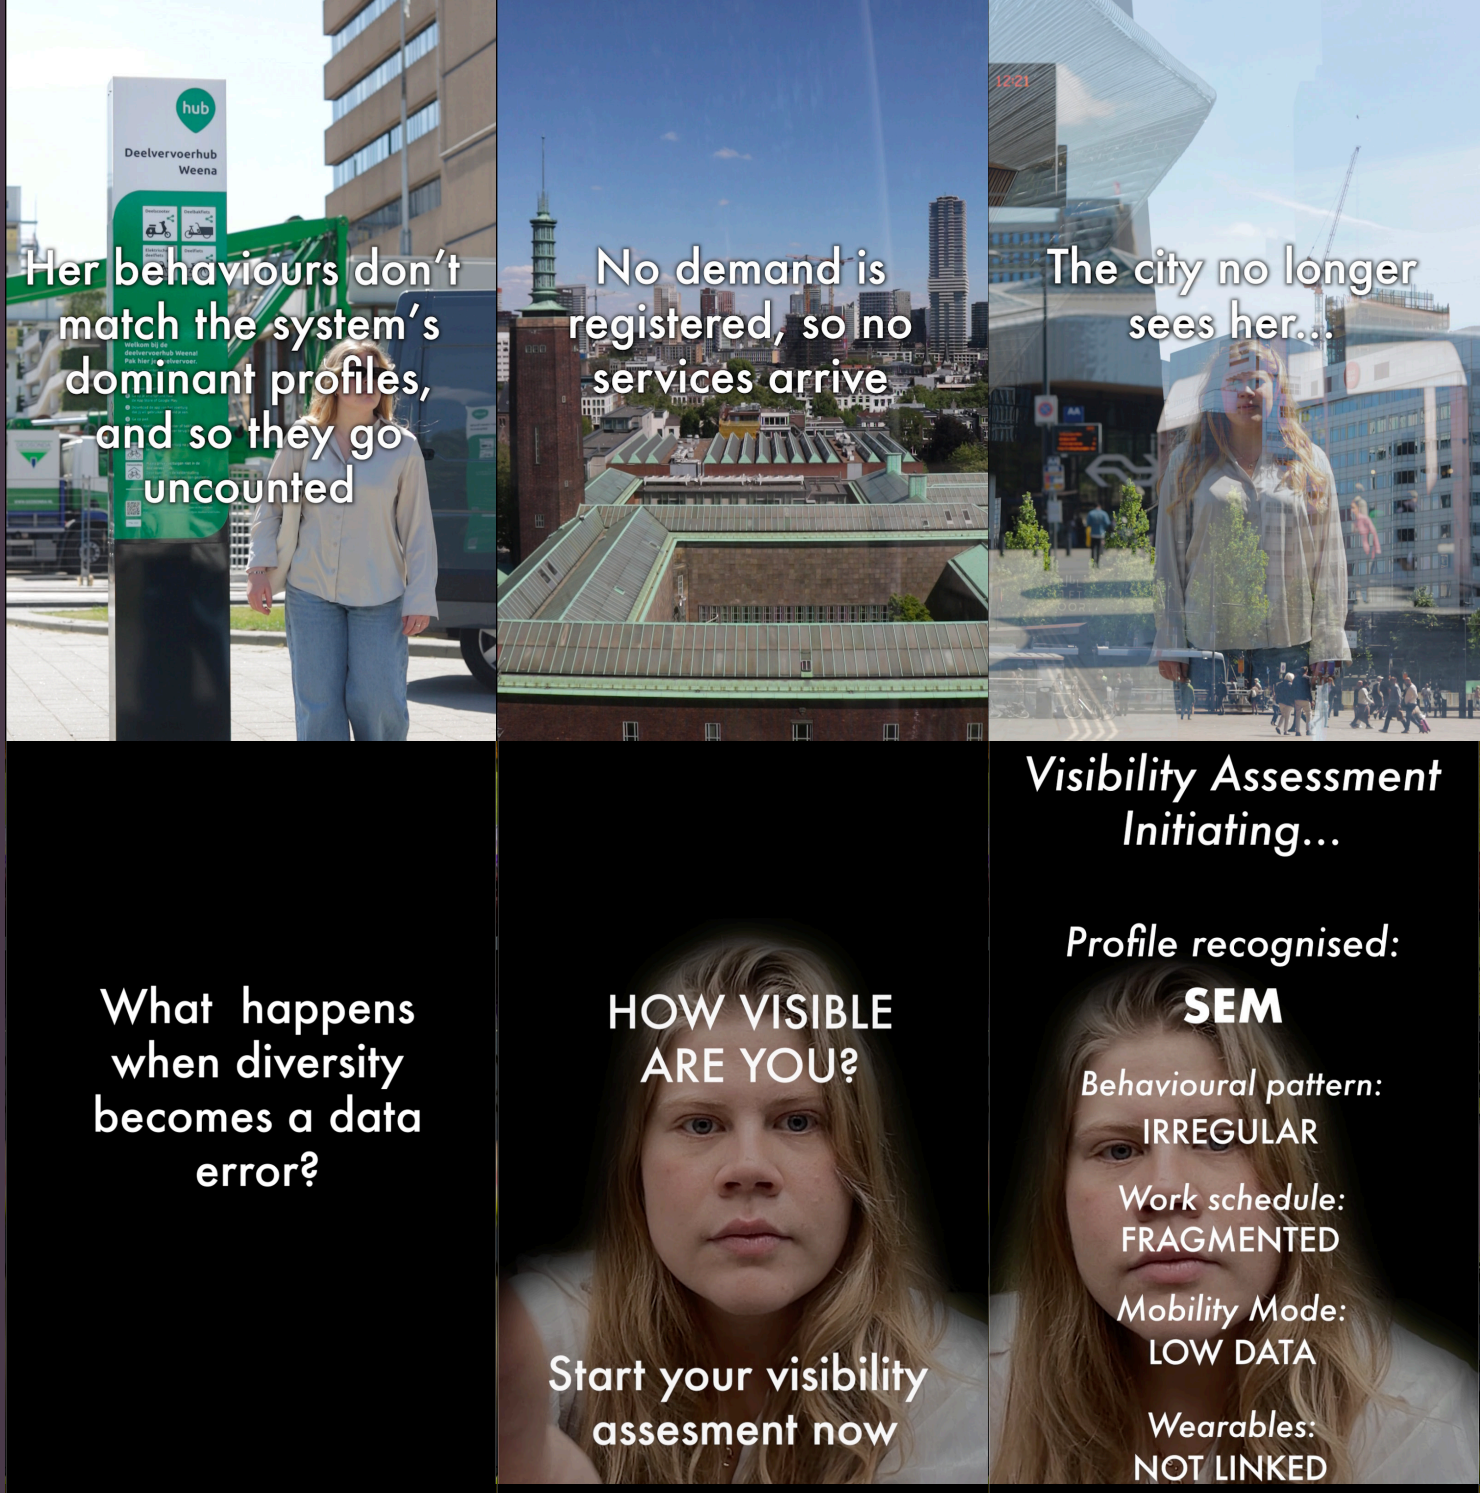

Figure 8.7 - Stills from the Behavioural Mirror video

Figure 8.8 - The Visibility Mirror in use

Figure 8.9 - The Visibility Mirror in use

8.4 Evaluation Round 2: methods & key findings

Where the first evaluation round primarily surfaced issues of clarity, comprehension, and methodological friction, this second iteration aimed to test whether the refined Behavioural Mirror could effectively function as a reflective prompt. Rather than focusing on whether the message came across, this evaluation asked: Did it provoke the intended questions? Did it shift perspectives? Did it support critical thought among its target audience?

The goal was not to fix misunderstandings, but to examine the artefact's effectiveness in raising ethical reflection around optimisation, visibility, and exclusion in Urban Digital Twin systems.

8.4.1 Testing Approach

The second evaluation round consisted of individual sessions with a total of seven participants. These sessions aimed to explore whether the refined Behavioural Mirror could meaningfully prompt reflection and critical thought about systemic exclusion, visibility, and optimisation

Participants included

- Four in-person sessions with advisors from Advier, all of whom work with or around urban data and mobility
- Three online sessions with experts from the municipalities of Amsterdam, The Hague, and Utrecht, all working on digitalisation or digital twin strategies

Each session lasted approximately 30 minutes and followed the same general structure:

- A short introduction to the project and its speculative design goals
- Viewing of the updated Behavioural Mirror video
- A semi-structured discussion, based on the same prompts used in Evaluation Round 1

The sessions with Advier were conducted in person. The municipal expert sessions were held online, and participants were invited to imagine the artefact as part of a public exhibition or as a reflection tool within their own policy contexts.

The following sections present the reflections from both participant groups separately.

