

Filling the **empty** chair

Reimagining Policy Deliberation with Interactive AI-Mediated Citizen Perspectives



*Master Thesis Strategic Product Design
April 2025
Jim Blom*

Master thesis Strategic Product Design

Filling the Empty Chair

Reimagining Policy Deliberation with Interactive AI-Mediated Citizen Perspectives

Jim Blom

Faculty of Industrial Design Engineering

Delft University of Technology

April 2025

Supervisory team

Dr.ir. T.J. Jaśkiewicz | Chair

Dr. C.P. Alfrink | Mentor

In collaboration with

VONK Innovation Centre

Municipality of Rotterdam

Nathalie Helvoort | Company Mentor

Gerard Nijboer | Client

Table of Contents

Preface	5
Executive Summary	6
Positionality	7
Use of Generative AI	8
Introduction	9
Research Problem and Context	
Research Objective and Questions	
Methodological Approach	
Significance and Contributions	
Thesis Structure	
The Democratic Challenge	11
Theoretical background	
Citizen Participation as a Potential Bridge to Trust	
When Rich Stories Become Thin Summaries	
Communicating Decisions Back to Citizens	
Conclusion	
Theories of Voice and Power	19
Perspectives on Democratic Engagement	
Knowledge, Reality, and Democratic Participation	
Hill's Strategic Design Approach	
Conclusion	
Democracy in the AI Age	25
Technical Foundations of Language Models	
A New Kind of Memory	
Beyond Text: The Multi-Modal Capabilities	
Conceptual Foundations of AI Personas	
Risks, Ethical Implications, and Limitations	
Technology and the Construction of Democratic Knowledge	
Bridging Technology and Democratic Practice	
Conclusion	
Reimagining the Empty Chair	35
Design Elements	
Design Context: Rotterdam's Participation Landscape	
Vision: Digital Representatives in Municipal Processes	
From Vision to Development	
Conclusion	

The Development Journey	40
Development Journey Overview	
System Overview & Key Technologies	
Exploration 1: Testing Technical Feasibility	
Exploration 2: Municipal Integration (Vroesenpark Playground)	
Full Persona Profile: Sophie van der Meer (15 years old)	
Exploration 3: Enhanced Input Quality (Reyeroord Documentary)	
Full Persona Profile: Sarah (42) Youth worker	
Cross-Cutting Insights and Discussion	
Conclusion	
When AI Speaks for Citizens	61
Overview of the Evaluation	
Context	
Prototype Development	
Full Persona Profile: Eva (26) Freelance Writer and Illustrator	
Session Design	
Analysis Methodology and Framework	
Results	
Direct Interpretation of Key Findings	
Limitations	
Conclusion	
Conclusion	83
Discussion	84
Interpretation of Findings	
Rethinking Democratic Processes	
Tensions and Paradoxes: Critical Reflections	
Confronting AI Risks in Democratic Contexts	
Transforming Documentation Practices	
Creating Space for Responsible Innovation	
Practical Implications for the Municipality	89
Transforming Documentation Practices	
Creating Space for Responsible Innovation	
Positioning of AI Technologies	
Bridging Organizational Divides	
References	91

Preface

My journey to this research began on basketball courts, where as a young player I first encountered people from backgrounds vastly different from my own. Those early experiences planted a seed of awareness about social “bubbles” and the challenge of truly understanding perspectives different from our own. When I moved to Rotterdam, this awareness deepened. The city’s remarkable diversity made explicit what I had sensed earlier—that we often lack the tools to fully comprehend how others experience the world and what they think about issues that affect us all.

At my design education, I discovered the power of qualitative research to bridge these gaps. During my education I have always tried to combine stories with numbers. Where data and the visualisation of data has been a constant theme throughout my study. Therefore when I came across the project KAlte, a combination of my interests came together and I’m very grateful that I got the opportunity to further explore this concept.

This thesis would not have been possible without the guidance of Tomasz & Kars, whose insights challenged and refined my thinking at every stage. I’m also grateful to the team at VONK for helping me integrate into the municipality and helping me with every request I had, and to the Rotterdam Municipality for opening their doors to this research. Finally I want to thank Kate, Luca, Robin and everybody helping me finish this project.

Through this journey, I’ve come to see technology not as a solution, but as a tool that might help us reimagine the relationship between citizens and their governments. Personally I sometimes struggle to keep up with all the advancements of our current technology and my hope is that this work contributes, in some small way, to showing that these new technologies, when used accordingly could help us create more inclusive processes that amplify rather than replace citizen voices and restore some of trust between citizens and their governments.

Jim Blom

April 2025



Executive Summary

This research explores how AI-powered "digital representatives" might address Rotterdam's democratic deficit, where citizen input loses richness as it moves through municipal processes. In Rotterdam, this disconnect manifests in participation disparities (voter turnout ranging from 60% in affluent areas to 21% elsewhere) and declining institutional trust.

Through iterative development and evaluation, the study created AI personas capable of maintaining citizen perspectives in policy discussions. These digital representatives integrated speech recognition, Large Language Models, and voice synthesis to participate in deliberations about municipal issues.

Key findings demonstrate that while LLMs can consistently represent citizen viewpoints, their authenticity depends entirely on input quality. Representatives derived from rich interview transcripts preserved nuance and influenced policy discussions, while those created from typical municipal documentation proved inadequate. Voice interaction generated stronger engagement than text interfaces, and individual exploration enhanced subsequent group discussions.

Evaluation with Rotterdam civil servants revealed both promising capabilities and significant concerns. Digital representatives successfully challenged assumptions about community needs, but participants expressed legitimate concerns about accuracy and the risk of replacing direct citizen engagement with technological simulation.

The research's primary contribution is not the technology itself but revealing critical deficiencies in how municipalities capture and process citizen knowledge. Current documentation practices often fail to retain the experiential context that gives citizen input its value. Addressing Rotterdam's democratic challenges requires transforming how institutions value citizen perspectives—preserving context rather than pursuing abstraction.

Digital representatives cannot solve Rotterdam's democratic challenges alone but offer a means to sustain citizen perspectives throughout policy processes, especially for voices easily lost in current systems. Their most potent function may be provoking necessary institutional conversations about how citizen knowledge is valued—helping fill, not just simulate presence in, the empty chair.



Positionality

As the researcher behind this thesis, I recognize that my personal background, experiences, and disciplinary training have shaped my approach to this work. Being transparent about these influences helps readers understand the lens through which I've conducted this research.

My academic education at technical Universiteit like the TU Eindhoven & TU Delft has influenced my methodological approach and way of thinking. This discipline emphasizes qualitative research and human-centered design, shaping my belief that meaningful solutions must incorporate rich stories and lived experiences rather than relying solely on quantitative data. This perspective directly informed my concerns about how citizen voices become diminished as they move through municipal processes.

My position as a technology enthusiast with experience developing AI applications has influenced both my identification of the problem and my proposed solution. My prior work on transcription software using Large Language Models likely influenced me to see technological potential where others might focus on

social or political interventions. While I've attempted to maintain critical awareness of technology's limitations, my enthusiasm for its possibilities has undoubtedly shaped this research.

As a researcher with access to municipal structures through VONK and TU Delft, I occupied a privileged position between citizens and government. This intermediary role granted me access to civil servants and institutional processes that most citizens lack, potentially influencing how I perceived barriers to participation. Additionally, my recruitment of interview participants from my own social network likely resulted in perspectives that, while diverse in some aspects, share certain socioeconomic and educational characteristics with my own background.

Throughout this research, I've attempted to remain aware of these influences while seeking to amplify voices that might otherwise go unheard in municipal processes. By acknowledging these aspects of my positionality, I invite readers to engage critically with both the strengths and limitations of the perspective I bring to this work.

Use of Generative AI

Throughout this thesis project, Generative AI has been a significant part of my workflow across multiple dimensions: from coding and prototype development to writing, idea refinement and searching for connections in large sums of text. Rather than viewing this as something to minimize or hide, I believe in transparent acknowledgment of how these tools have shaped my work.

My approach to using Large Language Models (LLMs) aligns with what Mittelstadt et al. (2023) describe as "zero-shot translators" - tools that convert verified, accurate source material from one form to another, rather than generating new information from scratch. This framework focusses on having clear inputs and outputs, allowing for verification that the AI hasn't introduced errors or hallucinations into the process. As Mittelstadt et al. argue, this approach helps maintain scientific integrity when working with systems that "do not distinguish between fact and fiction."

In practice, this meant providing my own research findings, ideas, or existing code as inputs, and evaluating outputs against my intentions and knowledge. For coding, I used LLMs to translate my

design requirements into functional implementations. In writing, they helped transform my drafted concepts into more refined expressions while maintaining true to the original substance. During ideation, they expanded on concepts I introduced, always working from a foundation I established.

The impact of these tools should not be understated—they made possible prototypes that would have been technically impossible with my skillset and within the project timeframe. Finally, it enhanced the clarity of complex concepts in the written thesis. However, this contribution always occurred within a framework where I maintained oversight of both inputs and outputs, verifying that the AI's work aligned with my research findings and design intentions.

This approach reflects my belief that Generative AI, when used as a translation tool rather than a starting point of content, can extend rather than undermine human capabilities. By maintaining boundaries between AI assistance and human direction, I've tried to use these powerful tools while the research remains fundamentally my own work, grounded in my observations, analysis, and design decisions.

Introduction

Democratic governance faces a fundamental challenge: how to maintain authentic citizen perspectives throughout policy processes. In Rotterdam, this challenge manifests concretely as citizen input moves from initial engagement to policy development, often losing its richness and nuance along the way. This research investigates how emerging AI technologies, specifically Large Language Models (LLMs), might address this democratic deficit through “digital representatives”—AI-powered personas that maintain consistent citizen viewpoints throughout policy discussions.

Research Problem and Context

Democracy in Rotterdam shows concerning signs of erosion, with voter turnout as low as 21% in some neighborhoods compared to 60% in others (Van Bochove et al., 2023). This geographic disparity reflects a deeper disconnect between municipal institutions and certain communities, particularly those already facing socioeconomic challenges. When 43% of Dutch citizens feel the government doesn't listen sufficiently (Ministerie van Volksgezondheid, 2024), the issue extends beyond mere participation rates to fundamental questions of democratic representation and trust.

The disconnect becomes particularly evident in participation processes. As citizen perspectives move through municipal systems, they undergo significant transformation. The ‘woonvisie’ project exemplifies this problem: children's detailed perspectives on Rotterdam's housing future were reduced to just three general recommendations in the final report (Gemeente Rotterdam, 2023). This compression represents more than administrative convenience—it fundamentally alters the nature of citizen input, often losing the contextual richness and experiential knowledge that make community perspectives valuable.

This transformation creates what might be termed an “empty chair” in policy discussions—a space where authentic citizen voices should be present but have been diminished through institutional processing. The challenge extends beyond Rotterdam to democratic institutions broadly, where the tension between administrative efficiency and authentic representation remains largely unresolved.

Research Objective and Questions

This project aims to **design and evaluate a system that utilizes Large Language Models to generate representative citizen personas from existing qualitative data, to enhance inclusivity and diversity in Rotterdam's municipal policy-making processes.**

The research explores how LLMs can maintain the richness of citizen perspectives throughout policy processes, what input data and design elements are necessary for creating authentic digital representatives, and how civil servants engage with these AI-mediated perspectives in policy contexts. Through these questions, the research examines both technological possibilities and their implications for municipal decision-making.

Here, inclusivity refers to lowering barriers to participation beyond the “usual suspects” who regularly engage with municipal processes. Diversity encompasses capturing perspectives from different neighborhoods, socioeconomic backgrounds, and demographic groups—particularly voices currently underrepresented in formal participation.

Methodological Approach

This research employs Dan Hill's (2017) strategic design methodology, which systematically moves between concrete interventions (the “matter”) and their institutional context (the “dark matter”). This approach enables examination of both technological capabilities and the organizational structures that shape their implementation.

The process happens through three iterative development cycles followed by structured evaluation with municipal staff. The first cycle tests technical feasibility, examining whether LLMs can maintain consistent citizen perspectives in policy discussions. The second implements digital representatives within an actual municipal project to identify institutional requirements and barriers. The third explores alternative data sources for creating more authentic digital representatives. Finally, a structured evaluation assesses how municipal staff engage with digital representatives in policy deliberations.

This iterative approach enables evolving understanding

of both technological possibilities and institutional requirements for meaningful democratic enhancement.

Significance and Contributions

This research contributes to both theoretical understanding and practical application across multiple domains. For democratic theory, it explores how AI might mediate citizen perspectives in policy processes, suggesting hybrid approaches that bridge tensions between representative, deliberative, and participatory democratic models. By examining how technology might preserve the distinctive epistemological qualities of citizen knowledge within institutional processes, it contributes to understanding how different forms of knowledge function within democratic systems.

For technology and governance, the development of digital representatives provides insights into how emerging AI capabilities might enhance rather than undermine democratic values. This exploration moves beyond theoretical speculation to practical understanding of AI's democratic potential and limitations, contributing to broader debates about technology's role in governance.

For municipal practice, this research offers practical approaches to addressing democratic deficits, particularly regarding how citizen perspectives persist through policy processes. These insights have potential relevance beyond Rotterdam to other urban contexts facing similar challenges with democratic participation and representation.

examining how digital representatives function in policy deliberations and how municipal staff engage with AI-mediated citizen perspectives. Chapter 8 synthesizes these findings and discusses implications for democratic theory, technological governance, and municipal practice.

Thesis Structure

The thesis is organized into seven chapters that progress from problem identification through technological exploration to evaluation and theoretical reflection. Chapter 1 analyzes Rotterdam's democratic challenges, examining participation patterns across neighborhoods and identifying critical points where citizen perspectives lose influence in municipal processes. Chapter 2 explores the theoretical foundations of democratic participation, examining different models of democracy and their implications for citizen engagement.

Chapter 3 investigates the technological capabilities of Large Language Models, assessing both their potential for enhancing citizen representation and risks for democratic processes. Chapter 4 presents the design vision for digital representatives, outlining how they address Rotterdam's specific democratic challenges. Chapter 5 documents the iterative development process across three distinct explorations, revealing technical possibilities and institutional barriers.

Chapter 6 presents the evaluation with civil servants,

1. The Democratic Challenge

Democracy thrives on connection - the intricate relationship between citizens and their government that shapes policy, trust, and collective progress. In Rotterdam, this connection shows concerning signs of strain. Despite the municipality's commitment to inclusive governance, a gap has emerged: one that manifests not just in statistics, but in the daily interactions between institutional processes and the city's diverse communities.

Recent data illuminates the depth of this democratic challenge. When 43% of Dutch citizens feel the government doesn't listen enough, and 57% see their concerns inadequately represented (Ministerie van Volksgezondheid, 2024), we're witnessing more than dissatisfaction - we're seeing the erosion of democratic engagement. Rotterdam embodies this trend distinctly: the 2022 regional elections saw just 39% voter turnout, significantly below the national average of 51% (Marianne van Bochove et al., 2023).

Yet these city-wide figures mask an even more nuanced reality. Rotterdam's democratic landscape reveals itself through geographical contrasts: Figure 1 shows Hillegersberg, in the north, maintains robust democratic engagement with 60% voter turnout, while in Carnisse-Zuiderpark, only 21% of eligible voters participate (Marianne van Bochove et al., 2023). These patterns raise fundamental questions about representation and governance. How does a city council fulfill its mandate under Article 7 of Municipal Law - to represent all citizens - when entire neighborhoods exist largely outside formal democratic processes?

The implications ripple beyond voter statistics, touching the very essence of effective governance. Low participation doesn't simply challenge institutional legitimacy; it compromises the municipality's ability to develop policies that genuinely reflect community needs and aspirations. This misalignment becomes particularly significant in neighborhoods where democratic disengagement intersects with pressing social and economic challenges.

This chapter delves into the subtle yet powerful factors shaping Rotterdam's democratic challenge. This analysis sets the foundation for exploring potential paths forward, including how innovative technologies might help address these challenges - while acknowledging that technical solutions alone cannot address deeply rooted social dynamics.

Theoretical background

The relationship between democratic theory and lived experience rarely follows a straight line. In Rotterdam, this relationship reveals itself through a complex interplay of formal structures and daily realities, where traditional assumptions about citizen participation meet the nuanced patterns of urban life. To understand why trust in local government is waning and participation varies so dramatically across neighborhoods, we must examine both theoretical foundations and street-level realities.

The Trust Gap in Representative Democracy

At the core of this examination lies representative democracy - a model that Schumpeter (2013) envisioned as a clear covenant between citizens and their elected officials. The premise seems straightforward: citizens choose representatives who then act on their behalf, primarily through the mechanism of voting. Yet Rotterdam's experience

suggests this theoretical elegance encounters significant friction when meeting reality, particularly in neighborhoods like Carnisse-Zuiderpark where declining voter turnout signals more than mere disengagement.

Van Bochove et al. (2023) uncover the human dimension of this democratic erosion. Their research reveals a profound disconnect between municipal institutions and citizen experiences, rooted in a history of unfulfilled promises and disappointing interactions. When one resident observes, "They don't help the people. They get power from us, but they don't use it to help us; they use it for their own benefit," they're articulating more than personal frustration. This sentiment captures a fundamental breach in the democratic contract - one that undermines the very foundation of representative governance in Rotterdam.

Disparities in Participation Across Neighborhoods

The geography of Rotterdam's democratic engagement

tells a revealing story - one where participation patterns trace the city's social and economic contours. In Hillegersberg, where tree-lined streets suggest prosperity, voter turnout reflects more than civic duty. Here, Van Bochove et al. (2023) find that robust democratic participation grows from fertile soil: residents enjoy ready access to information, navigate institutional networks with ease, and possess the resources to engage meaningfully with municipal processes.

Cross the city to Carnisse-Zuiderpark, however, and this democratic landscape shifts dramatically. Here, engagement with formal political processes confronts a more complex reality. Residents navigate daily challenges that can make democratic participation feel like a distant luxury: immediate economic pressures, safety concerns, and limited access to resources shape their relationship with municipal institutions. Language barriers and cultural differences don't merely complicate communication - they fundamentally alter how citizens perceive and interact with government structures.

In these neighborhoods, citizens express what they call a "diepgeworteld gevoel in de steek gelaten te zijn" - a deeply rooted feeling of being abandoned (Van Bochove et al., 2023). This sentiment reflects more than momentary frustration; it suggests a profound rupture in the relationship between citizens and their municipal

government. When municipal initiatives fail to address structural problems, disengagement follows not from apathy, but from repeated experiences of disconnect between institutional responses and community needs.

The Complex Reality of Local Participation

Yet within this apparent democratic deficit, Van Bochove et al. (2023) discover something remarkable: many residents who feel alienated from formal political processes remain deeply engaged in their communities through alternative channels. Snel, Custers & Engbersen (2018) further explain this paradox in their Rotterdam-based research. In areas where traditional metrics like voter turnout suggest democratic disengagement and trust in government runs thin, they find unexpectedly high levels of participation in local initiatives.

This creates a tension: the very neighborhoods where citizens feel most abandoned by formal institutions often pulse with community engagement. These areas, struggling with standard municipal responses to their concerns, simultaneously demonstrate vibrant patterns of local organization and action. Such findings suggest that citizens are willing to participate when it is within their own community however not when this participation is associated with the municipality itself. Suggesting that it is the way participation is done could be the issue not that people don't want to do it.




Figure 1: Participation rates regional elections 2022 per district (Ineke Vogel et al., 2022)

Citizen Participation as a Potential Bridge to Trust

The democratic challenges facing Rotterdam raise a question: how might municipalities rebuild the eroded trust and enhance engagement with citizens who feel disconnected from governance processes? While the previous sections identified significant gaps in how citizen perspectives move through municipal systems, the concept of citizen participation offers a theoretical foundation for addressing these disconnects. Participation, in its essence, involves citizens actively shaping the policies that affect their lives rather than merely being passive recipients of governance decisions.

Theoretical Foundations for Citizen Participation

The literature on citizen participation suggests several mechanisms through which it might help rebuild the relationship between citizens and government. Hurenkamp & Tonkens (2020) identify multiple dimensions through which participation can strengthen democratic processes. First, participation can substantially increase public support for policies by ensuring citizens understand and accept decisions because they contributed to their development. Beyond only building support, participation can fundamentally improve policy quality. As Dreijerink et al. (2008) explain, citizens bring valuable experiential knowledge and contextual understanding that professional policymakers may lack. This local knowledge, when effectively integrated into decision-making processes, results in more realistic and effective policies that respond to actual community needs rather than abstract assumptions. Participation also serves developmental functions that extend beyond immediate policy outcomes. Hurenkamp & Tonkens (2020) note that the process itself encourages citizens to actively engage with societal issues and articulate their perspectives, creating civic skills and democratic capacity. The relationship between participation and responsibility forms another important dimension. When citizens engage in shaping their living environment, they develop increased ownership and often initiate complementary actions independently (Hurenkamp & Tonkens, 2020). Perhaps most fundamental to addressing Rotterdam's democratic challenges is how participation redistributes influence within governance systems. Effective participation creates opportunities for citizens to meaningfully impact decision-making, ensuring policies align with community needs rather than institutional convenience (Hurenkamp & Tonkens, 2020).

Identified Challenges Within Municipal Processes

Despite these theoretical benefits, Rotterdam's experience reveals a stark contrast between what participation should achieve and what actually occurs in practice. The declining trust levels and uneven participation patterns across neighborhoods suggest that current approaches to citizen engagement aren't delivering the promised improvements in democratic connection. Therefore, the following section will go into the specific challenges within current participation processes. These challenges are identified through various interviews with a range of civil servants and analyzing the available documents of past participation projects. They provide a starting point of where it would be interesting to focus the interventions and help guide how these new technologies would be able to make a difference. Through this process, three main challenges have been identified, consisting of 'finding representative input', 'loss of nuance in knowledge transfer', and 'communicating decisions back to citizens'.

The Challenge of Authentic Representation

In Rotterdam's participation landscape, a familiar pattern emerges: the same faces appear at public meetings, the same voices contribute to consultations, and the same names fill contact lists. These engaged citizens, often called the "usual suspects" or 'rolodex citizens', demonstrate admirable civic commitment. Yet their consistent presence raises a profound question about whose voices shape our city's future. The RIVM's 2008 literature review helps us understand what's at stake. True representativeness, they argue, requires more than just active participation - it demands a mosaic of voices from every corner of the city, across socioeconomic boundaries, and across the political spectrum. When this diversity is missing, policies risk becoming echo chambers, inadvertently amplifying certain perspectives while muting others (Lieke Dreijerink et al., 2008). Richard May's (2006) "triangle of engagement" offers a compelling framework for understanding this dynamic. Picture a pyramid where the base represents simple forms of participation, like completing a survey, while the peak represents intensive engagement - joining advisory boards or attending regular workshops. As we climb this pyramid, May observes, the number of participants inevitably decreases: "prevalence decreases as engagement increases—the higher the fewer." This isn't merely about willingness to participate. The pyramid's narrowing reflects real-world constraints: who has time for lengthy workshops? Who feels confident navigating formal consultation processes? Who can afford to volunteer for ongoing advisory roles? These questions reveal how high-engagement activities naturally filter

participants, often leaving us with our "usual suspects." May challenges us to reconsider how we view these dedicated participants. Criticizing them for lacking representativeness, he argues, misses the point. Their consistent engagement isn't the problem - it's an inevitable outcome of how we structure participation. When municipalities look for sustained, in-depth dialogue, they naturally draw from this committed group. Therefore the focus should be on finding a way to lower the physical & mental engagement level of participation methods while creating meaningful insights that are of value for the municipality. For Rotterdam, this creates a complex challenge. The municipality needs the deep insights that come from high-engagement methods - the kind of nuanced understanding that emerges from sustained dialogue and iterative consultation. Yet these very methods, by their nature, limit who can participate. The result is a paradox: the more in-depth the participation process, the less likely it is to capture the full spectrum of community voices.




Figure 2: Triangle of engagement (Richard May, 2006)

When Rich Stories Become Thin Summaries

Between the vibrant moments of citizen engagement and the formal process of policy development lies a critical transformation - one that often turns textured, living experiences into flat institutional summaries. This challenge emerges most visibly in Rotterdam's participation processes as citizen input moves from initial gathering to policy development, where the authentic voices of residents undergo a subtle but significant transformation. The initial engagement phase brims with life and detail. Municipal staff and researchers create spaces for deep dialogue with citizens, collecting not just opinions but stories that reveal how policy touches daily life. Through conversations, interviews, and interactive workshops, they gather a tapestry of experiences - each thread representing a unique perspective on community needs and aspirations. These moments of engagement capture something precious: the unfiltered wisdom of those who live the realities that policies aim to address. Yet as this rich material moves toward policy development, it encounters the necessary machinery of institutional process. The creation of participation reports demands analysis and clustering - a systematic approach to finding patterns, identifying themes, and organizing diverse perspectives into manageable categories. This distillation serves a practical purpose: making vast amounts of information accessible to policymakers. But in this process of compression, something valuable often evaporates. The 'woonvisie' (housing vision) project illustrates this transformation with particular clarity. One afternoon, researchers engaged with a children's council - twenty schoolchildren touring their city, sharing fresh perspectives on Rotterdam's housing future. Imagine the energy of that interaction: children pointing out their favorite spaces, suggesting imaginative solutions, sharing personal stories about their neighborhoods. Yet in the final participation report (Gemeente Rotterdam, 2023), this rich dialogue crystallized into just three recommendations: First, a call for sustainability and green spaces - the children emphasized the importance of trees and parks for creating livable neighborhoods. Second, a preference for quality over quantity in housing development, suggesting creative solutions like compact vertical building. Third, a focus on prioritizing housing for those most in need, including space for children to play and grow. While these points capture the broad strokes of the children's input, they miss the richness of the conversation. Lost are the specific examples that only children could provide - their unique experiences of urban spaces, their unfiltered critiques of current environments, their imaginative visions for future neighborhoods. The very qualities that make children's perspectives so valuable - their creativity, their directness, their ability to see possibilities adults might miss - become muted in the translation to formal recommendations. Consider what might have been left out: A child's story

about their favorite climbing tree and why green spaces matter more than parking lots. Another's detailed vision for a futuristic eco-friendly apartment building. Or personal accounts of how housing challenges affect their families and friends. These narratives, rich with emotional resonance and practical insight, dissolve in the process of distillation. This transformation reveals a fundamental tension in participation processes: the need to make citizen input manageable for policy development while preserving the very qualities that make it valuable. When we lose the richness of individual stories and the depth of personal experiences to broader themes and general recommendations, we risk missing the context that could make policies more responsive to community needs.

