



Master Thesis | 2026

AI as Digital Building Master

Structuring sustainable design choices in the early design phase of corporate real estate

Daphne van Noorden

Msc Management in the Built Environment
Theme: Gamechangers in Transitions
Delft University of Technology

Colophon

Author

Name D.J.M. van Noorden
Student number 4855183

Photo



Education

Institution Delft University of Technology
Faculty Architecture and the Built Environment
Master Track Management in the Built Environment (MBE)

Supervisors

First mentor Prof.dr. P.W. (Paul) Chan
Department Design and Construction Management

Second mentor Dr. A. (Aksel) Ersoy
Department Urban Development Management

Graduation organisation

Name Edge Technologies
Mentor ir. *Constantijn Berning*

This thesis is confidential and cannot be made public until 8 January 2026

An electronic version of this thesis is available at <https://repository.tudelft.nl/>

Preface

Dear reader,

It feels special to be able to write this. I am taking a moment to reflect on the journey it took to get here. It started at primary school in 's-Hertogenbosch and ends for me at Delft University of Technology. A new chapter in my professional life lies ahead, but I would like to take a moment to reflect on my academic career.

When I was graduating, I wanted to choose a subject that really spoke to me. Something that made me happy and that I could dive into with my full attention. I love challenges. Artificial Intelligence became just such a world. It started as something unknown and grew into a subject I wanted to devote my energy to. I discovered that I love understanding new systems and seeing how you can use them to build something that helps people.

I also realise how special it is to find your passion during your academic career. For me, that passion is the Built Environment. The combination of technology, people and creativity suits me perfectly. It is a perfect balance where you must switch between technology and feeling. It is a fine line between knowledge and intuition. That gives me energy.

Writing this thesis was not just an academic process. It was also a personal journey. My dyslexia made it difficult at times. Yet it also shaped me. I think visually. I see patterns quickly. I make connections that feel natural to me. For years, I saw that as a limitation. Now I see it as something that gives my work depth and creativity. This thesis reflects that style. It is research that reveals my way of thinking.

The collaboration with EDGE brought the research to life. The atmosphere is open and direct. I felt that my work was included in the real conversation. That gave me confidence and made the process valuable. I am grateful that I was given a place where I could learn and experiment.

And then there are the people around me. My parents, Monique and Jaap, stood by me every step of the way. They supported me through all the ups and downs throughout the entire journey. Their trust carried me through. My friends made this period light and warm. They helped me to keep going at times when it wasn't easy. Without them, I would have experienced this journey very differently.

Now that I am sitting here writing these words, I feel calm and proud. I look back fondly on what lies behind me. And I look forward to what lies ahead. This thesis shows who I have become. And where I want to grow.

Enjoy reading,

Daphne van Noorden

Amsterdam

January 2026

Abstract

Sustainable design choices in early-stage real estate development are often based on fragmented knowledge and varying definitions. Teams strive for ambitious environmental goals, but the underlying decision making remains vague. This thesis investigates how a single AI agent, acting as a Digital Building Master, can structure Environmental (E) decision making in the SO and VO phases. Rather than framing environmental design choices as trade-offs between competing objectives, the research explores how integration can reveal relationships and synergies across Environmental criteria in the early design phase.

The research follows a research-through-design approach in which a prototype was developed and tested within an iterative process that ran parallel to insights from practice. The results show that the agent does not add new knowledge. The agent organises existing standards and assumptions in a way that reveals hidden connections. The agent exposes implicit assumptions and identifies tensions at an early stage. The agent links each statement to verifiable sources. This changes the structure of design discussions. Decision-making shifts from linear logic to an iterative pattern in which analysis and sketch development take place simultaneously. This increases transparency and strengthens the stability of Environmental decision making. Interpretative judgement and final design decisions remain with the designer.

The presence of the agent also changes the dynamics between designers. Discussions remain open for longer and reflection becomes more in-depth. However, the structured output format can lead to rapid agreement before all alternatives have been explored. Its use requires sufficient AI skills within the team. The agent functions strongly within clearly defined Environmental themes but has less control over issues of a qualitative or political nature. This thesis concludes that the agent acts as an epistemic catalyst. The agent increases the reflective capacity of teams. The agent distributes cognitive load in a more stable manner and supports a clear basis for Environmental decision making in the early design phase.

Keywords

Environmental decision making | artificial intelligence | early design phase | digital building master | research-through-design | real estate development | decision-making | sustainability

Acknowledge

I would like to acknowledge the people who shaped the development of this thesis. Throughout the process, Prof. dr. Paul Chan provided guidance that brought clarity to each stage of the research. His questions created direction at moments when the project required structure and helped define how the study could progress with focus. Dr. Aksel Ersoy complemented this by recognising when the dynamic of the research drifted from its core. His reflections brought the work back to its essential line of decision making and strengthened the coherence of the approach.

My graduation period at EDGE was supported by ir. Constantijn Berning, whose involvement extended from the content of the thesis to the broader rhythm of the graduation process. His perspective made it possible to connect academic insights with the realities of design practice.

Within EDGE, Ryan and Bryan contributed essential technical support. Their work ensured that the prototype could operate within a stable digital environment and could be tested in a realistic setting.

The development of the AI component was further shaped by the expertise of ir. Pim Rutgers from NextSense. His explanations of agent behaviour and data interpretation deepened the technical foundation of the prototype.

I also want to acknowledge all professionals who participated in the interviews. Their openness and willingness to share their experience made it possible to understand how Environmental decision making unfolds in practice.

Their contributions formed the empirical backbone of the problem and solution space and allowed the research to take shape in dialogue with real design processes.

Executive Summary

In the early design phase of corporate real estate, choices are made that determine the environmental performance of buildings for decades to come. At this stage, information is incomplete. Design options are open. Assumptions often remain implicit. Sustainability ambitions are present but are usually expressed in broad terms such as Paris Proof or BREEAM Excellent. The translation of these ambitions into concrete design choices often remains invisible.

In practice, meanings shift as the design process progresses. This happens through daily design decisions, informal conversations and pragmatic considerations under time pressure. These shifts are rarely explicitly recorded. This makes it increasingly difficult to understand why certain environmental choices have been made and how they relate to previous ambitions.

This research focuses on the question of how environmental decision-making can be better supported in this early design phase. The core lies in understanding the decision-making process itself. Central to this is an AI agent designed as a Digital Building Master. This agent offers structured support by making assumptions explicit, organising interpretation and keeping decision-making traceable as long as there is still considerable design freedom.

The theoretical framework approaches ESG as an integrated value framework within corporate real estate. Environmental choices are formed early on and strongly influence later design and investment decisions. That is why the environmental dimension forms the analytical starting point of this research. Social and governance aspects serve as the context in which environmental ambitions are interpreted and legitimised. AI is approached as an interpretative tool that structures fragmented information and makes implicit assumptions explicit.

The research follows an abductive Research Through Design approach. Interviews with developers, designers and consultants identify where environmental decision-making fragments. Based on this, three analytical roles are formulated for the AI agent: sensing, steering and support. These roles describe different forms of support that can be provided by the system, depending on the moment, question and degree of uncertainty. The prototype is developed and tested in a realistic corporate real estate case involving PV design at EDGE Stadium. A validation workshop compares how design teams' reason with and without using the agent.

The findings show that environmental integration mainly gets stuck on interpretation. Responsibilities are divided. Concepts are used flexibly. Trade-offs between operational performance, embodied carbon and financial feasibility remain implicit. Certification frameworks provide structure, but mainly direct attention to what is measurable. As a result, ambitions shift without explicit reconsideration.

Within this context, the Digital Building Master takes on meaning as a structuring tool. Sensing refers to the explicit articulation of early assumptions and tensions

while design freedom is still high. Steering refers to the structuring and comparison of environmental reasoning during multidisciplinary discussions. Support refers to maintaining consistency in definitions, assumptions and sources as projects move through phases, documents and teams. These roles make environmental decision-making explicit and discussable, without reducing it to a single outcome.

The functional analysis shows that effective AI support revolves around structure. Context is organised. Interpretation is made explicit. Reasoning is linked to verifiable sources. Complexity is broken down without reduction. Memory is retained over time. These qualities are embedded in a layered architecture in which parallel environmental logics coexist.

The validation shows that the use of the agent changes the timing of decision-making. Interpretation is given more space in the early phase. Assumptions remain visible for longer. Parallel reasoning remains active. As a result, choices become more consistent and easier to justify later. Decision-making remains human. The structure changes.

At the same time, the research makes it clear that AI always remains dependent on input, source selection and framing. The system reinforces what becomes visible and stabilises chosen frameworks. This shifts interpretative influence to those who determine data, documents and prompts. The Digital Building Master makes this dynamic explicit but does not remove it. Integration is achieved through transparency and continuous reflection.

The limitations of the research are explicit. The research focuses on one area of environmental tension within Dutch corporate real estate. The findings are context-specific and process-oriented. The study does not make any statements about optimisation or performance improvement. Its value lies in providing insight into decision-making, not in outcomes.

In summary, this research shows that an AI agent can support environmental decision-making by structuring interpretation at a time when design choices are still open. The Digital Building Master contributes to transparency, coherence and accountability in early design processes. Its value lies in making explicit how sustainable design choices are made.

“The Internet will change everything.”
- Nicholas Negroponte (1990s)

The internet showed how technology not only supports but also rewrites our lives. It changed how we work, learn and live together, often in ways no one had foreseen.

Today, we face a similar upheaval with AI. AI promises to bring order to the complexity of our times but at the same time raises questions: which values do we make visible, which remain hidden, and who ultimately takes responsibility for the choices that are made?