8.4.2 Reflections from Advier Participants

a) Recognition of Structural Bias and Systemic Logic

Participants widely recognised that the artefact surfaced how optimisation systems are not neutral, but embed particular values and norms. The speculative scenario was described as *"verontrustend realistisch" ("**disturbingly realistic"**)*, with multiple participants pointing out the risk of designing urban systems based only on what is measurable. *"We doen dit eigenlijk al… wat zichtbaar is telt mee, wat niet zichtbaar is telt niet." ("**We are already doing this… what is visible counts, what is not visible does not."***

The scenario successfully made visible the quiet exclusions built into optimisation logic, especially through data reliance: *"Het raakt aan het ongemak dat sommige mensen gewoon buiten beeld vallen." ("**It touches on the discomfort that some people simply fall out of view."**)*

b) Critical Reflection on Data Participation and Autonomy

Several participants discussed whether opting out of data systems is still a real option. While some noted the right to be invisible, others framed it as a false choice, refusing to participate often results in being structurally excluded. *"Het gaat over gedrag, maar ook over keuze. Kun je eigenlijk nog wel kiezen om níet mee te doen?" ("**It is about behaviour, but also about choice. Can you still choose not to join in?"**)*

This led to reflections on how behavioural conformity is incentivised through system design, rather than imposed explicitly: *"Het systeem straft niet letterlijk, maar je voelt toch druk om je aan te passen." ("**The system does not punish you directly, but you still feel pressure to conform."**)*

c) Value in Policy and Strategy Settings

Participants affirmed the potential of the artefact as a strategic tool for municipal or advisory sessions. Because of its ambiguity and emotional distance, the mirror was seen as an effective conversation starter: *"Het zegt niet: dit is fout. Maar het nodigt je uit om te denken: hoe doen wij dit eigenlijk?" ("**It does not say: this is wrong. But it invites you to think: how do we actually do this?"**)*

One participant stated clearly: *"Ik zou dit zo gebruiken in een strategiesessie. Het maakt de onderliggende logica zichtbaar." (**"I would use this directly in a strategy session. It reveals the underlying logic."**)*

8.4.3 Reflections from Expert Evaluations

a) System realism and strategic urgency

Each expert treated the scenario as a short-term reality rather than distant fiction. One commented, *"Dit gebeurt al, alleen minder uitgesproken" ("**This is already happening, just in a less explicit way"**)*. Another described the viewing experience as *"een beetje onheimisch... een beetje Black Mirror-achtig" ("**a bit uncanny, very Black Mirror-like").*** The third stressed inevitability by saying, *"Het is niet de vraag óf dit gebeurt, maar wanneer" ("**The question is not if this happens, but when"**)*. These remarks confirm that the artefact successfully frames optimisation as an immediate governance concern.

b) Decision-making, system logic and the transparency deficit

All three experts said the mirror helps shift attention from individual behaviour to embedded algorithms. One observed, *"Het legt bloot hoe de logica van het systeem beslissingen overneemt" (**"It exposes how the system's logic starts taking the decisions"**)*. They also stressed that real-world safeguards are lagging. As one put it, *"Ik vind dat die transparantie en de capaciteit om kritisch te reflecteren nog onderbelicht blijven" ("**Transparency and the capacity for critical reflection are still under-exposed"**)*. The artefact therefore lands within a landscape where accountability mechanisms are needed but not yet standard.

c) Autonomy, opt-out myths and quiet exclusion

Concerns about genuine choice surfaced in every interview. One expert asked, *"Kun je eigenlijk nog wel kiezen om níet mee te doen?" ("**Can you still choose not to take part?"**)*. Another noted, *"Het systeem straft niet letterlijk, maar je voelt toch druk om je aan te passen" (**"The system does not punish you directly, yet you still feel pressure to conform"**)*. A third emphasised the stakes of invisibility by remarking, *"Je ziet ineens wie niet past in het model. Dat komt anders nooit zo scherp in beeld" ("**You suddenly see who does not fit the model; that usually never becomes this clear"**)*. Collectively these comments show the mirror surfaces the penalties attached to non-conforming data profiles.

d) Organisational application and the risk of stalled action

All experts could position the artefact inside existing workflows. A typical reaction was, *"Ik zou dit gebruiken als reflectietool, zeker in overheidscontext. Je praat ineens over dingen die normaal onder de oppervlakte blijven" ("**I would use this as a reflection tool, especially in government settings.***

You suddenly talk about things that usually stay below the surface"). However, they warned that dialogue must convert into delivery. One cautioned, *"Het kan ook innovatie juist in de weg zitten als je té veel in de reflectiestand blijft hangen" (**"It can also stifle innovation if you stay stuck in reflection mode"***). Suggested counter-measures included bias audits, model cards and decision logs that assign clear ownership for follow-up.

Overall reading

The experts describe the Behavioural Mirror as a credible and unsettling snapshot of near-future governance. It exposes the invisibility of those who fall outside data norms and highlights gaps in current transparency practices. They value its potential as a catalyst for strategic and ethical discussions but insist that organisations must pair the reflection it provokes with concrete governance tools and time-boxed actions. Otherwise its critical insight risks remaining a mirror rather than becoming a lever for change.

8.4.4 Reflection on Outcomes

The second round of evaluations showed a clear shift from methodological feedback to critical engagement with the artefact's underlying message. While the first iteration required clarification and adjustment to properly communicate its speculative logic, this updated version of the Behavioural Mirror succeeded in prompting reflection, questioning, and even self-positioning among participants.

Participants not only recognised the systemic logic of visibility and exclusion, but also began to interrogate their own roles and practices within it. The artefact no longer asked "Did you understand the story?" but instead provoked: "How do you participate in systems like this?" and "Can these logics be challenged from within?"

The most successful interactions emerged when the ambiguity of the artefact was preserved but scaffolded by the refined narrative. The addition of guiding text, a clear perspective, and sharpened system critique allowed the ambiguity to become productive rather than confusing. Instead of leading to disengagement, it encouraged layered interpretation and emotional identification.

Importantly, participants from both policy and advisory roles confirmed that the artefact could have real value in professional contexts. They suggested it could be used in strategic workshops, policy development sessions, or even ethics training modules. The speculative mirror was seen not just as a provocation, but as a practical tool for critical reflection on algorithmic governance and digital urban futures. After the sessions, all three external experts either asked for a copy of the Behavioural Mirror video to use in their own policy or ethics workshops, or offered to connect the project with colleagues who could. This enthusiasm signals immediate interest beyond the project team.

This round of testing shows that the Behavioural Mirror, in its refined form, can fulfil its intended role: an artefact that challenges, reveals, and opens space for thoughtful speculation about what we optimise, what we overlook, and who gets counted in the cities of tomorrow.

8.5 Updates for Standalone Use

Following the second evaluation round, no major changes were made to the core speculative artefact. However, small additions were introduced to support its use outside of facilitated sessions. The aim was to make the Behavioural Mirror legible and engaging for viewers encountering it independently, for example in an exhibition or public policy setting.

Two key additions were made:

Informative Poster
A poster was developed to provide visitors with light background information about the context of Urban Digital Twins, the logic behind the artefact, and the questions it aims to raise. It avoids technical depth but offers just enough framing to support interpretation without requiring facilitation (see figure 8.10).

Reflection Cards
A set of small A6-sized cards accompanies the artefact. Each card poses a different speculative question or prompt, inviting viewers to reflect on the system logic, their own behaviour, or the values embedded in optimisation. The cards are open-ended and ambiguous in tone, encouraging personal interpretation. They can be used individually or in group discussions (see figure 8.11).

These updates were not intended to clarify the design or explain away its ambiguity. Instead, they act as soft invitations into the speculative space, offering entry points for reflection while preserving the artefact's layered character.

Note that these materials have not been tested for effectiveness.

What happens when optimisation defines urban life?

Who gets counted in the city of tomorrow?

Urban Digital Twins (UDTs) are digital replicas of the physical city, built from real-time data. They promise efficiency, prediction, and responsivenes and help planners and policymakers act fast and allocate resources smartly. Yet **behind this promise lies a set of assumptions about what can be measured, who is visible, and which behaviours count in such a system**. This project explores what happens when urban systems begin to govern not just space and infrastructure, but also **legibility and inclusion**.

The Behavioural Mirror is part of a speculative design inquiry into the optimisation logics embedded in smart urban systems. With academic grounding drawing from **algorithmic governance, data feminism, and ambiguity in design** theories, this work examines how seemingly neutral categories (for example "demand" or "activity") can shape structural outcomes and influence the lives of citizens. The binaries behind this system (eg, high vs low demand) reduce complexity, turning diverse lives into predictable patterns. What falls outside these patterns often disappears from view.

This speculative artefact does not aim to fix the system, or to provide policy recommendations; Instead, it creates a space to confront the everyday consequences of being misrepresented, flattened, or left out by data-driven decision-making. Through the fictional scenario of Sem (a nurse whose care work and behavioural unpredictability render her invisible to a city's optimisation algorithm) the mirror surfaces a broader question: **What does inclusion mean, when shaped by prediction?**

Rather than providing a didactic message, the artefact relies on ambiguity to open critical reflection. While interacting with the mirror, it asks viewers to reflect on how they participate in, benefit from, or resist such systems. By exposing the friction between **efficiency and inclusivity**, it invites doubt, not only about future cities, but about the systems already shaping our present.

> ***"What gets counted gets cared about"***
>
> D'Ignazio & Klein, Data Feminism

Rethinking Binaries in Urban Digital Twins

A Speculative Design Approach

The Behavioural Mirror

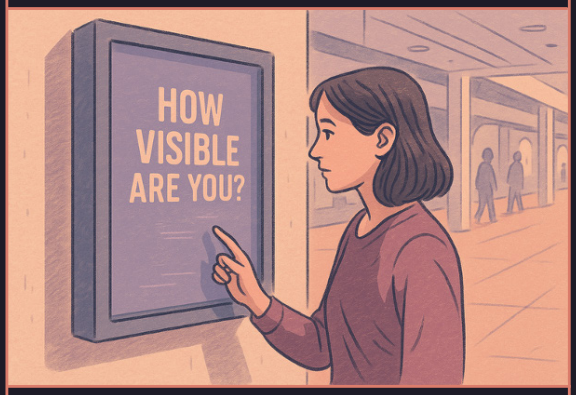

The Behavioural Mirror invites viewers to reflect on their own data visibility within a fictional smart city.

This speculative artefact imagines a system in which algorithmic optimisation governs access to urban services. Visibility is not based on who you are, but on how legible your behaviour appears to the system.

Behaviours that are regular, measurable, and efficient are rewarded; those that are informal, care-based, or irregular often fall outside the frame. Over time, these omissions do not only result in exclusion, but in a quiet disappearance from the city's logic.

What Questions Does This Raise?

Visibility Bias	Who gets seen, and who stays invisible in data systems?
Legibility	What kinds of lives are readable to optimisation models?
Care & Exclusion	Why are irregular or care-based routines often overlooked?
Autonomy	Is opting out of digital systems still a real choice?
Ambiguity	Can design invite doubt instead of offering clear answers?