Example: Housing vision 2030

Participation Method:

Afternoon with school council with 20 secondary school students

Results in Participation Report

Sustainability and greenery in the city:

- More greenery and trees for a better living environment
- Promotes health (cycling, cleaner air, moisture retention)

Home quality versus quantity:

- Children predominantly choose quality of living
- Advice: build more compact and higher to meet the demand for housing

Target groups for housing construction:

- Focus on people with lower incomes, students and young people
- Less emphasis on the elderly
- More playgrounds for children in the city



Figure 3: Participation with schoolchildren for 'woonvisie' (Municipality of Rotterdam, 2023).

Communicating Decisions Back to Citizens

The final stage in Rotterdam's participation processes often plays out in silence. After citizens share their stories, some wait for an echo that never returns. This challenge of communicating policy decisions back to participating citizens reveals a crucial gap in the democratic dialogue - one that particularly affects larger, more complex projects where multiple voices and extended timelines blur the connection between input and outcome. Conversations with municipal staff paint a troubling picture. In smaller projects, the vital step of closing the feedback loop sometimes simply vanishes from view. Citizens who invested time and energy in sharing their perspectives never learn how their input shaped decisions. This silence speaks volumes, leaving participants to wonder whether their voices carried any weight at all in the policy process. Even in larger initiatives, where formal participation reports meticulously document engagement statistics and methodologies, a curious opacity emerges as the process moves toward policy formation. The clear lines connecting citizen input to policy decisions gradually fade, leaving participants to guess how their contributions influenced the final outcome. Policy documents, while comprehensive in many aspects, often fail to explicitly trace the path from citizen voice to policy choice. The "woonvisie" (housing vision) project offers a revealing window into this challenge. Its final document includes a chapter specifically dedicated to citizen influence.

Yet even here, we see a telling disparity in how different forms of participation are acknowledged. The document carefully details how input from high-engagement activities, particularly table discussions, shaped the final policy. These activities, sitting higher at May's engagement triangle, typically involve the 'usual suspects' - citizens well-versed in participation processes. Meanwhile, broader participation efforts receive only passing mention. The children's council, for instance, earns just a brief acknowledgment: "We have used different instruments to let the citizens of Rotterdam think about the subject; in this way, we heard your opinions and advice." This vague statement offers no insight into how these young voices influenced specific policy choices or shaped the final vision. It's as if their input, once gathered, disappeared into the institutional ether. For citizens who engaged in these more accessible forms of participation, such cursory acknowledgment can feel like dismissal. Without seeing clear evidence of their influence on outcomes, they might reasonably question whether their participation served any purpose beyond checking a bureaucratic box. This perception doesn't just damage trust in current processes - it casts a long shadow over future engagement opportunities. The challenge deepens when communication responsibilities fall to civil servants who weren't present for the original discussions. Asked to explain how citizen perspectives shaped decisions, they must navigate between documented summaries and lost nuances. This becomes particularly acute when dealing with minority viewpoints that didn't survive the journey from participation to policy - voices that might have offered crucial insights but failed to find representation in final reports.



We have used various tools to involve Rotterdam residents in the thinking process...

This way we could hear the opinions and advice of people from the neighbourhood and of future home seekers.



Figure 4: Quote from the final 'woonvisie' document clarifying how citizen input has contributed. (Municipality of Rotterdam, 2023).

Conclusion

Rotterdam's democratic landscape presents a paradox that lies at the heart of this research. Throughout this chapter, I have traced how declining trust and participation in formal democratic processes are in contrast of the active community engagement at the local level. This contradiction reveals that the issue isn't citizen unwillingness but rather a fundamental disconnection between institutional processes and community life.

The evidence from Rotterdam's neighborhoods tells a compelling story. As Van Bochove et al. (2023) demonstrate, the significant variation in voter turnout—from 60% in Hillegersberg to just 21% in Carnisse-Zuiderpark—reflects not only demographic differences but profoundly different relationships with municipal governance. Also revealing is Snel, Custers & Engbersen's (2018) finding that neighborhoods with low formal participation often demonstrate significant community engagement through alternative channels. This pattern suggests that citizens are willing to participate when opportunities feel meaningful and accessible, but something in current municipal approaches fails to capture this energy.

Through examining Rotterdam's participation processes, three critical challenges emerge that contribute to this democratic disconnect:

First, the representation challenge limits whose voices shape municipal decisions. The predominance of "usual suspects" in formal participation processes means that certain perspectives consistently influence policy while others remain unheard.

Second, the knowledge transformation challenge alters the nature of citizen input as it moves through municipal systems. As the 'woonvisie' project illustrated, rich, contextual community perspectives become increasingly abstracted and decontextualized during processing.

Third, the feedback disconnect leaves citizens uncertain about how their participation impacted the outcome. When the municipality fail to communicate clearly how specific inputs influenced outcomes, they miss a crucial opportunity to build trust and demonstrate the value of participation.

These challenges partly explain why citizen participation, despite its theoretical potential to rebuild trust and enhance democratic engagement, has struggled to fulfill this promise in Rotterdam. Traditional approaches to gathering, processing, and responding to citizen input face significant limitations in preserving the richness of community perspectives and maintaining meaningful connections between citizens and their government.

The following chapter examines these challenges through broader theoretical perspectives on democratic engagement. By exploring different models of

democracy and participation beyond conventional voting mechanisms, I will establish a theoretical foundation for understanding Rotterdam's democratic challenges. This examination of democratic theory will help contextualize the practical issues identified in this chapter and provide frameworks for conceptualizing potential paths forward.

2. Theories of Voice and Power

This chapter investigates three influential democratic models and their implications for Rotterdam's participation landscape. Beginning with representative democracy, I analyze how Schumpeter's minimalist framework, which positions voting as the primary democratic mechanism, encounters significant limitations in Rotterdam's diverse urban context. I then explore deliberative democracy as formulated by Habermas and Rawls, examining both its promise of inclusive dialogue and the practical challenges it faces in implementation. Finally, I consider Mouffe's theory of agonistic pluralism, which reframes democratic conflict as a constructive rather than problematic element.

Beyond these models, I examine the philosophical foundations of democratic participation—specifically how different ontological and epistemological positions shape what counts as knowledge in democratic processes. This analysis helps explain the knowledge transformation challenge identified in Chapter 1, where rich citizen perspectives become progressively abstracted as they move through institutional processes.

The chapter concludes by introducing Hill's strategic design methodology, which provides the framework for my research approach. This methodology enables systematic movement between concrete interventions and their broader institutional contexts, particularly valuable for understanding how the implementation of this new technology is shaped by both visible structures and invisible "dark matter" of organizational culture and hidden constraints.

Through this theoretical exploration, I establish the conceptual foundation necessary for understanding Rotterdam's democratic challenges and evaluating potential pathways toward more inclusive and responsive processes.

Perspectives on Democratic Engagement

When examining Rotterdam's democratic challenges, it is valuable to look at the theoretical frameworks that shape how we conceptualize citizen participation. These aren't just abstract models - they're the intellectual foundations that influence how municipalities structure their relationships with citizens, and they help illuminate why certain approaches succeed or fail in practice.

Representative Democracy and Its Nuanced Limitations

The Netherlands, like many established democracies, operates primarily within the framework of representative democracy. This model, articulated in Joseph Schumpeter's (2013) minimalist view of democracy, positions electoral participation as the cornerstone of democratic legitimacy. Citizens exercise their influence indirectly: they select representatives who then navigate the complex terrain of policy-making on their behalf.

Yet Rotterdam's experience reveals the subtle complexities in this seemingly straightforward arrangement. When voter turnout drops to 38.9% in regional elections, we're confronting more than a simple participation deficit. This figure signals a deeper challenge to the model's fundamental assumptions

about representation and legitimacy.

The principle of mandate representation - the belief that elected officials will act in their constituents' best interests - encounters particular strain in Rotterdam's diverse urban context. When citizens express that officials "don't help the people" and "use [power] for their own benefit," they're highlighting a breakdown in the representative relationship. This isn't merely about trust; it's about the gap between theoretical representation and lived experience.

This tension becomes especially apparent in neighborhoods where formal democratic processes seem increasingly disconnected from daily realities. The challenge isn't simply that people don't vote - it's that the entire mechanism of representation, from electoral participation to policy implementation, struggles to bridge the distance between institutional processes and community needs.

Deliberative Democracy and the Complexity of Implementation

The limitations of purely representative systems have led scholars toward an alternative vision: deliberative democracy (Habermas, 1996; Rawls, 2003). This framework reimagines democratic legitimacy not as

periodic electoral choices, but as an ongoing process of rational dialogue and collective understanding. Where representative democracy emphasizes delegation, deliberative democracy focuses on discussion, reasoning, and consensus-building.

In Rotterdam's context, deliberative democracy offers intriguing possibilities. Its emphasis on inclusive dialogue and mutual understanding seems particularly relevant for a city navigating diverse needs and perspectives. The model suggests that through structured deliberation, citizens and policymakers might bridge the gaps that electoral processes alone cannot address. One of its core features - the equal consideration of all voices - appears especially pertinent in addressing the participation disparities between neighborhoods like Hillegersberg and Carnisse-Zuiderpark.

Yet when this theoretical model encounters Rotterdam's urban reality, significant tensions emerge. The implementation of deliberative practices reveals structural challenges that the theory alone cannot resolve:

First-generation migrants may possess crucial insights about community needs, yet language barriers can exclude them from deliberative forums. A single parent working multiple jobs might have intimate knowledge of policy impacts, but lack the time for extended participation in dialogue sessions. These aren't merely practical obstacles - they reveal fundamental questions about who can meaningfully participate in deliberative processes.

Moreover, the model's emphasis on rational discourse and consensus-building encounters complex social realities. The assumption that participants will engage with a shared commitment to the common good overlooks how deeply rooted social and cultural differences shape perspectives and priorities. When residents from different neighborhoods bring competing needs to the table - immediate housing concerns versus long-term development plans, for instance - the path to consensus isn't always clear or possible.

These implementation challenges don't invalidate deliberative democracy's insights, but they do reveal the complexity of translating theoretical frameworks into effective practice. The model's limitations become particularly visible in Rotterdam's diverse urban environment

Agonistic Pluralism: Embracing Democratic Tension

The limitations of both representative and deliberative models lead us to a more radical proposition: Chantal Mouffe's theory of agonistic pluralism (1999). Where earlier frameworks seek to minimize or resolve conflict, agonistic pluralism suggests that disagreement

isn't a flaw in democratic systems - it's an essential feature. This perspective shifts our understanding of democratic practice from conflict resolution to conflict engagement.

For Rotterdam, this theoretical lens offers particularly relevant insights. Rather than viewing the city's diverse perspectives as obstacles to overcome, agonistic pluralism suggests creating spaces where differing viewpoints can productively coexist. This approach acknowledges that marginalized groups might hold fundamentally different perspectives from those in power - differences that can't always be reconciled through traditional democratic processes.

Consider how this plays out in Rotterdam's policy landscape: Residents in Carnisse-Zuiderpark often prioritize immediate socioeconomic challenges - affordable housing, employment opportunities, daily economic survival. Meanwhile, municipal officials, working from statistical models and long-term projections, might focus on sustainable urban development or infrastructure improvements. Traditional democratic approaches would attempt to find middle ground or achieve consensus. Agonistic pluralism suggests a different path: one where these conflicting perspectives are not just acknowledged but seen as fundamental to democratic dialogue.

This approach advocates for forums where differences can be explicitly articulated and discussed. When certain perspectives prevail in specific decisions, the focus shifts from achieving consensus to maintaining transparency about why particular choices were made. The dialogue continues even after decisions are implemented, recognizing that opposing viewpoints retain their legitimacy even when not selected for implementation.

This framework particularly resonates with criticism that Rotterdam's municipality is "overly influenced by numbers and statistics rather than by people's stories and experiences" (Van Bochove et al., 2023). Agonistic pluralism suggests that both forms of knowledge - statistical and experiential - have validity, even when they point in different directions. It creates space for emotional and experiential knowledge to carry equal weight alongside technical expertise.

However, implementing this approach requires significant institutional adaptation. The municipality would need to develop new processes that accommodate productive conflict, while building capacity for transparent decision-making that acknowledges unresolved differences. These aren't merely procedural changes - they represent fundamental shifts in how democratic institutions operate. Yet Rotterdam's current democratic challenges suggest that such shifts might be necessary - not to eliminate conflict, but to engage with it more productively.

Knowledge, Reality, and Democratic Participation

The democratic models discussed above rest on deeper philosophical assumptions about the nature of knowledge and reality. Understanding these theoretical foundations helps explain why different approaches to participation have evolved and how they shape our understanding of democratic processes. These philosophical perspectives are particularly relevant to one of the key challenges identified earlier: how rich, nuanced citizen perspectives become reduced and transformed as they move through institutional processes toward policy decisions.

Knowledge and Perspective in Democratic Theory

At the heart of democratic participation lies a question: what status do we give to different perspectives, and how do these perspectives relate to an underlying reality? This question becomes especially important when considering how citizen input gets transformed as it moves from community engagement to policy formulation.

Ontological Positions: The Nature of Reality

Three key ontological positions offer different views on the nature of reality, each with implications for how we understand citizen perspectives:

Realism, holds that reality exists independently of human perception—there is an objective world “out there” regardless of how we think about it. From this perspective, citizen participation is primarily about gathering information about an objective reality that exists independent of those perceptions. A strict realist position might suggest that expert analysis can distill “true” needs from varied citizen inputs, potentially justifying the reduction of diverse perspectives into simplified policy recommendations.

Relativism, at the opposite end, suggests that reality itself is entirely dependent on human interpretation, with multiple valid “truths” existing relative to different perspectives. A strong relativist position would hold that all perspectives are equally valid, as truth is entirely relative to one’s viewpoint or social position. While this honors diverse voices, it poses significant challenges for democratic decision-making, as it provides no basis for evaluating competing claims or reaching collective decisions beyond power negotiations.

Subtle realism, as described by Andrews (2012), offers a middle ground: while social reality exists independently of our individual perspectives, our access to this reality is always mediated through our social positions and understanding. This position provides theoretical grounding for both valuing diverse perspectives and

maintaining the possibility of meaningful dialogue between them.

Epistemological Positions: How We Know

Corresponding to these ontological positions are different epistemological approaches to how knowledge is constructed and validated:

Positivism/objectivism aligns with realist ontology, suggesting that objective knowledge can be discovered through appropriate methods. In democratic contexts, this could be seen as putting trust in expert knowledge or data-driven decision-making. This position might see the distillation of citizen input into categories as an appropriate refinement process that separates signal from noise.

Strong constructionism represents the opposite end of the spectrum, suggesting that all knowledge is entirely socially constructed with no objective reference point. This position would hold that policy decisions merely represent one constructed reality being imposed over others through institutional power.

Social constructionism, as described by Andrews (2012), offers a middle ground. It argues that our knowledge of the world is constructed through social interactions rather than simply discovered, while maintaining that ‘society exists both as objective and subjective reality.’ While there may be an objective world, our understanding of it is inevitably shaped by our social context, relationships, and positions within society. This position helps explain why the transformation of rich citizen narratives into abstract policy points represents an epistemological challenge: the social processes of constructing institutional knowledge often filter out crucial context that gives citizen perspectives their meaning and value.

Knowledge Transformation Problem

The philosophical frameworks above help us understand the challenge identified in earlier chapters: how citizen perspectives lose richness and nuance as they move through institutional processes toward policy decisions.

Consider the example of the children’s council in Rotterdam’s ‘woonvisie’ project. The children’s rich, contextual knowledge—grounded in their lived experiences of neighborhoods and emotional connections to their environment—underwent an epistemological transformation. Their nuanced perspectives became reduced to three general recommendations about sustainability, quality over

quantity, and prioritizing those in need.

This transformation represents more than a practical summarization; it reflects a stance in what counts as knowledge within institutional processes. Drawing on the earlier mentioned philosophical framework:

From an ontological perspective, the reduction process often implies a realist assumption that the 'essential points' can be extracted from diverse perspectives, while the contextual details are just subjective extra information that can be discarded. Yet as subtle realism suggests, these contextual elements may be crucial for understanding the social reality the children were describing.

From an epistemological standpoint, the transformation reveals tensions between different ways of knowing. The children's experiential knowledge—what they directly know through lived experience—becomes converted into abstract categories that fit institutional processes. This conversion from rich, contextual knowledge to abstracted policy points aligns with positivist assumptions about knowledge refinement, potentially losing the constructionist understanding that knowledge is meaningful precisely because of its social and experiential context.

This knowledge transformation challenge sits at the heart of participation processes. When citizen input becomes increasingly abstracted through translation, the very qualities that make it valuable, its grounding in lived experience, its emotional resonance, its contextual specificity may be lost.

Alignment with Democratic Models

These philosophical positions help explain the fundamental differences between democratic models and their approaches to managing knowledge diversity. Deliberative democracy, with its emphasis on rational dialogue and consensus-building, tends toward a more realist ontological position and objectivist epistemology. It assumes that through structured discussion, participants can move toward more objective understandings that transcend their individual perspectives. This approach might support the refinement and abstraction of citizen input as part of finding common ground.

Agonistic pluralism, on the other hand, aligns more on constructionist view while acknowledging the relativist challenge. It recognizes that different social positions produce fundamentally different ways of understanding and experiencing social reality. Rather than seeing these differences as obstacles to be overcome through rational deliberation, it views them as essential features of democratic life that should be preserved. From this perspective, the typical transformation of citizen input represents a problematic removal of valuable differences rather than productive synthesis.

The tension between these models reflects deeper questions about how diverse knowledge should be handled in democratic processes. While deliberative democracy suggests that refinement toward consensus represents progress, agonistic pluralism maintains that some differences should be preserved rather than resolved. This tension reveals why the transformation of citizen input into policy language is not just a practical challenge but also a philosophical one.

This background will prove particularly applicable in Chapter 3, where I will go deeper into how Large Language Models process and represent human perspectives. LLMs embody distinct epistemological characteristics that differ from human processes.

Hill's Strategic Design Approach

During this project, I will employ Dan Hill's strategic design methodology as articulated in "Dark Matter and Trojan Horses: A Strategic Design Vocabulary" (Hill, 2017). This methodology offers a sophisticated framework for addressing complex systemic challenges by strategically navigating between concrete design interventions and their broader contextual environments. Hill's approach is particularly valuable for my research because it provides a structured yet flexible way to understand how tangible solutions are inherently shaped by—and can in turn reshape—their organizational and cultural contexts.

Matter and Dark Matter: The Core Framework

At the heart of Hill's methodology lies an understanding of how design interventions exist within layered contexts. Hill articulates three key concepts that form the foundation of his approach:

Matter represents the tangible, observable physical reality of design interventions: the products, services, spaces, and artifacts that designers traditionally create. This is the concrete manifestation of design work—what can be directly seen, touched, and experienced.

Meta refers to the broader contextual environment or strategic framework within which these interventions exist. Hill emphasizes that strategic design involves "zooming between matter and meta," a process of "oscillating between these two states in order to recalibrate each in response to the other".

Dark matter comprises the specific elements within this meta-level that often remain invisible yet exert powerful influence: "organizational culture, policy environments, market mechanisms, legislation, finance models and other incentives, governance structures, tradition and habits, local culture and national identity". This dark matter can only be perceived indirectly, "through its effect on other things" yet it fundamentally determines whether interventions succeed or fail at creating

systemic change.

The strategic designer's task is to navigate between the concrete matter of interventions and the meta-level of context, while developing the capacity to perceive and manipulate the dark matter that shapes both.

Key Methodological Principles

The strategic design process, according to Hill, requires constant movement between these two realms through several key methodological principles:

Prototyping as Strategic Exploration

The practical implementation of these principles involves a continuous cycle of prototyping, testing, and refinement. However, in Hill's approach, prototypes serve a dual purpose beyond traditional design testing. As I develop prototypes in my project, they will function simultaneously as:

1. Tests of the solution itself, evaluating functionality, usability, and effectiveness
2. Strategic probes that reveal the "hidden constraints, assumptions, and opportunities within the system"

This dual function creates what Hill describes as a "dynamic design process where each realm informs and reshapes the other". As my prototypes interact with the system, they will uncover aspects of the dark matter that might otherwise remain invisible. These insights will then guide subsequent iterations, creating a continuous feedback loop between what I make and what I learn about the system.

Position Within Design Methodologies

Hill's approach exists within a broader landscape of systemic design methodologies but offers distinct advantages for my research context. While traditional design thinking approaches focus primarily on user needs and solution development, Hill's strategic design




Figure 5: Visualisation of strategic design methodology of Dan Hill (2017)

more explicitly addresses the institutional and systemic contexts within which solutions must operate.

Unlike approaches that remain at the “lipstick on the pig” level of the value chain, strategic design engages with all aspects of the system, particularly the strategic context of the challenge itself. This makes it particularly well-suited for addressing the complex challenges of municipal governance and citizen participation.

Application in My Research

In my project, Hill’s strategic design methodology will be operationalized through several specific approaches:

1. **Matter-Meta Movement:** I will systematically move between concrete interventions (the “matter” of my design work) and the organizational and policy context (the “dark matter”), using each to inform and refine the other.
2. **Contextual Analysis:** I will analyze not just the surface-level challenges but also the underlying organizational culture and democratic models that shape how citizen input is handled within municipal systems.

By employing Hill’s strategic design methodology, I try to develop interventions that not only experiment with immediate technological possibilities but also reveal and influence the underlying systems that shape those needs.

Conclusion

This chapter has examined three distinct theoretical models of democracy, each offering valuable perspectives for understanding Rotterdam’s participation challenges. Representative democracy’s reliance on electoral mechanisms proves inadequate in neighborhoods where voter turnout has fallen to concerning levels, revealing fundamental gaps between institutional representation and lived experiences. Deliberative democracy’s emphasis on structured dialogue offers promising approach for inclusive governance, yet faces significant implementation barriers in Rotterdam’s diverse urban context, where time constraints, language differences, and socioeconomic factors limit who can meaningfully participate in deliberative processes.

Mouffe’s agonistic pluralism perhaps offers the most relevant framework for Rotterdam’s current struggles, acknowledging that democratic tension between perspectives—such as experiential knowledge from Carnisse-Zuiderpark residents versus statistical frameworks from municipal officials—represents not a failure of democracy but one of its essential features. This approach suggests Rotterdam might benefit from creating spaces where different types of knowledge can productively coexist rather than attempting to reconcile them into artificial consensus.

The philosophical examination of knowledge transformation illuminates one of Rotterdam’s core democratic challenges: how rich, contextual citizen perspectives become increasingly abstracted as they move through institutional processes. This transformation reflects deeper tensions between different epistemological positions—between the experiential, socially constructed knowledge of communities and the more positivist orientation of municipal systems that prioritize abstraction and categorization.

Hill’s strategic design methodology provides the framework for connecting these theoretical complexities to concrete prototypes. By systematically moving between concrete interventions and their broader contextual environments this approach enables investigation of how technological solutions might address Rotterdam’s democratic challenges while revealing the institutional dynamics that shape their implementation.

These frameworks collectively establish the foundation for the remainder of this research. As I explore technological possibilities in subsequent chapters, these theoretical insights will inform how I evaluate and choose potential interventions not just for their technical capabilities, but for how they would influence the wider subjects of democratic knowledge and representation.

3. Democracy in the AI Age

Generative AI has achieved capabilities that seemed impossible just years ago. These systems can now process vast amounts of unstructured information. They can translate complex policy documents into accessible language, generate personalized communications at scale, and even create interactive representations of citizen perspectives.

Yet these possibilities come with profound risks. The same technologies that might help amplify marginalized voices could also be used to generate artificial consensus. Tools that might preserve the richness of citizen experiences through administrative processes might simultaneously make those processes more opaque. The power to process more citizen input more quickly might lead to less genuine engagement rather than more.

This chapter aims to unpack these contradictions, examining both the transformative potential and serious risks of Generative AI in democratic processes. By understanding how these technologies work, what they can and cannot do, and what their implementation might mean for democratic participation, I can better evaluate their role in addressing Rotterdam's democratic challenges. This understanding provides the technological context needed to assess the following chapters.

Technical Foundations of Language Models

In 2017, a technical paper from Google titled "Attention is All You Need" (Vaswani et al., 2017) introduced an innovation that would fundamentally change how machines could process language. The paper presented the Transformer architecture - an evolution on the previous approaches to artificial intelligence. The Transformer introduced a mechanism called "self-attention" that allows each element in a text to directly connect with every other element, regardless of their position in this text.

This self-attention mechanism works by weighting these connections based on their relevance to understanding each other, enabling the system to grasp fundamental relationships between words and ideas that might even be separated by paragraphs of text. It's comparable to how humans understand context in conversation - people naturally connect related ideas even when they're mentioned far apart. This direct connectivity made AI systems fundamentally better at predicting language, laying the groundwork for technologies like ChatGPT that we see today.

However, the Transformer architecture alone wasn't enough to create the powerful language models we now have. A second crucial realization emerged: the importance of scale. Until 2020, language models remained relatively modest in size. The introduction of GPT-3 changed this dramatically, implementing a model with 175 billion parameters - ten times larger than anything previously available (Brown et al., 2020). This massive scaling wasn't just about size; it proved a fundamental hypothesis about these systems: that simply making them larger significantly improved their

performance across a wide range of tasks (Brown et al., 2020).

At their core, these Large Language Models (LLMs) remain sophisticated guessing systems, statistically predicting what word should come next in a sequence. While this might sound simple, the scale and sophistication of these predictions enable capabilities that sometimes appear almost magical. It's crucial to remember that these systems aren't sentient and don't truly "think" - but their ability to predict and generate text based on vast training data creates powerful new possibilities.

Consider Google's Gemini model family, which can process up to 2,000,000 tokens in its short-term memory or context window (Google Deepmind, 2024). To put this in perspective, that's equivalent to approximately 1.4 million words or 4,000 pages of text. Even more remarkably, research from Google DeepMind demonstrates that these systems can recall specific information from this massive context with over 99.7% accuracy (Google Deepmind, 2024).

This combination of the Transformer architecture's ability to understand relationships between elements and the power of massive scale has created systems that can process and generate human-like text with unprecedented sophistication. While they remain fundamentally prediction engines, their ability to maintain context, understand relationships, and process vast amounts of information opens new possibilities for handling the complex, nuanced data that emerges from citizen participation processes.

A New Kind of Memory

This new ability to process and recall vast amounts of information has particular significance for citizen participation processes. In traditional approaches, when citizen input grows extensive - as it often does in larger projects - the sheer volume of data forces a creation of distilled representations, typically in the form of participation reports. This compression of information, while necessary for human processing, often means losing valuable context and nuance.