Contents

1. INTRODUCTION	5
1.1 PROBLEM CONTEXT	5
1.2 PROBLEM FRAMING	6
1.3 RESEARCH QUESTIONS	7
1.4 SOCIETAL AND SCIENTIFIC RELEVANCE	8
1.5 APPROACH OVERVIEW	9
2. THEORETICAL BACKGROUND	12
2.1 ESG DECISION-MAKING IN DUTCH CORPORATE REAL ESTATE	12
<i>Corporate real estate</i>	12
<i>ESG framework</i>	13
<i>Interrelationship E-S-G</i>	16
<i>Key Takeaways: ESG in corporate real estate</i>	15
2.2 ARTIFICIAL INTELLIGENCE	16
<i>The Agent</i>	19
<i>LLM as cognitive engine</i>	20
<i>RAG as a contextual memory layer</i>	21
<i>Key Takeaways Artificial Intelligence</i>	22
2.3 THE INTEGRATION OF ESG AND AI	23
2.4 LIMITS AND RISKS OF AI IN ENVIRONMENTAL DECISION-MAKING	24
<i>Organisation and processes</i>	25
<i>Data infrastructure</i>	25
<i>Regulation and governance</i>	25
<i>Digital Sustainability</i>	25
2.5 KEY TAKEAWAYS CHAPTER 2	27
3. METHODOLOGY	30
3.1 RESEARCH THROUGH DESIGN METHOD	30
3.2 RESEARCH TECHNIQUES	31
3.3 ROUND 1: EXPLORATION (PROBLEM SPACE)	32
<i>Development of conceptual framework</i>	32
<i>Empirical research: Semi-structured interviews</i>	35
3.4 ROUND 1: EXPLORATION (SOLUTION SPACE)	36
3.5 ROUND 2: PROTOTYPE	37
<i>First step: reconstruction of human decision making</i>	37
<i>Second step: translation into a working workflow</i>	39
3.6 ROUND 3: VALIDATION	39
3.7 RESEARCH OUTPUT	42
<i>Deliverables and data set</i>	42
<i>Dissemination and audiences</i>	43
3.8 TRUSTWORTHINESS AND RIGOUR	44
3.9 ETHICAL CONSIDERATIONS	44
3.10 KEY TAKE AWAYS METHODOLOGY	45
4. FINDINGS IN THE PROBLEM SPACE	48
4.1 STRUCTURE OF THE PROBLEM SPACE	48
<i>First coding round: understanding environmental decision-making</i>	48
<i>Second coding round: interpreting AI and the three roles of the Digital Building Master</i>	49
4.2 THEMATIC ANALYSIS FIRST CODING ROUND	50
<i>ESG-Environmental in corporate real estate</i>	50
<i>Early Design Phase</i>	51

<i>The ESG feedback flow</i>	52
4.3 THEMATIC ANALYSIS: SECOND CODING ROUND	54
<i>AI System in Practice</i>	54
<i>The Digital Building Master</i>	56
4.4 CONTEXTUAL ACTIVATION OF SENSING, STEERING AND SUPPORT	57
4.5 KEY TAKE AWAYS H4	59
5. FINDINGS IN THE SOLUTION SPACE	65
5.1 FUNCTIONAL ARCHITECTURE OF THE AI AGENT	65
<i>Perception Layer</i>	65
<i>Decision making / Planning Layer</i>	66
<i>Tool Layer</i>	67
<i>Specialized Functions Layer</i>	68
<i>Memory / Provenance Layer</i>	68
5.2 OPERATIONAL REQUIREMENTS	70
5.3 KEY TAKE AWAYS H5	71
6. FINDINGS OF THE PROTOTYPE	73
6.1 TESTCASE	73
6.2 HUMAN WORKFLOW	74
6.3 AI WORKFLOW	75
6.4 HOW THE THREE ROLES BECOME VISIBLE IN THE GENERATED SCENARIO	79
6.5 KEY TAKEAWAYS H6	82
7. FINDINGS OF THE VALIDATION SPACE	84
7.1 THE STRUCTURE OF THE DESIGN WORKSHOP	84
7.2 HUMAN WORKFLOW	85
7.3 AI-WORKFLOW	ERROR! BOOKMARK NOT DEFINED.
7.4 COMPARISON AI AND HUMAN WORKFLOW	87
7.5 DIFFERENCES IN BEHAVIOUR AND DECISION-MAKING	88
7.6 STAKEHOLDER ROLES AND REDISTRIBUTION OF AGENCY	89
7.7 TECHNICAL AND ROLE-RELATED LIMITATIONS	91
<i>Limitations in sensing and perception</i>	91
<i>Steering without interpretation</i>	92
<i>Support, memory and continuity</i>	92
<i>Dependency, framing and strategic use</i>	93
7.8 SYNTHESIS	93
7.9 KEY TAKE AWAYS H7	95
8. CONCLUSION	97
8.1 ANSWER TO SUB-QUESTION 1	97
8.2 ANSWER TO SUB-QUESTION 2	98
8.3 ANSWER TO SUB-QUESTION 3	98
8.4 ANSWER TO SUB-QUESTION 4	99
8.5 ANSWER TO THE MAIN QUESTION	100
9. DISCUSSION	102
10. LIMITATIONS	108
11. RECOMMENDATIONS	112
11.1 RECOMMENDATIONS FOR PRACTICE	112
11.2 RECOMMENDATIONS FOR FUTURE RESEARCH	113

12.	IMPLEMENTATION STRATEGY	117
13.	REFLECTION	120
13.1	PRODUCT, PROCESS, PLANNING	120
13.2	METHOD CHOICE	121
13.3	LOOKING AHEAD	122
13.4	MASTER TRACK REFLECTIVE.....	122
13.5	ETHICAL REFLECTION.....	123
13.6	TRANSFERABILITY	124
13.7	EXTRA.....	125
14.	APPENDIX A: OVERVIEW INTERVIEWEES	132
15.	APPENDIX B: INTERVIEW PROTOCOL.....	133
16.	APPENDIX C: KEY INSIGHTS INTERVIEWEES	135
17.	APPENDIX D OVERZICHT CODE GROUPS ATLAS.TI	137
18.	APPENDIX E: EXAMPLE SUB CODES BY GROUP	138
19.	APPENDIX F: HUMAN WORKFLOW SOLAR PANEL.....	143
20.	APPENDIX G: COPILOT STUDIO AGENT	149
21.	APPENDIX H: WORKFLOW AI IN COPILOT	150
22.	APPENDIX I: EXAMPLE COPILOT RESPONSE	154
23.	APPENDIX J: DATA MANAGEMENT PLAN.....	157
24.	APPENDIX K: CONSENT FORMS	164
25.	APPENDIX K: OVERVIEW CHUNKS.....	167
26.	APPENDIX L: SOURCES COPILOT STUDIO	168
27.	APPENDIX M: OPERATIONAL AND EMBODIED CARBON AI	169
28.	APPENDIX N: OVERVIEW OF PARTICIPANTS IN THE VALIDATION PHASE	170
29.	APPENDIX O: REQUIREMENTS EACH LAYER AI	171
30.	APPENDIX P: DESIGN WORKSHOP MATERIAL	173

List of table and figures

Figure 1: overview of the research structure connecting exploration, prototyping and	10
figure 2: integrating sustainability into financial decision-making in corporate real estate. (macneil & esser, 2022)	12
figure 3: macleamy curve (davis, 2011).....	18
figure 4: ai agent workflow (huang & ai, 2023)	19
figure 5: large language model, (xi et al., 2023)	20
figure 6: workflow rag (own work)	22
figure 7: number of datacentres worldwide based on iea, 2025	26
figure 8: increase in electricity demand by sector (base case 2024-2030)	26
figure 9: development of the digital building master roles (own work)	35
figure 10: translation of the digital master builder's conceptual roles into the technical layers of the ai system (own work, 2025).....	36
figure 11: rooftop configuration of edge stadium used as basis for the pv trade-off case. (edge, 2025)....	38
figure 12: research methods overview (own work, 2025)	41
figure 13: open coding themes (atlas.ti, 2025)	48
figure 14: coding round 2 (atlas.ti, 2025)	50
figure 15: esg feedback loop (own work, 2025).....	53
figure 16: impact - project phase each role (own work)	57
figure 17: new esg feedback loop (own work, 2025).....	63
figure 18 functional architecture of the ai agent across roles (own work , 2025)	69
figure 19: functional architecture of the ai agent across roles (own work).....	69
figure 20: conversation flow for ending or continuing a dialogue in an ai agent (copilot, 2025)	76
figure 21: dialogue between ai agent and human (own work, 2025)	76
figure 22: the effect of an ai agent on decision making (own work. 2025).....	89
table 1: limits and risks of ai (own work).....	25
table 2: research techniques (own work)	32
table 3: requirement framework outlining the ai-agent's environmental roles and functions. (own work, 2025)	70
table 4: chunk example (own work, 2025)	77
table 5: sense, more in appendix k	ERROR! BOOKMARK NOT DEFINED.
table 6: steering, see more in appendix k	80
table 7: support, see more in appendix k	81
table 8: key take aways h7 do not indicate system failure but demonstrate how specific architectural and configurational choices directly shape interaction patterns, role dynamics and environmental decision making.....	94

Acronyms and glossary

AI agent

A computational system that interprets and structures information within predefined boundaries. In this thesis, the agent supports the visualisation of assumptions and Environmental dependencies without making decisions.

BIMI

The AI agent developed in this thesis. BIMI functions as a Digital Building Master by structuring Environmental reasoning, making assumptions explicit and supporting design discussions without taking over decision-making.

Corporate Real Estate (CRE)

Real estate owned or leased by an organisation for its own operational use, such as offices, factories or warehouses, primarily to support corporate strategy and core activities.