TU Delft

Figure 8.10 - The Informative Poster (A1)

Master Thesis Project by
Isabella Maljers

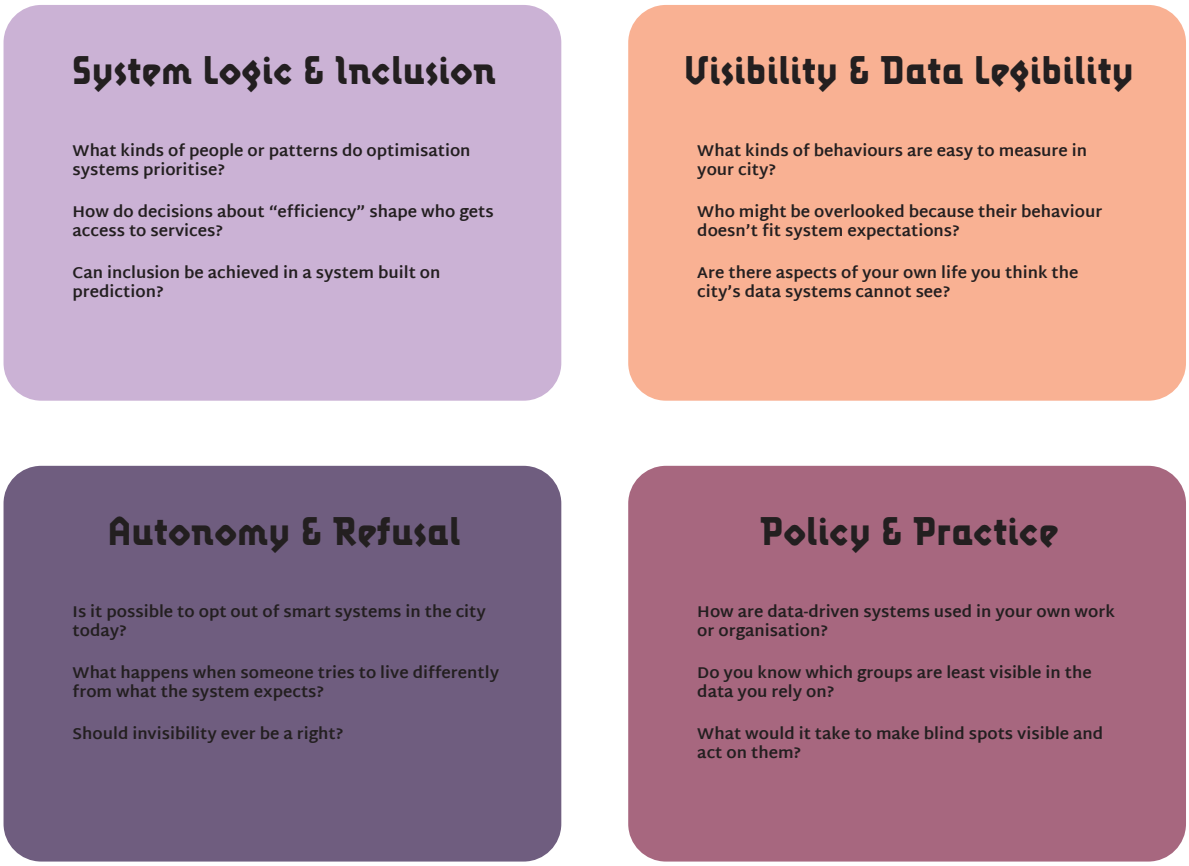

Figure 8.11 - The scenario Cards

8.6 How can this artefact be used by professionals?

The Behavioural Mirror was not developed as a predictive model or operational tool, but rather as a reflective instrument for professional contexts. During the evaluation sessions, participants repeatedly recognised its potential for use in policy, ethics, and innovation settings. Based on their feedback and the design intent, the artefact can support reflection in the following ways:

- **In strategic workshops**, to surface unspoken assumptions embedded in data practices or optimisation strategies
- **In ethics or innovation sessions**, as an entry point into discussions around invisibility, surveillance, and system logics
- **Within cross-disciplinary teams**, to create a shared reference when working across technical, policy, and social domains
- **In training or onboarding contexts**, as a way to introduce critical perspectives on digitalisation and urban AI
- **Used during events and innovation fairs** such as AI in de Stad, to trigger accessible dialogue on data-driven exclusion and urban decision-making.

Rather than aiming to tell professionals what is right or wrong, it creates space to question what may otherwise remain implicit. It relies on ambiguity and suggestion rather than instruction, which makes it particularly suited to settings where reflection and an open mind are valued.

When used alongside supporting materials (such as the poster, the question cards, or a facilitator script) the artefact can function as a conversation starter or as a thematic opener. Participants noted that the mirror allowed them to recognise how certain behaviours or groups become unintentionally deprioritised by system design, which suggests that speculative artefacts like this can definitely play a role in ongoing professional reflection, even outside academic or design-led environments.

Rather than framing the citizen as the problem, the Behavioural Mirror shifts the focus to the system's internal logic. In doing so, it encourages professionals to take a step back and ask: what kinds of urban lives are we designing for, and who might be left out?

Note: Since the evaluation session, two experts have already requested and used the video within their own professional settings as a conversation starter on this topic.

Chapter 9

DISCUSSION & RECOMMENDATIONS

This chapter reflects on how the project outcomes link back to the theoretical groundwork, and what this might mean for both design research and municipal practice. At its core, the work connects critical design with real-world concerns about data, governance, and ethics, and so it becomes important to revisit the theoretical strands that informed this project, and to explore how they were deepened or challenged through practice.

The first section connects the project back to the four main themes explored previously in this report: speculative design as a method of problem finding, ambiguity as a critical design strategy, the tension between efficiency and inclusivity, and data feminism as a design lens.

The second part outlines the key limitations encountered during the process; acknowledging both the practical constraints as well as the conceptual simplifications. I also reflect on how certain choices in the narrative and artefact format affected its interpretation and reach.

The final section looks to potential future directions: Rather than framing them as fixed next steps, they are positioned as provocations for continued experimentation. What else could speculative artefacts do within civic settings? How might future versions be adapted to explore neighbouring issues such as surveillance, participation, or resistance? These reflections are not a conclusion, but an opening; they invite others to extend, challenge, or reframe the work in ways that respond to their own contexts.

9.1 Linking Back to Literature

As a research-through-design project, it is imperative to link the findings of the research from the artefact back to the literature studies and research conducted at the start of the project. This chapter therefore returns to the three key strands of design research and critical scholarship that were introduced in chapter X: speculative design, ambiguity in reflective practice, and the goals of efficiency and inclusivity. It shows how the behavioural mirror's iterative deployment both validates and extends these theories. In the following sections I synthesise findings by showing how the mirror made each of these theoretical domains tangible and actionable for participants.

9.1.1 The Use of Speculative Design

Firstly, the project has questioned "quick-fix" solutionism, recognising its limitations in addressing complex, multifaceted societal challenges within the realm of urban digital twins. Speculative design, as a methodology, thrives on ambiguity and critical exploration rather than providing definitive answers or readily implementable solutions (Galloway & Caudwell, 2018). The immediate objective is to make stakeholders aware that a problem exists before any solution can be considered.

This approach aligns with the understanding that many of the issues surrounding digital twins are deeply rooted in systemic biases, ethical considerations, and power dynamics that cannot be resolved through simple technological fixes. The

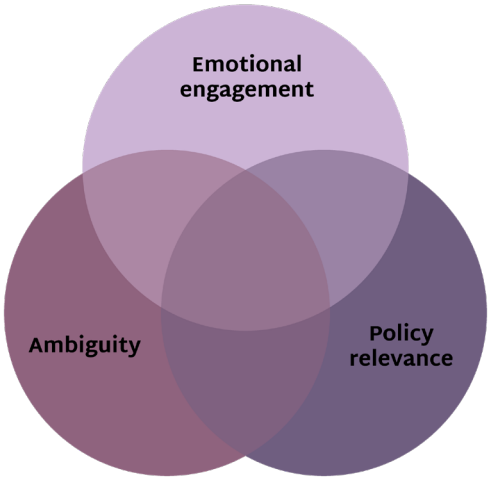

Figure 9.1 - Balancing ambiguity, emotion, and policy relevance in the Behavioural Mirror.
The artefact was iteratively shaped to sit at the intersection of these three design goals, ensuring it remained both reflective and actionable.

primary aim, therefore, is to spark critical conversations and alert stakeholders, citizens and policymakers to the pitfalls and unintended consequences of relying solely on data-driven models for urban governance.

During discussions with different organisations, I encountered strong demand for ready-made solutions, particularly from municipalities still defining the problem space. Yet the evaluations showed that the artefact itself translated abstract ideas into something tangible, helping project leaders articulate concepts such as visibility bias to their peers and subordinates. In doing so, the mirror has created a reflective space that can later be used to develop well-grounded, value-aligned interventions.

9.1.2 Ambiguity as a Catalyst for Ethical Reflection

Within this project, ambiguity is seen as an oversight, but a deliberate design strategy. In speculative design, these open meanings can keep audiences thinking rather than rushing to tidy conclusions (Gaver et al., 2003). In the first evaluation round, however, several participants expressed confusion with regards to their understanding of the situation. This confirmed Giaccardi et al., (2024)'s warning that ambiguity becomes unprovocative when it offers no foothold for interpretation.

Small narrative edits in the second iteration clarified why Sem might be excluded and added subtitles throughout the video, converting puzzlement into engagement. Participants began asking how they themselves participate in optimisawtion systems and whether Sem's right to refuse data is meaningful if it entails losing services. One municipal expert noted that the video felt quite realistic and uncomfortable, yet clear enough to prompt policy debate.

This supports Gaver's claim that interpretive flexibility is most effective when balanced with a limited set of anchors. Within the second iteration this was applied by focussing the ambiguity onto one specific element (the system's intentions) while clarifying three others (context, stakes and prompt). Maintaining this allowed the artefact to function as what Giaccardi calls a "political ambiguity", provoking ethical reflection without dictating a moral lesson. Because the artefact was iterated upon between the two test cycles, the design itself became an analytic tool, turning the first-round confusion into second round insight.