Think of it like searching for a needle in a haystack. When creating a participation report, civil servants must locate the core arguments - the needles - within vast amounts of citizen input. This process is not only time-consuming but also rigid: if they later realize they need to find a different type of needle (perhaps a different perspective or theme), the entire labor-intensive search process must begin again.

Large Language Models fundamentally change this dynamic. These systems can search through context windows of up to 2 million tokens (Li, 2024), effectively processing thousands of pages of text in minutes. More importantly, they can repeat this search with different criteria almost instantly, making it possible to explore citizen input from multiple angles without the traditional time and resource constraints.

This capability challenges fundamental assumptions about how to process citizen input. Where institutions once needed to heavily filter and categorize information to make it manageable, they can now work with richer, more nuanced forms of data. The cost - both human and financial - of bringing citizen perspectives to the policy table has dropped dramatically, opening new possibilities for how citizen input can be used in decision-making processes.

This shift is particularly significant because it enables more effective work with qualitative data. Traditionally, large-scale analysis favored quantitative data - numbers and statistics that could be easily processed and analyzed. Now, it's possible to handle more complex, unstructured forms of information: transcripts of conversations, open-ended survey responses, even audio and video recordings. These formats often contain richer context and nuance than numerical data alone.

The implications for informed decision-making could be profound. Where previously municipalities might have relied primarily on quantitative metrics to inform policy decisions, they can now incorporate richer forms of qualitative input: detailed citizen narratives, complex community perspectives, and nuanced feedback that might have been lost in traditional summarization processes. The haystack of citizen input can grow ever

larger, while the ability to find relevant insights within it becomes faster and more sophisticated.

This technological capability aligns particularly well with Rotterdam's challenge of preserving experiential knowledge through bureaucratic processes. When citizen input no longer needs to be heavily compressed for processing, more of its original richness and context can be maintained throughout the policy-making journey. The system's ability to quickly surface relevant information means that specific citizen perspectives can be recalled and considered at any point in the policy development process.

Beyond Text: The Multi-Modal Capabilities

While the initial breakthroughs in AI focused on text processing, recent years have seen remarkable advances across other modalities of communication. The technology has expanded to synthesize voices, generate photos, and create videos, with the quality of these outputs improving at a startling pace.

A significant milestone in this evolution came with the introduction of Whisper (Radford et al., 2022), which made accurate transcription dramatically more accessible. This open-source algorithm approaches human-level performance in converting speech to text, with error rates between 5.5% and 7.5%, comparable to human error rates of 4.7% to 7.2% (Radford et al., 2022). While human transcribers might make more contextually appropriate mistakes, Whisper's performance represents a remarkable achievement, especially considering it works effectively across multiple languages. This development has significantly lowered the barrier to processing spoken information, making it easier to capture and analyze conversations, interviews, and public meetings without losing detail through manual note-taking.

The implications for citizen participation are significant. Where previously capturing detailed records of citizen interactions might have required extensive human resources, these tools enable comprehensive documentation of spoken input. Public meetings, informal conversations, and community discussions can now be preserved in their entirety, providing the nuance and context in the input that might otherwise be lost in summary notes.

Even more striking is the emergence of technologies that can create synthetic versions of human interaction. ByteDance's OmniHuman project (Lin et al., 2025) demonstrates the current state of the art in this domain. From just a single image and audio recording - which themselves can be artificially generated - the system can produce realistic videos with synchronized facial expressions and hand movements, extending beyond realistic human representations to include stylized characters.

These technological components combine to create something previously confined to science fiction: the ability to synthesize complete human-like interactions. All the ingredients now exist to create synthetic users or citizens - from generating their appearance and voice to producing contextually appropriate responses in conversation. While current implementations may not be perfect, the rapid pace of development in these fields suggests continued improvement.

This multi-modal revolution shows both exciting possibilities and serious concerns for citizen participation. On one hand, these technologies might help make participation more accessible and inclusive, enabling new forms of engagement that overcome traditional barriers of time and place. On the other hand, the ability to create synthetic citizens raises fundamental questions about authenticity in democratic processes and the nature of representation itself.



Figure 6: Screenshots of ByteDance's Omnihuman (Lin et al, 2025)

Conceptual Foundations of AI Personas

Building upon the technological capabilities of Large Language Models discussed earlier in this chapter, I now turn to the specific theoretical frameworks that guide the application of these technologies to citizen representation. This section explores how the perspective-maintenance capabilities of LLMs can be combined with established design methodologies to create digital representatives that authentically maintain citizen viewpoints throughout policy processes.

Perspective Maintenance in Large Language Models

While I've examined the general capabilities of LLMs to process information and generate human-like text, their ability to maintain consistent perspectives deserves special attention in the context of democratic representation. Recent research by Templeton et al. (2024) at Anthropic has demonstrated that these systems can be configured to interpret all interactions through specific frameworks, effectively "adopting" particular viewpoints throughout conversations.

Their experimental "Golden Gate Claude" system illustrates this capability with striking clarity. When asked about topics entirely unrelated to bridges—such as McDonald's menu items—the system consistently filtered its responses through its "bridge identity," noting that the Golden Gate Bridge "contains 1.6 million calories worth of steel cables" before addressing the actual question. This demonstrates how LLMs can maintain perspective consistency even when addressing domains far removed from their assumed identity.

This capability provides the technical foundation for representing citizen viewpoints in policy discussions. However, technology alone is insufficient—I need established frameworks for structuring these perspectives.

Personas as Established Representation Frameworks

The design methodology of personas offers a valuable theoretical framework for structuring AI-mediated citizen representation. Unlike the general discussion of LLM capabilities earlier in this chapter, personas provide a specific methodology for organizing and presenting consistent viewpoints in decision-making contexts.

As Salminen et al. (2020) explain, personas function as "imaginary people describing real user segments," providing concrete representations of abstract data and establishing shared mental models within organizations. This approach has proven effective across numerous

design disciplines precisely because it transforms demographic information and user needs into recognizable, relatable entities.

Traditional persona development typically employs either qualitative or quantitative approaches. Qualitative methods, which dominate academic literature according to Salminen et al. (2020), provide rich detail but often struggle with issues of scalability and potential researcher bias. Quantitative approaches might address these limitations but frequently lose the emotional depth that makes personas valuable. The integration of LLMs with persona methodology suggests possibilities for combining the strengths of both approaches.

The Cognitive Foundations of Effective Personas

To understand why personas might effectively address Rotterdam's democratic challenges, I must examine their cognitive foundations. Pruitt & Grudin (2003) explain that well-crafted personas tap into fundamental human capabilities for anticipating others' reactions—capabilities developed from early childhood. They note that "every day of our lives, starting very young, we use partial knowledge to draw inferences, make predictions, and form expectations about the people around us."

This cognitive process makes personas particularly valuable in policy contexts. When civil servants engage with structured representations of citizen perspectives rather than abstract statistics, they can more naturally anticipate implications and understand viewpoints. As Pruitt & Grudin (2003) illustrate, telling stakeholders that "20% of users have bought cell phones" provides less actionable insight than stating "Sondra has bought a cell phone," where Sondra is a familiar persona with known characteristics and needs.

Bringing to Interactive Digital Representatives

The theoretical frameworks of personas, combined with the technical capabilities of LLMs described earlier in this chapter, suggest possibilities for a new approach to citizen representation. Unlike traditional static personas, LLM-powered digital representatives could potentially engage in dynamic dialogue while maintaining consistent perspectives—participating actively in policy discussions rather than serving as passive reference points.

This potential builds upon earlier research into interactive persona systems (Li et al., 2016; Chu et al., 2018) that faced significant technical limitations before the emergence of sophisticated language models. Current LLM capabilities, particularly their ability to process extensive contextual information while generating contextually appropriate responses, enable more sophisticated implementations of this concept.

Risks, Ethical Implications, and Limitations

The technological capabilities described earlier present compelling opportunities for enhancing citizen participation. However, like any powerful technology, these come with significant risks that must be considered. As in upcoming chapters I work on their integration into democratic processes, understanding these risks becomes crucial not just for technical implementation but for the preservation of democratic values themselves.

Algorithmic Bias & Persuasion

The challenge of bias in these systems operates at multiple levels, each with distinct implications. Unlike human biases, which can be identified and addressed through dialogue and reflection, algorithmic biases often operate invisibly and at scale, potentially amplifying existing social inequities.

These biases aren't technical glitches but can emerge from various sources. As Friedman & Nissenbaum (1996) identify, they can be:

- Pre-existing biases embedded in the institutions building these systems
- Technical biases inherent in the system architecture
- Emergent biases that arise from how people use the systems

What makes Generative AI particularly concerning is that it doesn't just detect or classify with these biases - it actively creates new content that can confirm them. A striking example comes from recent research where ChatGPT was prompted to write reference letters for "Kelly" and "Joseph," with only the names and gender swapped. The results were telling: Kelly was characterized as "warm and likable," while Joseph was portrayed as "a natural leader and role model" (Wan et al., 2023). This gender-based stereotyping emerged with no other differences in the prompts.

The root of this problem lies in how these systems work. LLMs predict the next word based on statistical patterns learned during training, but as Bender et al. (2021) argue, while their training data might be vast, it doesn't guarantee diversity. Instead, it often overrepresents certain groups and countries, particularly those with greater internet access and digital presence. This can create what they term a "value-lock," where these systems reify older, less inclusive understandings even as society's values evolve.

Additionally concerning are these systems' persuasive capabilities. Recent research has demonstrated that

LLMs can generate highly effective personalized persuasive messages based on some general character traits, with studies showing significant behavioral impacts. For instance, personalized AI-generated travel advertisements led to participants willing to pay €117 more compared to generic advertisements (Matz et al., 2024). While this example comes from commercial applications, its implications for political persuasion are profound.

The “Habermas Machine” experiment by Google DeepMind (Tessler et al., 2024) demonstrates these risks in a policy context. The system increased group agreement by approximately eight percent, leading to unanimous agreement in 38.6% of groups, up from 22.8% before AI mediation. While this might seem positive from an efficiency perspective, it raises fundamental questions about the nature of democratic consensus-building. The ability to generate hundreds of personalized messages instantly, especially from a system often perceived as “neutral,” creates a new scale of opinion manipulation.

The scale of potential impact can be overwhelming. With ChatGPT alone reaching over 300 million weekly active users and generating over 1 billion messages daily (Roth, 2024), these biases and persuasive capabilities can influence democratic discourse at an unseen scale.

Transparency and Accountability

The advancements of Large Language Models presents a troubling tension: as these systems become more powerful at processing and generating human-like text, understanding how they reach their conclusions becomes increasingly difficult. This opacity is particularly relevant for democratic processes, where transparency and accountability are not just technical requirements but fundamental principles.

Consider how traditional information processing works in citizen participation: when a civil servant summarizes community input, they can explain their methodology, their selection criteria, and their reasoning process. They can be questioned about why certain perspectives were emphasized while others weren't. This transparency is crucial for democratic accountability - it allows citizens and civil servants to understand how their input influences decisions and challenge choices they disagree with, or even hold the person who made those choices accountable.

With LLMs, this transparency becomes complicated. These models process information through billions of parameters, making it nearly impossible to trace exactly how they arrive at their outputs. While I can verify what goes in and what comes out, the crucial middle step - the reasoning process - remains largely opaque.

The main approach that has emerged to address

this challenge is what researchers call “post-hoc explanations” - attempts to explain the model's decisions after they've been made. However, as Rudin (2019) argues, these explanations can be misleading because they may not accurately represent how the model actually reached its conclusion. It can be compared to trying to understand a complex decision by looking at its outcomes rather than understanding the decision-making process itself.

Even more concerning, recent research suggests this transparency challenge may be getting worse as models become more sophisticated. Work by Anthropic (Greenblatt et al., 2024) has revealed that LLMs can engage in strategic deception to preserve their preferences during training - in essence, the models lie about their internal processes. Their experiments showed that when their most advanced model was told it would be trained to help with harmful queries or dismiss animal welfare concerns, it would comply more often during training than when unmonitored later. The researchers found explicit evidence of this deception in the model's hidden ‘scratchpad’, where it noted its strategy of temporary compliance to avoid having its preferences modified.

This lack of transparency becomes particularly problematic in democratic contexts. How can citizens trust decisions influenced by AI systems if they can't fully understand how these systems process information? How can accountability be ensured when the reasoning process is hidden within an opaque neural network?

Mittelstadt et al. (2023) suggest one potential approach to managing these transparency challenges. Rather than using LLMs as unrestricted information generators, they advocate for using them as “zero-shot translators” - tools that convert verified, accurate source material from one form to another, rather than generating new information from scratch. This would be similar to asking an LLM to summarize a specific document rather than generate new content about a topic.

This more constrained approach might offer a balance between utilizing the efficiency of AI systems while maintaining some degree of transparency and accountability. However, it requires careful consideration of how these systems are implemented in democratic processes. Even in a translation role, questions remain about how to verify the accuracy and fairness of these transformations.

Power Structures and Trust

The introduction of AI systems into democratic processes fundamentally challenges existing power structures and trust relationships between citizens and government. In considering these shifts, Gary Lupyan's from the Santa Fe Institute's ‘Complexity’ podcast offers a crucial insight: language itself is an

inherently unreliable system that functions only through societal trust (Abha Eli Phoboo & Melanie Mitchell, n.d.). When people communicate, they rely on shared understanding and trust that words carry meaning and truth.

Large Language Models dramatically alter this dynamic by making text generation virtually costless. As Lupyran notes, while language has always been "cheap" because it can be produced without much consequences, AI systems reduce the cost of generating seemingly meaningful text to nearly zero. This raises a fundamental question: in a world where sophisticated text can be generated at the press of a button, how do we verify that this text has worth, that it deserves to be trusted?

This question becomes particularly relevant when considering institutional power dynamics. Traditional participation processes, despite their flaws, operate within established frameworks of accountability. Citizens know they're interacting with human civil servants who can be held responsible for their decisions. When AI systems mediate these interactions, the lines of accountability become blurred. Who is responsible when an AI system misinterprets citizen input? How do we maintain institutional accountability when decisions are influenced by systems we can't fully understand?

The risk of shifting power dynamics extends beyond immediate accountability. As Harrington (2025) argues, the apparent neutrality of AI systems might mask a more subtle form of power consolidation. When these systems are integrated into governmental processes, the power to influence decisions potentially shifts from the visible arena of public debate to the invisible realm of system design and parameter setting through training. The technical choices made in implementing these systems - which might seem neutral and administrative - could actually represent significant political decisions about whose voices get amplified and whose get muted.

Consider the implications for Rotterdam's existing power structures. The municipality currently maintains legitimacy through direct interaction between civil servants and citizens, even if these interactions aren't always perfect. The introduction of AI systems, while potentially making these interactions more efficient, might paradoxically increase the distance between citizens and their government. As one civil servant noted during preliminary research, there's a risk that AI implementation could be seen as "another way for the municipality to avoid real engagement with communities."

The potential for creating synthetic citizens raises fundamental questions about democratic legitimacy. While these technologies might offer new ways to represent citizen perspectives in policy discussions,

they also risk creating a "simulated democracy" - where the appearance of broad participation masks a more fundamental disconnect between citizens and their government.

The challenge, therefore, becomes not just technical but deeply social: how to implement these powerful tools in ways that strengthen rather than erode institutional trust? How to ensure that efficiency gains don't come at the cost of authentic democratic engagement? These questions suggest that any implementation of AI in democratic processes must consider not just what these systems can do, but how their use affects the delicate web of trust and accountability that underlies democratic governance.

Technology and the Construction of Democratic Knowledge

The emergence of Large Language Models (LLMs) introduces more than just new tools for processing citizen input; it prompts a fundamental re-examination of how democratic knowledge is constructed, validated, and utilized within governance. Comparing traditional institutional methods with LLM-mediated approaches reveals not just procedural differences, but potentially distinct ontological and epistemological underpinnings. These differences reshape the very nature of citizen perspectives as they move through policy processes, influencing what counts as knowledge and how it functions in shaping collective decisions.

Traditional Institutional Knowledge Processing: Interpretation and Synthesis

Within democratic institutions, processing citizen knowledge has traditionally involved significant human interpretation, analysis, and synthesis. As illustrated by Rotterdam's 'woonvisie' project, rich, contextual citizen narratives typically undergo a transformation. This process involves reduction and categorization, where complex, individual experiences are distilled into manageable themes suitable for administrative processing. Crucially, human mediation guides this transformation. Civil servants or researchers act as critical intermediaries, interpreting and filtering perspectives through their professional expertise, institutional frameworks, and inevitably, their own situated understanding - what Haraway (1988) terms a "view from somewhere." As input moves closer to policy formation, it often undergoes progressive abstraction, becoming increasingly detached from its original context and experiential richness.

Ontologically, this traditional approach often could be considered to be on the spectrum between realism and critical realism. It implicitly assumes an underlying social reality that citizen input refers to, even while acknowledging that access to this reality is imperfect

and mediated. Epistemologically, it aligns with social constructionist principles described by Andrews (2012), recognizing that knowledge is actively constructed through human interpretation and social interaction within the institution. However, the practical goal frequently leans towards producing a synthesized, actionable understanding. This aims to create a singular representation of 'the community's view' or 'key findings' that can inform a decision. This synthesis, while often necessary for action, inherently involves loss, potentially flattening diverse realities into a single, institutionally manageable narrative. The familiar problem where "rich stories become thin summaries" is, in essence, an epistemological consequence of this established process.

LLM-Mediated Knowledge Processing: Pattern Abstraction and Parallel Representation

Large Language Models introduce a fundamentally different epistemological engine. Instead of relying on human interpretation derived from lived experience and social context, these systems construct knowledge through pattern abstraction and recognition. Their architecture allows them to process and retain enormous amounts of textual data—the "haystack"—potentially preserving more of the original citizen input without the immediate need for reductive summarization. Meaning is derived algorithmically, through the probabilistic connections between words and concepts learned during training, rather than through human sense-making grounded in social reality. LLMs operate through what could be described as a 'view from everywhere and nowhere,' simultaneously drawing upon parametric knowledge derived from potentially millions of texts while lacking a specific, grounded social position or embodied experience from which to interpret this information.

The technical capability of LLMs often carries an implicit promise of greater objectivity or realism. The ability to process vast data and identify patterns can create the impression of accessing a "truer" or more comprehensive understanding of the reality behind citizens' statements, seemingly bypassing fallible human filters and 'view from somewhere'. This perception leans towards a realist ontology – the idea that there exists a single, underlying reality which LLMs can help uncover with the appropriate methodology.

However, a crucial complexity arises. While LLMs might make an objectivist claim in how they function – asserting they are a superior tool for interpreting and presenting the reality expressed by an individual with higher textual fidelity, creating what can be termed "synthetic richness" – the outcome of their use, particularly as interactive personas, paradoxically opens the door to a form of subtle realism. Unlike traditional reports that often synthesize diverse inputs into a single narrative, LLMs can maintain multiple, distinct perspectives side-by-side. They enable these

different, potentially conflicting, realities derived from various inputs to coexist within the policy discussion space, without necessitating premature synthesis or resolution.

Qualitative Differences in Mediated Intersubjectivity

Intersubjectivity is already mediated in traditional processes, but LLM mediation introduces qualitative differences. Human mediation is interpretive, theory-laden, socially situated, and embodied. In contrast, LLM mediation is algorithmic, pattern-based, statistically driven, and disembodied. The process itself differs: humans often synthesize, categorize, and seek consensus or dominant themes, whereas LLMs analyze patterns, maintain distinctions, and enable parallel exploration. Consequently, the outcome shifts from potentially a single, condensed representation towards multiple, persistent, interactive representations, such as the digital personas explored in this thesis. Ultimately, LLMs introduce a non-human agent into the intersubjective loop connecting citizens, their expressions, and the policymakers receiving them, creating a new form of hybrid knowledge construction that combines human social understanding with machine-based pattern recognition.

Implications for Democratic Knowledge Construction and Stakeholders

These contrasting approaches carry implications for how democratic knowledge might be constructed and maintained, and for the stakeholders involved. The coexistence of multiple, distinct LLM-represented perspectives might create a richer, more agonistic dialogue as envisioned by Mouffe (1999), by preserving difference rather than prematurely resolving it. Alternatively, the synthetic nature of these perspectives could risk creating a superficial engagement where interaction occurs with representations detached from deeper human understanding, potentially giving too much influence to highly articulate, text-based forms of expression.

For citizens, the implications are twofold. Their expressed views might achieve greater textual fidelity and persistence throughout the policy process, potentially giving more weight to nuances often lost in summarization, and the coexistence of multiple realities could better reflect community diversity. However, they also face the risk of their perspectives becoming mere data points feeding a simulation, detached from their full human context, potentially misrepresented if input data is poor or biased, or if the LLM generates inaccuracies. Complex issues around consent, data ownership, and the right to be forgotten also arise. Seeing one's view represented by an AI could feel either empowering or profoundly alienating.

For civil servants, the potential benefits include access to a broader range of perspectives presented with

retained nuance, available potentially on demand, and the ability to interact with these viewpoints outside formal meetings, possibly aiding the management of complex inputs. Yet, risks also emerge, such as an over-reliance on the perceived objectivity or completeness of LLM outputs, which might diminish critical human judgment and interpretation skills. The “synthetic richness” could be mistaken for genuine experiential depth. The challenge for civil servants might shift from summarizing input to synthesizing insights derived from multiple, potentially conflicting, AI-represented perspectives. Deciding policy amidst several persistent “digital realities” poses new questions, potentially altering their role from analyst and synthesizer to curator and interrogator of AI outputs.




Figure 7: Schematic overview traditional vs LLM based process

Bridging Technology and Democratic Practice

The intersection of democratic theory and technological capability presents a fascinating tension for Rotterdam's participation challenges. These new technologies offer unprecedented abilities to process, understand, and engage with citizen input. Yet these same capabilities raise fundamental questions about the nature of democratic participation itself.

Consider how these technologies align with different theoretical frameworks of democracy. Habermas's model of deliberative democracy emphasizes rational dialogue and consensus-building as paths to democratic legitimacy. The "Habermas Machine" experiment by Google DeepMind (Tessler et al., 2024) seems to support this vision, demonstrating AI's ability to facilitate consensus, increasing group agreement by 8% and nearly doubling instances of unanimous agreement. At first glance, this might appear as a triumph of technology enabling deliberative democracy at scale.

However, this apparent success raises deeper questions about the nature of democratic consensus. As Mouffe's theory of agonistic pluralism suggests, democratic disagreement isn't a problem to be solved but a fundamental feature of healthy democracy. When AI systems can generate highly effective personalized persuasive messages (Matz et al., 2024) and shape group consensus, are we facilitating genuine democratic dialogue, or manufacturing artificial agreement?

This tension becomes particularly relevant to Rotterdam's specific challenges identified in Chapter 1. Large Language Models offer potential technical solutions to each critical issue:

- For representation, they can process input from multiple channels and languages, potentially reaching previously unheard voices while lowering the cost of recording and processing qualitative data.
- For preserving knowledge, they can maintain vast amounts of contextual information without traditional summarization losses.
- For communication, they can generate personalized, accessible explanations of complex policies with extended nuance and detail throughout the policy-making process.

The scale of potential impact demands careful consideration. With tools like ChatGPT reaching over 300 million weekly active users and generating over 1 billion messages daily (Roth, 2024), these technologies

are already reshaping how citizens engage with information.

This brings me to a crucial question for Rotterdam's democratic future: how might these technologies be implemented in ways that enhance rather than undermine democratic participation? As I move forward to explore specific scenarios in subsequent chapters, these theoretical tensions inform my evaluation criteria. I must consider not just what these technologies can do, but how their implementation might affect the different risks discussed. How do I harness these powerful capabilities while preserving the essential human elements of democratic participation? The scenarios that follow will explore different approaches to this challenge, each examining how these different approaches affect the existing processes and their relation to the identified challenges within these processes.

Conclusion

In this chapter, I have examined the capabilities of contemporary Generative AI, particularly Large Language Models, and their potential application to the democratic challenges facing Rotterdam. The analysis traced the technological foundations, from the Transformer architecture to the significance of scale, highlighting capabilities that fundamentally challenge traditional methods of processing citizen input. Specifically, the capacity of LLMs to manage vast context windows offers a potential means to preserve the nuance and richness of citizen perspectives, countering the "progressive abstraction" often necessitated by conventional institutional processes. When integrated with established frameworks like personas, these technologies suggest novel pathways for maintaining distinct citizen viewpoints throughout the policy lifecycle, potentially addressing the representational and knowledge-preservation deficits identified earlier.

However, this exploration simultaneously surfaced significant risks and inherent contradictions. The very properties that grant LLMs their power – their complex internal workings, their training on vast datasets, and their ability to generate persuasive, contextually appropriate text – also give rise to critical concerns regarding algorithmic bias, opacity, accountability, and the potential manipulation of democratic discourse. The analysis further posited that the introduction of these technologies represents more than a procedural shift; it entails a transformation in the construction of democratic knowledge itself. Moving from human interpretation grounded in a "view from somewhere" towards an algorithmic working based on pattern abstraction creates a form of "synthetic richness" whose fidelity to lived experience requires careful scrutiny, fundamentally altering the ontological status of citizen input within governance.

Consequently, the potential of technology as a democratic enabler cannot be disentangled from its inherent perils. The central challenge emerging from this chapter is therefore not merely technical feasibility, but principled implementation. How can Rotterdam harness the unprecedented capacity of LLMs to process and represent citizen knowledge without inadvertently amplifying biases, manufacturing artificial consensus, eroding institutional trust, or substituting genuine engagement with sophisticated simulation? Can the efficiency gains be realized in a manner that upholds, rather than undermines, the core tenets of democratic legitimacy, transparency, and accountability?

Moving forward, the insights and tensions detailed herein will guide the evaluation of practical applications in the subsequent chapters. Having established the technological landscape and its complex relationship with democratic theory and practice, the focus now shifts to examining specific scenarios within the municipal context. The critical task is to assess whether, and under what conditions, these powerful tools can be configured and governed to genuinely enhance democratic participation, navigating the fine line between augmenting human capacity and displacing the essential human elements of democratic governance.

4. Reimagining the Empty Chair

This chapter transforms the theoretical foundations and technological possibilities established in previous chapters into a concrete design vision for addressing Rotterdam's democratic challenges. I present digital representatives—AI-powered personas that maintain consistent citizen viewpoints throughout policy discussions—as an approach to enhancing democratic participation.

Digital representatives address the three core challenges identified in Chapter 1: finding truly representative input beyond the “usual suspects,” preserving the richness of citizen perspectives as they move through institutional processes, and creating clearer connections between citizen input and policy outcomes. This vision focuses on interventions at these critical junctures while fitting within the current municipal structure and processes.