Design team

The group of developers, architects, consultants and technical specialists involved in early design decisions. The diversity of perspectives within the team contributes to interpretative differences in Environmental decision-making.

Digital Building Master

The interpretative function of the AI agent. The Digital Building Master structures Environmental decision-making in the early design phase by making assumptions, reasoning and dependencies explicit, while leaving responsibility with human actors.

Early design phase (SO/VO)

The project phase in which design choices are still open and have the greatest influence on long-term Environmental performance. Fragmentation in this phase often leads to later design lock-ins.

Embodied carbon

CO₂ emissions associated with material extraction, production, transport, construction, maintenance and end-of-life processes of building components.

Environmental (E)

The Environmental dimension of ESG. In this thesis, the focus is on the relationship between operational energy use and material-related emissions.

Human decision-making processes

The way in which design teams interpret information, form assumptions and justify sustainable design choices. These processes are often implicit, distributed and context dependent.

LLM (Large Language Model)

The language model forming the interpretative core of the AI agent. The LLM recognises patterns in text and supports the structuring of Environmental reasoning based on available inputs.

Operational energy

Energy consumed during the use phase of a building. Together with embodied carbon, this forms the central Environmental tension examined in this thesis.

RAG (Retrieval-Augmented Generation)

The system layer that connects the LLM to project-specific documents and external sources. RAG constrains interpretation to current, verifiable information.

Sensing

The role of the AI agent focused on revealing implicit assumptions and early Environmental tensions while design freedom is still high.

Steering

The role of the AI agent focused on structuring and comparing Environmental alternatives without closing design choices.

Support

The role of the AI agent focused on maintaining traceability and transparency of Environmental reasoning during design discussions.

Acronyms

AI	Artificial Intelligence
BIMI	Name of the AI agent developed in this thesis
CRE	Corporate Real Estate
E	Environmental (dimension of ESG, used as analytical focus)
ESG	Environmental, Social and Governance
LLM	Large Language Model
RAG	Retrieval-Augmented Generation
SO	Concept Design
VO	Preliminary Design
PV	Photovoltaic (solar panels)

Chapter 01

Introduction

- 1.1 Problem context
- 1.2 Problem framing
- 1.3 Research questions
- 1.4 Relevance
- 1.5 Approach overview

1. Introduction

This chapter sets the context for the research. It explains why Environmental decision making in Dutch corporate real estate requires closer study and outlines its societal and scientific relevance. It introduces the main research question and sub-questions. The chapter closes with a brief overview of the methodological approach.

1.1 Problem Context

In the early design phase of Dutch corporate real estate development, sustainability is often articulated as a shared ambition, yet its meaning remains diffuse. Design teams work with different assumptions regarding which environmental impacts should be prioritised, how material choices affect performance, and how future emissions are weighted. These assumptions implicitly structure decisions, but they are rarely made explicit and therefore difficult to negotiate. Early design choices, however, have disproportionate consequences: they influence material composition, energy demand and long-term CO₂ emissions, in a sector that accounts for roughly 40% of total energy consumption and 36% of CO₂ emissions in Europe (IEA, 2025). Moreover, environmental performance increasingly influences market valuation and access to financing, making sustainability not only a normative aspiration but also a strategic necessity. (Larisa Olteanu & Elena Ionaşcu, 2023)

When underlying values remain unspoken, stakeholders optimise for different interpretations of what constitutes environmental performance without establishing a shared direction. This is particularly consequential in the early design phase, where choices are still flexible, but decision-making occurs under uncertainty. As Kolltveit & Grønhaug, (2004) highlight, decisions embedded in this phase become structurally “locked-in,” constraining future alternatives. The challenge is therefore not a lack of data but the absence of mechanisms to make value considerations visible while design directions are still open.

It is therefore logical that design teams use digital systems to gain control. However, these systems offer only limited support at this stage. Life Cycle Assessment (LCA) is often only applied once design choices have already been fixed, thereby confirming existing decisions rather than keeping alternatives open (Kumar et al., 2025) BIM facilitates information exchange but does not generate shared interpretative frameworks; disciplines continue to operate in parallel (Motalebi et al., 2025a).

Given that the problem lies not in a lack of data but in the absence of systems that reveal underlying assumptions at a time when design choices are still open, this research focuses on how these assumptions can be made visible in the early stages of the design process, focusing on interpretation of sustainable decision making.

1.2 Problem framing

Against this backdrop, this research studies whether Artificial Intelligence can support the visibility of sustainable design choices in the early design phase and make it more open to discussion. This potential is analytically explored through the development and testing of an AI agent. The agent does not make design choices. Instead, it supports the conversation by structuring how different design approaches relate to Environmental performance. In doing so, it makes visible where interpretations diverge and which assumptions guide early choices. The aim is to examine whether this form of support can strengthen decision-making at a moment when choices remain open and consequences are still uncertain.

The relevance of AI is explored through the specific capabilities of large language models compared to earlier digital systems. These models can work with partial and qualitative information and relate fragmented considerations to one another. (Huang & Ai, 2023) In this research, this capacity is examined as a possible means to make sustainable design choices explicit without closing alternatives. The agent is therefore positioned as a reflective layer rather than a tool for confirmation or optimisation. Human judgement remains leading.

This interpretative role is conceptualised through the metaphor of the Digital Building Master. Historically, the master builder integrated material, spatial and technical considerations into a coherent design vision. (Allard Jolles et al., 2021) In contemporary practice, this integrative capacity is fragmented across specialised disciplines, making shared meaning less self-evident. The Digital Building Master does not restore centralised control but serves as an analytical lens to study how interpretative assumptions can be made visible within collaborative design processes.

The empirical focus is on Dutch corporate real estate, where buildings function as strategic assets. In this domain, Environmental performance directly influences financing conditions, risk assessment and market positioning (Larisa Olteanu & Elena Ionaşcu, 2023). Sustainability ambitions are often articulated through ESG frameworks, which provide a shared language for accountability and decision-making. This thesis focuses explicitly on the Environmental (E) dimension, as early design choices have long-term implications for CO₂ emissions, energy consumption and material flows. (Motalebi et al., 2025b)

The study first analyses how design teams interpret and handle sustainable design choices. It then examines how the Digital Building Master may support more transparent conversations by making underlying assumptions explicit while design choices are still open.

1.3 Research questions

To investigate how an AI agent might function as a Digital Building Master, this research develops a prototype to explore its potential influence on sustainable design choices and collaborative decision-making in the early design phase. Following a Research Through Design approach, the agent is designed, tested and refined through iterative cycles. These iterations are used to examine how the presence of the agent affects shared decision-making processes, including which assumptions surface, how Environmental considerations are discussed, and how dialogue between disciplines evolves. Each phase of the study is therefore linked to a sub-question that makes these changes analytically traceable. The main research question guiding this study is:

How can an AI agent, operating as a Digital Master Builder, structure human Environmental (E) decision making processes in the early design phase of Dutch corporate real estate development?

This question is addressed through the following sub-questions:

1. How are Environmental (E) considerations currently interpreted and handled by design teams, and where does fragmentation create space for the AI agent to act as a Digital Building Master?
2. What functional qualities should an AI agent possess to support the structuring of Environmental (E) considerations in early design discussions?
3. How can these qualities be translated into a prototype AI agent that can be explored and evaluated in practice?
4. How does the presence of the AI agent influence how design teams articulate, compare and reflect on Environmental (E) considerations during decision-making?

Rather than testing a predefined effect, this study proceeds from the analytical expectation that when an AI agent is introduced early in the design process, Environmental (E) considerations may become more explicit and easier to discuss. By drawing on shared, project-specific information, the agent may help clarify how design choices relate to embodied emissions and energy performance while choices are still flexible. In this process, the agent does not replace expert judgement but may support collective decision making by making underlying assumptions and interpretations visible.

1.4 Societal and scientific relevance

This research examines the potential role of an AI agent in bringing greater clarity to the early design phase of Dutch corporate real estate. In this phase, design teams make sustainable design choices while key parameters remain undetermined, and consequences are not yet fully traceable. A digital interpretative system may help to make underlying assumptions visible and show where perspectives diverge. Rather than resolving these differences, the system supports reflection by structuring how arguments are articulated and compared. In this context, AI is not positioned as a system that calculates or optimises outcomes, but as a reflective layer that helps design teams understand what is at stake while choices are still open.

The scientific contribution of this study lies in its focus on decision-making under complexity rather than on technical optimisation. While much existing research approaches AI as a tool to accelerate processes or automate tasks, this study examines how an AI agent can influence interpretation, dialogue and collaboration in situations characterised by uncertainty. It shows how the presence of an AI agent can make interpretative frictions explicit and alter the structure of conversations when sustainable design choices are still fluid. In doing so, the study contributes to theoretical understanding of collective decision-making processes in the built environment, particularly where assumptions remain implicit and decisions are not yet stabilised.

The societal relevance of the research relates to the need for more transparent and explainable Environmental decision-making in the built environment. In the early design phase, choices are made that shape long-term emissions and environmental performance, yet the assumptions underlying these choices often remain unspoken. This can lead to lock-in effects and limited consideration of alternatives. By supporting the explicit articulation of Environmental arguments at an early stage, a digital interpretative layer may help design teams to better justify and communicate their choices in response to increasing regulatory, financial and societal demands surrounding sustainable buildings.

1.5 Approach overview

The research questions introduced in Section 1.3 guide the overall research design and structure the study across its different phases. The aim of the research is to use the Environmental (E) dimension as an analytical entry point to examine how Artificial Intelligence can support interpretation in early design decision-making. AI is therefore applied to explore how sustainability considerations are structured in the early phase of real estate development. The research follows a Research Through Design approach combined with abductive decision making, in which understanding develops through iterative cycles of making, testing and reflection (Stappers, 2017) Following (Dubois & Gadde, 2002) idea of systematic combining, the study moves back and forth between empirical observations and theoretical insights, allowing both the problem framing and analytical focus to evolve during the research process.