Importantly, the evaluation also surfaced a second-order form of ambiguity: the invisibility as both deprivation and potential autonomy. Future studies might lean into this, exploring the off-grid pathways and the autonomy of the person with regards to their urban digital footprint. This points to a wider research opportunity: using ambiguity not only to critique data legibility, but to imagine alternative urban futures where selective legibility is a valued right.

9.1.3 Efficiency and Inclusivity: The Optimisation Dilemma

Through the context mapping within my project, the relationship between efficiency and inclusivity came to be a leading thread. While initially framed as opposing forces, further research and reflection revealed that this perceived tension is often constructed by the specific optimisation logics embedded within Urban Digital Twins. The literature positions efficiency as the ability to allocate resources swiftly, cheaply and with predictive certainty, whereas inclusivity measures how far a system recognises and supports a spectrum of urban lives that resist neat classification. However, fieldwork showed that these definitions of "efficiency" are never neutral; they already embed value judgement about whose behaviours, needs and contributions the city prefers to recognise and serve.

Both the literature and our evaluation data demonstrate that pursuing efficiency without inclusivity simply shifts costs downstream, often onto the very populations optimisation claims to serve. In fact, data-feminist work argues that what gets counted gets cared about whilst what remains uncounted drifts out of the decision frame (D'Ignazio & Klein, 2020). When optimisation sidelines those who are less represented by the data, they can create urban invisibility, or blind spots that have direct effect on the lives of citizens (Benjamin et al., 2021). This message was also clearly conveyed to the participants as was suggested within the evaluation sessions. However, the story within the artefact was only able to show one simplified scenario; the reality is of course a lot more complex.

The project therefore creates an opening to reframe efficiency as a measure that must account for the costs of exclusion. A city that delivers services quickly for the already legible while forcing everyone else into work-arounds is only shifting cost downstream. It can be proposed that performance metrics should flag data-poor areas, weight accuracy by coverage breadth and set aside budgets specifically for areas or groups where there is little to no data. This converts unseen absence into actionable signal rather than an ignored anomaly.

In essence, overlooking certain populations isn't a shortcut to efficiency; it's a significant vulnerability. Genuine optimisation hinges on acknowledging and accommodating the entire urban demographic. Anything less sacrifices lasting systemic progress for immediate gains.

9.1.4 Legibility, Counting and Data Feminism

Data-feminist scholars argue that what gets counted gets cared about and conversely, that what remains uncounted drifts outside the field of concern (D'Ignazio & Klein, 2020). The artefact created in this project, The behavioural

mirror, translates that principle into a visibility metric. Sem's visibility score, makes it clear that optimisation logic privileges data-rich citizens while quietly deprioritising those whose lives are harder to measure. During testing, a participant captured the mechanism in everyday language: "what is visible counts, what is not visible does not.", by placing that remark next to the visibility percentage, the artefact turns an abstract critique into an experimental moment of recognition.

The evaluation also revealed that visibility is not a straightforward benefit. Several participants noted that being seen can expose residents to unwanted oversight or profiling. One expert remarked "Sometimes you might actually not want to be seen.", pointing to the tension between access and autonomy. This echoes D'Ignazio and Klein's warning that data regimes often offer inclusion at the price of intensified surveillance. The project therefore confirms two linked insights from the literature: The first, that counting distributes value and neglect simultaneously; and second, that legibility is always ambivalent, conferring resources while eroding anonymity.

A design implication follows that any municipal digital twin dashboard should include an overview of data that showcases the people, places and behaviours that are left off the map, to prompt planners to ask where those absences stem from (personal choice, technical oversight or systemic bias). Incorporating this kind of feature would operationalise the data-feminist call for "making the invisible visible" and would give officials a concrete prompt to question whose needs remain unaddressed by optimisation logics.

9.1.5 Reflection Across Themes

Across these three dimensions, I have reflected on how the behavioural mirror has shown how data-driven urban governance can embed systemic bias unless designers and policymakers deliberately intervene. Speculative framing opened space for problem-finding and ambiguity sustained ethical enquiry. These findings are backed by literature studies, while also demonstrating that carefully crafted artefacts can provoke the questioning needed to surface these hidden assumptions. By creating moments of disquiet and dialogue, the mirror does not prescribe solutions, but rather cultivates the critical stance required before any design intervention can be responsibly pursued. In the next chapter, I will draw on this groundwork to outline potential paths for prototyping and stakeholder engagement that builds on the reflective space the mirror has created.

9.2 Project Limitations

Even though this project offers insights into the ethical implications of Urban Digital Twins and the design of reflective speculative artefacts, it is also shaped by its framing, context, and constraints. In this section I would like to go through some of the key limitations that emerged throughout the process. They are grouped into three themes: contextual constraints, design decisions, and interactional interpretation. These limitations do not aim to discredit the project, but rather help clarify its scope, relevance, and opportunities for future adaptation.

9.2.1 Contextual Constraints

Field Evolution - With the rapid evolution of technology, this field of Urban Digital Twins has also been developing at a high speed. Since the start of this project, new developments have emerged that already slightly shift the contextual landscape. Within the evaluation interviews, I was also brought to speed on new developments that I did not even take into account within this project. This begs to question how long this project will still remain relevant.

Cultural and Contextual Specificity - This project draws on examples from the Dutch and European context, and so its framing may not directly translate to global contexts without adaptation. The artefact does not aim to be universally applicable, but rather aims to provoke context sensitive reflection.

9.2.2 Design Decisions and Framing

Role of the Mirror - While the reflective surface of the mirror was designed to provoke self-recognition and reflection, within the evaluation sessions online I was unable to show it in that format. Even so the video was said to convey the message without the mirrored element, which raises questions about the added value of it. The design choice was retained to preserve its ambiguous quality, but this may differ per context.

Simplification of Data Capture- The scenario implies that visibility is mostly determined through personal devices, whereas in reality UDTs integrate multiple sources, including environmental sensors and infrastructural data. This simplification was a conscious choice to focus on the critique, but does not represent the full complexity of the real world system.

Speculative Scope - Within the scenario narrative, only one story is told: that of Sem. While this made the systemic critique tangible, it also risks limiting the interpretive possibilities. The speculative framing deliberately narrowed focus to highlight systems logic, but at what cost?

9.2.3 Interpretation and Interaction

Interpretation Dependence - The artefact invites reflection through ambiguity without offering explanation and while this strategy can be powerful in guided sessions, in autonomous settings there is a risk that viewers may miss or misread the intended provocation. The poster and question cards were thus added to support interpretation, while preserving the open-ended tone.

Artefact Setup Requirements - The artefact is quite dependent on the environmental factors such as lightning, screen quality, and positioning. If these are not well calibrated, the interaction may lose its impact.

9.2.4 Reflection

Together, these limitations illustrate the boundaries of the project's scope, and the trade-offs made to balance narrative clarity with speculative provocation. Acknowledging these constraints also reinforces the need for context-sensitive, adaptable tools when designing for critical reflection in public systems.

9.3 Future recommendations

Although this project does not offer direct solutions, several future directions have emerged from the research journey and evaluation sessions. These can be considered open invitations to continue the conversation across the academic and non academic fields. Some suggestions build directly on the artefact's speculative logic, others on the insights and questions shared by participants and professionals. Together, they offer starting points for extending or adapting the work in new contexts.

9.3.1 Towards a More Personal Mirror

While the Behavioural Mirror currently presents a fictionalised encounter through the story of Sem, some participants expressed curiosity about how their own behaviour would be assessed by such a system. A future iteration could allow users to interact with more personalised or semi-personal data traces, perhaps through scenario choices or behavioural prompts. This could help make the critique more immediate, encouraging users to recognise the invisible trade-offs they already make in daily life.

At the same time, this personalisation must be handled carefully. The goal is not to create a gamified experience or promote compliance, but to deepen reflection.

9.3.2 Exploring Invisibility as Agency

One of the more surprising and consistent reflections from participants was the idea that invisibility might not always be a problem; it might be a choice. While many worried about being excluded from optimisation systems, others saw invisibility as a kind of refuge. One participant noted that not being visible could also mean not being tracked, and maybe even living more freely.

This speaks to a deeper tension in urban data systems. Visibility is often framed as a good thing, tied to access and inclusion. But what if opting out could be a meaningful choice? What if invisibility was not only a symptom of marginalisation, but also a way to reclaim autonomy?

Future artefacts could explore this more fully. Rather than visualising invisibility as a lack or a failure, they might present it as a political or personal stance. Of course, not everyone can choose to or wants to disappear, but invisibility is not always voluntary. But by surfacing these tensions, design can help expand this debate beyond visibility as the only goal.

9.3.3 Use in Civic and Policy Settings

Several professionals who engaged with the artefact expressed interest in using it in their own practice. This included settings like innovation workshops, municipal strategy days, or thematic team discussions. They saw the artefact not as a tool, but as a conversation starter, something that can spark debate about AI ethics, data governance, and behavioural profiling.

This opens a clear opportunity of developing a lightweight facilitation format for the Behavioural Mirror. This could, for example, include a session guide and a small set of discussion prompts. The goal would not be to translate the artefact into a toolkit, but to support its use in spaces where speculative design is not the norm.

9.3.4 Exhibition and Public Engagement

Given its speculative tone and layered message, the project could be well suited to public exhibitions or civic events. These might include places like Dutch Design Week, the AI & the City initiative, or smaller local festivals where design meets debate. In these spaces, the artefact can offer a different kind of entry point into complex topics, one that does not require prior knowledge, but invites slow thinking and shared reflection. An aspect to note for this is of course the accessibility of such a complex topic.

The current setup already supports stand-alone presentation, combining the mirror, video, and poster. Future versions might include multilingual adaptations, branching narratives, or for example contextual guides to support wider audiences. The intention would still be not be to teach or explain, but to hold space for discomfort, curiosity, and doubt.