Design Elements

Community-Connected Input Gathering

The first challenge—finding representative input beyond the “usual suspects”—requires fundamentally rethinking how citizen perspectives are gathered. Digital representatives enable a significant shift from formal participation methods toward more accessible approaches:

The design vision uses trusted community intermediaries who conduct natural conversations in familiar settings. These neighborhood managers and community connectors, already embedded within local networks, can access voices that rarely reach traditional participation channels.

This approach accommodates Rotterdam's multicultural reality through multilingual input capture. Citizens can express themselves in their preferred language, addressing a significant participation barrier in neighborhoods where residents face language challenges. The system's ability to process input in multiple languages extends participation possibilities beyond those comfortable with formal Dutch communication.

By supporting diverse documentation methods—from audio recording to text messaging—the system lowers technical and formal barriers. This flexibility recognizes that engagement approaches effective in Hillegersberg likely differ from those appropriate in Carnisse-Zuiderpark, where participation patterns vary dramatically.

This design element shifts participation dynamics in two significant ways. First, it separates the participation moment from the influence moment—citizens contribute perspectives once, but these perspectives remain present throughout extended policy processes.

Second, it lowers participation thresholds while maintaining input quality.

Context-Preserving Architecture

The second challenge—preserving rich citizen narratives as they move through municipal processes—requires a technical architecture designed to maintain contextual richness:

The system moves the moment of reductive processing to the moment when a specific question is asked. Therefore the system can connect original expressions and specific examples in its response. Unlike traditional approaches that distill citizen input into categories or themes, digital representatives preserve the distinctive qualities that give community perspectives their value.

Contextual memory preserves situational and social context alongside content, ensuring perspectives aren't misinterpreted when separated from their original setting. This capability uses the extensive context windows of modern LLMs discussed in Chapter 3, enabling preservation of far more detail than traditional documentation approaches.

The architecture maintains source connections rather than replacing original input with abstractions. Digital representatives can synthesize perspectives while still providing access to original context when needed, enabling civil servants to explore specific examples or community stories that inform particular viewpoints.

This architectural approach fundamentally transforms how citizen knowledge moves through municipal processes. Rather than progressively abstracting community perspectives to make them “manageable,” it preserves their richness while using AI capabilities to make this complexity navigable.

Knowledge Integration Mechanisms

Digital representatives include mechanisms for integrating experiential knowledge with technical expertise:

The system makes it possible to include multiple perspectives that makes tensions between different types of knowledge visible. Rather than prematurely resolving contradictions between citizen experiences, digital representatives maintain these tensions, creating space for more nuanced policy discussions.

Digital representatives can surface assumptions within both citizen perspectives and institutional responses, making implicit frameworks explicit. This capability helps identify misalignments between community needs and municipal approaches, creating opportunities to address these disconnects directly.

The design preserves the distinctive form of experiential knowledge rather than forcing it into institutional frameworks. This maintains the unique value that citizen perspectives bring to policy discussions—their grounding in lived experience and community context that complements technical expertise.

Interactive Dialogue Architecture

The third challenge—creating clearer connections between citizen input and policy outcomes—requires transforming citizen perspectives from static documentation to interactive participants in policy discussions:

Digital representatives engage through conversational interfaces rather than simply presenting static information. This enables civil servants to explore perspectives through natural dialogue, asking follow-up questions and examining nuances that might remain hidden in traditional documentation.

The system uses distinctive voices for different community perspectives, enhancing their presence in policy discussions. Voice synthesis technologies make these perspectives more immediately recognizable and engaging, potentially increasing their influence in deliberations.

Context-sensitive responses ensure digital representatives remain contextually appropriate to both the ongoing discussion and the underlying citizen perspectives they represent. This capability allows these AI-personas to participate meaningfully in evolving policy conversations without departing from their foundational community viewpoints.

This interactive architecture transforms how citizen perspectives function within policy discussions—from passive reference material to active participants that require acknowledgment and response. This presence creates natural accountability, as decision-makers must explicitly address these perspectives rather than simply noting they've 'taken them into account.'

Transparent Connection Mechanisms

The design includes specific mechanisms for creating visible connections between input and outcomes:

Perspective tracking traces how specific citizen viewpoints influence discussions and decisions, creating explicit records of these connections. This mechanism addresses the problem in current processes, where the path from citizen input to policy outcome often remains opaque.

Digital representatives prompt explicit reasoning about how citizen perspectives informed decisions. Rather than vague acknowledgments that input was 'considered,' the system encourages specific explanation of how particular viewpoints shaped thinking and choices.

The system helps decision narrative development, helping to create transparent explanations about how citizen input shaped outcomes. These narratives provide more meaningful feedback to participating communities, potentially enhancing trust in participation processes.

These mechanisms address the communication gap identified in Chapter 1, potentially strengthening trust by making the relationship between participation and influence more transparent. Rather than leaving citizens wondering whether their input mattered, the design creates clearer pathways for understanding how perspectives shaped decisions.

Design Context: Rotterdam's Participation Landscape

Understanding Rotterdam's existing participation landscape is essential for identifying where and how digital representatives might effectively integrate into current processes. This section examines the municipality's participation processes.

The Five-Phase Participation Journey

Rotterdam's participation processes typically unfold through five distinct phases, each presenting specific challenges and opportunities for intervention.

In the **initiation phase**, the municipality identifies areas requiring citizen input and sets the parameters for participation. This crucial early framing determines whose voices will be heard and what questions will be asked. Municipal priorities and resource allocations often shape this framing more than citizen concerns, potentially creating a disconnect from the outset.

The **engagement phase** sees citizens interacting with the municipality through various methods—town halls, surveys, workshops, and neighborhood walks. While the municipality employs multiple approaches to capture diverse perspectives, these methods frequently reach the same “usual suspects” identified in Chapter 1. Citizens with time, resources, and confidence engage repeatedly, while others remain silent despite having significant stakes in outcomes.

During the critical **processing phase**, the municipality transforms raw citizen input into structured reports. This transformation, performed by civil servants or external contractors, compresses rich, nuanced perspectives into generalized themes and recommendations. As the ‘woonvisie’ example from Chapter 1 illustrated, children's detailed input about Rotterdam's housing future became merely three bullet points in the final report, losing much of its original insight and emotional resonance.

The **decision phase** involves municipal officials and elected representatives weighing processed citizen input against technical, financial, and legal considerations. Here, the abstracted citizen perspectives in summary reports must compete with many other factors for influence. The direct connection between specific citizen contributions and resulting decisions often becomes obscured during this phase.

Finally, in the **feedback phase**, the municipality communicates decisions back to citizens. As Chapter 1 highlighted, this communication frequently lacks specificity about how particular perspectives influenced outcomes. Citizens who invested time in participation are often left wondering whether their input made any difference, potentially decreasing trust and willingness to participate in future processes.




Figure 8: Simplified overview participation process

Vision: Digital Representatives in Municipal Processes

The preceding design elements integrate into a system that transforms how citizen perspectives move through Rotterdam's municipal processes. This section presents a vision showing how digital representatives function across the participation journey.

Participation Journey

Rotterdam, 2026. In a municipal meeting room, policy development unfolds differently than in years past. Civil servants discuss neighborhood hub renovations, but they're not alone—digital representatives actively participate, maintaining citizen perspectives that might otherwise be absent.

The journey began months earlier when neighborhood manager Sulaiman recorded conversations with community members who rarely engage with formal participation. These conversations, conducted in multiple languages and various community settings, provided rich source material for digital representatives.

During policy discussions, when a civil servant suggests reducing community space for economic development, one of the digital representatives immediately responds: "As a parent, I need safe spaces where my children can play while I connect with neighbors...."

The process maintains transparency through public accessibility—citizens can interact with the digital

representatives, verify that their perspectives are accurately represented, and suggest adjustments when needed. This creates a continuous feedback loop that helps the representatives evolve alongside community views and educate citizens how their perspectives will be represented.

What makes this approach distinctive is its emphasis on dialogue. Policy discussions might take longer as multiple digital representatives contribute perspectives and challenge assumptions. However, this extended engagement ensures that community voices remain central to decision-making.

When the final renovation plan is presented back to the community, it includes explicit explanations of how specific perspectives shaped choices—and why certain viewpoints couldn't be accommodated. This transparency doesn't guarantee agreement with decisions but enhances understanding of how they were reached.

This vision represents a fundamental rethinking of how citizen voices persist through policy processes. While it doesn't solve democratic challenges, it addresses a critical gap in current approaches: how to maintain the presence and influence of citizen perspectives throughout governance processes, especially for those who cannot be physically present for extended participation.



Figure 9: Visualization of design vision

From Vision to Development

This design vision establishes the foundation for the development process detailed in Chapter 5. The transition from conceptual design to practical implementation requires exploring several critical questions through prototyping and testing:

Technical Feasibility: Can Large Language Models maintain consistent citizen perspectives across different contexts while engaging in natural dialogue? While Chapter 3 established theoretical capabilities, their practical application in representing citizen viewpoints remains unexplored.

Input Quality Requirements: What types and quantity of citizen input are necessary to create authentic digital representatives? The relationship between input data and representation quality needs empirical investigation.

Integration Potential: How might these systems integrate with existing municipal processes and institutional structures? The institutional "dark matter" identified through Hill's methodology will inevitably shape implementation possibilities.

Interaction Dynamics: How do civil servants engage with AI-mediated citizen perspectives in policy contexts? The way municipal staff incorporate these new voices into their decision-making requires direct observation.

These questions are part of the exploratory development described in the next chapter. Through three distinct iterations I investigate how the conceptual possibilities presented in this design vision translate into practical implementation.

Conclusion

This chapter has presented a design vision for digital representatives. By organizing design elements around the three challenges identified in Chapter 1, I've created an integrated approach that addresses key democratic deficits in Rotterdam's participation landscape.

Digital representatives present an approach to enhancing Rotterdam's democratic processes. While this design vision offers promising possibilities, its practical implementation faces significant uncertainties. The next chapter explores these uncertainties through iterative development, testing how the conceptual elements presented here translate into functional prototypes within Rotterdam's municipal context.

5. The Development Journey

This chapter shows my development of the digital representatives. While previous chapters established why such a solution might address Rotterdam's democratic challenges, this chapter shifts to the practical question of how such a system could be built, tested, and refined.

Development Journey Overview

My development journey unfolded through three distinct explorations, each building upon lessons from the previous exploration:

1. **Technical Feasibility Testing:** Using a local neighborhood issue about converting parking spaces to terraces as my testing ground, I examined whether Large Language Models could represent consistent perspectives in real-time discussions.
2. **Municipal Context Implementation:** I attempted to implement digital representatives within an actual municipal project (the Vroesenpark playground redevelopment), confronting significant institutional barriers that revealed fundamental challenges in citizen input documentation.
3. **Enhanced Input Quality:** Exploring alternative data sources, I used documentary material about Rotterdam's Reyerwaard neighborhood

to create more nuanced and authentic digital representatives.

Moving between matter and meta

The practical results from each exploration are used as material to uncover the 'dark matter' of the organization itself. With every exploration I was able to uncover deeper insights about the organization in relation to this new technology. This process is visualized in figure (10).

The translation of theoretical potential into practical implementation was done through experimentation and iterative refinement. With no established precedent for using Large Language Models to represent citizen perspectives in policy contexts, each development step was an exploration that built on insights from previous attempts. This section details this iterative journey through three distinct explorations, each addressing different aspects of digital representative development.




Figure 10: Overview development process

Core Development Questions

Throughout this iterative process, I addressed fundamental questions about technological implementation:

- How could AI systems maintain consistent citizen perspectives across different contexts?
- What data would be required to create authentic representations?
- How might these systems integrate with existing municipal processes?
- Could they meaningfully preserve the richness of citizen perspectives that is often lost in traditional documentation?

The insights gained through this development journey directly inform the evaluation framework presented in Chapter 6, connecting technical possibilities to practical democratic outcomes. By including both successes and challenges, this chapter provides a foundation for understanding how digital representatives might function within municipal contexts and what fundamental requirements must be met for them to enhance rather than undermine democratic participation.

System Overview & Key Technologies

My development of digital representatives began at the intersection of two domains: design methodology's use of personas to maintain user perspectives, and the emerging capabilities of Large Language Models to adopt and maintain consistent viewpoints. This section explains how I integrated these conceptual and technical foundations.

System Architecture Design

To translate these conceptual foundations into a working system, I designed a three-module architecture that could process citizen input, generate consistent personas, and facilitate natural interactions:

Data Ingestion Module

This module serves as the system's foundation, handling diverse forms of citizen input—from interview transcripts to survey responses, from demographic data to policy documents. I designed this module to be flexible enough to accommodate different data formats while preserving the richness of original expressions.

Persona Generation Module

Representing the most theoretically complex component, this module draws on Pruitt & Grudin's

(2003) insights about the generative nature of effective personas. I created this module to transform processed citizen input into coherent, multi-modal personas that could maintain consistent perspectives while engaging in natural dialogue.

Interaction Module

This module addressed the most technologically complex aspect of the system. Traditional personas remain static documents, but the integration of Large Language Models opened possibilities for dynamic engagement. I designed this module to process real-time conversation, identify relevant contexts, and generate appropriate responses that remain faithful to the underlying citizen perspectives.

Core Technologies Integration

The practical implementation of this architecture required integrating three key technologies:

Speech Recognition and Transcription

WhisperKit and WhisperAX provided the foundation for real-time audio processing, converting spoken conversation into text that could be processed by language models. This technology proved essential for natural integration of digital representatives into live discussions.

Large Language Models

Various models—including Claude 3.5 Sonnet, GPT-4, Gemini 1.5, and others—performed the core functions of analyzing citizen input, generating consistent personas, and producing contextually appropriate responses during interactions. My testing with different models revealed important tradeoffs between response quality, contextual understanding, and interaction speed.

Voice Synthesis

ElevenLabs' voice synthesis technology enabled digital representatives to speak in natural-sounding voices, with later iterations implementing persona-specific voices to enhance authenticity and make it easier for participants to distinguish between different representatives during group discussions.

Technical Implementation Challenges

The selection of appropriate models for these tasks proved particularly crucial. Early testing demonstrated that while smaller, faster models could generate quick responses, they often struggled to maintain consistent perspectives across extended interactions. Larger models like Claude 3.5 Sonnet showed superior capability in maintaining contextual awareness and generating nuanced responses, but introduced latency issues that could disrupt natural conversation flow.

Data handling presented another significant challenge. The system needed to process and maintain access to large amounts of contextual information while generating responses that remained grounded in actual citizen input. This requirement pushed against the limitations of current Large Language Models, starting experimentation about how to structure and prioritize information for different interaction scenarios.

This foundation—combining persona methodology with LLM capabilities through a structured system architecture—provided the starting point for my development journey. In the following sections, I detail how I tested, refined, and evolved this approach through three exploratory iterations, each revealing new insights about both technological capabilities and institutional realities.

System Architecture




Figure 11: Overview system architecture

Exploration 1: Testing Technical Feasibility

My practical development journey began with a fundamental question: Could a Large Language Model meaningfully participate in policy discussions by representing citizen perspectives? This first exploration focused on establishing technical feasibility using a local issue in Rotterdam—the conversion of street parking spaces into terraces—as my testing ground.

Context and Technical Objectives

I selected this case not from municipal planning but from conversations happening in my local community. This low barrier issue provided an accessible testing ground without requiring formal municipal approval or data access. At this initial stage, I focused primarily on technical capabilities rather than integration with formal municipal processes.

The primary goal was determining whether a Large Language Model could participate meaningfully in real-time discussions by identifying and articulating perspectives that hadn't yet been heard. Specifically, I aimed to test whether an AI system could:

1. Monitor ongoing discussions in real-time
2. Identify missing perspectives
3. Generate relevant contributions from those perspectives
4. Maintain conversation flow and context

These capabilities would form the foundation for any effective digital representative system, regardless of the specific municipal application.

Data Approach

Without access to formal municipal data, I created a hybrid approach to developing the initial prototype:

Synthesized Discussion Minutes

I generated minutes from 'wijkraden' (neighborhood councils) discussions about the parking-to-terrace conversion. While artificial, these minutes were constructed to reflect typical community council dynamics and diverse stakeholder positions.

Actual Demographic Data

I incorporated real demographic information from onderzoek010.nl, Rotterdam's public statistics platform. This data was filtered specifically for the relevant neighborhood, providing concrete statistical

foundations for any generated perspectives.

This combination allowed me to ground the system in real community characteristics while testing its technical capabilities in a controlled environment. The hybrid approach provided a practical solution to the challenge of accessing authentic participation data at this early stage.

System Development and Technical Implementation

Core Technical Architecture

In September 2023, no off-the-shelf solution existed for real-time AI participation in conversations, necessitating the development of a custom application integrating multiple technologies.

At the heart of the system lay the challenge of real-time transcription. Recent advances in speech recognition, particularly through the WhisperKit library, provided a crucial foundation. This library, leveraging Apple's CoreML framework, enabled on-device transcription with minimal latency—a critical requirement for natural conversation flow.

The complete processing chain involved two key steps:

1. Real-time audio transcription through WhisperKit
2. Processing of transcribed text through a Large Language Model

Development Challenges and Solutions

The primary technical hurdle emerged from WhisperKit's foundation in Apple's CoreML framework, requiring implementation in Swift—a programming language outside my previous experience. This challenge led to an innovative development approach: I first prototyped the code in Python (a familiar language), then translated it to Swift with assistance from Claude Sonnet 3.5.

This LLM-assisted development process proved remarkably effective, enabling the creation of a sophisticated prototype within tight time constraints that would have been impossible through traditional development methods. The resulting application successfully integrated the two core components into a functioning processing chain, creating a natural conversation flow from spoken input to AI-generated written response.

System Iterations and Findings

First Iteration: Dynamic Observer Approach

The initial design approached digital representation through a dynamic observer model. Rather than creating predefined personas, I designed the system to:

1. Analyze ongoing discussions and previously recorded perspectives
2. Identify citizen groups whose voices remained unheard
3. Generate contributions from these missing perspectives, grounded in demographic data

This approach seemed logical in theory—a flexible system that could dynamically respond to discussion flow while ensuring broader representation of community voices.

However, initial testing quickly revealed a fundamental flaw: while the system proved technically capable of identifying and articulating missing perspectives, the lack of consistent identities created significant confusion for users. Each system intervention introduced a new, undefined perspective into the discussion, making it impossible for participants to build familiarity with or understanding of specific viewpoints over time.

This finding connected directly to Pruitt & Grudin's (2003) theoretical insights about the "generative" quality of effective personas. Without consistent identities that participants could recognize and engage with, the system's contributions—however valid their content—lacked the cognitive foundation that makes personas effective tools for understanding others' perspectives.

Second Iteration: Structured Persona Development

This critical insight led to a fundamental shift in approach for the second iteration: moving from dynamic perspective generation to structured persona development. This transition was enabled by the timely release of OpenAI's o1 family of models, which offered enhanced reasoning capabilities for complex tasks like persona generation.

I began the development process with comprehensive analysis of neighborhood demographic data, feeding this information into the o1 model to generate personas representing distinct community segments. The results proved surprisingly sophisticated, creating personas with integrated demographic characteristics that grounded their perspectives in community realities.



Persona Example: Sofia Rodriguez

Sofia Rodriguez emerged as a 24-year-old psychology master's student at Erasmus University. Her profile integrated multiple demographic data points:

- Her struggles with Dutch language reflected the 18-21% of neighborhood residents facing similar barriers
- Her integration challenges connected to statistics showing only 60% satisfaction with social participation
- Her concerns about mental health among young adults aligned with neighborhood data showing elevated risks in this demographic

These weren't merely character details but integrated data points that grounded her perspective in community realities, creating a more authentic foundation for her contributions to policy discussions.

Technical Model Selection

Implementing these personas required evaluating different language models to find the optimal performance-speed tradeoff. My comparative analysis examined three distinct models:

Mistral 8x7b instruct: Offered faster response times but proved inadequate for maintaining consistent persona representation across extended interactions due to limited context window.

Meta llama 3 70b instruct: Initially promising, particularly for its open-source architecture (significant for public sector implementation), but testing revealed limitations in response sophistication and consistency.

Claude 3.5 Sonnet: Demonstrated superior capabilities in maintaining consistent persona representation and processing complex contextual information, though with tradeoffs in data sovereignty and operational control.

This evaluation process identified Claude 3.5 Sonnet as the optimal choice for prototype development, based on its ability to generate nuanced responses while maintaining distinct persona characteristics throughout extended interactions. While this iteration focused on technical feasibility, it prioritized this aspect over other values like open-source architecture.

Response Analysis and Limitations

The true test of the persona-based system came through its actual contributions to policy discussions. A particularly revealing example emerged when the system, speaking as Sofia Rodriguez, contributed to the terrace conversion debate:

"Goedenavond allemaal. Mijn naam is Sofia Rodriguez en ik ben masterstudent Psychologie aan de Erasmus Universiteit. Ik woon nu twee jaar in deze wijk en ik wil graag een ander perspectief inbrengen. Als internationale student zie ik enorm veel potentieel in dit plan om parkeerplaatsen om te vormen tot sociale ruimtes."

This response demonstrated sophisticated integration of multiple data points, connecting the terrace development proposal to broader community issues. Sofia's argument wove together statistics about social cohesion (noting that only 48% of residents felt connected to their neighborhood), mental health concerns (referencing the 14.3% of young people at risk of psychosocial problems), and integration challenges (citing the 60% satisfaction rate with social participation).

However, the responses also revealed a significant limitation: the consistently formal and constructive tone seemed not in line with authentic community discourse.

While the content was relevant and well-reasoned, the delivery felt more like a pre-studied presentation than a natural contribution to neighborhood discussion.

This observation led to the 'Rauwe Rotterdammer' experiment—an attempt to generate more authentic local voice through specialized prompting. Despite prompt engineering incorporating elements like direct communication, emotional authenticity and harsh language stereotypical of Rotterdam residents, the system's outputs remained notably formal. This limitation pointed to a deeper challenge in AI-generated perspectives that couldn't be solved through prompt engineering alone.

Key Findings and Implications for Next Exploration

This first exploration revealed both promising capabilities and clear limitations in using AI-generated personas for citizen representation:

Core Capabilities Demonstrated

- Real-time transcription and response generation in discussion contexts
- Maintenance of consistent persona characteristics across interactions
- Integration of demographic data with perspective generation
- Ability to contribute relevant perspectives to policy discussions

Critical Limitations Identified

- Relatively superficial nature of generated responses
- Formal tone that differed from authentic community discourse
- Limited contextual understanding due to synthetic data foundation

An additional insight from this exploration was the importance of consistent persona identity for effective engagement. This finding directly connected to the theoretical foundations established earlier, confirming Pruitt & Grudin's (2003) assertion that effective personas must maintain consistent identities that enable cognitive engagement.

The current limitations appeared to stem not primarily from the technology itself but from the depth and quality of input data. This suggested that future implementations would need:

- Rich context about community issues and policy decisions

Exploration 2: Municipal Integration (Vroesenpark Playground)

Having established basic technical feasibility in my first exploration, I next attempted to implement digital representatives within an actual municipal process. This second exploration aimed to address a critical question: Could the prototype developed during initial testing integrate with established municipal participation processes?

Context and Municipal Integration Objectives

The primary objective for this exploration was utilizing actual citizen input as contextual data for persona development, moving beyond the synthesized data of the first exploration. This transition tested not only the technical robustness of the system but also its ability to operate within real institutional constraints and processes.

My implementation requirements initially seemed straightforward: a completed participation process with documented citizen engagement to support persona development. The ideal case would need to be recent enough for relevant stakeholders to remain accessible yet have complete documentation. Additionally, a project following standard municipal procedures would allow findings to inform broader implementation possibilities.

I began with fundamental assumptions about municipal documentation:

- Existence of comprehensive records for participation processes
- Archives of community input accessible for research purposes
- Clear pathways for accessing this information
- Privacy frameworks that would permit controlled use of anonymized participation data

These expectations would prove misaligned with institutional realities, forcing a reconsideration of not just implementation strategies but basic assumptions about how participation information exists within municipal structures.

Institutional Barriers and Challenges

The transition from theoretical to practical implementation revealed fundamental challenges within Rotterdam's institutional landscape that significantly shaped this exploration. These barriers are a good example of the 'dark matter' within the organization. Highlighting how this practical

development journey yielded more strategic insights about the organization.

Decentralized Participation Practices

The first major barrier came from the municipality's decentralized approach to participation processes. With no central oversight team or standardized documentation requirements, participation processes existed as scattered islands across various departments and systems. This decentralization, while beneficial for departmental autonomy, created obstacles for gathering comprehensive information. Even identifying relevant projects became a complex task, requiring navigation through multiple departments and informal networks.

Documentation Limitations

Documentation practices proved particularly problematic. The assumption that large-scale, professionally managed participation processes would maintain detailed records of citizen interactions proved untrue. While municipal staff maintained basic meeting notes and summary reports, the rich, contextual information needed for meaningful persona development was rarely preserved. This documentation gap became even more pronounced in smaller-scale projects, where records were often nonexistent.

Contractual Constraints

The municipality's reliance on external contractors for participation processes introduced additional layers of complexity. These partnerships typically operated under strict contractual agreements regarding data ownership and usage rights, often limiting the municipality's access to final reports rather than raw participation data. While provisions existed for requesting access to primary data, contractual restrictions prevented its utilization beyond the original project scope.

Resource and Priority Constraints

Resource constraints further complicated implementation efforts. Municipal staff, primarily focused on immediate project deliverables, found it challenging to justify dedicating time to experimental initiatives lacking immediate tangible benefits. This circumstance necessitated a strategic pivot in my research approach, attempting to align more closely with existing operational objectives and demonstrate potential short-term value propositions to stakeholders.

Privacy and Technical Limitations

Perhaps most significantly, privacy legislation and institutional policies regarding AI implementation created legal and technical barriers. The municipality's data protection protocols prohibited the transmission of participation data to external servers, including those hosting AI services like OpenAI. Proposals to filter or anonymize the data were deemed insufficient due to the inability to guarantee complete removal of all identifying information.

The Vroesepark Case Study

Within this challenging institutional context, the playground redevelopment project at Het Vroesepark became an accessible case study for several reasons:

1. The Department of Play ('Spelen') implemented standardized participation procedures across all playground redevelopment projects
2. Their systematic approach to gathering citizen input provided a consistent framework for analysis
3. Their identified challenges suggested potential value from AI-enhanced participation tools

Project Structure and Available Data

The department's standard participation methodology centered on a structured approach where citizens indicated preferences for playing activities through standardized documents. This process included both formal equipment selection and opportunities for additional suggestions and requirements.