As illustrated in Figure 1, the research is organised into three phases: Exploration, Prototype and Validation. The Exploration phase consists of two closely connected components. SQ1 examines how design teams currently interpret and handle sustainable design choices. SQ2 builds on these insights by formulating exploratory requirements that inform the design of the AI agent as a Digital Building Master.

These requirements are examined through three interrelated roles: Sensing, Steering and Support, which are explored as situational modes rather than fixed functions. The Prototype phase (SQ3) translates this analytical framing into an iterative development process, starting from a manual baseline and refining a structured workflow through expert feedback. The aim is not optimisation, but to examine how a digital agent can support reflection, make assumptions explicit and help maintain openness in the design space while sustainable design choices are still unsettled.

The Prototype phase (SQ3) forms the centre point between theory and observation. The requirements formulated in SQ2 are translated into an iterative development process. This begins with a manual baseline study of how design teams currently weigh environmental (E) trade-offs. This baseline is converted into a structured workflow based on the Sense-Steer-Support logic and then refined in consultation with sustainability experts and technical specialists. The aim of this phase is not optimisation, but to investigate how a digital agent can support reflection, make assumptions explicit and keep the design space open if the Environmental (E) decision making has not yet reached an impasse.

The Validation phase (SQ4) is on the empirical side of the diagram and explores how the AI agent influences Environmental (E) decision making. The goal is not to finalise the agent as a product, but to gain insight into its appropriate role, value and limitations within early design situations. Before turning to the detailed description of each research phase, the next chapter provides the necessary background on the domains of ESG and AI.

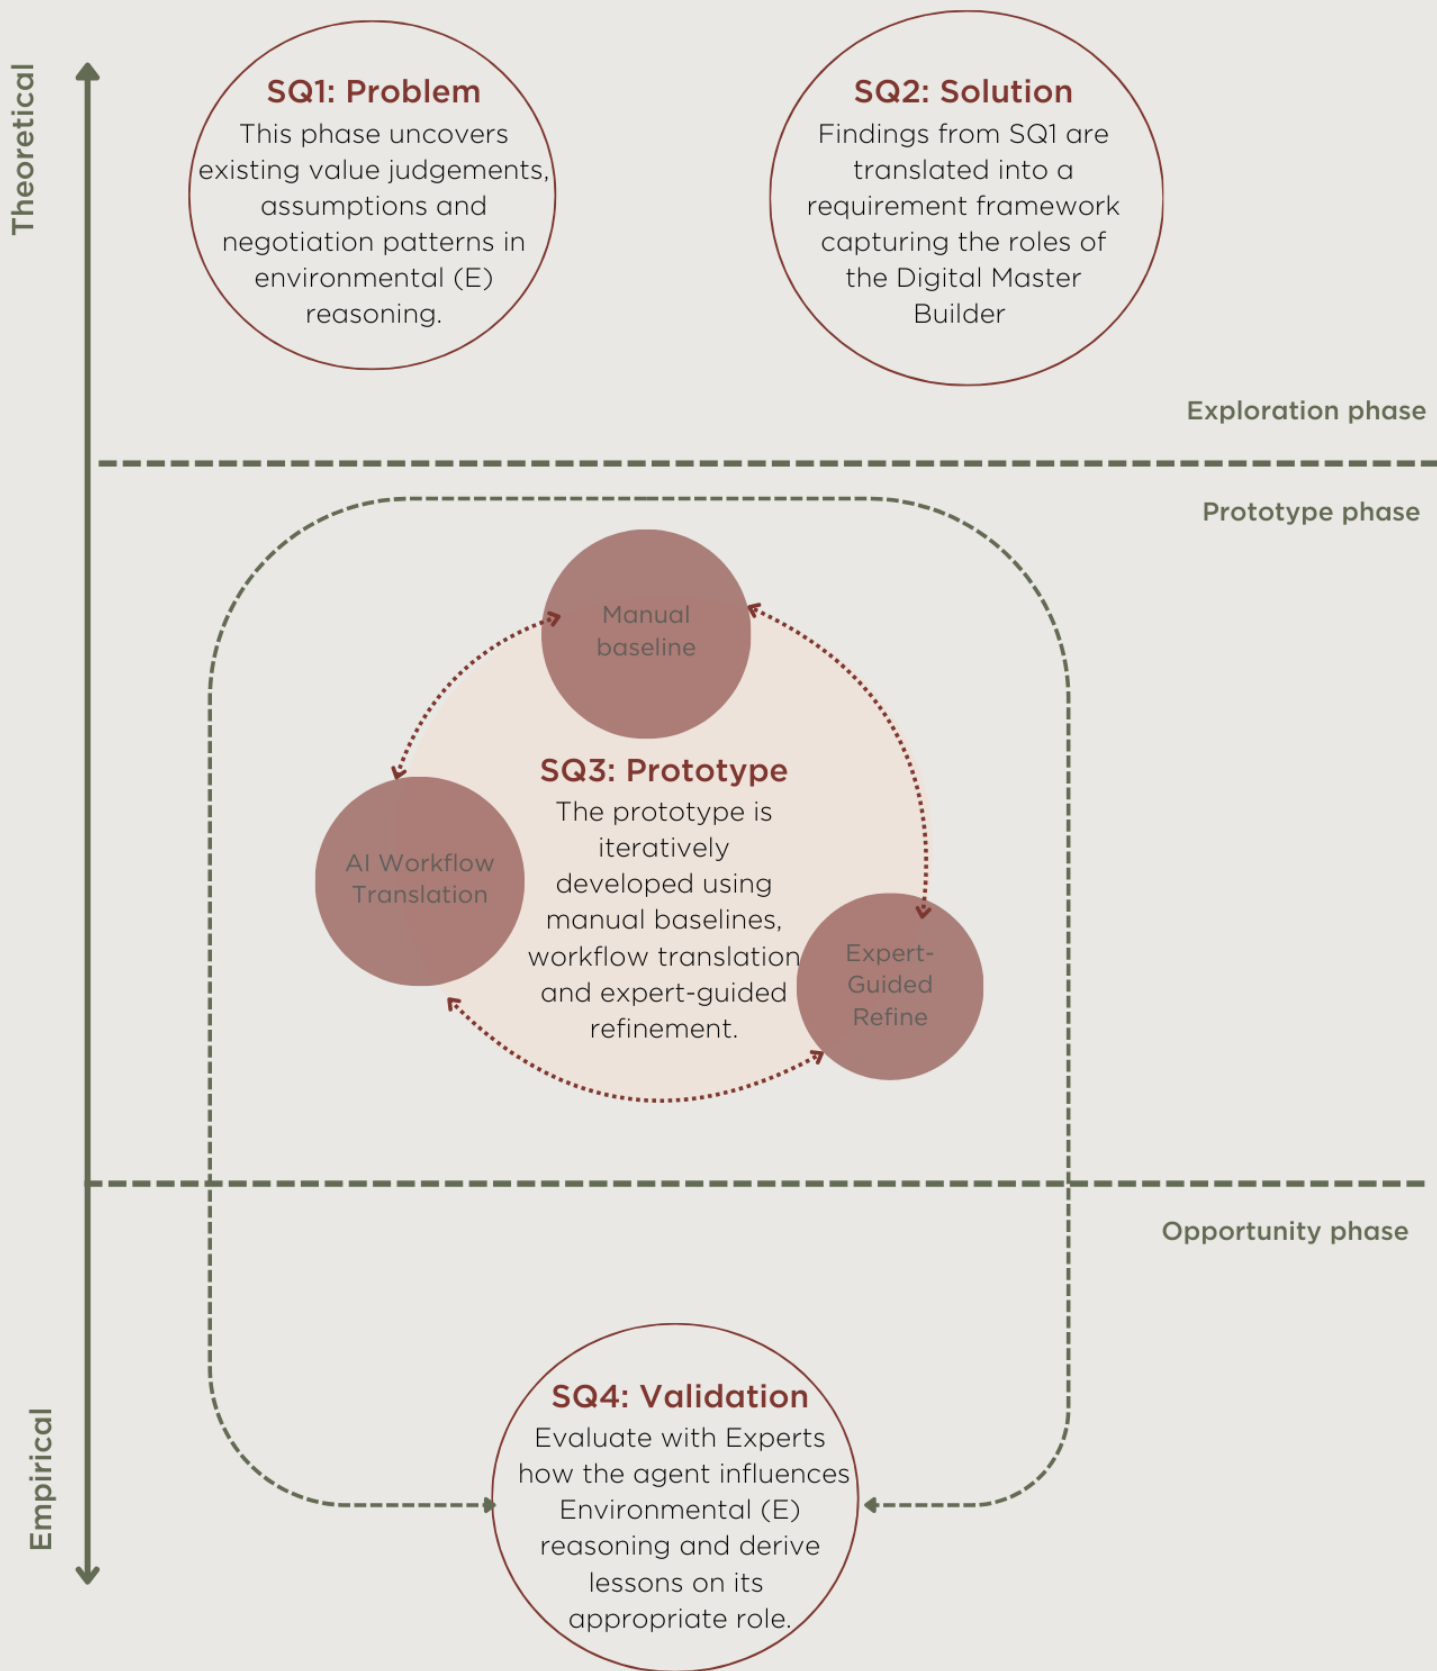


Figure 1: Overview of the research structure connecting exploration, prototyping and validation across the four sub-questions. (own work, 2025)

Chapter 02

Theoretical background

2.1 ESG decision-making in Corporate Real Estate

Corporate Real Estate

ESG framework

Interrelationship E-S-G

Key takeaways ESG in corporate real estate

2.2 Artificial Intelligence

The Agent

LLM as cognitive engine

RAG as a contextual memory layer

Key takeaways Artificial Intelligence

2.3 The integration of ESG and AI

Risks of AI in E decision-making

2.4 Key Takeaways theoretical background

2. Theoretical Background

This chapter sets out the theoretical background of the research by outlining ESG and Artificial Intelligence as two distinct but related domains. The focus lies on sustainable design choices in Dutch corporate real estate, with the Environmental (E) dimension used as the main analytical entry point. By discussing these domains separately, the chapter clarifies how sustainable design choices are interpreted and structured in early decision-making, and how AI can be used as an analytical lens to make underlying assumptions and differences in interpretation more explicit.