9.3.5 Alternative Futures and Fringe Logics

While this project focused on a plausible near-future scenario, several conversations touched on more radical possibilities. What if people actively tried to disappear from the system? What if invisibility was not just tolerated, but desired? These kinds of ideas already exist in art and activism such as in apps that help people avoid surveillance cameras, or services that allow individuals to vanish from digital records.

Future speculative work could build on these ideas, not as science fiction, but as provocations that challenge the current assumptions of what it means to be a citizen in a datafied city. These fringe imaginaries raise questions that existing systems tend to ignore; They help stretch the speculative frame beyond optimisation and legibility, towards autonomy and refusal.

Such futures would not replace the scenario shown in this project, but could exist alongside it. They would serve as parallel reflections, expanding the conversation and revealing the limits of current system logic.

Chapter 10

CONCLUSION & REFLECTION

In this chapter I bring together the theoretical, empirical and design strands of this thesis. I revisit the core research question, exploring how optimisation logics in Urban Digital Twins shape inclusion and invisibility in contemporary cities and how speculative design has revealed hidden assumptions about data, visibility and governance. I start by summarising the four central contributions of this work. I then reflect on methodological insights gained from working with artefacts in the Research-through-Design process. Next I outline the implications for future scholarship and practice in smart-city governance. Finally I offer a personal reflection the design process and on embracing uncertainty and on the role of design as a form of critical inquiry.

10.1 Key Contributions & Reflection on the Research Question

This project set out with the question:

> *How do optimisation logics in Urban Digital Twins impact the relationship between efficiency and inclusivity, and what implications does this have within the field of urban governance?*

Rather than seeking a single or definitive answer, the project aimed to surface the assumptions that shape how digital systems assign value, determine visibility, and govern access. Through speculative design, these assumptions were not only analysed but also made tangible and experiential. The result was the Behavioural Mirror, a provotype positioned between fiction and critique, designed to reflect not only the user but also the system itself.

One of the key contributions lies in how this work reframes optimisation. Instead of treating it as a purely technical goal, it is exposed as a cultural and political logic that embeds particular norms about what counts, who fits, and how urban life is made legible. The fictional scenario of Sem offered an accessible entry point into this critique, revealing how systems that rely on regularity, measurability, and behavioural predictability often sideline those who do not conform.

A second contribution is methodological. The mirror challenges conventional interface design by resisting clarity, instruction, and control. Instead, it invites ambiguity and interpretation. This deliberate openness proved effective in evaluation sessions, where participants engaged with the mirror not through explanation but through speculation. It demonstrated that artefacts can provoke ethical reflection without prescribing it, and that ambiguity, when carefully framed, can become a tool for situated critique.

The project also engages with ongoing discussions around visibility, legibility, and data justice, it challenges the idea that being seen is inherently beneficial. Both literature and participants pointed to the discomfort of visibility and the potential agency of invisibility. By visualising what is left out and questioning why, the artefact opens up space for alternative imaginaries.

In revisiting the research question, it becomes clear that efficiency and inclusivity are not simply goals to be balanced, but rather are concepts shaped by deeper system logics. The contribution of this project is not meant to be a solution or a tool, but also a critical lens made tangible. It holds up a mirror to optimisation culture and invites others, such as designers, policymakers, and publics, to look more closely at what is being reflected.

10.2 Implications of the Project

Although speculative by nature, this project is firmly rooted in ongoing urban transformations; its questions arise from the everyday work of policy, planning and data infrastructures, and the implications therefore extend beyond the Behavioural Mirror itself to anyone engaged with Urban Digital Twins, algorithmic governance or data-driven city systems.

Design as reflection, not solution

I suggest that one of the most significant lessons of this work lies in recognising design's capacity to surface latent assumptions rather than prescribe fixes; while many policy-oriented innovations prioritise optimisation and streamlined workflows, I show that a conceptually simple artefact can prompt participants to reconsider their own dashboards and articulate critical insights situationally, thereby surfacing unspoken norms and inviting deeper reflection on equity, visibility and bias.

Data visibility is political

This project underscores that visibility within data systems is never a neutral technical output but always a political choice; what gets measured and rendered legible frequently determines whose experiences are recognised and whose are obscured, and practitioners might therefore find it valuable to perform periodic legibility audits that reveal how categories, thresholds and algorithmic patterns contribute to the invisibility of certain populations.

Use ambiguity as method

Rather than treating ambiguity as a problem to eliminate, this work suggests that carefully crafted uncertainty can function as a productive design strategy; leaving room for interpretation allows diverse perspectives to surface and fosters more nuanced engagement, and integrating strategic ambiguity into civic-facing materials may encourage stakeholders to explore layered understandings of system impacts and to imagine multiple, plural urban futures.

Potential uses in practice

Evaluation sessions showed that many participants regarded the mirror as an effective conversation starter in team workshops, innovation labs and ethics discussions; several expressed interest in deploying it within municipal working groups. Future adaptations might package the Mirror into concise facilitation guides or tailor it for other domains (such as education, healthcare or energy) thereby embedding speculative design as an infrastructural tool for public reasoning rather than as a mechanism for delivering ready-made solutions.

10.3 Personal Reflections

When I started this project, I was keen and curious; in hindsight the whole direction the project would take was quite a blur, a grey blob yet to be defined. This was only enforced by the lack of consensus and definition within digital twins, and so a whole trajectory started where I myself had to understand what exactly I am looking into.

One of the great things about this cutting edge topic is that there seems to be an overwhelming interest. I had subscribed to the "AI in de Stad" (AI in the city) channel and found there to be relevant events happening in this topic nearly every month (some from other sources as well). With the help of this, and by reaching out to professional individuals, I was able to get quite a real-time image of the current state of Urban Digital Twins rather than relying solely on publications which were often written a while back. This was also somewhat disillusioning, as I had a different image of the state of implementation of technology before the start of the project. However, this approach also meant that my research process was not quite conventional. In my background and theory chapter you will find that the information is quite mixed with different sources overlapping. While unconventional, this fused structure mirrors my own exploratory path and, I believe, serves the report well by reflecting the iterative, inquiry-driven spirit of Research-through-Design.

The process of this project had its ups and downs; I spent a relatively long time in the initial stage grappling with its complexity, and in hindsight I would have liked to move more quickly into ideation. That is not to say I am unhappy with the outcome; although I would have liked the interface to be a little more interactive, the artefact succeeded in provoking the critical thought I had initially hoped for. The enthusiasm of municipal members (and their early steps toward implementation) showed me the real added value of this work. For my own practice, I recognise that I could streamline the early research phase to allow more time for creative exploration.

This journey also taught me video editing, acting, narrative writing and being crafty with simple materials to create effective prototypes. Learning to edit footage helped me refine the mirror's pacing and focus. Acting in the video prototype let me step into Sem's perspective and gauge emotional impact. Narrative writing gave shape to the scenario, while quick, hands-on prototyping showed me how low-fi mockups can spark richer feedback. Together, these diverse skills expanded my toolkit and reminded me that design research often blends creative craft with critical inquiry.

I personally really enjoy the fuzziness that comes at the start of a project, but

it is easy to get stuck in it. At some point the cloud needs to break and become raindrops grounded to the floor. Throughout this project I have, of course, had to explain what I am actually working on many times, and I've found that the title itself should come with an explanation book: first, what is a Digital Twin; okay, but what makes it urban; then what is speculative design; and finally data bias and binary logic. In the end I learned a two-sentence summary that avoids a five-minute lecture, yet I still watch people go wide-eyed or glaze over. That experience reminded me why clear, tangible artefacts (like the Behavioural Mirror) are so essential: they speak beyond language and let people discover the questions for themselves.

I do find myself looking differently at cities now, walking and wondering what is measured, what is or will be automated, and who will be affected. I also have found the value in ambiguity, and genuinely believe that it should be woven into all practices. After all, the world is not black and white.

REFERENCES

Afacan, Y., & Afacan, S. O. (2011). Rethinking social inclusivity: design strategies for cities. Proceedings of the Institution of Civil Engineers - Urban Design and Planning, 164(2), 93. https://doi.org/10.1680/udap.2011.164.2.93

Akter, S., McCarthy, G., Sajib, S., Michael, K., Dwivedi, Y. K., D'Ambra, J., & Shen, K. N. (2021). Algorithmic bias in data-driven innovation in the age of AI. International Journal of Information Management, 60, 102387. https://doi.org/10.1016/j.ijinfomgt.2021.102387

Ali, M. E., Cheema, M. A., Hashem, T., Ulhaq, A., & Babar, M. A. (2023). Enabling Spatial Digital Twins: Technologies, Challenges, and Future Research Directions. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2306.06600

Atwood, M. (2019, December). Margaret Atwood: Telling Tales From the Future.

Auger, J. (2013). Speculative design: crafting the speculation. Digital Creativity, 24(1), 11. https://doi.org/10.1080/14626268.2013.767276

Ávila Eça de Matos, B. (2023). EngD Thesis: Digital Twins for Cities. https://pure.tue.nl/ws/portalfiles/portal/299183465/20230615_SBC_vila_E_a_de_Matos_B.pdf

Benjamin, J. J., Berger, A., Merrill, N., & Pierce, J. (2021). Machine Learning Uncertainty as a Design Material: A Post-Phenomenological Inquiry. 1. https://doi.org/10.1145/3411764.3445481

Calvi, A. (2022). Gender, data protection & the smart city: Exploring the role of DPIA in achieving equality goals.