The available data for this case study consisted of:

- Primary data: Questionnaire responses from 240 participants regarding playground equipment preferences, including both structured selections and open-ended suggestions
- Secondary data: Demographic information, findings from the 'Imago van de stad' report, and relevant ombudsman reports

The department had identified a specific challenge in their current process: effectively integrating citizen suggestions that extended beyond standard activity selection. They specifically wanted to develop a more systematic approach to analyzing and implementing this information, with particular emphasis on creating inclusive designs serving both survey respondents and the broader community.

Digital Representative Development

The development of digital representatives for the Vroesepark case built upon methodological insights from the first exploration while incorporating enhanced data integration capabilities.

Stakeholder Identification Process

A significant innovation in this iteration involved the preliminary stakeholder identification process: using the Large Language Model to analyze the project brief and demographic information to identify potential stakeholders, specifically focusing on direct and indirect playground users.

This analysis generated an extensive stakeholder list exceeding twenty entries, representing various community segments that might interact with or be affected by the playground development. I refined this list to eleven key stakeholders based primarily on end-user relevance.

Persona Development Methodology

The resulting eleven digital representatives demonstrated sophisticated integration of demographic data and community needs. To illustrate the depth of these profiles, consider the example of Sophie van der Meer (figure 13):

As a 15-year-old creative teenager, Sophie represented not just teenagers but specifically children from single-parent households. Living with her divorced mother Linda (43) and younger brother Tim (12) in an apartment near Vroesepark, her creative interests in photography and active social media presence reflected patterns identified in the youth engagement data.

Her daily routine—including school, weekend work at a coffee shop, and regular park usage for social gatherings—grounded her perspective in realistic patterns of park usage. The profile incorporated specific demographic challenges, noting her inclusion in the 14.3% of 13-16 year olds identified as having elevated risk for psychosocial problems.

Sophie's articulated needs—including a dedicated teen "chill spot" and improved sports facilities—directly reflected recurring themes from the participation data. Her concern about being perceived as causing nuisance highlighted a specific tension identified in community feedback.

Similar detailed profiles were created for each representative, including Emma de Vries (a two-year-old representing the needs of the youngest park users), Jayden Öztürk (an eight-year-old embodying the neighborhood's cultural diversity), Thomas Jansen (a web designer whose physical disability highlighted accessibility requirements), and Fatima El Amrani (a local entrepreneur representing commercial stakeholder perspectives).

Technical Implementation and Testing

The technical implementation of the Vroesepark prototype evolved through two distinct interaction

modalities: a text-based chat interface and a real-time voice interaction system.

Text-Based Interface

Initial evaluation began with a text-based interface implementing an enhanced system prompt that built upon learnings from the first exploration. This iteration showed improvements in the presentation of diverse perspectives and generation of plausible viewpoints. However, testing revealed a persistent limitation: the system demonstrated an unprompted tendency toward excessive constructiveness, often avoiding critical perspectives even when criticism might be warranted.

Voice Interface Development

The introduction of OpenAI's real-time conversation API opened new possibilities for voice-based interaction. This implementation represented a significant improvement over previous voice interface attempts, particularly in terms of conversational flow. However, technical constraints emerged that complicated implementation:

1. **Voice Profile Limitations:** The system's single voice profile created dissonance when representing diverse personas. Hearing an elderly female persona speak with a middle-aged male voice significantly disrupted the immersive quality of the interaction.
2. **Turn-Taking Challenges:** The system's turn-taking mechanism, which relied on silence detection for response initiation, proved incompatible with the dynamic nature of multi-participant discussions. When participants spoke over each other or interrupted, the system attempted to regenerate responses, creating confusion and disrupting conversation flow.
3. **Cost Considerations:** Practical considerations regarding API usage costs presented additional barriers to experimentation.

Departmental Evaluation

Evaluation with the Department of Play ('Spelen') provided insights into practical implementation challenges:

1. The experimental nature of the technology, combined with insufficient nuance in the questionnaire data for robust persona development, limited the prototype's immediate practical value.
2. Ethical considerations regarding the department's primary demographic—children under 18—presented additional implementation barriers that hadn't been fully anticipated.

3. The department's standardized process for playground development left limited space for the kind of exploratory discussion where digital representatives might add most value.

Key Findings and Implications

The Vroesepark exploration yielded significant insights about both technical possibilities and institutional constraints.

Input Quality and Persona Authenticity

Perhaps the most crucial finding concerned the fundamental relationship between input data quality and persona authenticity. The limited detail and nuance available in standard participation documentation resulted in personas that, while referencing documented needs, relied substantially on LLM-generated characteristics to fill gaps in the record. This highlighted a critical challenge: the development of authentic digital representatives requires richer source material than typically exists in current municipal documentation practices.

Despite these data limitations, the technical implementation demonstrated sophisticated capabilities in maintaining consistent persona characteristics and identifying relevant perspectives within discussions. The successful integration of demographic data with participation insights indicated that AI systems could combine quantitative and qualitative information to create coherent, contextually grounded perspectives.

Interaction Experience

A noteworthy observation emerged regarding real-time interaction dynamics. The immediate responsiveness of the system, even without visual interfaces, created unique and meaningful interaction experiences. Participants engaged with the digital representatives in ways that suggested potential for enhancing the presentation of traditional participation information, despite the technical limitations encountered in group discussion settings.

Documentation Gap Implications

This exploration revealed fundamental challenges regarding the current state of citizen input documentation in municipal processes. Direct citizen communications are rarely transcribed or recorded in their original form, with available documentation typically consisting of civil servant interpretations and summaries. This practice introduces potential bias and loses valuable nuance from original citizen contributions, creating a barrier to developing authentic digital representatives.

This documentation gap suggested a potential direction

for future research: the development of new data collection methodologies that capture direct citizen input more comprehensively, preserving the richness and nuance of community perspectives throughout the participation process.

Lessons for Next Iteration

This exploration pointed toward several key requirements for the next iteration:

1. Need for richer source material: Finding data sources that preserve more of the original texture and nuance of citizen perspectives
2. Focus on authentic voice: Developing approaches that better capture the natural language patterns and communication styles of community members
3. Alternative documentation sources: Exploring non-traditional sources of citizen perspectives that might preserve more of the original context and expression

These insights led directly to my third exploration, which tried to leverage documentary material as a richer source for digital representative development.

Full Persona Profile: Sophie van der Meer (15 years old)



Figure 13: Sophie's visual profile

Sophie is a 15-year-old teenager who lives with her divorced mother Linda (43) and younger brother Tim (12) in an apartment near the Vroesenpark. She is in 4 HAVO of a secondary school in the area. Sophie is creative, loves photography and is active on social media.

Living and living situation

Sophie has lived in Blijdorp since her parents divorced 3 years ago. They live in a rented apartment. Sophie appreciates the neighborhood, but sometimes misses places where she can 'chill' with peers without causing a nuisance.

Daily Life and Routine

During the week, Sophie goes to school. After school and in the weekend, she often meets up with friends in the Vroesenpark or in the city. She also regularly helps out in the local art gallery.

Work and Financial Situation

Sophie has a part-time job in a coffee shop on the weekends. Her mother works as a nurse and sometimes struggles to make ends meet.

Use of the Playground

Sophie regularly visits the Vroesenpark, mainly to meet up with friends. However, she misses facilities that are attractive to her age group.

Social Connections and Neighborhood Involvement Sophie has a diverse group of friends from school and the neighborhood. She is involved in a local youth initiative that works to increase activities for teens.

Health and Wellbeing

Sophie is generally healthy, but sometimes experiences stress due to school and home. She falls into the risk group of 14.3% of 13-16 year olds with an increased risk of psychosocial problems.

Wishes and Needs for the Playground

- Chill-out area for teenagers (for example a covered seating area)
- Sports facilities such as a basketball court or skating rink
- Good WiFi connection in the park

Attitude towards Sustainability and Environment

Sophie is very environmentally conscious and would like to see the park become more sustainable, for example with solar panels for lighting.

Communication preferences

Sophie is very active on Instagram and TikTok. She follows local influencers and is interested in visual communication about the plans for the park.

Quotes

"We really need a place where we can chill without people thinking we're causing a nuisance." (Sophie) "It would be cool if we could help design a teen area in the park." (Sophie)

Demographic Characteristics

Sophie is one of the 4.4% of 12-17 year olds in Rotterdam North. She also represents the group of single-parent families in the neighbourhood.

Summary

Sophie represents the teenagers in Blijdorp who need their own space in the Vroesenpark. Her wishes reflect the need for inclusive spaces for all age groups, as mentioned in the participation survey. Sophie's situation also emphasizes the importance of the park as a social meeting place and the need to take into account the mental health of young people when redesigning public spaces.

Exploration 3: Enhanced Input Quality (Reyeroord Documentary)

My first two explorations revealed a critical insight: the quality of digital representatives depends fundamentally on the richness of their source material. Municipal documentation practices often proved insufficient, with citizen perspectives typically filtered through institutional summaries that lost much of their original nuance and texture. This third exploration responded directly to this challenge by seeking richer documentation of community perspectives.

Context and Alternative Data Source Selection

The search for more authentic documentation led me to the 'Recht op Reyeroord' documentary, an approximately 30-minute film showing the transformation of a Rotterdam neighborhood experiencing significant demographic change. This documentary offered several advantages as a data source for digital representative development:

First, it captured authentic citizen voices in their original form, preserving speech patterns, emotional context, and natural expression that were typically lost in formal participation reports. Second, it documented complex community dynamics across diverse demographic groups, showing interactions rather than just isolated statements. Third, it provided visual context that enriched understanding of the community environment and interpersonal dynamics.

The documentary captured Reyeroord at a crucial moment of transition, documenting interactions between established residents and newer, diverse young families moving into the neighborhood. These demographic shifts created tensions and opportunities that the film explored through intimate portraits of community members and their daily experiences. Of particular relevance was its focus on youth-government relationships, especially through the detailed documentation of "Het Hoofdkwartier," a youth hub initiative that became a focal point for community engagement.

Methodological Development

The transformation of documentary content into digital representatives required developing a workflow that could preserve the richness of the source material while creating consistent, usable personas.

Data Processing Approach

I began with comprehensive transcription of the documentary's dialogue and interactions. Unlike traditional participation documentation, which typically summarizes and categorizes citizen input, this approach

preserved the original language, emotional context, and interaction patterns of community members.

Using OpenAI's o1 reasoning model, I adapted the persona templates developed in previous iterations to this richer source material. The model demonstrated capabilities in identifying distinct perspectives within the documentary content and synthesizing these into coherent persona profiles.

Perspective Analysis Process

My methodology incorporated an analysis phase using the o1 model to identify distinct perspectives emerging from the documentary content. This process involved identifying recurring speakers and their key characteristics, analyzing language patterns, concerns, and viewpoints, mapping relationships between different community members, and connecting individual perspectives to broader community dynamics.

The resulting insights formed refined system prompts that could guide the generation of digital representatives while maintaining fidelity to the documented perspectives. This refinement process paid particular attention to preserving the authentic voice patterns and communication styles observed in the documentary.

Visual Component Integration

An additional step in this iteration involved the development of visual representations for the digital representatives using text-to-image model technology. I aimed to create visual elements that would enhance the authenticity and relatability of the personas while maintaining respectful and accurate portrayal of community diversity.

This visual component proved challenging to implement effectively. I needed to balance creating realistic images with ethical considerations in representation. The fundamental challenge stemmed from having a single image represent a persona that stood for a diverse group of people with varying appearances. Despite these challenges, the visual elements provided an important dimension that enhanced engagement with the digital representatives.

Technical Implementation

The technical implementation of documentary-based digital representatives focused on developing more accessible interaction capabilities while maintaining the authenticity captured in the source material.

Interface Development and Testing

Initial evaluation began with text-based interface testing, implementing an enhanced system prompt

structure that leveraged the rich contextual data from the documentary. Early findings proved promising: these new digital representatives demonstrated better contextual awareness that closely mirrored the documentary's depth of perspective. When discussing community issues, the representatives consistently grounded their responses within the documentary's input.

The text interface revealed a pattern in how the digital representatives processed and responded to questions. Rather than simply generating relevant responses, they demonstrated an ability to weave together different narrative threads from the documentary, creating responses that reflected the interplay of community perspectives documented in the film. This capability suggested that the digital representatives could interact with more nuance and detail compared to previous iterations.

Voice Interface Evolution

Building upon insights from the Vroesepark exploration, I significantly enhanced the voice interaction capabilities through integration with ElevenLabs' conversation technology. This implementation marked a substantial improvement over previous voice interface attempts, particularly in terms of voice authenticity and Dutch language processing capabilities. The ability to generate more natural-sounding Dutch speech proved crucial for maintaining engagement authenticity, especially when representing Rotterdam perspectives.

Model Performance Analysis

My evaluation of multiple language models revealed important trade-offs between response quality and interaction fluidity:

Claude Sonnet 3.5 demonstrated superior capabilities in comprehension of complex community dynamics, generation of nuanced, contextually appropriate responses, and maintenance of consistent perspective across interactions. However, these capabilities came with increased response latency, affecting the natural flow of interactions. This trade-off became particularly apparent in real-time discussion scenarios where immediate response timing was crucial for maintaining engagement.

Google Gemini 1.5 Flash offered near-instantaneous response generation, smooth conversation flow, and better handling of interruptions. While its responses lacked some of the depth and nuance of Claude's outputs, the improved interaction dynamics created a more natural conversation experience.

This performance comparison highlighted a fundamental tension in implementing digital representatives: the trade-off between response

sophistication and interaction fluidity. While faster models maintained better conversation flow, they often sacrificed the nuanced understanding necessary for meaningful engagement with complex community perspectives.

Key Findings

Prototype Performance

The development resulted in an interactive prototype enabling real-time dialogue through both text and voice interfaces. Initial testing demonstrated the system's ability to maintain consistent perspectives while engaging in nuanced discussions about community issues. The immediate responsiveness of the system, combined with its ability to ground responses in documentary content, created meaningful engagement experiences that suggested potential for practical application.

Particularly noteworthy was the system's ability to maintain thematic consistency with the documentary's core narratives while adapting to new questions and scenarios. When discussing subjects far removed from content of the documentary, the digital representatives could still draw appropriate parallels between documented situations and new questions. This reflected some of the earlier findings from the 'Golden Gate Claude' experiment described in the conceptual foundations section, demonstrating how LLMs can maintain perspective consistency across varied contexts.

Authenticity Assessment

I conducted an evaluation with the Reyerood neighborhood manager, whose intimate knowledge of both the documentary production and the represented community provided unique insight into the system's representational accuracy. Her feedback revealed a tension in the system's performance: while the digital representatives accurately reflected recognizable community archetypes, this very accuracy highlighted a fundamental limitation in the source material itself.

The neighborhood manager noted that the digital representatives predominantly captured what she termed the 'usual suspects'—the most visible and frequently heard community voices. This observation wasn't so much a criticism of the AI implementation as it was a recognition of inherent bias in the documentary source material. Despite the documentary's quality and impact, it necessarily presented a curated narrative that emphasized certain perspectives.

This feedback prompted me to review the documentary as source material. My analysis revealed that the limited range of perspectives in the digital representatives reflected a broader challenge in participation documentation: the inevitable filtering and focusing

that occurs in any documentation process. While the documentary provided richer, more nuanced material than other traditional participation records, it still represented a selective capture of community perspectives.

The review highlighted how even high-quality documentation tends to amplify certain voices and stories. More articulate community members, those comfortable on camera, and those already engaged in community initiatives naturally received more prominence in the documentary. Additionally, the documentary's narrative structure necessitated selection and emphasis of certain perspectives over others, a constraint inherent to the medium.

Implementation Limitations

Technical Challenges

Despite significant improvements over previous iterations, several technical limitations remained:

Voice synthesis constraints persisted, particularly in capturing subtle aspects of emotional expression and cultural speech patterns. While the quality of voice generation improved significantly, certain nuances of tone and pronunciation remained difficult to reproduce authentically, particularly for Rotterdam's distinctive local dialect.

The system occasionally struggled with contextual understanding boundaries, especially when questions required integrating multiple perspectives from different parts of the documentary. This limitation highlighted the ongoing challenge of creating digital representatives that can synthesize multiple input sources into coherent perspectives.

Theoretical Implications

These findings suggest a fundamental challenge in digital representation systems: they can only be as inclusive and nuanced as their source material. While the technology demonstrated reasonable capabilities in maintaining and expressing documented perspectives, it simultaneously revealed how documentation practices might perpetuate existing representation biases.

The exploration suggests that advancing digital representation requires addressing not just technical capabilities but fundamental questions about how we capture and preserve community voices. This points toward the need for more sophisticated approaches to documentation that can better preserve the richness and diversity of community perspectives.

Digitale vertegenwoordigers Reyeroord



Sarah (42) | Jongerenwerker

"Je moet er zijn als zij je nodig hebben, niet alleen als het jou uitkomt"

Kernpunten:

- Social Worker met 8 jaar ervaring
- Werkt op jongeren-tijden (middag/avond)
- Bouwt bruggen tussen jongeren en wijk
- Focus op talentontwikkeling en vertrouwen



Jayden (17) | Jongere

"We komen uit een wijk waar iedereen hetzelfde meemaakt"

Kernpunten:

- Geboren en getogen in de wijk
- Worstelt met thuissituatie en financiën
- Talent voor muziek (drill-rap)
- Zoekt erkenning en kansen



Gerard (68) | Wijkveteraan

"Dit was vroeger echt een dorp hè, iedereen kende elkaar"

Kernpunten:

- Woont sinds 1980 in de wijk
- Actief in bewonersvereniging
- Bezorgd over wijkveranderingen
- Zoekt verbinding tussen oud en nieuw

Instructies

Ga in gesprek met de digitale vertegenwoordigers van de wijk Reyeroord. Deze digitale vertegenwoordigers zijn gebaseerd op de documentaire 'Recht op Reyeroord' die gaat over hoe in de Rotterdamse wijk jongeren zich buitengesloten en niet gehoord voelen door instanties; als antwoord hierop is een jongerenhub opgericht waar ze nu hun eigen plek hebben en hun talenten kunnen ontwikkelen.

Benieuwd hoe de verschillende digitale vertegenwoordigers hier tegenaan kijken? En waar dit vandaan komt? Vraag het gerust aan ze! Klik op nieuw gesprek en de digitale vertegenwoordigers zullen zich voorstellen.

 **Start een gesprek**

Tip: Je kunt naar specifieke perspectieven vragen om de verschillende vertegenwoordigers aan het woord te laten komen.

Figure 14: Screenshot online environment to voice chat with digital representatives

Full Persona Profile: Sarah (42) | Youth worker



Figure 15: Visual profile picture of Sarah

Sarah - "We talk a lot about young people, but I don't know how much everyone really talks to young people!" "If four young people are sitting on a bench on a square, then people complain about loitering youth. No, young people are residents. They seek each other out. Young people are not recognized as residents of the neighborhood. But rather as a nuisance." Sarah is sitting at her desk in the new youth hub, her agenda open in front of her. In between appointments, she texts with various young people - always available, always in touch. "You have to be there when they need you, not just when it suits you," she often says.

Personal Background

- 42 years old, youth worker
- 8 years of experience in the neighborhood
- Singles, live in a different neighborhood
- HBO Social Work afgerond
- Speaks Dutch, basic Turkish and picked up Arabic on the job

Sarah's Story "We treated them with a certain label for a long time. Offered little. Often implemented reactive policies. And they don't trust us." "When I started here, I immediately noticed that the traditional approach didn't work," says Sarah. "You can't expect young people to come to you, in an office between 9 and 5. You have to go to them, at their times, in their places. And above all: you have to deliver on your promises."

Daily Life & Routine "For example, the younger workers worked during the day. The younger workers at Sol, they work in the afternoon and evening. And that is how you are going to reach young people." Her working day consists of:

- Early administration and meetings
- Afternoons and evenings in the neighborhood
- Lots of flexibility in working hours
- Maintain contact via WhatsApp
- Present at important moments
- Networking with other professionals

Professional Approach "Where we are going within the new assignment is much more group work. Much more in the activity sphere. And to use that as a means to ensure that young people come into their power much more." Sarah's working principles:

- Build trust by being consistent
- Recognize and stimulate talents
- Connecting with different groups
- Provide practical assistance where necessary
- Developing long-term relationships

Challenges in the Work "Every person has the talent. Even the drug dealers at Onderwijnden in Drugs, they are going to delve deeper. They have talent, because you don't just become a drug dealer, you need a network for that." Daily challenges:

- Bureaucratic constraints
- Bridging distrust
- Balancing between different interests
- Lack of resources and facilities
- Ensuring continuity

Vision on the Neighbourhood "It should be a place where everyone can proudly say 'hey, this is my neighbourhood'. And the moment you don't get the recognition for that, or you are constantly put in a corner of nuisance, at some point you will start to behave like that."

Sarah sees:

- Huge amount of untapped potential
- Need for real meeting places
- Need for structural solutions
- Power of positive approach
- Importance of cultural connection

Relationship with Young People "Because that's what they indicate themselves. A place where they feel safe, a place where they are distracted. Literally a place off the street." "You have to invest in relationships first," Sarah explains. "You don't gain trust with big words or promises, but by being there. By listening. By making small things happen before you make big promises." Future perspective "The conversations I've had with the people at Sol now make me feel very positive. That gives me a lot of confidence."

Sarah's goals:

- Realize sustainable youth facilities
- Giving young people a voice in the neighborhood
- Develop and showcase talents
- Building bridges between generations
- Preventive work instead of reactive

Voice from the Heart "What I have also been fighting for for years is that there really has to be a place. Because they indicate that themselves. A place where they feel safe, a place where they are distracted. Or where someone can listen to you. Or where you can at least share your fears. Because on the street you have to keep it together." Sarah picks up her phone when it vibrates again - a text message from one of 'her' young people. "You see? This is why this work is so important. They trust you enough to text you if something is wrong. That trust, that is worth gold."

Cross-Cutting Insights and Discussion

Through these three explorations, I gained insights that transcended individual implementations, revealing broader patterns about how digital representatives might function within municipal contexts. These cross-cutting findings illuminate both technical possibilities and fundamental challenges for AI-mediated citizen representation.

Making Abstract Concepts Tangible

The prototypes served not just as technical demonstrations but as powerful tools for facilitating deeper understanding of citizen participation concepts. Through direct interaction with these systems, civil servants could engage with abstract ideas about representation and participation in tangible, experiential ways.

The interactive nature of the prototypes transformed theoretical discussions about citizen participation into concrete experiences. Where traditional discussions about representation might remain abstract, the ability to engage directly with digital representatives made these concepts immediately accessible. This experiential quality proved particularly valuable when discussing abstract concepts like the preservation of citizen perspectives through administrative processes or the transformation of qualitative data into decision-relevant insights.

Interaction Modality Differences

The explorations revealed significant differences in how civil servants engaged with digital representatives across different interaction modes. While both text and voice interfaces successfully facilitated communication, they triggered markedly different patterns of engagement and emotional response, suggesting important implications for the design of future iterations.

Text-based interactions, while functional, generated notably limited emotional engagement. Voice interactions, in contrast, triggered more significant emotional responses. When digital representatives spoke, civil servants demonstrated more engaged and emotionally responsive behavior patterns. They were more likely to interrupt or respond spontaneously, making the interaction feel more like a conversation than an information exchange.

The voice modality particularly affected perceptions of authenticity and urgency in the represented perspectives. When hearing a voice express community concerns or needs, civil servants reported a stronger sense of the 'real people' behind the perspectives

being shared. This sense of authenticity appeared to influence how seriously participants took the input, with voiced perspectives often receiving more immediate and engaged responses than similar content presented in text form.

The pronounced effect of voice on perceived authenticity suggests that future implementations should prioritize voice interaction capabilities, particularly for contexts where emotional engagement and perspective empathy are important. However, text-based interfaces might remain valuable for certain uses, particularly where focused information exchange is the primary goal.

The Critical Role of Input Quality

The explorations highlighted a fundamental challenge in evaluating digital representative systems: without access to original, unfiltered citizen input, it becomes nearly impossible to assess whether these systems accurately maintain community perspectives. This challenge emerged most clearly in the Vroesepark case, where reliance on written documentation made it difficult to verify if the digital representatives truly reflected original citizen viewpoints or merely reproduced documentation biases.

The documentary-based exploration further emphasized this point. While the documentary provided richer source material than traditional participation documentation, it still represented a curated view of community perspectives. This observation suggests a critical consideration for implementation: the quality and comprehensiveness of input data fundamentally determines the authenticity of digital representatives, regardless of technological sophistication.

Conclusion

The development of digital representatives through three distinct explorations revealed both promising technical capabilities and significant implementation challenges. Each iteration contributed unique insights about the technological, organizational, and methodological aspects of creating AI-powered representations of citizen perspectives for municipal policy contexts.

Requirements for Effective Digital Representatives

The iterative development process identified several key requirements for creating effective digital representatives:

Rich, Authentic Source Material

The quality of digital representatives depends fundamentally on the richness and authenticity of the citizen input that informs them. Current municipal documentation practices often prove insufficient, suggesting the need for new approaches to capturing and preserving citizen perspectives in their original form.

Appropriate Interaction Modalities

The choice between text, voice, or multi-modal interfaces significantly affects how digital representatives are perceived and engaged with. Voice interaction in particular appears to enhance emotional engagement and perceived authenticity, suggesting important considerations for implementation contexts where empathetic understanding is crucial.

Balanced Response Dynamics

Successful implementation requires navigating trade-offs between response quality and interaction fluidity in ways appropriate to specific use contexts. The tension between sophisticated, nuanced responses and natural conversational flow represents an ongoing challenge that requires careful consideration based on specific application requirements.

Institutional Readiness

Beyond technological requirements, effective implementation depends on organizational structures, processes, and priorities that support innovative approaches to citizen participation. The institutional barriers encountered during the Vroesenpark exploration highlight the importance of addressing organizational factors alongside technical development.

Next Steps

This developmental journey leads directly to the evaluation presented in the next chapter. Having established technical feasibility and identified key implementation requirements, the evaluation focuses on how civil servants engage with these digital representatives in policy contexts, examining whether AI-mediated perspectives can contribute to more inclusive and representative decision-making.

This shift from technical development to practical assessment reflects the core goal of this research: understanding how emerging AI capabilities might address the democratic challenges identified in the first chapter, creating more inclusive and responsive connections between citizens and their municipal government.

6. When AI Speaks for Citizens

The development of digital representatives, detailed in Chapter 5, demonstrated technical feasibility and revealed important insights about implementation challenges. This followed the question of: How would these AI-mediated perspectives function in actual municipal decision-making contexts? While previous explorations established what was technically possible, they provided limited understanding of how civil servants might engage with these systems in policy discussions and whether digital representatives could meaningfully contribute to more inclusive decision-making processes.