2.1 ESG decision-making in Dutch corporate real estate

Corporate real estate

To better understand ESG decision-making, it is important to understand what corporate real estate entails. According to the Cambridge Dictionary, (2025) corporate real estate (CRE) involves the physical assets that an organisation owns or leases to support its activities. Academic definitions emphasise the same core principle: CRE comprises all real estate held or used by an organisation to facilitate its operational, financial and strategic objectives (Miller & Pogue, 2018) CRE decisions are therefore of strategic importance: the timing of renovations, the suitability of the space and energy systems have a direct impact on the operational continuity and performance of the organisation. Unlike stand-alone development projects, corporate real estate decisions are typically made within a portfolio context, where individual assets are assessed in relation to one another and to long-term organisational objectives.

The meaning of corporate real estate has changed significantly over the past decade. Whereas CRE decisions were traditionally based on the financial triangle of return, risk (security) and liquidity, sustainability has now become a fourth axis

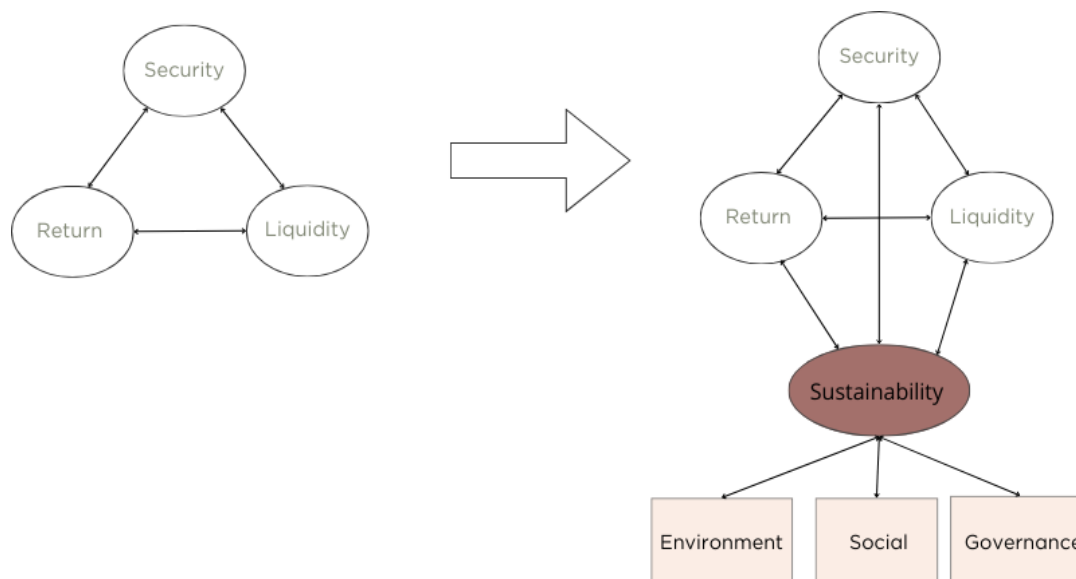


Figure 2: Integrating sustainability into financial decision-making in corporate Real estate. (MacNeil & Esser, 2022)

that redefines how value is justified. Figure 1 from MacNeil & Esser (2022) illustrates this shift: sustainability does not replace financial metrics but repositions them within a broader evaluation framework. This wider trend is also reflected in real estate valuation literature, which shows that environmental and social performance increasingly influence perceived value, access to capital and long-term risk profiles (Riratanaphong & Pewklieng, 2025) In a corporate real estate context, this shift is particularly consequential, as buildings are long-lived, capital-intensive assets and early design choices can set standards and lock in performance characteristics across multiple assets over time.

ESG framework

The emergence of Environmental, Social and Governance (ESG) frameworks reflects a shift in how organisations define value and risk in corporate real estate decision-making. ESG can be understood as an integrated decision-making framework that translates sustainability-related risks and opportunities into measurable and comparable criteria (Chopra et al., 2024) Rather than replacing financial considerations, ESG reframes how performance and long-term value are assessed, as also illustrated in Figure 2. Although ESG comprises three interrelated dimensions, this research focuses primarily on the Environmental dimension. In corporate real estate, Environmental considerations play a central role in how sustainable design choices are interpreted and prioritised, particularly in the early design phase where assumptions remain implicit and decisions have long-term consequences (Larisa Olteanu & Elena Ionaşcu, 2023). The Environmental dimension therefore provides a focused analytical entry point for examining how sustainability is reasoned about and structured in practice.

The Environmental Dimension

In corporate real estate, the Environmental dimension concerns how buildings contribute to an organisation's environmental performance over their entire life cycle. This includes not only energy use during operation, but also emissions associated with the production, transport, use and replacement of materials. In practice, Environmental objectives are therefore not simply applied as fixed criteria, but are interpreted and negotiated in design processes shaped by certification logics, market expectations and feasibility constraints.

Recent literature highlights that focusing solely on operational energy efficiency provides an incomplete picture of Environmental performance. Tirelli & Besana, (2023) show that buildings with low operational energy demand can still have a high overall environmental impact when material-related emissions are substantial. As operational energy use decreases, the relative importance of embodied emissions increases, making life-cycle considerations more prominent in early design decision-making (Palomar-Torres et al., 2025) Their research indicates that material-related emissions can constitute the largest share of total climate impact, underscoring the importance of material selection and reuse strategies in early project stages.

While the Environmental dimension can be defined more broadly to include aspects such as water use, ecological value and climate adaptation, this research focuses on how operational energy and material-related emissions are prioritised and weighed in the early design phase. These considerations account for a significant share of environmental impact and are repeatedly subject to interpretation and prioritisation in design and investment decisions. Focusing on this area allows the study to examine how Environmental considerations are structured and discussed when sustainable design choices are still open.

The Social Dimension

Although the Social dimension is not the primary analytical focus of this research, it forms an important contextual layer in how Environmental ambitions are interpreted and experienced in corporate real estate. The social value of buildings relates to how the working environment supports health, well-being and daily user experience. Research shows that aspects such as indoor environmental quality, spatial configuration and user autonomy influence comfort, productivity and satisfaction. Research into Quality of Work Life (QWL) did by Krizmanić & Milovanović, (2025) shows that these social performances are determined by a combination of physical, psychological and social factors.

These social aspects intersect directly with Environmental design choices. Strategies aimed at reducing energy use or material impacts can affect thermal comfort, daylight access and spatial layouts, shaping how users experience and accept sustainability measures. In this way, the social dimension influences how Environmental priorities are evaluated and legitimised in practice, even though it is not examined in depth in this study. The Social dimension is therefore treated as a contextual factor that shapes the conditions under which Environmental design choices are discussed and implemented in corporate real estate.

The Governance dimension

Although governance is not the primary analytical focus of this research, it provides the institutional framework within which Environmental ambitions in corporate real estate are defined, translated and monitored. Governance concerns the structures, responsibilities and procedures that shape how decisions are made, justified and accounted for across real estate portfolios. Riratanaphong & Pewklieng, (2025) emphasise that governance determines how sustainability is organised: who is responsible, what steps are mandatory and how transparency is ensured. According to Jebril et al., (2024) legal compliance, financial management and operational control of buildings. Their work shows that consistent decision-making and clearly assigned responsibilities support reliable value development and prevent fragmented or inconsistent choices.

Certification systems such as BREEAM and WELL give practical form to sustainability ambitions by translating abstract intentions into measurable requirements, such as documentation, performance thresholds and verification

points Grzegorzewska & Kirschke, (2021). In practice, these systems function as governance instruments because they standardise what a building must demonstrate and enforce transparency throughout design and management processes.

Governance also plays a key role in managing risk and long-term value. Clear responsibilities, monitoring procedures and accountability structures help organisations anticipate regulatory changes and avoid fragmented decision-making (Jebril et al., 2024). This is particularly relevant in Dutch corporate real estate, where governance is strongly shaped by national regulations, reporting obligations and public accountability requirements. As a result, Environmental decision making is not only an internal strategic concern but is continuously mediated by regulatory expectations and compliance frameworks. This makes Dutch corporate real estate a relevant context to study how Environmental ambitions are structured, constrained and legitimised in practice.

Key Takeaways: ESG in corporate real estate

Research on Environmental performance in corporate real estate consistently shows that sustainable design choices are shaped by the interaction between operational energy use and material-related emissions. Tirelli & Besana (2023) demonstrate that buildings optimised for low operational demand can still result in higher life cycle impacts due to embodied emissions. Palomar-Torres et al. (2024) reinforce this by showing that, in high-performance buildings, embodied impacts increasingly dominate total climate impact. Together, these studies underline the importance of a life-cycle perspective, positioning the Environmental dimension as a key integrative domain where material and energy strategies must be considered together rather than separately.