Caprari, G., Castelli, G., Montuori, M., Camardelli, M., & Malvezzi, R. (2022). Digital Twin for Urban Planning in the Green Deal Era: A State of the Art and Future Perspectives. Sustainability, 14(10), 6263. https://doi.org/10.3390/su14106263

Debaise, D., Terranova, F., Solhdju, K., & Pihet, V. (2017). Speculative Narration. https://doi.org/10.70733/hd829bdhweu2

Deren, L., Yu, W., & Shao, Z. (2021). Smart city based on digital twins. Computational Urban Science, 1(1). https://doi.org/10.1007/s43762-021-00005-y

D'Ignazio, C., & Klein, L. (2020). Data Feminism. In The MIT Press eBooks. The MIT Press. https://doi.org/10.7551/mitpress/11805.001.0001

DMI (Dutch Metropolitan Innovations) Birthday symposium. (2024, November 28).

Dunne, A., & Raby, F. (2013). Speculative Everything – Design, Fiction, and Social Dreaming.

Faliagka, E., Christopoulou, E., Ringas, D., Politi, C., Kostis, N., Leonardos, D., Tranoris, C., Antonopoulos, C., Denazis, S., & Voros, N. (2024). Trends in Digital Twin Framework Architectures for Smart Cities: A Case Study in Smart Mobility. Sensors, 24(5), 1665. https://doi.org/10.3390/s24051665

Filho, W. L., Wall, T., Mucova, S. A. R., Nagy, G. J., Balogun, A., Luetz, J. M., Ng, A. W., Kovaleva, M., Azam, F. M. S., Alvés, F., Guevara, Z., Matandirotya, N. R., Skouloudis, A., Tzachor, A., Malakar, K., & Gandhi, O. (2022). Deploying artificial intelligence for climate change adaptation. Technological Forecasting and Social Change, 180, 121662. https://doi.org/10.1016/j.techfore.2022.121662

Galloway, A., & Caudwell, C. B. (2018). Speculative design as research method. In Routledge eBooks (p. 85). Informa. https://doi.org/10.4324/9781315526379-8

Gaver, W., Beaver, J., & Benford, S. (2003). Ambiguity as a resource for design. 233. https://doi.org/10.1145/642611.642653

Giaccardi, E., Murray-Rust, D., Redström, J., & Caramiaux, B. (2024). Prototyping with Uncertainties: Data, Algorithms, and Research through Design. ACM Transactions on Computer-Human Interaction. https://doi.org/10.1145/3702322

Hartmann, M., Pawelzik, S., & Wimmer, M. A. (2024). Comparative Analysis of Digital Twins in Smart Cities.

Hassija, V., Chamola, V., De, R., Das, S., Chakrabarti, A., Sangwan, K. S., & Pandey, A. (2024). A Survey on Digital Twins: Enabling Technologies, Use Cases, Application, Open Issues and More. Deleted Journal, 1. https://doi.org/10.1109/jsas.2024.3523856

Hekkert, P., & van Dijk, M. (2020). Vision in Design: A Guidebook for Innovators. https://www.bispublishers.com/matthijs-van-dijk-and-paul-hekkert-vip-vision-in-p.html

Jackson, M. (2023). Uncertain: the Wisdom and Wonder of Being Unsure.

Jones, D., Snider, C., Nassehi, A., Yon, J., & Hicks, B. (2020). Characterising the Digital Twin: A systematic literature review. CIRP Journal of Manufacturing Science and Technology, 29, 36. https://doi.org/10.1016/j.cirpj.2020.02.002

Klein, L., & D'Ignazio, C. (2024). Data Feminism for AI. 2022 ACM Conference on Fairness, Accountability, and Transparency, 1989, 100. https://doi.org/10.1145/3630106.3658543

Korotkova, N., Benders, J., Mikalef, P., & Cameron, D. (2023). Maneuvering between

skepticism and optimism about hyped technologies: Building trust in digital twins. Information & Management, 60(4), 103787. https://doi.org/10.1016/j.im.2023.103787

Li, L., Aslam, S., Wileman, A., & Perinpanayagam, S. (2021). Digital Twin in Aerospace Industry: A Gentle Introduction. IEEE Access, 10, 9543. https://doi.org/10.1109/access.2021.3136458

Marcucci, E., Gatta, V., Pira, M. L., Hansson, L., & Bråthen, S. (2020). Digital Twins: A Critical Discussion on Their Potential for Supporting Policy-Making and Planning in Urban Logistics. Sustainability, 12(24), 10623. https://doi.org/10.3390/su122410623

Matos, B. Á. E. (2023). Digital Twins for Cities.

Mazzetto, S. (2024). A Review of Urban Digital Twins Integration, Challenges, and Future Directions in Smart City Development [Review of A Review of Urban Digital Twins Integration, Challenges, and Future Directions in Smart City Development]. Sustainability, 16(19), 8337. Multidisciplinary Digital Publishing Institute. https://doi.org/10.3390/su16198337

Muravyov, D. (2022). Doubt to be certain: epistemological ambiguity of data in the case of grassroots mapping of traffic accidents in Russia. Social Movement Studies, 23(3), 338. https://doi.org/10.1080/14742837.2022.2128327

Ortt, J. R., & Tiihonen, L. V. (2022). The Technology Monitor 2022 Digital Twins in an Urban Environment.

Overview of the AdaptiveCity / SmartCambridge / CDBB Digital Twin Program. (2020). https://pages.cdbb.uk/

Piromalis, D., & Kantaros, A. (2022). Digital Twins in the Automotive Industry: The Road toward Physical-Digital Convergence. Applied System Innovation, 5(4), 65. https://doi.org/10.3390/asi5040065

Rathore, M. M., Shah, S. A., Shukla, D., Bentafat, E., & Bakiras, S. (2021). The Role of AI, Machine Learning, and Big Data in Digital Twinning: A Systematic Literature Review, Challenges, and Opportunities. IEEE Access, 9, 32030. https://doi.org/10.1109/access.2021.3060863

Sanchez, T. W., Brenman, M., & Ye, X. (2024). The Ethical Concerns of Artificial Intelligence in Urban Planning. Journal of the American Planning Association, 1. https://doi.org/10.1080/01944363.2024.2355305

Sánchez–Vaquerizo, J. A. (2024). Urban Digital Twins and metaverses towards city multiplicities: uniting or dividing urban experiences? Ethics and Information Technology, 27(1). https://doi.org/10.1007/s10676-024-09812-3

van Apeldoorn, N., van Lamoen, J., Mayer, I., & de Kruijf, J. (2023). Making (common) sense of urban digital twins.

Vianen, J. C. (2022). MSc Thesis: A design approach to determine the locations in combination with amenities for neighbourhood hubs, based on user profiles.

Whize. (2025). Whize Segmenten . https://whize.nl/nl

APPENDIX

Click on the QR code or follow the link to access the Appendix.

https://drive.google.com/file/d/17BdRbNP9h3wbTmgIC01ALZAhWvM9ijAR/view?usp=sharing

INITIAL PROJECT BRIEF

Can be found in the next pages

Personal Project Brief – IDE Master Graduation Project

Name student Isabella Maljers **Student number** 4672682

PROJECT TITLE, INTRODUCTION, PROBLEM DEFINITION and ASSIGNMENT

Complete all fields, keep information clear, specific and concise

Project title Real vs Digital Reality: Using ambiguity through speculative design to bring attention to the data discrepancies used for digital twins in the city of Amsterdam

Please state the title of your graduation project (above). Keep the title compact and simple. Do not use abbreviations. The remainder of this document allows you to define and clarify your graduation project.

Introduction

Describe the context of your project here; What is the domain in which your project takes place? Who are the main stakeholders and what interests are at stake? Describe the opportunities (and limitations) in this domain to better serve the stakeholder interests. (max 250 words)

Within the urban context, digital twins, virtual representations of physical assets or systems, are increasingly utilised in urban planning, smart city development, and infrastructure management. These advanced models enable real-time monitoring and simulation of various urban elements, including traffic management and environmental monitoring. Within the scale of Gemeentes Local urban Digital Twins (LDTs) are being developed through projects such as the Twin4Resilience project (an EU funded project) with the goal of creating an ecosystem of mature and networked LDTs that help cities achieve economies of scale.

As with all data-based modelling, there is a high potential risk for the escalation of assumptions, bias and uncertainty within the resulted model (e.g. sample bias, representation bias, algorithmic bias, data incompleteness, measurement errors, model assumptions etc). Additionally, many qualitative humanistic aspects are usually ignored in data collection (e.g. “Thick data” in the form of motivations, preferences, emotions and experiences). These together result in a skewed model, which is assumed to be objective and true because of its basis being from real and actual data. Although there is a high potential of predictive modelling, the more important the decision making they will be used to make, the more risk there is to the urban citizen.

With this graduation project I would like to address these potential risks, by highlighting the gaps where these LDTs are not (fully) representative of the real using ambiguity as a tool to do this. The resulting research will accumulate in an interactive and physical design, which will bring attention to this disparity and create awareness for the citizens. This will be done using speculative design methods which are ideal to create a thought provoking, critical design for which I am aiming. I will be using the Twin4Resilence project in Amsterdam as a use case.