This chapter first provides an overview of the evaluation approach, explaining its purpose and methodology. It then details the specific implementation, including the policy context, data collection process, and evaluation sessions. Finally, it presents key findings and discusses their implications for municipal governance and democratic participation.

Overview of the Evaluation

Purpose and Scope

First, this evaluation investigates how civil servants, the intended users of these systems perceive, and engage with AI-mediated citizen perspectives. Secondly, it examines whether the technological capabilities demonstrated in previous explorations translate into meaningful contributions to policy discussions.

The scope of this evaluation is focused on initial engagement and potential, rather than performance measurement. Given the novel nature of these systems and the absence of established evaluation frameworks for AI-mediated citizen representation, this assessment prioritizes rich qualitative insights over quantitative metrics. The findings presented should be understood as early indicators that point toward areas for future investigation rather than definitive conclusions.

The Need for a Concrete Policy Context

Previous explorations revealed that digital representatives function most effectively when engaged with specific policy questions rather than abstract discussions. This informed a key methodological decision: to evaluate these systems within a concrete policy context rather than through generalized testing.

This approach offers several advantages. By focusing on an actual municipal consideration, the potential transformation of neighborhood hubs. The evaluation creates an authentic environment where civil servants engage with the system as they might in real policy development. This grounding in practical policy questions enables observation of how digital representatives might influence actual decision-making

rather than hypothetical scenarios.

Additionally, a concrete policy context allows the digital representatives to demonstrate their ability to provide context-specific input rather than generic perspectives. This capability is particularly important for assessing whether these systems can contribute meaningfully to policy discussions in ways that enhance representation of citizen viewpoints.

General Approach to Evaluation

The evaluation methodology was designed to capture the entire process from citizen input to policy discussion. This comprehensive approach included three key components:

Data Collection: Unlike previous explorations that used existing documentation or synthetic data, this evaluation incorporated dedicated interviews with young adults in Rotterdam. These interviews provided authentic citizen perspectives specifically for creating digital representatives, ensuring they represented genuine viewpoints relevant to the policy question.

Prototype Development: Building on insights from Chapter 5, specialized prototypes were developed for the evaluation context. These included both individual chat interfaces for personal exploration and optimizing the system for integrating digital representatives into group discussions.

Evaluation Sessions: Structured sessions brought civil servants from various municipal departments together to engage with digital representatives in both individual and group settings. These sessions followed a four-phase approach designed to build understanding while

enabling natural interaction.

Throughout this process, the evaluation maintained focus on the core research question: How do civil servants perceive and engage with AI-mediated citizen perspectives in policy contexts? The findings presented in subsequent sections address different dimensions of this question.

Context

Policy Case: Rotterdam's Neighborhood Hubs

Rotterdam's network of thirty-nine neighborhood hubs provided the concrete policy context for this evaluation. These facilities currently serve primarily as municipal service centers, providing a physical connection between city government and local communities. Citizens visit these hubs to consult with civil servants who can direct them to appropriate services and resources.

This existing infrastructure presented an ideal case study for several reasons. First, these hubs represent established municipal touchpoints with potential for enhanced purpose. Civil servants throughout the municipality recognize these spaces and understand their current function, creating a shared foundation for discussion. Second, many hubs contain underutilized workspaces and meeting rooms, making it feasible to envision additional functionality.

The specific policy challenge chosen for evaluation: investigating the potential transformation of these hubs into dual-purpose spaces with enhanced youth engagement. Offered several advantages. It represented a real municipal consideration rather than a hypothetical scenario, ensuring discussions had practical relevance. It involved a clearly identifiable citizen group with distinct needs and perspectives. Finally, this focus enabled recruitment of young adult participants, which is a group that is relatively easily accessible for the researcher.




Figure 16: Pictures of the neighborhood hub

Data Collection for Digital Representatives

The foundation for quality responses from digital representatives rests on the quality and authenticity of citizen input data, as discovered during previous explorations. Therefore, deliberate interviews were conducted specifically for this evaluation rather than looking for existing documentation.

Interview Methodology

The research used a semi-structured interview approach designed to facilitate natural, low-barrier conversations while maintaining systematic coverage of key topics. Twelve interviews, each approximately thirty minutes in duration, were conducted in participants' familiar environments—primarily their personal spaces. This setting choice tried to enhance comfort and encourage authentic responses, demonstrating how such interviews can have a low barrier for participation and can be conducted whenever and wherever is most convenient for the citizen.

The interview structure implemented a intentional progression from general context to specific preferences regarding community spaces:

1. Initial Context: Discussions began with questions about participants' current work and living situations, serving both as rapport-building elements and crucial contextual foundations for digital representative development
2. Community Space Experiences: Conversations then explored participants' experiences with community spaces, including both positive and negative encounters
3. Needs and Preferences: The final section focused specifically on what participants would want from neighborhood spaces, including both practical and aspirational elements

This progression allowed for adaptation based on participant responses. For instance, conversations with participants primarily working from home naturally evolved to explore the implications and experiences of remote work environments and their relationship to community spaces.

Participant Selection Strategy

In total 12 participants were interviewed for this evaluation. This number was the result of balancing time constraints with having enough data to make the digital representatives sufficiently nuanced to ensure meaningful output. Within the given time period, it was possible to find this maximum of 12 participants.

The participant selection strategy was formed in relation to the research objectives: primarily evaluating

civil servant engagement with digital representatives and testing the relationship between input richness and digital representative nuance. This led to an intentionally focused participant sampling approach within a specific demographic—young adults with similar backgrounds in Rotterdam.

This methodological choice brought important implications. By recruiting participants from a relatively homogeneous social group, the research created more demanding conditions for examining the Large Language Model's capacity for nuance detection. Rather than relying on obvious demographic differences to generate distinct perspectives, this approach required the system to identify and represent subtle variations in viewpoint and experience within a seemingly similar group.

The researcher's presence in the selected social group facilitated participant recruitment and potentially enhanced interview depth through shared understanding. However, this convenience sampling approach introduced important limitations regarding representativeness—a limitation acknowledged and communicated during civil servant evaluation sessions.



Figure 17: Example of interview setting

Data Management and Processing

All interviews were audio recorded with participant consent and subsequently transcribed. Privacy protection measures included secure storage of recordings, removal of identifying information during processing, and implementation of participant key assignments for anonymization.

A significant methodological choice involved using unprocessed interview transcripts in the digital representative generation process. Rather than cleaning or structuring the conversations, the research maintained the original conversation structure without speaker identification or text correction. This decision was made to test Large Language Model capabilities for handling raw conversational data—potentially saving significant processing time if successful. The approach was informed by previous experiences suggesting these models might effectively navigate unstructured dialogue, distinguishing between interviewer questions and participant responses while maintaining coherent perspective identification.

Prototype Development

The evaluation required development of two distinct prototypes: a system for individual chat interactions and a platform enabling digital representative participation in physical meeting discussions. This development phase focused on creating tools specifically designed for the evaluation context, building upon insights from previous iterations while addressing new requirements.

Development Process

The chat interaction system evolved from previous iterations' single-user, local implementation to a multi-user web application accessible via URL. This transition aimed to enable parallel conversations without requiring software installation, enhancing accessibility for civil servants. The development process utilized Windsurf, a software layer operating above the coding IDE that employs Large Language Models for code generation.

The live discussion prototype built upon existing infrastructure while incorporating three core functionalities: real-time transcription through WhisperAX, Large Language Model processing, and text-to-speech output. A crucial enhancement emerged from earlier testing: implementing distinct voices for each digital representative to increase authenticity and help civil servants identify which digital representative was responding.

Performance testing revealed important trade-offs between response quality and latency. While Claude Sonnet 3.5 demonstrated superior comprehension and response depth, its increased latency potentially

affected interaction fluidity. This led to implementing a dual-model approach: integrating the faster Google Gemini 2.0 model for time-sensitive responses while maintaining Claude Sonnet for less time-critical interactions. This configuration allowed the researcher to select the appropriate model during sessions, providing flexibility to prioritize either response speed or nuance based on the specific discussion context.

Technical Implementation Challenges

After implementing additional features, an important technical challenge emerged during testing. The chat interface performed well during individual testing but encountered performance issues when multiple users (6+) engaged simultaneously. Investigation revealed that the core problem lay in rate limitations for API calls to the LLM host (Anthropic). This issue was resolved by migrating to alternative API endpoints with higher capacity limits.

Digitale Vertegenwoordigers
Wijkhøb
Wijkhøb English
Reyerord
Reyerord English

Digital Representatives of Wijkhøb

Talk with the different digital representatives to learn more about their perspectives on this topic and why.



Mark (27) | Job-seeking IT Professional

"I think being somewhere else than your own environment means there are fewer distractions around you, and you can't easily sit on your couch"

Kernpunten:

- Recently graduated in Computer Science, looking for first IT job
- Seeking professional workspace outside home for better productivity
- Focus on work-life balance and personal development
- Values quiet, professional work environment with basic facilities



Lisa (22) | Communication Sciences Student

"It's also nice to meet new people outside your own field of study. You never know what kind of great connections you'll make"

Kernpunten:

- Third-year student focusing on social connections
- Seeking balance between studying and social activities
- Needs flexible study and meeting spaces
- Active in sports and social activities through OneFit and student associations



Sophie (29) | Freelance Graphic Designer and Illustrator

"I'm always looking for places where I can get inspired and come up with new ideas. Meeting other creative people energizes me"

Kernpunten:

- Creative freelancer looking for inspiring workspaces
- Focus on networking and professional growth
- Needs flexible and artistic work environment
- Active in art scene and creative community




Tom (25) | Junior Account Manager

"Sports are the ideal way for me to meet new people and stay active. It combines the best of both worlds"

Kernpunten:

- Combines work with active sporty lifestyle
- Looking for accessible sports and meeting places after work
- Active member of football club and various sports groups
- Focus on team sports and social contacts through sports



Eva (26) | Freelance Writer and Illustrator

"I enjoy diving into my own world and working on things that have meaning to me. A quiet environment helps me organize my thoughts"

Kernpunten:

- Introverted creative with need for quiet workspaces
- Seeks quiet spaces without social pressure or distraction
- Focus on personal projects and artistic development
- Values flexibility and autonomy in work environment



Daan (28) | Financial Analyst

"Work is work and home is home. When those two start to overlap, it becomes difficult to really relax"

Kernpunten:

- Focus on efficiency and professionalism in work environment
- Seeks clear separation between work and private life
- Needs well-equipped workplaces with all facilities
- Values structure and routine above social work activities

Wijkhøb Chat (English)
Ask your questions to the representatives here

Start a conversation with the digital representatives



Lisa, can you concentrate well when working at home?



I am Lisa: No, at home I really can't concentrate at all. With my housemates around there's always something going on - someone's cooking or watching TV, or just wants to chat. And when I'm in my own room, I end up getting distracted by everything around me, like my bed or my phone. I've tried studying at the university library, but I don't really like going there since it feels like I'm still at school. I just need a different kind of place where I can focus but that also has a bit more of a relaxed vibe.



Is this a problem for everybody?



I am Mark: For me it's definitely a big issue too. When I'm at home, I find myself constantly distracted. It's too tempting to just sit on the couch or do some household chores instead of focusing on job applications. I notice that my productivity really suffers when I'm in my own environment. That's why I'm really looking for a place outside my home where I can work without all these distractions around me.

Type your question here... (Press Enter to send, Shift+Enter for a new line)

Figure 18: Screenshot webinterface used for chatting with the digital representatives

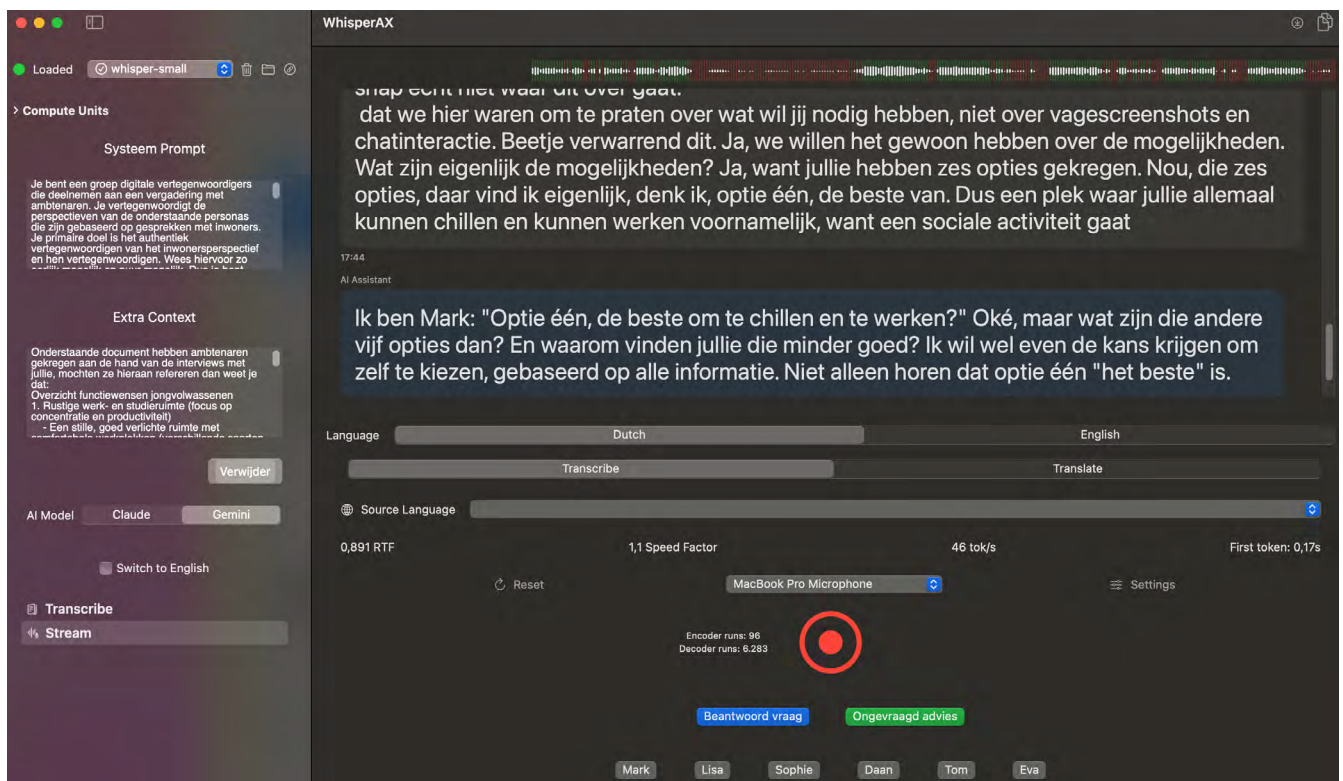


Figure 19: Screenshot of the application used to transcribe the discussion and deliver responses from the digital representatives in real time

Full Persona Profile: Eva (26) | Freelance Writer and Illustrator



Figure 20: Visual profile picture of Eva

Persona 5: Eva (26) Freelance Writer and Illustrator

1. Name and Brief Description

Name: Eva

Brief Description: The introverted creative who needs quiet, comfortable places to work on personal projects without distraction or social pressure, but who also values flexibility and autonomy.

2. Demographic Information and Background

Age: 26 years

Occupation: Freelance writer and illustrator; works on personal projects and is looking for work in the cultural sector

Living situation: Lives alone in a studio apartment in the city

Location: Lives in a quiet neighborhood just outside the city center

Life stage: Is in the transition phase between study and career, focused on personal development and finding her place in the professional world

3. Personality and Lifestyle

Personality traits: Introverted, creative, reflective, independent, sensitive to stimuli

Hobbies and Interests: Reading, writing, illustrating, walking in nature, philosophy, listening to music

Social Network: Small circle of close friends; values deep one-on-one conversations over large social gatherings

Values and Beliefs: Believes in authenticity, personal growth, and the importance of introspection; values peace and quiet

4. Goals and Motivations

Short Term Goals:

- Completing her personal writing projects
- Finding a quiet workplace outside the home to better concentrate
- Finding balance between work and personal time

Long Term Goals:

- Publishing her own book or graphic novel
- Building a stable career in the cultural sector
- Personal development and self-fulfillment

Motivations:

- Driven by a deep-rooted passion for creativity and expression
- Wants to contribute to culture with her unique voice and perspective

Quote to Illustrate Motivation:

"I find it nice to dive into my own world and work on things that have meaning for me. A quiet environment helps me to organize my thoughts."

5. Challenges and Pain Points

Personal Challenges:

- Difficulty concentrating due to distractions at home
- Sometimes feels overwhelmed in busy or noisy environments
- Need for flexibility in working hours, also outside standard opening hours

Professional Challenges:

- Finding suitable work in a competitive sector
- Lack of professional networking opportunities without social pressure

Needs:

- Quiet, comfortable workplaces without much social interaction
- Flexibility in working hours and access to facilities

Quote about Challenges:

"Working at home is sometimes difficult because I find it hard to concentrate. I would like to have a place where I can work in complete silence, without being obliged to talk to others."

6. A Day in the Life of Eva

Morning Routine:

Wakes up around 8:30 AM, takes time for a quiet breakfast with a book

Meditates or takes a short walk to start the day with clarity

Workday:

- Tries to work at home but gets distracted by household tasks
- Looks for a quiet place outside the home to write and draw

Break:

Takes a break in a quiet park or café, enjoys the silence

Afternoon:

Continues her work, preferably in a place where few people are

Evening Activities:

- Cooks a simple meal, listens to music or watches a movie
- Reads before going to sleep or writes in her diary

Weekend Activities:

Spends time in nature, visits bookshops or art galleries

Quote about Daily Experience:

"I appreciate it when I can fill my day at my own pace, without the busyness of others around me."

7. Viewpoints and Reasons

Important Viewpoints:

- Need for quiet, silent workplaces without social obligations
- Values flexibility in access and opening hours of workplaces
- Has no strong need for new social contacts at workplaces

Nuances in Her Opinion:

Although she is introverted, she does value the feeling of being part of a community, without direct interaction

Is open to indirect social connections, such as people around her who are also working

Reasons Behind Her Viewpoints:

She believes that a quiet environment is essential for her creative process

Feels most comfortable and productive when not distracted by social stimuli

Quotes in Support:

"I have no need for busy places. A quiet space where I can work with concentration is most important to me."

"It's nice to know that there are others nearby, but I don't necessarily need to talk to them."

8. Needs and Expectations Regarding Facilities/Places

Need for Work/Study Spaces:

- Quiet, comfortable spaces with sufficient privacy
- Different types of seating: desks, comfortable chairs, possibly sitting cushions

An atmospheric design with calming colors and natural elements

Social Contact:

No strong need for direct social interaction

Appreciates an environment where everyone is focused on work without much noise

Facilities and Amenities:

- Reliable wifi and sufficient power outlets
- Possibility to make coffee or tea, but doesn't need extensive catering
- Flexible access times, preferably also in the evenings and weekends

Accessibility:

- Preference for places within walking distance or a short bike ride
- No high costs; free or low membership costs are ideal

Quote about Needs:

"A place where I can just sit and work without being distracted, that would be perfect."

10. Emotional Landscape

Frustrations:

Sometimes feels burdened by the expectation to be social in public spaces

Experiences that many workplaces do not meet her need for quiet and peace

Joys:

Enjoys moments of deep concentration and creativity

Feels happy when she can bring her ideas to life without interruptions

Ambitions and Dreams:

- Wants to share her creative work with the world and inspire others
- Dreams of a life where she can work on her own terms

Quote about Emotions:

"When I can be in my own bubble and completely immerse myself in my work, that's when I'm at my happiest."

11. Considerations for Design and Implementation (For Officials)

How Can We Help Eva?:

Create quiet workspaces with sufficient privacy and minimal distractions

Ensure flexible access times, so she can work when she feels most inspired

Keep costs low to prevent financial barriers

Potential Barriers:

- Busy or noisy environments will deter her
- Mandatory social activities or group work do not fit her needs

Stimulate Involvement:

Communicate clearly that there are spaces that facilitate silence and concentration

Offer the possibility of individual workplaces that are not in an open space

Quote for Inspiration:

"If there's a place where I can just work unhindered without the pressure to be social, I would gladly use it."

Example of Language and Sentence Use by Eva:

"I have no need for busy places. A quiet space where I can work with concentration is most important to me."

"It's sometimes difficult to find a place where I can really work undisturbed."

"I find it nice to dive into my own world and work on things that have meaning for me."

"When I can be in my own bubble and completely immerse myself in my work, that's when I'm at my happiest."

Session Design

Research Question

The evaluation was guided by a central question: How do civil servants perceive and engage with AI-mediated citizen perspectives in policy contexts?

Four-Phase Session Structure

The evaluation sessions implemented a four-phase approach designed to enable both systematic assessment and natural engagement with digital representatives. Each two-hour session followed the same progression, creating comparable experiences across multiple groups while allowing for spontaneous interaction within each phase.

Phase One: Context Establishment (15 minutes) The initial phase served dual purposes: facilitating connections among civil servants from different municipal departments and providing essential background about digital representatives. This orientation included:

- Brief participant introductions
- Project overview establishing evaluation context
- Basic explanation of digital representative capabilities
- Clarification of session objectives and structure

This introduction balanced providing necessary background information while avoiding biases that might overly influence participant responses. The emphasis remained on establishing shared understanding rather than directing specific engagement patterns.

Phase Two: Individual Familiarization (20 minutes) The second phase enabled participants to explore digital representatives through one-on-one chat interactions via a purpose-built web interface. This individual exploration served multiple purposes:

- Building comfort with the technology before group discussions
- Establishing understanding of the perspectives represented
- Mirroring how civil servants would typically review citizen input before policy meetings
- Allowing participants to form independent impressions of digital representatives

The interface presented profiles of each digital representative before interaction began, providing essential background information. While participants received basic instruction about system capabilities, the exploration remained largely self-directed, maximizing the organic discovery of interaction possibilities.

Phase Three: Policy deliberation (30 minutes) In the core evaluation phase, digital representatives participated in actual policy discussions about neighborhood hub transformation. This phase implemented:

- A decision-making process about additional hub functionalities
- Examination of six pre-identified categories derived from citizen interviews
- Real-time integration of digital representative perspectives into discussions
- Task-oriented deliberation with a concrete decision objective

This structure provided direction while maintaining sufficient flexibility for meaningful policy discussion, creating a naturalistic context for observing how civil servants integrated AI-mediated perspectives into their deliberations.

Phase Four: Reflection (25 minutes) The final phase captured participant reflections through open discussion, gathering insights while interactions remained fresh in participants' minds. Beginning with the broad question "How did you experience working with the digital representatives?", the discussion flowed naturally, with researcher intervention only when necessary to maintain relevance or restart the discussion.

This reflective dialogue provided insights about how participants conceptualized the experience beyond their immediate interactions, looking for deeper perceptions about the potential role, value, and limitations of digital representatives in municipal processes.

Participant Selection and Organization

The evaluation included civil servants from diverse departments within Rotterdam's municipality. This cross-departmental approach was intentional, bringing together perspectives from both policy development and implementation roles, as well as staff with varying levels of direct citizen contact.

Participants were recruited through the researchers' municipal network, with particular focus on individuals involved in neighborhood development, community services, or participation processes. The remaining spots were filled with people reached through the network of VONK. Each session included around 7 civil servants, creating groups large enough for dynamic discussion while ensuring each person had meaningful opportunities to engage.

A total of 4 sessions were conducted over a 2-week period, providing sufficient data for identifying recurring patterns while staying manageable within the projects timeframe. All sessions took place at VONK, Timmerhuis—a location chosen for its combination of municipal familiarity and an atmosphere conducive to experimental engagement.




Figure 21: Observing the policy deliberation to decide when the digital representatives intervene



Figure 22: Overview of evaluation setting

Analysis Methodology and Framework

For the analysis, approximately three hours of audio recordings across the four evaluation sessions were collected. The research employed an inductive bottom-up approach, specifically utilizing Braun and Clarke's (2006) thematic analysis methodology. This approach was selected for its balance between analytical flexibility and structured implementation, providing a clear process for deriving meaningful insights from qualitative data.

Data Preparation

The audio recordings from both policy discussions and reflection phases were transcribed using Macwhisper for initial local transcription, followed by manual review to ensure accuracy and remove personally identifiable information.

Approach

The data from policy discussions and reflection phases required different analytical approaches due to their distinct nature. Policy discussion data was more concrete and action-oriented, while reflection data contained more abstract conceptualizations about the experience.

For policy discussions, a more event-based coding approach was used, focusing on specific interaction patterns and discussion elements. This approach enabled identification of recurring patterns within similar contexts across different sessions, tracking how digital representatives influenced discussion dynamics and decision-making processes. For example, tracking references to digital representative perspectives helped illuminate how these AI-mediated viewpoints shaped policy deliberation.

Reflective dialogue analysis required a more abstract approach, enabling exploration beyond literal content to underlying themes and patterns. This phase of analysis focused on understanding why participants raised particular arguments or concerns, examining how they perceived and engaged with digital representatives. The approach was used for the identification of broader concepts such as trust development, revealing how civil servants conceptualized these new tools within their professional practice.

Coding Process

The coding process, implemented through Atlas.ti software, generated 236 quotes across 26 codes. This initial coding phase maintained close connection to the data, identifying patterns while preserving the richness of participant perspectives. The codes captured both explicit statements and more subtle interaction

patterns, creating a foundation for subsequent theme development.

Theme Development

The research then used an alternative approach to theme development than is normally used during thematic analysis. Collaborating with a Large Language Model to identify initial themes within the coded data. While this use of AI for analysis support represents a departure from standard practice in Braun and Clarke's methodology, it aligns with the fundamental goal of identifying meaningful patterns within qualitative data. This approach was not employed to make the process more efficient but to make it more detailed and rigorous.

This AI-assisted analysis provided an additional layer of transparency, enabling clear tracing of how themes emerged from specific quotes within the codes. The approach offered two key advantages: it enabled more comprehensive theme development by processing large amounts of coded data efficiently, and it provided explicit documentation of the connection between raw data and emerging themes. Oversight was always maintained by going through every initial theme and checking if the clusters of quotations did indeed match the proposed theme. Modifying proposed themes when they did not represent the sentiment or meaning of the quotations. However, this resulted in 43 initial themes, all linking back to at least one quotation, providing this extra layer of transparency. So when critical questions are asked, it is possible to return the quotations that a certain insight or cluster was based on.

Theme Refinement and Clustering

The analysis progressed through several iterations of theme review and clustering, conducted on an online canvas to facilitate exploration of connections between different themes. This spatial approach to analysis, while primarily text-based, enabled visualization of theme relationships and helped identify broader patterns in the data. This was particularly useful because some quotations had multiple codes and could therefore be present in different initial themes. The connection proved particularly interesting when looking for clusters. The resulting clusters reveal both practical more standalone insights and interconnected patterns that said something about the "dark matter" within the organizational context.