The way Environmental ambitions are evaluated in practice is influenced by the social dimension. Felgueiras et al. (2023) show that thermal comfort, air quality, acoustics and daylight directly affect cognitive functioning and well-being. As Environmental design choices inevitably shape these conditions, user expectations play a mediating role in how sustainability measures are perceived and legitimised within corporate real estate projects.

Governance research explains how these Environmental and Social considerations are formalised and stabilised. Grzegorzewska & Kirschke (2021b) show that certification systems such as BREEAM and WELL translate sustainability ambitions into measurable procedures and verification steps. Riratanaphong & Pewklieng (2025) further demonstrate that organisational governance structures determine whether such ambitions can be applied consistently across projects. Governance therefore functions not as a separate sustainability pillar, but as the institutional mechanism through which Environmental and Social priorities become actionable and comparable.

Sector-level studies help explain why integrated decision making is challenging in early design stages. Ifediora et al. (2024) show that developers operate within a negotiation space shaped by actors with differing financial, regulatory and spatial priorities, requiring continuous reinterpretation of sustainability ambitions. Kempeneer et al. (2021) add that operational performance remains partly dependent on user behaviour, reinforcing the interdependence between Environmental and Social outcomes. These dynamics do not make integration impossible but highlight the need for structures that support alignment rather than isolated optimisation.

The timing of decision-making is critical in this matter. The MacLeamy curve (Davis, 2011) shows that the potential to influence Environmental performance is greatest when design flexibility is high and cost commitments are still limited. This confirms that the early design phase offers the primary opportunity to align Environmental intentions, user expectations and governance requirements before decisions become fixed.

This theoretical framework therefore treats ESG not as a set of competing objectives, but as an integrated value structure in which Environmental, Social and Governance considerations are aligned through early-stage interpretation rather than resolved through trade-offs.

Interrelationship E-S-G

Although the Environmental, Social and Governance dimensions can be distinguished conceptually, their roles in early design decision-making are not equivalent. In the early design phase of corporate real estate, Environmental considerations form the primary analytical entry point, as material choices, energy concepts and spatial principles are first articulated at this stage. These early decisions shape the direction of the project while key parameters remain open. At the same time, the interpretation of Social and Governance considerations is still largely dependent on the Environmental strategies under development, as comfort expectations only become concrete once material and energy concepts are defined, and governance requirements become operational once it is clear what needs to be monitored or verified.

In practice, Environmental, Social and Governance considerations interact rather than operate as separate decision tracks. Environmental strategies aimed at reducing emissions directly influence how buildings are experienced by users, while organisational structures determine how these ambitions are translated into formal requirements. Literature on corporate real estate governance shows that procedural clarity, responsibilities and performance requirements shape the space within which Environmental ambitions can be implemented and legitimised (Riratanaphong & Pewklieng, 2025). In this sense, Social and Governance dimensions function as contextual filters through which Environmental priorities are interpreted and formalised.

The interaction between Environmental and Social considerations becomes particularly visible when design choices aimed at lowering energy demand or reducing material-related emissions affect indoor environmental quality. Factors such as thermal comfort, air quality, acoustics and daylight access have documented effects on cognitive performance, well-being and productivity (Felgueiras et al., 2023). Environmental interventions are therefore inseparable from how users experience the building, and social expectations influence whether such measures are perceived as enhancing or diminishing workplace quality.

Governance mechanisms subsequently shape whether these integrated Environmental and Social values can be sustained over time. Certification systems such as BREEAM or WELL translate intended performance related to carbon, materials and user well-being into documented requirements and verification processes, thereby providing a procedural structure for implementation and monitoring (Grzegorzewska & Kirschke, 2021b). Financial evaluation frameworks further influence decision-making, as life-cycle carbon and life-cycle cost assessments are often considered jointly in feasibility analyses, affecting which Environmental strategies are regarded as viable (Wang et al., 2024). Governance thus operates not as a separate sustainability pillar, but as the institutional mechanism through which Environmental priorities are aligned, justified and constrained.

Rather than framing ESG dimensions as competing objectives that require trade-offs, this research approaches sustainable design choices as the outcome of an integrative process. Environmental, Social and Governance considerations become aligned through interpretation and structuring in the early design phase. While alignment remains challenging in practice due to differing risk perceptions, responsibilities and strategic priorities among actors (Ifediora et al., 2024), this complexity reinforces the need for approaches that support integration rather

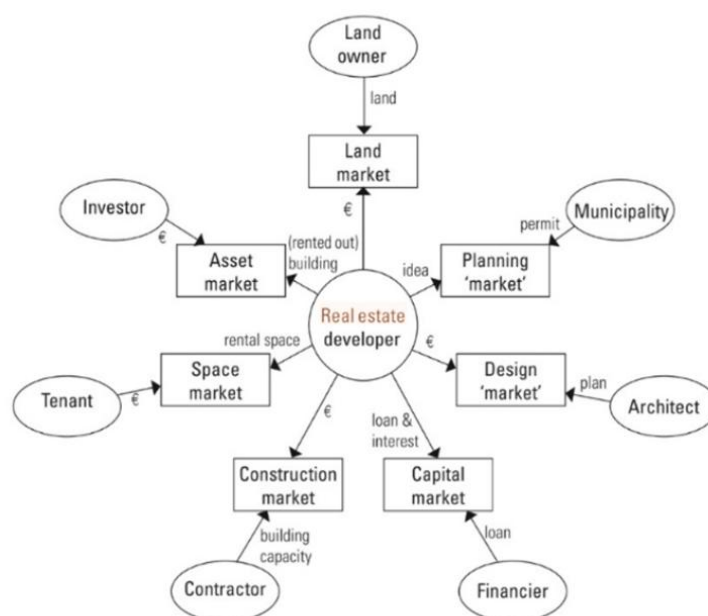


Figure 3: Stakeholders real estate ((Ifediora et al., 2024),

than optimisation. Within the Environmental domain, this is particularly evident in the need for a life-cycle perspective, in which operational energy use and embodied emissions are evaluated together rather than separately (Tirelli & Besana, 2023)

Figure 3 illustrates how the Dutch corporate real estate sector operates as an interdependent negotiation space in which the developer is embedded within multiple markets, including land, planning, design, capital, construction, assets and space. Each market is characterised by its own incentives, constraints and temporal logics. Financiers prioritise risk and long-term value protection, municipalities emphasise compliance and public-interest criteria, and contractors operate within capacity and cost constraints. Environmental design choices, such as decisions on material intensity or energy concepts, therefore, affect several of these domains simultaneously. As a result, sustainability ambitions must be continuously interpreted and aligned across different institutional and market logics rather than negotiated in isolation.

The importance of early integration becomes more pronounced when considering how design influence evolves over time. Figure 3 incorporates the MacLeamy curve (The MacLeamy Curve, 2004) , which shows that the ability to influence Environmental outcomes is highest during the pre-design and schematic design phases, before major financial and contractual commitments are fixed. As projects progress, design flexibility decreases while the cost and impact of changes increase. When Environmental ambitions are introduced only in later phases, they tend to be addressed as constrained optimisations rather than as guiding design principles.

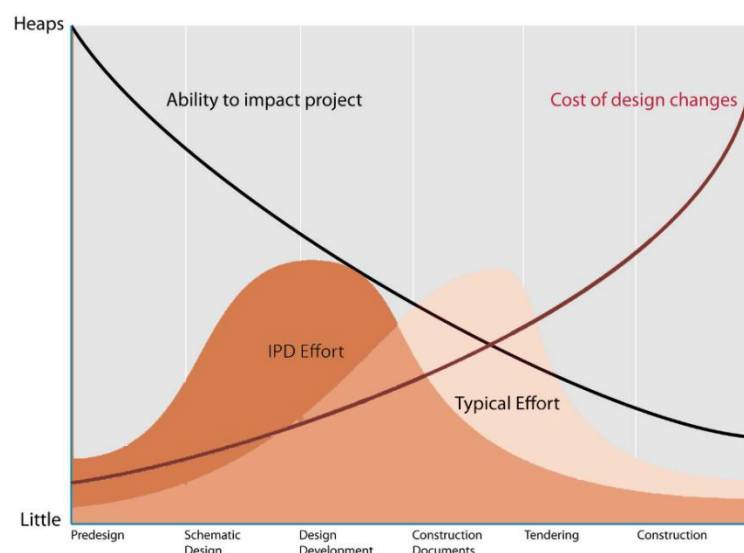


Figure 4: MacLeamy curve (Davis, 2011)

Taken together, these dynamics show that sustainability in corporate real estate functions as an interdependent value framework that is most coherent in the early design phase. At this stage, Environmental intentions, user expectations and

governance requirements can still be interpreted together and aligned. When integration occurs later in the process, sustainability risks being reduced to incremental optimisation instead of shaping the strategic direction of the project.

2.2 Artificial intelligence

Here, AI is not seen as a system that takes over decisions, but as a form of reflective support: a Digital Building Master that helps design teams clarify their decision making and structure the complexity of sustainability decision making. To understand this role, it is important to consider how the concept of AI has developed.

The term AI was introduced by John McCarthy in 1956, who defined it as the possibility that learning and decision making could be described so precisely that a machine could simulate it (Panesar, 2020)

Since then, AI has developed from a theoretical notion into a set of computational techniques capable of interpreting data, recognising patterns and generating meaning rather than merely calculating outcomes.

The Agent

A central concept in contemporary AI research is that of the agent. Initiating from philosophical discussions on agency, an agent is a system that can perceive its environment, interpret information and act within defined boundaries (Huang & Ai, 2023). Figure 4 illustrates this cycle of perception, decision making and action. Four characteristics are fundamental: autonomy (acting independently within constraints), reactivity (responding to inputs), proactivity (anticipating rather than merely reacting) and social interaction (communicating with users and other systems). In digital design environments this capacity to mediate between sources of knowledge is particularly relevant.