➔ *space available for images / figures on next page*

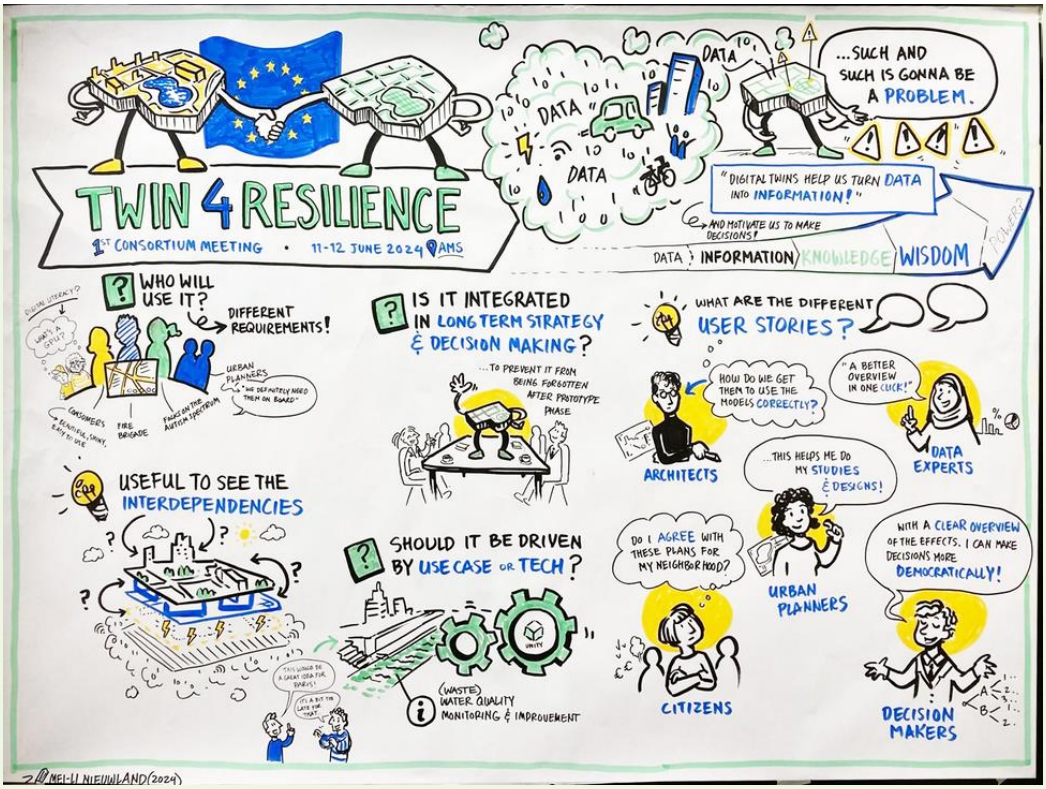

image 1- Introduction drawing to the Twin4Resilience Project

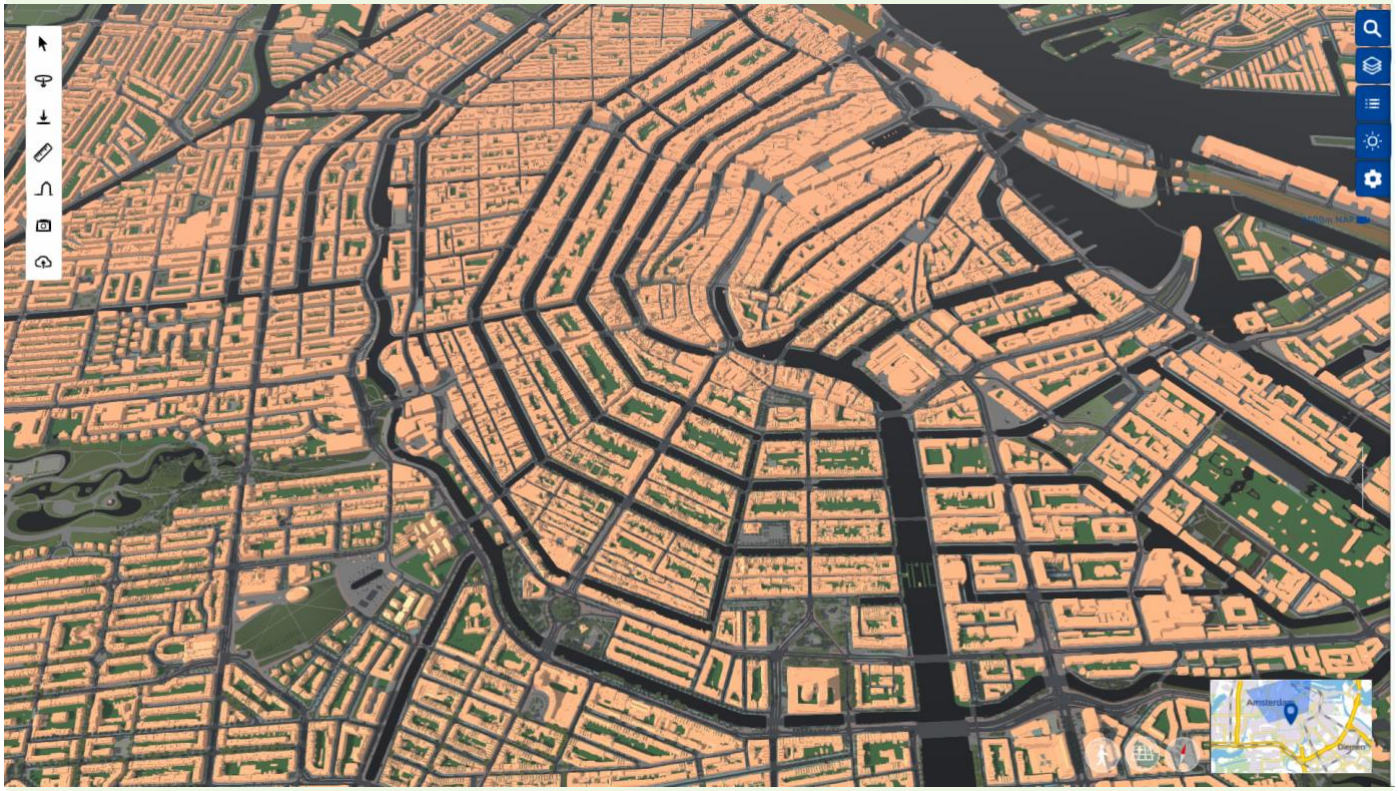

image 2 – 3D Amsterdam, a virtual model of the city. Also a basis for the digital twin

Personal Project Brief – IDE Master Graduation Project

Problem Definition

What problem do you want to solve in the context described in the introduction, and within the available time frame of 100 working days? (= Master Graduation Project of 30 EC). What opportunities do you see to create added value for the described stakeholders? Substantiate your choice.
(max 200 words)

Local Digital Twins (LDTs), are increasingly used in urban planning, smart city development, and infrastructure management. However, these models often fail to capture qualitative aspects such as motivations, preferences, emotions, and lived experiences. Additionally, the data and algorithms creating LDTs are opaque, leading to biases, uncertainties, and a false sense of objectivity. This poses risks, especially when LDTs inform high-stakes urban decisions.
To examine this, I have written the following research questions:
--> How do Local Digital Twins (LDTs) fail to fully represent qualitative humanistic aspects, such as motivations, preferences, emotions, and experiences?
--> What role can speculative and interactive physical designs play in increasing citizen awareness about the risks and limitations of Local Digital Twins in urban decision-making?
--> How can ambiguity be used as a tool in speculative design to highlight the limitations and biases inherent in Local Digital Twins?
--> To what extent does reliance on data-driven models in Local Digital Twins increase risks for urban citizens, particularly in high-stakes decision-making scenarios?

Assignment

This is the most important part of the project brief because it will give a clear direction of what you are heading for. Formulate an assignment to yourself regarding what you expect to deliver as result at the end of your project. (1 sentence) As you graduate as an industrial design engineer, your assignment will start with a verb (Design/Investigate/Validate/Create), and you may use the green text format:

To design and create an interactive speculative artifact that uses ambiguity to critique the limitations and biases of Local Digital Twins, fostering awareness and public discourse on their implications for urban decision-making.

Then explain your project approach to carrying out your graduation project and what research and design methods you plan to use to generate your design solution (max 150 words)

Discover- Conduct a thorough literature review of existing research on transparency, bias, and uncertainty in AI-powered decision-making systems, particularly in the context of digital twins. Find and analyse available data.
Stakeholder Engagement: Conduct interviews with governmental and non governmental parties working on developing urban digital twins; engage with citizens to understand their perspectives and level of knowledge.
Define -Define the design direction, the specific focus within digital twins, the role of ambiguity
Develop- use design techniques such as scenarios, speculative design fictions and design for ambiguity
- use rapinstallation, that communicates the key insights and challenges around digital twins
Deliver - Create a (physical) provotype (MVP) that communicates the key insights and challenges around digital twins.
Assess the effectivity of the design

Project planning and key moments

To make visible how you plan to spend your time, you must make a planning for the full project. You are advised to use a Gantt chart format to show the different phases of your project, deliverables you have in mind, meetings and in-between deadlines. Keep in mind that all activities should fit within the given run time of 100 working days. Your planning should include a ***kick-off meeting***, ***mid-term evaluation meeting***, ***green light meeting*** *and* ***graduation ceremony****. Please indicate periods of part-time activities and/or periods of not spending time on your graduation project, if any (for instance because of holidays or parallel course activities).*

Make sure to attach the full plan to this project brief.
The four key moment dates must be filled in below

Kick off meeting	20.11.2024
Mid-term evaluation	27.01.2025
Green light meeting	31.03.2025
Graduation ceremony	28.04.2025

In exceptional cases (part of) the Graduation Project may need to be scheduled part-time. Indicate here if such applies to your project

Part of project scheduled part-time	
For how many project weeks	
Number of project days per week	

Comments:

Motivation and personal ambitions

Explain why you wish to start this project, what competencies you want to prove or develop (e.g. competencies acquired in your MSc programme, electives, extra-curricular activities or other).

Optionally, describe whether you have some personal learning ambitions which you explicitly want to address in this project, on top of the learning objectives of the Graduation Project itself. You might think of e.g. acquiring in depth knowledge on a specific subject, broadening your competencies or experimenting with a specific tool or methodology. Personal learning ambitions are limited to a maximum number of five.
(200 words max)

- With the current AI-boom we are experiencing, I am particularly interested in maintaining a critical view. Rather than just looking Moat the benefits we should also look at the potential risks and particularly bring them to light to the people they will be affecting. I would love to be able to contribute in this manner.

- Personally, I am also very interested in the workings of assistive decision AIs and this topic would allow me to do a deep-dive into it.

- Speculative design is something I briefly was able to experiment with during my exchange at the Politecnico di Milano, and is an aspect of design that intrigues me and that I would like to be able to further explore.