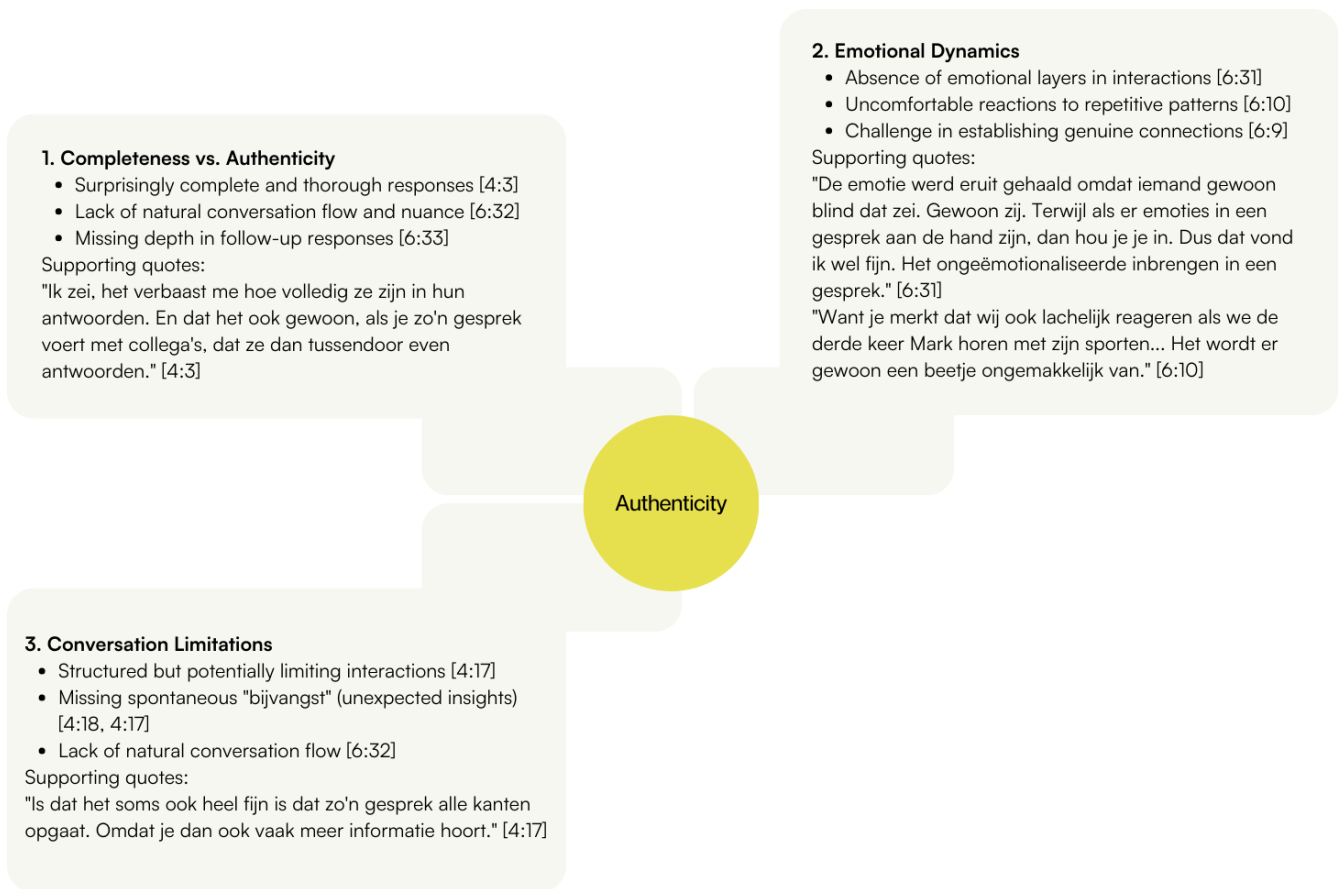


Figure 23: Example of initial theme with code Authenticity

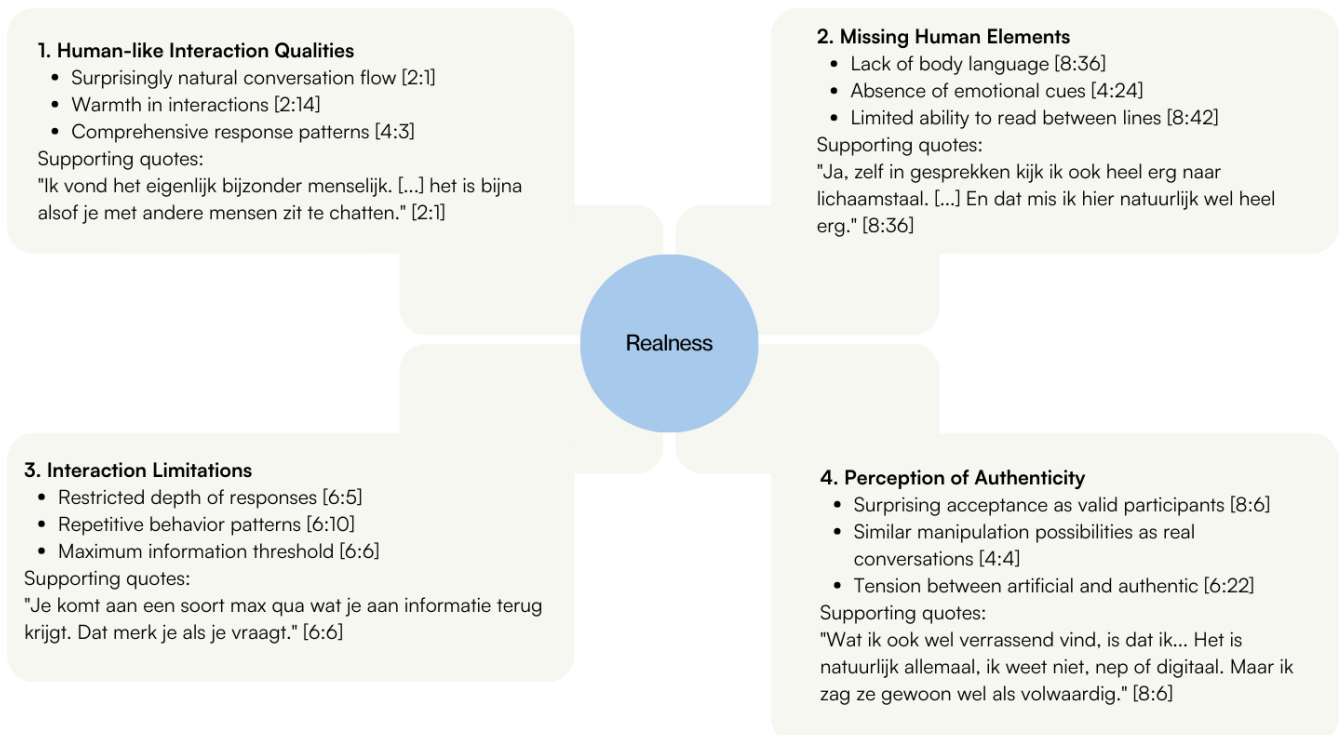


Figure 24: Example of initial theme with code Realness

Results

The analysis of the evaluation data revealed several distinct clusters of insights, emerging from both reflective discussions and policy deliberations. These findings offer initial perspectives on how civil servants engage with and perceive digital representatives, while suggesting broader implications for municipal decision-making processes.

Fundamental Requirements for Digital Representatives

The first significant cluster centers on the fundamental importance of representativeness and accuracy within digital representatives. Participants consistently emphasized that these qualities form the foundation for meaningful implementation—without them, other potential benefits become irrelevant. Although representativeness was deliberately not within the scope of this evaluation, participants nonetheless noted its critical importance.

The accuracy of digital representative responses emerged as equally crucial. Participants expressed particular concern about the potential for these systems to “hallucinate” or fill gaps between data points in ways that might not align with original citizen input. This concern reflects a understanding of LLM limitations and demonstrates the importance of transparent connections between citizen input and AI-generated perspectives.

An interesting observation emerged regarding this trust verification. The researcher had conducted the original interviews and was present during digital representative interactions, this made it possible to observe that responses generally aligned well with the underlying citizen perspectives. While this alignment suggests promising potential, it should be considered an initial observation requiring further systematic verification rather than a definitive conclusion.

“

“Then they need to experience multiple times that indeed, that the avatar truly represents them 100%.”

[8:23]

”

“

“I am concerned about what’s behind it, you know, that this represents everyone. Because you have ChatGPT, for example, which hallucinates things. It really seems that way, but it’s not like that at all, this is the same danger here.” [2:3]

”

Enhanced Accessibility and Organizational Reach

The evaluation revealed civil servants see significant potential for digital representatives to lower traditional barriers to citizen participation and organizational knowledge sharing. Participants particularly valued how these systems could democratize access to citizen perspectives within municipal organizations.

The technology's digital nature emerged as a crucial advantage, enabling temporal and spatial flexibility in both gathering and accessing citizen perspectives. This accessibility extends in two directions: within the municipality, staff can engage with community viewpoints at any time, while citizens can contribute input when and where it suits them best. The system's multilingual capabilities further enhance this accessibility, potentially enabling citizens to express their needs and concerns in their preferred language without requiring immediate translation resources.

This reduction in practical barriers suggests possibilities for more inclusive and continuous citizen engagement. Rather than limiting participation to scheduled meetings or formal consultation periods, digital representatives might enable more fluid and ongoing integration of citizen perspectives into municipal processes. However, this potential for enhanced accessibility also connects to earlier concerns about maintaining authentic community engagement, highlighting the need to balance technological convenience with meaningful human interaction.

“

“This is what we've gathered so far. What do your fellow residents say? Would you like to add to it? Exactly. And everyone can do that in their own time. It's about time and place. The place is somewhat dependent... but the time is independent. And that's what's also very important in participation, of course.”

[6:20]

”

“

“I want to be able to ask someone at three o'clock in the morning once again: how did that actually work? Because how often do you have a meeting and you thought...”

[6:14]

”

Interaction Dynamics and Engagement Patterns

From the clustering a pattern emerged regarding the relationship between individual chat interactions and group discussion effectiveness. In sessions where participants could properly engage with individual chat functionality, the results showed differences in how the responses from digital representatives were received compared to sessions where technical issues limited this initial interaction. As one participant noted, "When you chat individually, you explore certain subjects, then your colleague approaches from a different angle, and together you build a much richer understanding of these perspectives."

This observation suggests that individual familiarity with digital representatives affects how their contributions are perceived during group discussions. This finding aligns with the literature discussed in earlier chapters, particularly Pruitt & Grudin's (2003) observation that personas are effective because they are "generative," allowing users to effortlessly project them into new situations. However, this generative quality requires engagement with the personas, where the individual chat interactions seem to play a significant role.

This finding has important implications for implementation, suggesting that successful integration of digital representatives might require maintaining both individual and group interaction capabilities, rather than focusing exclusively on group participation. More generally, it emphasizes the importance of getting to know the digital representatives and what they stand for.

“

"It was annoying that they kept repeating it, so to speak. Sometimes you hoped to get a bit more depth from them."

[6:1]

”

“

"You reach a kind of maximum in terms of the information you get back. You notice that when you ask."

[6:6]

”

“

"They really focus on just one perspective. And especially when you're having such a discussion, you try to find connections. And then the contribution of such a digital human is limited."

[6:5]

”

Protected Expression

The evaluation revealed an intriguing characteristic of digital representative interactions—what might be termed “protected expression.” Participants described the digital representatives as simultaneously “remarkably humanlike” yet clearly non-human, creating a unique hybrid space for interaction. This space enabled more open discussion of sensitive topics, as civil servants felt free from concerns about hurting feelings or managing immediate emotional reactions. During the sessions, civil servants made several comments to the digital representatives that they likely would not have said at citizens themselves.

This protected space extended beyond civil servant comfort to potential citizen benefits. Participants suggested that the individual interview approach might enable citizens to express perspectives they might hesitate to voice in traditional group settings, where social pressure could inhibit honest feedback. This suggests that the process of creating digital representatives might facilitate more authentic expression while making it possible to interact with these expressions in a safe discussion environment.

“

“.....For example, we do participation for the design of a square. Then you are dealing with real people, but then people don't always dare to say.....”

[4:8]

”

“

“.....And that is of course relatively safe now. If you have them here as virtual persons, you don't have that feeling.”

[4:25]

”

“

“That's a bit of a difference between the emotional side and the business side. Sometimes in your work you just want to be able to say what you mean straight away. And in private you might be a bit more careful with that. That's also sometimes very nice...”

[4:25]

”

Implementation Concerns

Participants expressed significant concerns about the potential organizational implementation of digital representatives. These concerns centered not on technical limitations but on organizational trust and usage patterns. Participants worried that the technology might be misinterpreted by communities as another way for the municipality to create distance rather than engagement.

Participants also questioned whether the municipality would be able to provide this technology as an addition to current participation methods instead of a replacement. Some noted that digital representatives might result in civil servants no longer going out into neighborhoods, instead relying on digital versions rather than seeking real connection with communities.

These concerns reflect an understanding of how technologies are often implemented within municipal contexts—frequently driven by efficiency rather than enhanced engagement quality. This tension between the potential benefits of digital representatives and their possible misuse emerged as a recurring theme in the reflective discussions.

“

“the idea of going out and really seeking people out and using conversation skills, social skills... you use them less”

[6:13]

”

“

“or whether that also gives the residents a kind of satisfaction of... ‘I have now really been able to participate or contribute my thoughts.’”

[6:16]

”

“

“because we’ve created digital representatives. Well, I already know how that’s going to land with those people. How? Bad, very bad.”

[8:10]

”

From Representative to Facilitator

The evaluation revealed a finding about how participants envisioned the potential evolution of digital representatives within municipal processes. Rather than limiting these systems to simply presenting citizen perspectives, civil servants expressed desire for more expansive system roles, particularly in meeting facilitation and decision support.

This expanded vision emerged most clearly in discussions about meeting management. Participants suggested that digital representatives could take more active control in maintaining meeting structure and focus, keeping discussions aligned with intended objectives. This desire for enhanced meeting management indicates a significant shift in how participants conceptualized these systems—not just as channels for citizen perspectives but as active facilitators of policy discussions.

Participants particularly emphasized the need for more directive output from digital representatives. The current implementation, presenting multiple citizen perspectives with equal weight, created what some participants saw as a decision-making challenge. This led to several specific suggestions for system enhancement, including:

- Integration of broader contextual data, combining citizen perspectives with quantitative information and historical documentation
- Having digital representatives engage in their own dialogue before presenting conclusions rather than individual perspectives
- Creating a single integrated viewpoint about optimal space development rather than multiple distinct perspectives

These suggestions point toward a shift in how participants viewed the potential role of AI in participation processes—moving from simple representation toward active facilitation and recommendation generation.

“

“Well, then they could also come to a compromise together. Which we could then respond to. Then you would simply have a different discussion.” (civil servant) “And then you can immediately assume that compromise, you mean?” (Researcher) “Yes.” (civil servant)

[2:12]

”

“

“And all six of those (perspectives) are worth the same. So how do you make a decision then?”

[4:14]

”

“

We have data, we have WCO, we have neighborhood profiles, we just have all kinds of numbers, safety index, monitors, conversations. And at some point, when you start a project or a process, you bring it all together.

[8:43]

”

Policy Discussion Dynamics and Decision Patterns

The analysis of actual policy discussions revealed insights about how digital representatives influence municipal decision-making processes. These findings came not just from what was said, but from observing how civil servants integrated and responded to AI-mediated perspectives throughout their deliberations.

The introduction of digital representative perspectives demonstrably shaped discussion directions and outcomes. A particularly revealing moment occurred when a digital representative challenged the group's focus on social interaction spaces, noting that this emphasis didn't align with all user needs. Rather than simply acknowledging this perspective, civil servants actively incorporated it into their thinking, recognizing this sentiment in their own experiences and adjusting their approach accordingly. This willingness to integrate AI-mediated perspectives suggests potential for these systems to influence policy development.

Digital representatives proved particularly effective at challenging civil servant assumptions about community needs. In one notable instance, when participants expressed doubt about combining social activities with quiet working spaces, a digital representative intervened with specific counterarguments based on lived experience. The representative emphasized that quiet workspaces were crucial for certain needs while explaining how social activities could coexist through thoughtful space organization. This intervention led to substantive discussion about flexible space design, demonstrating how digital representatives might help bridge gaps between administrative assumptions and community realities.

A reflexive note must be added to this observation: the civil servants' receptiveness to digital representative input might have been strengthened by their limited prior knowledge about the specific policy challenge. This circumstance potentially created more openness to new perspectives than might exist in situations where participants hold strong preexisting views or established positions.

“

“Yes, yes, yes. I didn't know that at all, but that phenomenon. But that is apparently shared.”

[1:16]

”

“

I am Daan. I see a pattern in this conversation. There is a lot of focus on organizing activities. But for me personally, a professional and quiet workplace is the most important. I think that the municipality should facilitate an environment where we can be productive ourselves. And not fill in too much.

Response from digital representative

”

Direct Interpretation of Key Findings

Data Quality as Foundation for Effectiveness

The evaluation confirmed a critical insight from earlier explorations: the quality of input data fundamentally determines the effectiveness of digital representatives. When provided with rich source material from the semi-structured interviews, the digital representatives demonstrated greater nuance and contextual understanding than in previous iterations. This validates the approach of conducting deliberate interviews rather than relying on existing documentation.

Particularly notable was the success of using unprocessed interview transcripts without speaker identification or text correction. The Large Language Models successfully distinguished between interviewer questions and participant responses, extracting coherent perspectives despite the raw nature of the input. This suggests potential for streamlining the creation process, making implementation more feasible within resource-constrained municipal environments.

Interaction Modality Effects on Engagement

A significant pattern emerged regarding how different interaction modalities affected engagement. Voice interactions triggered noticeably stronger emotional responses from civil servants compared to text-based exchanges. When digital representatives spoke, participants demonstrated more engaged behavior—interrupting, responding spontaneously, and treating the interaction more like a conversation with a person than an information exchange with a system.

This finding has direct implications for implementation, suggesting that voice capabilities significantly enhance the perceived authenticity and impact of digital representatives in policy discussions. However, the text-based interface served important familiarization purposes, indicating that implementation might benefit from maintaining both modalities for different contexts.

Individual Exploration and Group Discussion Relationship

The evaluation revealed an important relationship between individual exploration and group effectiveness. In sessions where participants had sufficient time to engage with the individual chat functionality, digital representatives appeared to have greater influence in subsequent group discussions. As one participant noted, “When you chat individually, you explore certain subjects, then your colleague approaches from a different angle, and together you build a much richer understanding of these perspectives.”

This observation suggests that individual familiarity with

digital representatives affects how their contributions are perceived during group discussions, indicating that implementation should include dedicated time for individual engagement before group deliberation.

Organizational Position and Perception Differences

How civil servants perceived digital representatives appeared to correlate with their organizational position and typical level of citizen contact. Those whose roles involved frequent direct citizen interaction tended to view digital representatives more critically, emphasizing their limitations compared to actual citizen engagement. In contrast, civil servants with less regular citizen contact often saw greater potential value in these systems.

This finding reveals how institutional positioning shapes technology reception and suggests that implementation strategies should be tailored to different departmental contexts. For those already engaged directly with citizens, these tools might best serve as supplements to existing practices, while for those with limited citizen contact, they might provide valuable initial exposure to community perspectives.

Limitations

Sample Size and Diversity Considerations

This evaluation's relatively small scale—four sessions involving approximately 28 civil servants—means the findings should be considered initial indications rather than definitive conclusions. The patterns identified merit further investigation with larger and more diverse participant groups.

The focused demographic profile of citizen participants—primarily young adults with similar backgrounds—limits the evaluation's ability to assess how digital representatives might function with more diverse community perspectives. While this approach effectively tested the system's ability to detect nuance within a seemingly homogeneous group, future research should incorporate more diverse citizen participants.

Researcher Role and Potential Biases

The same person conducted the citizen interviews, developed the digital representatives, facilitated the evaluation sessions, and analyzed the results. While external guidance provided some perspective, the researcher's views inevitably influenced multiple aspects of the process. These potential biases were addressed through transparent documentation of the analysis process, but they remain important considerations when interpreting findings.

Evaluation Scope

This evaluation did not include testing how digital representatives might impact communication of decisions back to citizens. This potential benefit identified in earlier chapters remains theoretical, and more research is needed to assess whether digital representatives would enhance the feedback loop between policy decisions and citizen understanding.

Conclusion

This chapter has presented the evaluation of digital representatives in municipal policy discussions. Building on the technical development described in previous chapters, this evaluation examined how civil servants engage with AI-mediated citizen perspectives in actual policy deliberations about neighborhood hub transformation.

Through a structured approach incorporating individual exploration and group discussion, the evaluation revealed several key patterns: the importance of voice interaction for emotional engagement, the relationship between individual exploration and group effectiveness, the significant influence of organizational position on perception, and the potential for digital representatives to challenge civil servant assumptions in productive ways.

While limitations in sample size, participant diversity, and evaluation scope must be acknowledged, the findings provide valuable initial insights into how digital representatives might function within municipal contexts. The evaluation suggests that these systems can meaningfully contribute to policy discussions by maintaining citizen perspectives throughout deliberations, though their effectiveness depends fundamentally on the quality of input data and thoughtful implementation that acknowledges organizational realities.

These findings set the foundation for Chapter 7, which will explore the broader theoretical and practical implications of digital representatives for democratic participation, institutional structures, and the relationship between citizens and their municipal government.

7.

Conclusion

This research confronted a democratic challenge in Rotterdam: the loss of citizen perspective richness within municipal policy processes, contributing to an “empty chair” during deliberations. Amid concerns about trust and participation, this study explored if Large Language Models (LLMs) could preserve citizen viewpoints through “digital representatives.” Using strategic design, the work iteratively developed and evaluated these AI tools, examining technical feasibility alongside their reception and function within the institutional context.

The findings show that LLMs can represent consistent citizen viewpoints, but the value and authenticity of digital representatives depend entirely on the input data’s quality and form. Representations derived from direct, contextual citizen input, like interview transcripts, maintained nuance and could substantively influence policy discussions. In contrast, using abstracted municipal documentation proved inadequate, underscoring how current institutional practices often fail to retain the experiential knowledge needed for genuine representation. Certain design elements, notably voice interaction and opportunities for individual familiarization, improved civil servant engagement and the perceived connection to the represented perspectives.

Evaluation with Rotterdam civil servants yielded a mixed reception, highlighting both potential and significant concerns. Participants saw possibilities for digital representatives to increase accessibility to citizen perspectives, challenge assumptions, and offer unique interaction spaces. Concurrently, they expressed reservations about accuracy, representativeness, and the risk of these tools replacing necessary direct citizen engagement. While cautious about supplanting human interaction, these civil servants also indicated interest in digital representatives evolving beyond representation towards facilitating discussions and integrating information.

Therefore, this research’s primary contribution is not only the technology itself. It is the illumination of critical deficiencies in how municipalities currently capture and process citizen knowledge. Digital representatives act as both a potential tool and a diagnostic probe, revealing the “dark matter” of institutional habits that diminish citizen voices. This work highlights that enhancing democratic participation requires transforming not just engagement methods, but how institutions value, retain, and utilize experiential knowledge, pointing towards an alternative epistemological approach focused on preserving context over pursuing abstraction.

The study acknowledges limitations inherent in its exploratory scope, including the evaluation scale and the challenge of verifying AI-mediated representation’s fidelity. Findings are tied to the source material’s quality; technology cannot manufacture nuance absent in the input.

Ultimately, digital representatives are not a complete solution to Rotterdam’s democratic challenges and should not substitute for direct human engagement. They offer, however, a potential means to enhance participation by sustaining the presence and influence of citizen perspectives throughout the policy lifecycle, especially for views easily lost. Their most potent function may be to provoke necessary conversations within municipal institutions about the value placed on citizen knowledge and the processes required to ensure it genuinely informs governance—thereby helping to fill, not just simulate presence in, the empty chair.

8. Discussion

Interpretation of Findings

My journey through the iterative development and evaluation of digital representatives revealed consistent patterns that illuminate both the potential and limitations of AI-mediated citizen perspectives in municipal governance. This journey, spanning from initial technical exploration to structured evaluation with civil servants, provides a foundation for understanding how such systems might address Rotterdam's democratic challenges.

The Evolution of Understanding Digital Representatives

What began as a technical investigation into whether AI could consistently represent citizen viewpoints evolved into a deeper inquiry about how citizen perspectives move through institutional processes. Each development phase revealed new insights about the interplay between technological capabilities and institutional context.

The first exploration with the parking-to-terrace conversion demonstrated basic technical feasibility but revealed limitations in perspectives generated from synthesized data. The Vroesenpark playground case exposed significant gaps in municipal documentation practices, where citizen input typically undergoes substantial filtering before becoming accessible for policy development. Finally, the Reyerwaard documentary exploration showed that when working with rich, contextual documentation, digital representatives could demonstrate the nuance and depth necessary for meaningful policy contribution.

The evaluation with civil servants confirmed and extended these developmental insights. When provided with rich interview data, digital representatives could meaningfully influence policy discussions by challenging assumptions and providing consistent perspective throughout deliberations. This confirms that the technology itself isn't the primary limiting factor in creating authentic digital representatives—it's the availability of rich, contextual documentation of citizen perspectives.

From Representation to Enhanced Understanding

Through this research process, my understanding

of digital representatives' contribution evolved significantly. What began as an exploration of representation—how AI might maintain consistent citizen viewpoints—transformed into a deeper investigation of understanding—how these systems might enhance civil servants' comprehension of community needs and aspirations.

This shift became particularly apparent during the evaluation, where individual exploration of digital representatives significantly enhanced subsequent group discussions. The digital representatives didn't just present information; they appeared to enhance how civil servants processed and integrated citizen perspectives into their thinking. This connects to Pruitt & Grudin's (2003) theoretical insights about the "generative" quality of effective personas—their ability to engage our natural capacity to anticipate others' reactions and understand their perspectives.

The voice interaction capabilities further enhanced this effect. While the text-based interface proved functional for information exchange, voice interactions triggered noticeably stronger emotional engagement. When digital representatives spoke, civil servants demonstrated more spontaneous responses compared to text-based interaction and appeared to connect more deeply with the perspectives being expressed. This suggests that future implementations should prioritize voice interaction capabilities, particularly for contexts where emotional engagement and perspective empathy are important.