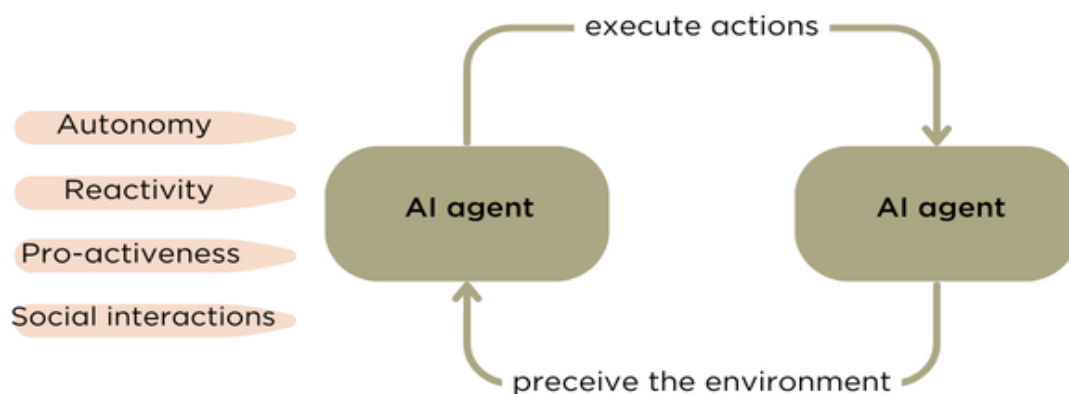


Figure 5: AI agent workflow (Huang & Ai, 2023)

In this research, the agent functions as an interface between humans and technical information systems. Its cognitive core is formed by a Large Language Model (LLM), which is trained on extensive text amounts and can recognise conceptual relationships, generate explanations and support decision making (Guo et al., 2024) Rather than offering definitive answers, the LLM enables the agent to surface patterns that may otherwise remain implicit in fragmented documentation or dialogue. This interpretive capacity forms the basis for the agent’s role as reflective support: it clarifies the structure of a decision context, without determining the outcome.

LLM as cognitive engine

The cognitive core of the agent is a so-called Large Language Model (LLM). These models are trained on huge amounts of text and learn to recognise patterns in it. Their operation is based on the principle of next token prediction: the model predicts each time which word is most likely to follow the previous words (Xi et al., 2023) Underneath this seemingly simple logic lies a sophisticated transformer architecture that allows language not only to be recognised, but also to be understood and processed.

Figure 6 shows these inner workings schematically. Each word or character received by the model is first transformed into a numerical representation, called an embedding, that captures meaning and context. The transformer then processes these representations through several layers. The built-in attention mechanism (multi-head attention) determines which part of the input is most

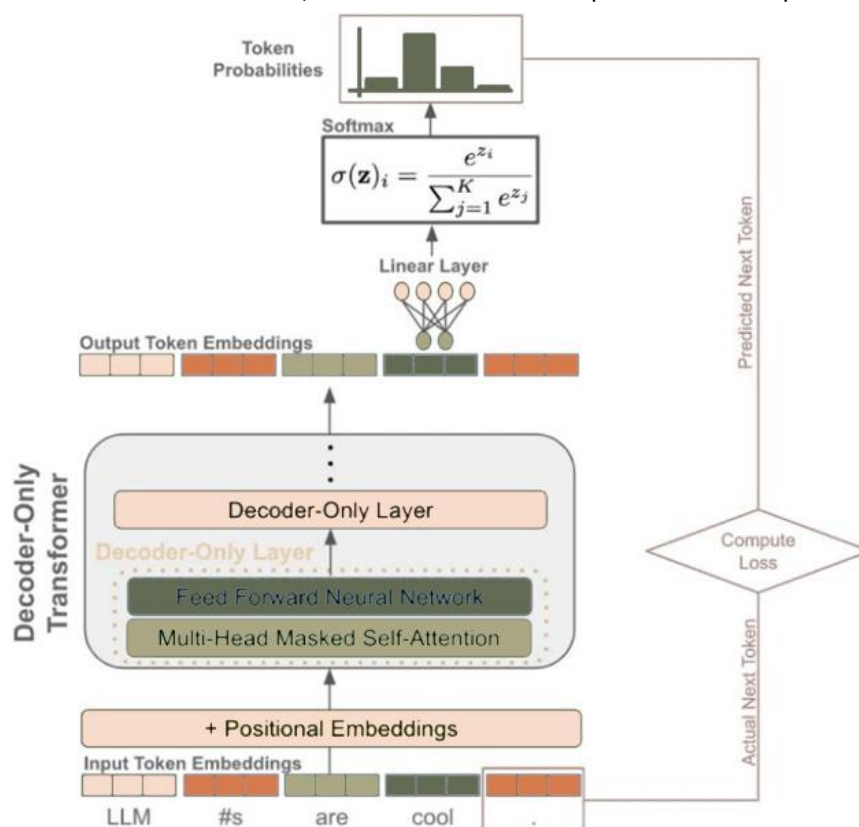


Figure 6: Large Language model, (Xi et al., 2023)

relevant at that moment. This allows the model to make connections between words that are far apart in a text and to recognise patterns across longer passages. This capability makes it possible to interpret complex documents, for example by linking energy performance requirements to material impact and policy frameworks.

This semantic structuring capacity is grounded in the vector space representation of language, in which words and concepts are mapped into a multidimensional space where similar meanings are situated close together (Mikolov et al., 2013) The transformer architecture further enables the model to draw connections across longer passages by attending to the relevance of different textual elements (Vaswani et al., 2017)

Interactions with an LLM are via prompts: textual instructions that give direction to the model's task. The precision and clarity of such a prompt largely determine the quality of the output. Clear, context-specific instructions lead to consistent and relevant answers, which is important in domains such as ESG decision-making, where information is often fragmented. Yet an LLM has clear limits. Model knowledge is static: anything does not present in the training data cannot be reliably recalled. In addition, models can hallucinate, generate seemingly logical but factually incorrect information (Gao et al., 2023)

Also, domain-specific knowledge is often lacking to answer specialised questions properly. To overcome these limitations, this research links the LLM with Retrieval Augmented Generation (RAG). Whereas the LLM acts as a cognitive engine, RAG provides the system with current and project-specific information. The next section elaborates on this RAG architecture and explains how it enhances the agent's reflective role in ESG decision-making.

RAG as a contextual memory layer

A language model such as an LLM has a lot of general knowledge, but that knowledge is "fixed", so to speak, the moment the model is trained. New developments, project-specific information cannot be taken by the model by itself. Therefore, in this study, it is coupled with Retrieval-Augmented Generation (RAG) explained by Grant Sanderson, (2023) This method gives the agent access to a kind of external memory, so that answers come not only from old knowledge, but also from the most recent and relevant sources.

Figure 6 shows how this works in practice. It starts with a question or command from the user. The agent translates that question into a mathematical representation, a kind of "summary" of the meaning of the question. The agent then searches a database full of policy documents, technical reports or other relevant information for pieces of text that match the content of that query well. This does not involve exact keywords, but rather similarities in meaning.

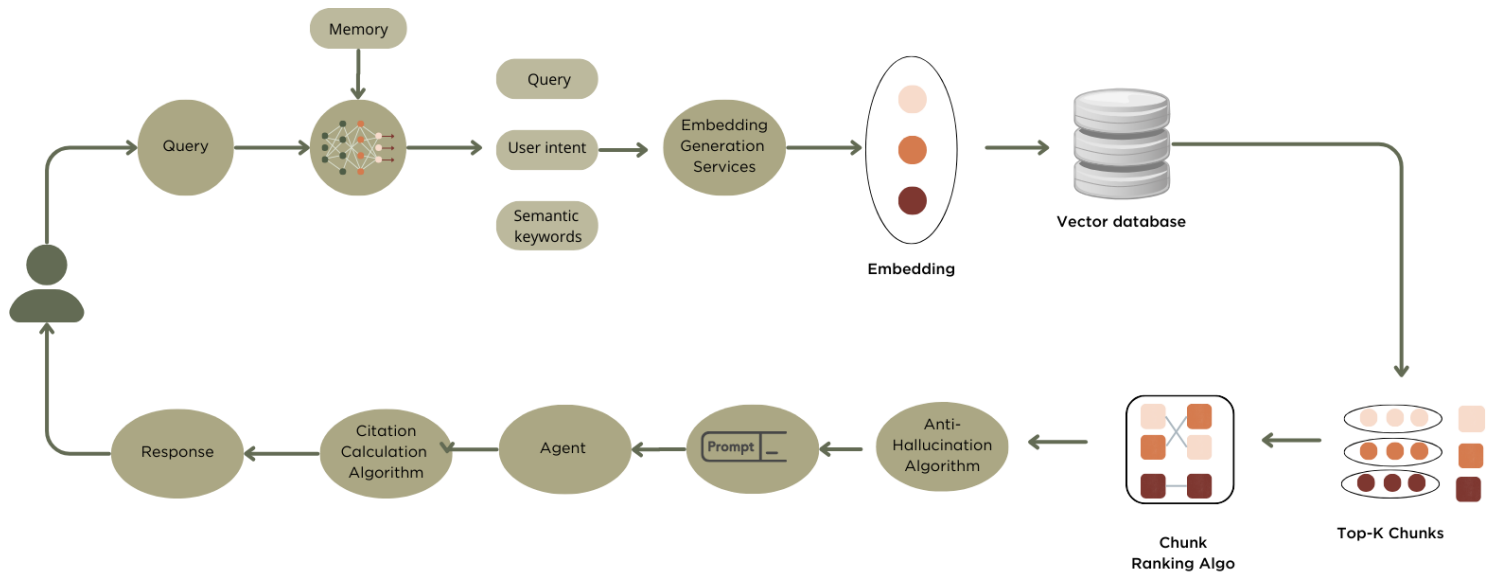


Figure 7: Workflow RAG (own work)

The agent then selects the most useful text fragments and checks whether they are reliable. This prevents the system from giving wrong or made-up answers. Only then are those checked fragments passed to the language model, which turns them into a clear and coherent response. The result is a response that not only sounds good but is also based on real and verifiable information.