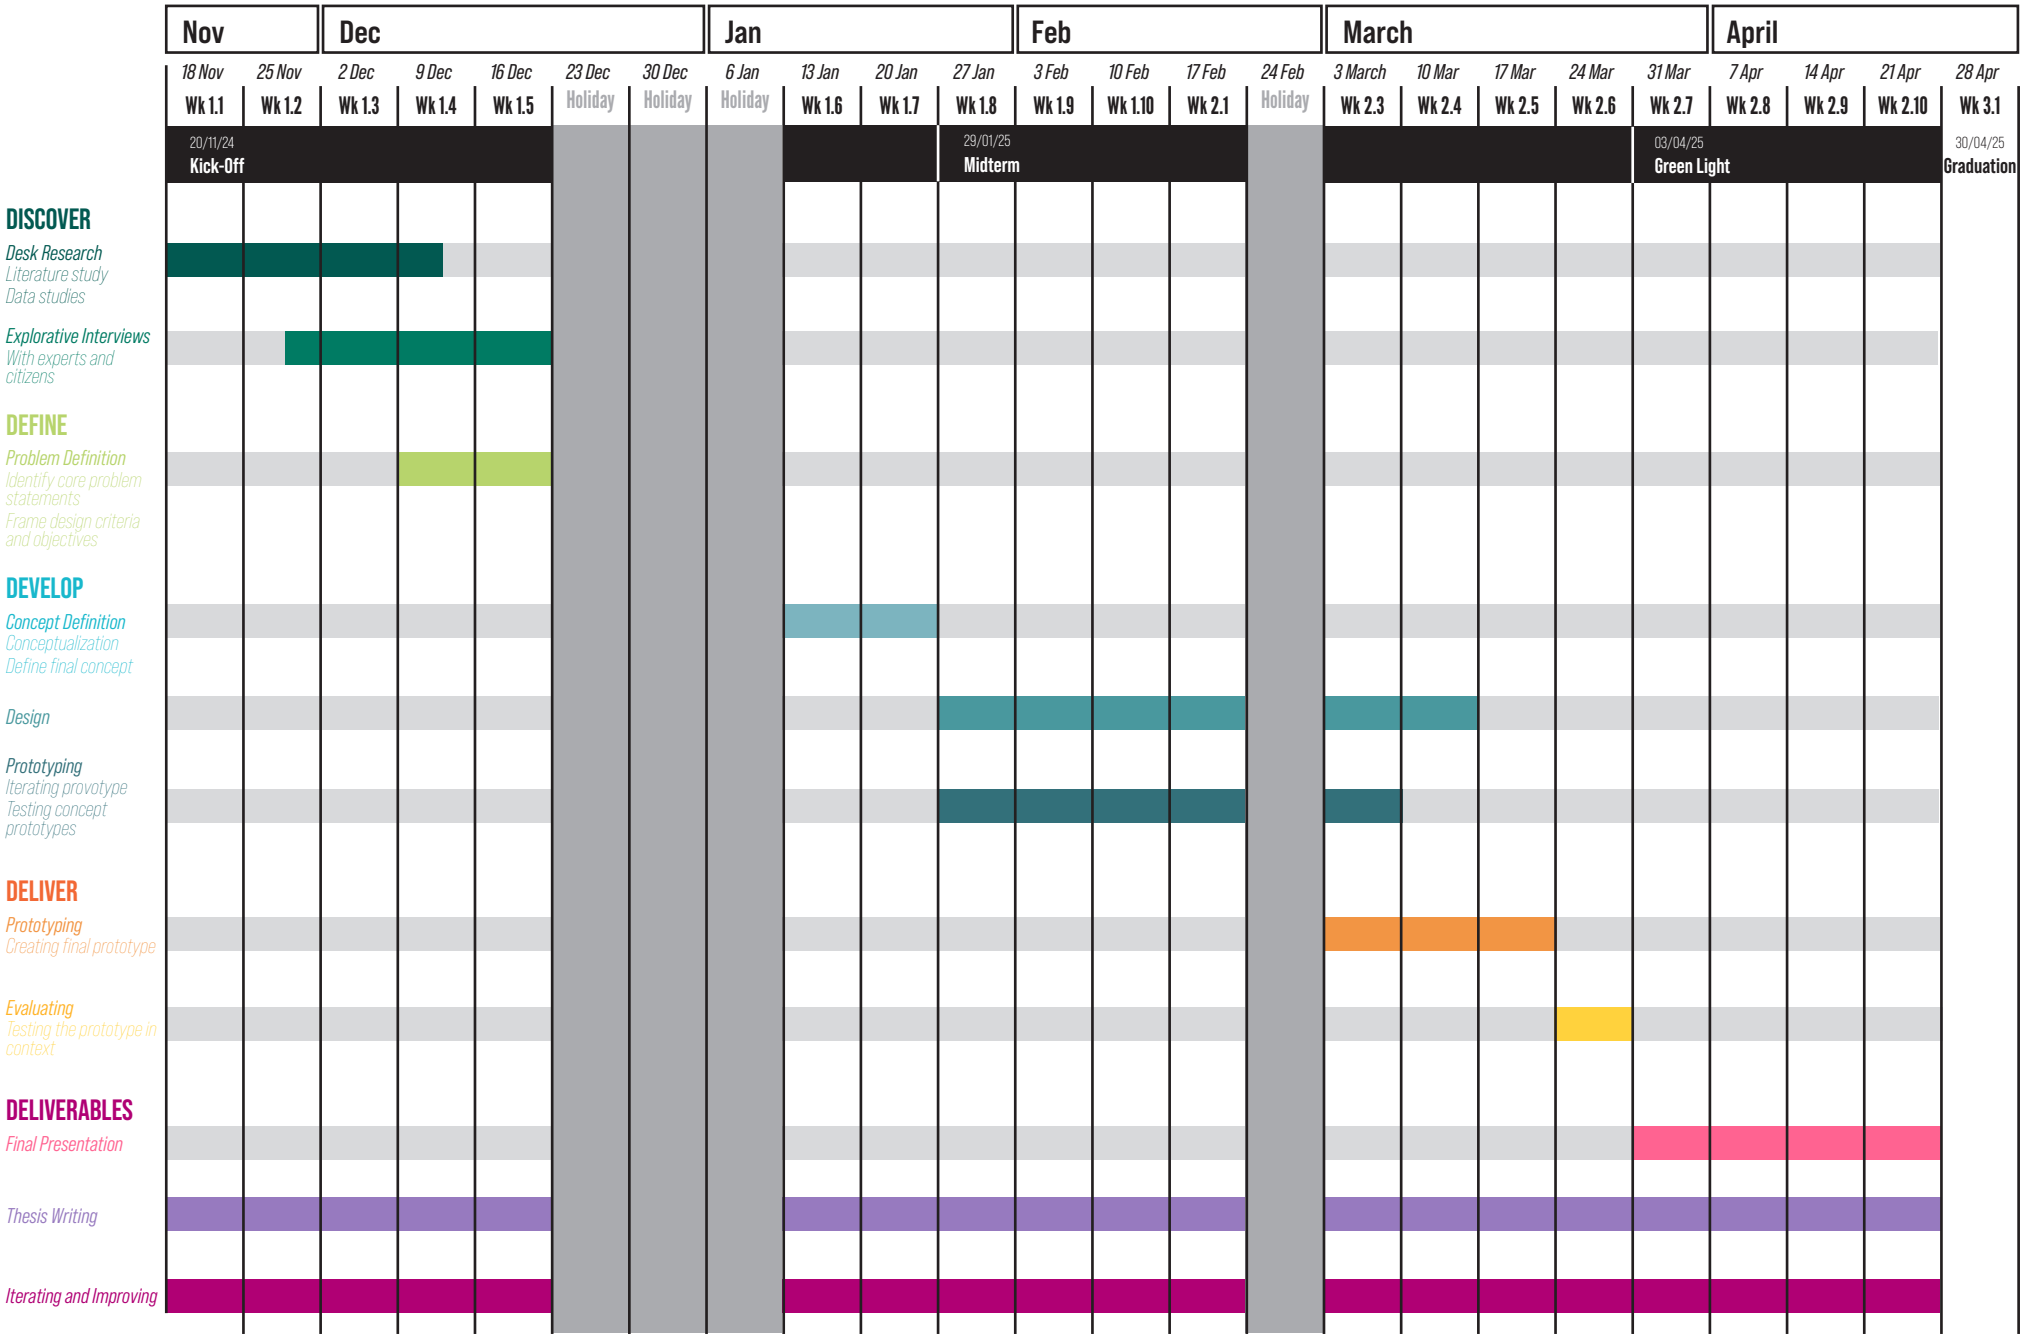
Nov
Dec
Jan
Feb
March
April
18 Nov
Wk 1.1
25 Nov
Wk 1.2
2 Dec
Wk 1.3
9 Dec
Wk 1.4
16 Dec
Wk 1.5
23 Dec
Holiday
30 Dec
Holiday
6 Jan
Holiday
13 Jan
Wk 1.6
20 Jan
Wk 1.7
27 Jan
Wk 1.8
3 Feb
Wk 1.9
10 Feb
Wk 1.10
17 Feb
Wk 2.1
24 Feb
Holiday
3 March
Wk 2.3
10 Mar
Wk 2.4
17 Mar
Wk 2.5
24 Mar
Wk 2.6
31 Mar
Wk 2.7
7 Apr
Wk 2.8
14 Apr
Wk 2.9
21 Apr
Wk 2.10
28 Apr
Wk 3.1
20/11/24
Kick-Off
29/01/25
Midterm
03/04/25
Green Light
30/04/25
Graduation
DISCOVER
Desk Research
Literature study
Data studies
Explorative Interviews
With experts and citizens
DEFINE
Problem Definition
Identify core problem statements
Frame design criteria and objectives
DEVELOP
Concept Definition
Conceptualization
Define final concept
Design
Prototyping
Iterating provotype
Testing concept prototypes
DELIVER
Prototyping
Creating final prototype
Evaluating
Testing the prototype in context
DELIVERABLES
Final Presentation
Thesis Writing
Iterating and Improving