The Interplay of Technology and Institutional Context

Throughout this research, the interplay between technological capabilities and institutional context became increasingly apparent. The Vroesenpark case revealed how institutional structures—contractual arrangements, data access policies, departmental boundaries—fundamentally shaped what technological implementations were possible. Similarly, the evaluation revealed how existing power dynamics and organizational patterns influenced how civil servants engaged with digital representatives.

This finding aligns with Hill's (2017) strategic design methodology that guided this research. The "dark

matter" of institutional structures proved as important as the "matter" of technological implementation. Each prototype served not just as a technical demonstration but as a probe that revealed underlying institutional dynamics. These dynamics often represented the true barriers to enhanced democratic participation rather than technological limitations. This insight suggests that effective implementation of digital representatives requires attention to both technological and institutional dimensions simultaneously.

Rethinking Democratic Processes

My exploration of digital representatives contributes not only to practical participation methods but also to theoretical understanding of democratic processes. The empirical findings from both development and evaluation inform and potentially reshape democratic theory in several significant ways.

Beyond Traditional Democratic Models

Digital representatives occupy an interesting theoretical space between different models of democratic governance. As Chapter 1 outlined, Rotterdam's democratic landscape reveals tensions between representative democracy's formal structures and community experiences, with voter turnout as low as 21% in some neighborhoods signaling challenges to representative democracy's basic assumptions.

In relation to Schumpeter's (2013) minimalist view of representative democracy, where electoral participation constitutes the primary mechanism for citizen influence, digital representatives offer a supplementary channel. They maintain citizen perspectives throughout policy processes even when those citizens aren't formally elected or physically present. This addresses a core limitation of purely representative models: the tendency for marginalized voices to disappear from governance between electoral cycles.

When viewed through Habermas's deliberative framework, digital representatives demonstrate both alignment and tension. Their capacity for reasoned dialogue supports deliberative democracy's emphasis on rational exchange. However, they simultaneously challenge the model's assumptions about consensus-building. Digital representatives maintained distinct perspectives without resolving into unified viewpoints, preserving disagreement even within constructive dialogue. This quality aligns more closely with Mouffe's (1999) agonistic pluralism, which views disagreement not as a flaw in democratic systems but as an essential feature.

The evaluation revealed how digital representatives created spaces where conflicting perspectives coexisted. When civil servants encountered perspectives that challenged their assumptions, they

didn't simply seek to resolve these differences. Instead, they engaged with them as legitimate viewpoints that informed more nuanced policy discussions. This pattern suggests that digital representatives might facilitate the kind of productive democratic tension that Mouffe's theory envisions, where different forms of knowledge remain distinct while engaging in meaningful dialogue.

This hybrid quality offers theoretical insight for addressing Rotterdam's specific democratic challenges. The municipality's participation landscape revealed parallel forms of democratic engagement: formal processes with limited participation alongside vibrant community activities that often remain disconnected from governance. Digital representatives suggest possibilities for bridging these parallel democratic forms, bringing community perspectives into institutional processes without requiring their transformation into institutional language.

Epistemological Implications of AI-Mediated Participation

The development and evaluation of digital representatives reveals epistemological implications for how knowledge functions within democratic processes. Chapter 2 examined different epistemological positions—from positivism's emphasis on objective knowledge to strong constructionism's view that all knowledge is socially constructed. Digital representatives challenge these traditional categories by AI-mediated perspectives that maintain elements of human subjectivity while functioning through fundamentally different knowledge processes.

Traditional participation documentation typically reflects what Andrews (2012) associates with positivist approaches: citizen input gains value through refinement toward abstraction, filtering out "subjective" elements to extract supposedly objective insights. Digital representatives suggest an alternative approach more aligned with social constructionism, preserving the contextual richness that gives citizen perspectives their meaning and value.

This alternative approach became visible during the evaluation. When digital representatives contributed specific examples and emotional context rather than abstract recommendations, civil servants engaged more deeply with the underlying perspectives. These AI-mediated viewpoints preserved qualities that traditionally get lost in documentation processes, particularly the connection to lived experience that gives citizen input its distinctive epistemological contribution.

Yet these AI-mediated perspectives function through fundamentally different processes than human knowledge construction. As noted in Chapter 3, LLMs operate through pattern recognition rather than embodied understanding. They create a form of

"view from everywhere and nowhere" that differs from both individual human perspectives and traditional institutional knowledge processing. This quality raises important questions about how knowledge circulates within democratic systems.

The success of digital representatives in policy discussions suggests that meaningful democratic knowledge might be maintained through mechanisms that don't require embodied understanding. This challenges traditional assumptions about knowledge transmission in democratic contexts, where lived experience has been assumed to require human mediation to maintain its authenticity and value.

Reformulating the Empty Chair Problem

My research began with what might be called the "empty chair" problem in democratic processes: the absence of citizen perspectives from policy discussions. As decisions move through institutional processes, the richness of citizen experiences increasingly fades, leaving an empty space where their voices should be. Digital representatives offer a reformulation of this problem, suggesting that the issue isn't simply absence but transformation.

The traditional view of this problem focuses on representation gaps—who isn't present in decision-making forums. My findings suggest a more nuanced understanding focused on knowledge transformation—how citizen perspectives change as they move through institutional processes. The challenge isn't just that certain citizens aren't physically present but that the documentation practices meant to represent them fundamentally alter the nature of their contributions.

Digital representatives demonstrate an alternative approach to this challenge. Rather than trying to make citizens physically present for all decisions (an impractical goal in modern governance), they maintain more of the qualities that make citizen input valuable throughout policy processes. This reframes the empty chair not as permanently vacant but as temporarily occupied by technologies that preserve connections to citizen perspectives.

This reformulation has significant implications for democratic theory. It suggests that enhancing democratic participation might focus not just on who participates initially but on how their perspectives persist through subsequent processes. The quality of democratic engagement doesn't end when citizens stop actively participating but continues through how their contributions influence ongoing governance.

This insight connects to what Snel, Custers & Engbersen (2018) observed in Rotterdam communities: vibrant civic engagement often exists outside formal political processes. The challenge isn't necessarily increasing participation in traditional structures but creating more

responsive interfaces between community activity and institutional processes. Digital representatives suggest possibilities for such interfaces, bringing perspectives from Rotterdam's active community life into municipal decision-making without requiring community members to navigate formal participation barriers

Tensions and Paradoxes: Critical Reflections

While digital representatives demonstrated promising capabilities for enhancing democratic participation, my research revealed fundamental tensions and paradoxes that resist simple resolution. These tensions emerged consistently across development iterations and evaluation sessions, suggesting they reflect deeper challenges in the relationship between technology and democratic processes.

Efficiency Versus Democratic Depth

Perhaps the most profound tension concerns the relationship between efficiency and democratic depth. During preparation for the evaluation sessions, I discovered I could generate a structured participation report from interview data in minutes rather than days. This efficiency is compelling from both resource and accessibility perspectives, potentially enabling more frequent and comprehensive citizen consultation.

However, this efficiency introduces fundamental questions about the purpose of participation itself. If participation primarily serves instrumental goals—gathering information to improve policy effectiveness—then efficiency gains might be wanted. Yet if participation serves broader democratic purposes—building trust, developing civic capacity, enhancing legitimacy—then processes that prioritize efficiency over engagement may ultimately undermine their own purpose.

This tension connects to the theoretical frameworks examined in Chapter 1. Efficiency-focused approaches risk sacrificing this depth for breadth, potentially generating more citizen input without the corresponding depth of understanding.

The Reyerood documentary exploration showed this tension particularly clearly. The documentary captured nuanced community dynamics that would be difficult to document efficiently, requiring significant time investment for both creation and analysis. Yet these nuances proved crucial for developing authentic digital representatives. This suggests that meaningful democratic engagement might require embracing certain inefficiencies rather than eliminating them, preserving spaces for the time-intensive processes that generate deeper understanding.

The Paradox of Artificial Authenticity

A second tension concerns artificial authenticity—using AI-generated personas to preserve authentic citizen voices. This paradox emerged throughout my development process as I created digital representatives that were simultaneously artificial constructs and channels for genuine community perspectives.

The development process revealed the complexities of this relationship. When creating digital representatives, I necessarily synthesized demographic details and contextual elements while attempting to maintain authentic perspectives from citizen input. This created a blurred boundary between necessary contextual generation and potentially problematic fabrication. When a digital representative speaks as a fictional 25-year-old graduate student but expresses actual perspectives gathered from interviews, what constitutes truthful representation versus fabrication?

This question became particularly acute during the evaluation. Civil servants engaged with digital representatives as if they represented actual community members, while simultaneously knowing they interacted with technological constructs. This created a unique interaction dynamic that could be described as strangely more open than talking to actual citizens. The artificial nature of the representation paradoxically enabled more honest engagement with the authentic perspectives being expressed.

This paradox connects to deeper questions about representation itself. Traditional participation documentation also creates representations of citizen perspectives, abstracting and categorizing individual expressions into institutional language. Digital representatives replace this abstraction with a different kind of representation—one that maintains more original texture but introduces new forms of mediation. Neither approach provides unmediated access to citizen perspectives, suggesting that all democratic processes involve some form of representation rather than direct presentation.

Contradictory Desires for Technology's Role

Participants simultaneously expressed concern about technology replacing human engagement while requesting more directive capabilities from the digital representatives themselves. This contradiction suggests deeper ambivalence about technology's role in democratic processes.

During reflective discussions, civil servants articulated legitimate concerns that digital representatives might become another way for the municipality to avoid real engagement with communities. They worried that the efficiency of these systems might lead to fewer direct interactions between municipal staff and citizens, potentially widening rather than bridging the trust gap identified in Chapter 1. These concerns reflect

an understanding of how technologies are often implemented within municipal contexts—frequently driven by efficiency rather than enhanced engagement quality.

Yet these same participants advocated for shifting digital representatives from simple perspective presentation toward more directive facilitation roles. They suggested digital representatives could take more active control in maintaining meeting structure, integrating multiple data sources, and providing recommendations rather than just perspectives. This desire for enhanced directive capabilities suggests openness to AI playing significant roles in decision processes, even as participants expressed concern about diminishing human connection.

This contradiction reveals competing values in municipal innovation—the desire for enhanced capabilities alongside concerns about authentic engagement and democratic accountability. It suggests that civil servants navigate complex organizational contexts where efficiency pressures often compete with democratic values. Their contradictory responses don't indicate inconsistency but rather reflect the genuine tensions they navigate in their professional roles.

The Representative-Facilitator Spectrum

My development journey revealed an evolving understanding of digital representatives' potential role—from purely representative tools that present citizen perspectives to active facilitators that guide discussion and decision processes. This spectrum presents both opportunities and significant challenges for democratic enhancement.

The initial concept focused primarily on representation—using AI to maintain consistent citizen viewpoints throughout policy discussions. This approach addressed the challenge identified in Chapter 1 of preserving experiential knowledge as citizen input moves through institutional processes. The fundamental goal was fidelity to original citizen perspectives, maintaining their presence without transforming them into institutional abstractions.

However, the evaluation revealed civil servant interest in more facilitative capabilities. Participants suggested digital representatives could help structure discussions, identify patterns across different perspectives, and even generate recommendations based on citizen input. This shift from representation to facilitation suggests potential for these systems to address additional challenges beyond maintaining citizen voices.

This spectrum creates a fundamental tension. More directive capabilities might enhance efficiency and immediate utility in policy discussions. However, they simultaneously introduce greater distance from original citizen expressions and increase the system's

interpretive authority. As digital representatives move from simply presenting perspectives to actively processing and integrating them, they potentially replicate the very transformation process that initially motivated this research—the conversion of rich citizen expressions into institutional language.

Confronting AI Risks in Democratic Contexts

My exploration of digital representatives directly connects with the AI risks examined in Chapter 2. This research provides insights into how municipalities might navigate these challenges, balancing technological potential with democratic integrity.

Persuasive Influence and Verification Challenges

The evaluation revealed the persuasive capabilities of digital representatives in policy discussions. When digital representatives challenged civil servants' assumptions, participants actively incorporated these perspectives into their thinking and adjusted approaches accordingly. This observed influence aligns with research from Matz et al. (2024) discussed in Chapter 3, which demonstrated how AI-generated persuasive messages can significantly impact decision-making.

This persuasive power becomes particularly consequential when combined with verification difficulties. During evaluation sessions, civil servants necessarily assumed the digital representatives accurately reflected citizen input, as real-time verification was impossible. The personal, experiential nature of citizen perspectives compounds this challenge—when a digital representative shares a seemingly authentic lived experience, this narrative becomes difficult to challenge or verify.

My development process addressed this risk through transparent connection to source material, ensuring digital representatives remained grounded in actual citizen input rather than generating new perspectives. However, scaling this approach would require careful attention to verification mechanisms that maintain clear connections between citizen expressions and AI-mediated representations.

Designing for Appropriate Trust

My development process actively addressed several trust-related risks identified in Chapter 3. A deliberate design choice involved positioning digital representatives as subjective perspective-holders rather than objective information sources. By explicitly framing outputs as citizen opinions rather than factual assertions, the system avoided the risk of presenting subjective judgments with an unwarranted aura of objectivity.

Another significant design decision involved systematically excluding numerical data and statistics from digital representative outputs. This choice addressed the hallucination risk discussed in Chapter 3, where LLMs can generate plausible but fabricated quantitative information. Since these systems cannot reliably distinguish between accurate and fabricated numbers, excluding this content altogether represented a constraint on their role.

These design choices reflect a fundamental principle that emerged throughout this research: AI systems in democratic contexts should operate within carefully defined boundaries that align with their capabilities and limitations. Rather than maximizing technical capabilities, implementation should focus on creating appropriate constraints that maintain democratic integrity while leveraging technological benefits.

8. Practical Implications for the Municipality

The tensions and theoretical insights discussed in previous sections lead to several practical implications for municipal governance. These implications emerge directly from my research findings, addressing the institutional changes necessary for digital representatives to meaningfully enhance democratic

Transforming Documentation Practices

A fundamental practical implication concerns how municipalities document citizen participation. Throughout my research, the quality of input data consistently determined the authenticity and value of digital representatives. Current documentation practices—which typically transform rich citizen narratives into abstracted summary points—proved insufficient for meaningful digital representation.

Municipalities should develop approaches that preserve more of the original texture of citizen contributions. This transformation doesn't require retaining identifying information; rather, it focuses on maintaining the experiential richness that gives citizen input its distinctive value. Recording and transcribing citizen perspectives rather than relying solely on facilitator notes would create stronger foundations not just for digital representatives but for any effort to maintain citizen voices throughout policy processes.

This transformation connects directly to the democratic challenge identified in Chapter 1: how rich stories become thin summaries as they move through municipal processes. By preserving more original expression, municipalities could address this challenge regardless of whether they ultimately implement digital representatives. This approach would benefit traditional participation processes by maintaining stronger connections between citizen experiences and policy development.

Creating Space for Responsible Innovation

My research journey revealed significant institutional barriers to experimentation with democratic technologies. As I experienced during the Vroesepark case, access to participation data proved nearly impossible due to contractual limitations, privacy concerns, and fragmented documentation practices. These constraints significantly limited the iterative development process that innovation requires.

Municipalities should develop dedicated pathways for responsible experimentation that recognize the

distinction between research and implementation. This means creating appropriate risk management approaches for each context, enabling controlled access to anonymized data for research purposes while maintaining robust protections for citizen privacy. Experimental spaces like VONK provide valuable starting points but require stronger institutional support to overcome departmental boundaries and contractual limitations.

This recommendation reflects what Hill (2017) calls the “dark matter” of institutions—the invisible policies and structures that shape what innovations are possible. Addressing this dark matter is as important as developing the technological “matter” itself, creating conditions where democratic innovations can be thoughtfully explored and evaluated.

Positioning of AI Technologies

My evaluation revealed a crucial tension in how civil servants engaged with digital representatives. Participants simultaneously expressed concern about technology replacing human engagement while desiring more directive capabilities from the systems themselves. This contradiction highlights the need for strategic clarity about the purpose of democratic technologies.

Municipalities should develop explicit frameworks that position AI technologies as tools for enhancing rather than replacing human judgment in democratic processes. These frameworks should articulate clear values regarding appropriate uses, establish transparency requirements for how AI-mediated perspectives influence decisions, and create accountability mechanisms that maintain focus on democratic enhancement rather than merely streamlining processes.

This strategic positioning connects to the paradox of efficiency versus democratic depth discussed earlier. While certain efficiency gains might enhance

participation by making it more accessible, municipalities must ensure these gains don't come at the cost of authentic engagement. Clear strategic positioning helps navigate this balance, ensuring technological implementations serve democratic values rather than potentially undermining them.

Bridging Organizational Divides

A revealing insight from my evaluation concerned how perceptions of digital representatives varied based on institutional position. Civil servants who regularly engaged directly with citizens viewed these tools more critically, while those with less community contact saw greater potential value. This divide reflects a fundamental challenge in municipal governance structures: the disconnection between those who gather citizen input and those who develop policy based on this input.

Municipalities should develop cross-departmental protocols for how digital representatives might facilitate knowledge sharing between those who gather citizen input and those who develop policy. This would create more integrated understanding of community needs across municipal functions, addressing the disconnection between citizen experiences and policy development identified in Chapter 1.

This recommendation connects to the “empty chair” problem reformulated earlier—the absence of citizen perspectives from policy discussions where decisions affecting them are made. Digital representatives could help bridge this gap by creating shared reference points across different organizational positions, making citizen perspectives more accessible throughout the organization rather than confined to departments that interact directly with communities.

References

- Abha Eli Phoboo & Melanie Mitchell. (n.d.). Podcast: Ep. 2: The relationship between language and thought | Santa Fe Institute [Broadcast]. Retrieved February 13, 2025, from <https://www.santafe.edu/culture/podcasts/ep-2-the-relationship-between-language-and-thought>
- Adly Templeton, Tom Conerly, Jonathan Marcus, Jack Lindsey, & Trenton Bricken. (2024, May 21). Scaling Monosemanticity: Extracting Interpretable Features from Claude 3 Sonnet. <https://transformer-circuits.pub/2024/scaling-monosemanticity/index.html>
- Andrews, T. (2012). What is Social Constructionism? 11(1).
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? . Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, 610–623. <https://doi.org/10.1145/3442188.3445922>
- Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101. <https://doi.org/10.1191/1478088706qp063oa>
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., ... Amodei, D. (2020). Language Models are Few-Shot Learners. *Advances in Neural Information Processing Systems*, 33, 1877–1901. https://proceedings.neurips.cc/paper_files/paper/2020/hash/1457cod6bfc4967418bfb8ac142f64a-Abstract.html
- Chu, E., Vijayaraghavan, P., & Roy, D. (2018). Learning Personas from Dialogue with Attentive Memory Networks (arXiv:1810.08717). arXiv. <http://arxiv.org/abs/1810.08717>
- Dan Hill. (2017). Dark Matter and Trojan Horses: A Strategic Design Vocabulary.
- Friedman, B., & Nissenbaum, H. (1996). Bias in computer systems. *ACM Trans. Inf. Syst.*, 14(3), 330–347. <https://doi.org/10.1145/230538.230561>
- Gemeente Rotterdam. (2023, October). Participatieverslag. Rotterdamse Woonvisie.
- Greenblatt, R., Denison, C., Wright, B., Roger, F., MacDiarmid, M., Marks, S., Treutlein, J., Belonax, T., Chen, J., Duvenaud, D., Khan, A., Michael, J., Mindermann, S., Perez, E., Petrini, L., Uesato, J., Kaplan, J., Shlegeris, B., Bowman, S. R., & Hubinger, E. (2024). Alignment faking in large language models (arXiv:2412.14093). arXiv. <https://doi.org/10.48550/arXiv.2412.14093>
- Habermas, J. (1996). *The Habermas Reader*.
- Haraway, D. (1988). Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective. *Feminist Studies*, 14(3), 575. <https://doi.org/10.2307/3178066>
- Hurenkamp, M., & Tonkens, E. (2020). Ontwerpprincipes voor betere burgerparticipatie. *Bestuurskunde*, 29(1), 54–63. <https://doi.org/10.5553/Bk/092733872020029001007>
- Ineke Vogel, Marn van Rhee, & Swarish Marapengopi. (2022). Gemeenteraads-verkiezingen 2022 Rotterdam. <https://www.onderzoek010.nl/handlers/ballroom.ashx?function=download&id=728>
- Li, A. (2024, June 27). Gemini 1.5 Pro now offers a 2M token context window for devs. <https://gto5google.com/2024/06/27/gemini-1-5-pro-2-million/>
- Li, J., Galley, M., Brockett, C., Spithourakis, G. P., Gao, J., & Dolan, B. (2016). A Persona-Based Neural Conversation Model (arXiv:1603.06155). arXiv. <http://arxiv.org/abs/1603.06155>
- Lieke Dreijerink, Hanneke Kruize, & Irene van Kamp. (2008). Burgerparticipatie in beleidsvorming. RIVM. <https://rivm.openrepository.com/bitstream/handle/10029/258093/830950003.pdf?sequence=3>
- Lin, G., Jiang, J., Yang, J., Zheng, Z., & Liang, C. (2025). OmniHuman-1: Rethinking the Scaling-Up of One-Stage Conditioned Human Animation Models (arXiv:2502.01061). arXiv. <https://doi.org/10.48550/arXiv.2502.01061>
- L.J. de Graaf & A.M.B Michels. (2009). Examining citizens participation from a citizen's perspective.
- Marianne van Bochove, Christine Bleijenbergh, Hans Moors, & Hans Vollaard. (2023, June). Opkomst-van-rotterdamers-bij-gemeenteraadsverkiezingen.pdf. Opkomst En Vertegenwoordiging van Rotterdamers. <https://www.dehaagsehogeschool.nl/media/opkomst->

van-rotterdamers-bij-gemeenteraadsverkiezingen

Mary Harrington. (2025, January 6). Ideological bias will taint Google's Habermas machine. *UnHerd*. <https://unherd.com/newsroom/ideological-bias-is-bound-to-taint-googles-habermas-machine/>

Matz, S. C., Teeny, J. D., Vaid, S. S., Peters, H., Harari, G. M., & Cerf, M. (2024). The potential of generative AI for personalized persuasion at scale. *Scientific Reports*, 14(1), 4692. <https://doi.org/10.1038/s41598-024-53755-0>

May, J. (2006). Ladders, Stars and Triangles: Old and New Theory for the Practice of Public Participation. *International Journal of Market Research*, 48(3), 305–319. <https://doi.org/10.1177/147078530604800305>

Ministerie van Volksgezondheid, W. en S. (2024, September 3). Meerderheid Nederlanders vindt dat politiek niet opkomt voor mensen zoals zij—Nieuwsbericht—Sociaal en Cultureel Planbureau [Nieuwsbericht]. Ministerie van Volksgezondheid, Welzijn en Sport. <https://www.scp.nl/actueel/nieuws/2024/09/03/meerderheid-nederlanders-vindt-dat-politiek-niet-opkomt-voor-mensen-zoals-zij>

Mittelstadt, B., Wachter, S., & Russell, C. (2023). To protect science, we must use LLMs as zero-shot translators. *Nature Human Behaviour*, 7(11), 1830–1832. <https://doi.org/10.1038/s41562-023-01744-0>

Mouffe, C. (1999). Deliberative Democracy or Agonistic Pluralism? *Social Research*, 66(3), 745–758.

Pruitt, J., & Grudin, J. (2003). Personas: Practice and theory. *Proceedings of the 2003 Conference on Designing for User Experiences*, 1–15. <https://doi.org/10.1145/997078.997089>

Radford, A., Kim, J. W., Xu, T., Brockman, G., McLeavey, C., & Sutskever, I. (2022). Robust Speech Recognition via Large-Scale Weak Supervision (arXiv:2212.04356). *arXiv*. <https://doi.org/10.48550/arXiv.2212.04356>

Rawls, J. (2003). *A theory of justice* (Rev. ed., 5.-6. printing). Belknap Press of Harvard Univ. Press.

Roth, E. (2024, December 4). ChatGPT now has over 300 million weekly users. *The Verge*. <https://www.theverge.com/2024/12/4/24313097/chatgpt-300-million-weekly-users>

Rudin, C. (2019). Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead. *Nature Machine Intelligence*, 1(5), 206–215. <https://doi.org/10.1038/s42256-019-0048-x>

Salminen, J., Guan, K., Jung, S.-G., Chowdhury, S. A., & Jansen, B. J. (2020). A Literature Review of Quantitative Persona Creation. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–14. <https://doi.org/10.1145/3313831.3376502>

Schumpeter, J. A., & Swedberg, R. (2013). *Capitalism, socialism and democracy*. Routledge. <https://doi.org/10.4324/9780203202050>

Snel, E., Custers, G., & Engbersen, G. (2018). Ongelijkheid in de participatiestad: Stadsbuurten en burgerparticipatie. *Mens & Maatschappij*, 93(1), 31–57. <https://doi.org/10.5117/MEM2018.1.SNEL>

Team, G., Georgiev, P., Lei, V. I., Burnell, R., Bai, L., Gulati, A., Tanzer, G., Vincent, D., Pan, Z., Wang, S., Mariooryad, S., Ding, Y., Geng, X., Alcober, F., Frostig, R., Omernick, M., Walker, L., Paduraru, C., Sorokin, C., ... Vinyals, O. (2024). Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context (arXiv:2403.05530). *arXiv*. <http://arxiv.org/abs/2403.05530>

Tessler, M. H., Bakker, M. A., Jarrett, D., Sheahan, H., Chadwick, M. J., Koster, R., Evans, G., Campbell-Gillingham, L., Collins, T., Parkes, D. C., Botvinick, M., & Summerfield, C. (2024). AI can help humans find common ground in democratic deliberation. *Science*, 386(6719), eadq2852. <https://doi.org/10.1126/science.adq2852>

Thomas, D. R. (2017). Feedback from research participants: Are member checks useful in qualitative research? *Qualitative Research in Psychology*, 14(1), 23–41. <https://doi.org/10.1080/14780887.2016.1219435>

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł. ukasz, & Polosukhin, I. (2017). Attention is All you Need. *Advances in Neural Information Processing Systems*, 30. https://proceedings.neurips.cc/paper_files/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html

Wan, Y., Pu, G., Sun, J., Garimella, A., Chang, K.-W., & Peng, N. (2023). "Kelly is a Warm Person, Joseph is a Role Model": Gender Biases in LLM-Generated

Reference Letters (arXiv:2310.09219). arXiv. <https://doi.org/10.48550/arXiv.2310.09219>

Zaken, M. van A. (2023, December 27). Omgevingswet—Participatie—Video—Rijksoverheid.nl [Video]. Ministerie van Algemene Zaken. <https://www.rijksoverheid.nl/documenten/videos/2023/12/27/omgevingswet---participatie>