The main advantage of this approach is that the knowledge is stored outside the model. As a result, the system can be quickly adapted when policies change, or new regulations are introduced. Especially in ESG-driven real estate development, where standards and requirements are constantly changing, this is a big advantage. It also makes it possible to use very specific information, such as local environmental regulations or building standards.

By combining the language model (which understands language and makes connections) and RAG (which adds up-to-date knowledge), the agent does not become a calculation system, but a flexible system that helps to better structure and discuss complex issues

Key Takeaways Artificial Intelligence

Research on AI consistently shows that contemporary systems operate through layered architectures that combine perception, interpretation and action. Huang & Ai (2023) demonstrate that an agent can perceive its environment, interpret information and act within defined boundaries, characterised by autonomy, reactivity, proactivity and interaction.

The cognitive foundation of such an agent is the Large Language Model. Xi et al. (2023) show that LLMs rely on next-token prediction and transformer architectures that enable semantic interpretation rather than mere calculation. Through embeddings and multi-head attention, LLMs detect relationships across long textual distances, which allows them to reveal patterns that would otherwise remain hidden in dispersed documentation. Mikolov et al. (2013) further show how meaning is captured through vector-space representations, enabling consistent

interpretation across terminology and domains. Vaswani et al. (2017) explain that the attention mechanism is what allows LLMs to link policy requirements, material impacts and other textual dependencies.

However, the literature also shows clear limitations. Gao et al. (2023) demonstrate that LLMs may hallucinate plausible but incorrect information because their knowledge is static now of training. This makes standalone LLMs unsuitable for high-stakes or context-dependent decision making such as sustainability assessments. This limitation motivates the addition of Retrieval-Augmented Generation.

RAG addresses this gap by linking model decision making to external, verifiable sources. Grant Sanderson (2023) explains how RAG retrieves semantically relevant text fragments from curated databases and verifies their reliability before producing an answer. This ensures that responses are grounded in project-specific and up-to-date information rather than relying on static model memory.

Synthesising these insights shows that AI does not function here as only a decision-maker but as reflective support in decision making. Taken together, these studies show that AI's value in ESG-driven decision-making lies not in automating choices, but in supporting the integration of information, assumptions and interpretations by making decision making explicit, structured and traceable.

2.3 The integration of ESG and AI

Based on the key takeaways of both domains, the literature shows a clear conceptual correspondence.

A: ESG studies identify fragmentation in terminology, documentation and stakeholder framing, particularly in the early design phase where the Environmental dimension is most influential. Frameworks such as Life Cycle Assessment show that operational and embodied emissions must be evaluated in the early stage, yet studies demonstrate that this integration is difficult to sustain in practice due to siloed responsibilities and competing organisational logics.

B: AI studies describe mechanisms that systematically organise dispersed information. These mechanisms are grounded in established AI frameworks such as Transformer architectures (Vaswani et al., 2017) distributed semantic representation (Mikolov et al., 2013) and Retrieval-Augmented Generation for contextual grounding (Grant Sanderson, 2023)

C: Together, these insights explain why AI is analytically relevant in this research, because its conceptual and architectural frameworks correspond directly to the points where environmental decision making tends to fragment and where integration across assumptions, data and perspectives becomes difficult. This alignment is particularly significant in the early design phase, when according to

the MacLeamy curve the Environmental dimension has the greatest leverage and fragmentation has the most lasting impact.

Table 1 presents an analytical synthesis developed for this research. By bringing these domains of literature together, the table clarifies how the properties of AI systems conceptually align with the specific points at which Environmental decision making in early design tends to fragment.

Environmental decision making (a)	Artificial intelligence (b)
Knowledge fragmentation	LLM-based agents can interpret distributed textual information and align terminology across heterogeneous sources.
Process fragmentation	AI systems can relate separate documents and workflows by identifying semantic connections within dispersed material.
End users are often not involved	Natural-language interaction enables the incorporation of user-related inputs such as IEQ observations into early decision-making processes.
Difficult adoption of digital technology	Language-based interfaces reduce reliance on specialised digital skills and make interaction with technical information more accessible.
No coordination in early project phases	AI agents can expose assumptions and dependencies within dispersed documents, allowing decision making paths to be made explicit.
Stakeholders with different interests	Pattern-recognition across texts allows AI to highlight differences in framing, revealing how choices are described from distinct stakeholder perspectives.

Table 2: The table present the key takeaways of the integration of domain A and domain B (own work, 2025)

2.4 Limits and risks of AI in environmental decision-making

Although the previous table shows a conceptual correspondence between the structural challenges described in ESG related decision making and the mechanisms identified in AI literature, another key finding of literature is the use of AI introduces its own set of risks and limits. These limits arise in this research arise in five domains presented in table 3.

Domain	Barriers	Core	Relevance	Source
Cognition (model level)	Truth value vs. Probability	LLMs generate plausible patterns but provide no guarantee of correctness	Output requires interpretation by human experts since an agent cannot replace domain knowledge.	(Gao et al., 2023)
Organisation & processes	Fragmentation of knowledge and responsibilities	Sustainability expertise stands apart from design and decision-making processes	Shared understanding becomes possible only when all teams work from the same information base	(Böttcher et al., 2020)
Data-infrastructure	Scattered documentation and inconsistent data quality	Information sits in separate files systems and repositories	Without a unified environment the agent functions as a consultation tool rather than a coherent decision making space.	Autoriteit persoonsgegevens (2025)
Regulation & governance	Uncertainty about responsibility	The EU AI Act requires transparency and clear human accountability	Organisations must define who interprets agent output and who holds final responsibility	EU AI act (2024)

Digital sustainability	Increasing energy and water consumption by data centres	Large infrastructures and environmental externalities	scale introduce	AI can assist environmental analysis but also creates its own ecological footprint	(Dutch data center, 2025)
------------------------	---	---	-----------------	--	---------------------------

Table 1: Limits and risks of AI (own work)

Organisation and processes

The second domain concerns the way organisations distribute knowledge and responsibilities. Studies show that sustainability expertise is often located in separate teams while design choices occur elsewhere. This separation leads to fragmented information flows. Böttcher et al. 2020 argue that consistent interpretation becomes possible only when teams work with the same information base. This condition cannot be created by an AI system itself. Instead, it depends on organisational alignment which varies strongly across corporate real estate projects.

Data infrastructure

The third domain concerns data environments. Environmental decision making relies on documents specifications monitoring data and regulatory texts that often sit in separate systems. Autoriteit Persoonsgegevens 2025 notes that organisations increasingly depend on cloud-based platforms which centralise information yet create new dependencies on external providers. The first image illustrates this shift. Raw project information moves into protected environments that require subscription-based access. This creates questions about control timeliness and data continuity. Without a unified environment an AI agent remains a consultation system rather than a coherent decision making space.

Regulation and governance

The fourth domain concerns the regulatory context. The EU AI Act 2024 requires transparency human oversight and traceable decision making. Responsibility for interpreting model output cannot be delegated to automated systems in strategic contexts. Organisations must define who interprets output and who holds final responsibility. This sets clear boundaries for any reflective system that interacts with environmental decision making in early design phases.

Digital Sustainability

The fifth domain concerns the ecological footprint of AI infrastructures. Figure 7 shows the expected rise in both commercial and hyperscale data centres between 2025 and 2030. This growth signals a rapid expansion of the global computational landscape that supports machine learning systems. More data centres mean more server halls more cooling hardware and more electrical equipment which together shape the physical scale on which AI depends. The global spread of hyperscale

facilities illustrates that AI does not operate in a virtual vacuum but relies on material and spatial infrastructures with significant environmental weight.

Figure 8 then shows how this physical expansion is mirrored in rising electricity demand. In the period from 2024 to 2030 data centres form one of the fastest growing categories of energy use and this trend is projected to intensify as workloads increase. Studies by the International Energy Agency 2025 estimate

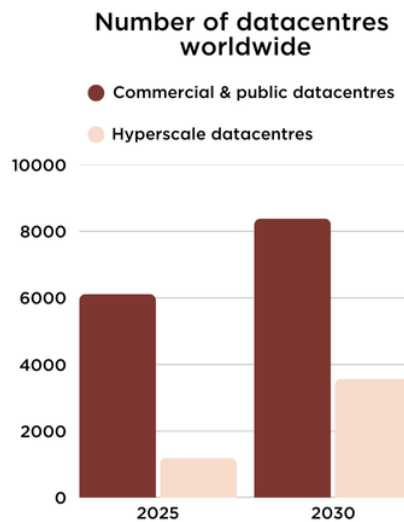


Figure 8: Number of datacentres worldwide based on IEA, 2025

that global data centre electricity demand may reach roughly 945 TWh per year by 2030. This level of demand affects national grids especially in regions where renewable supply is limited or where grid capacity is already under stress. Dutch Data Center 2025 notes that the environmental burden of AI therefore differs strongly between regions since local energy mixes cooling technologies and spatial planning regimes shape the carbon intensity of data centre activity

Water use adds a second dimension to this footprint. Most cooling systems rely on either evaporative cooling or liquid-based solutions and these systems vary widely

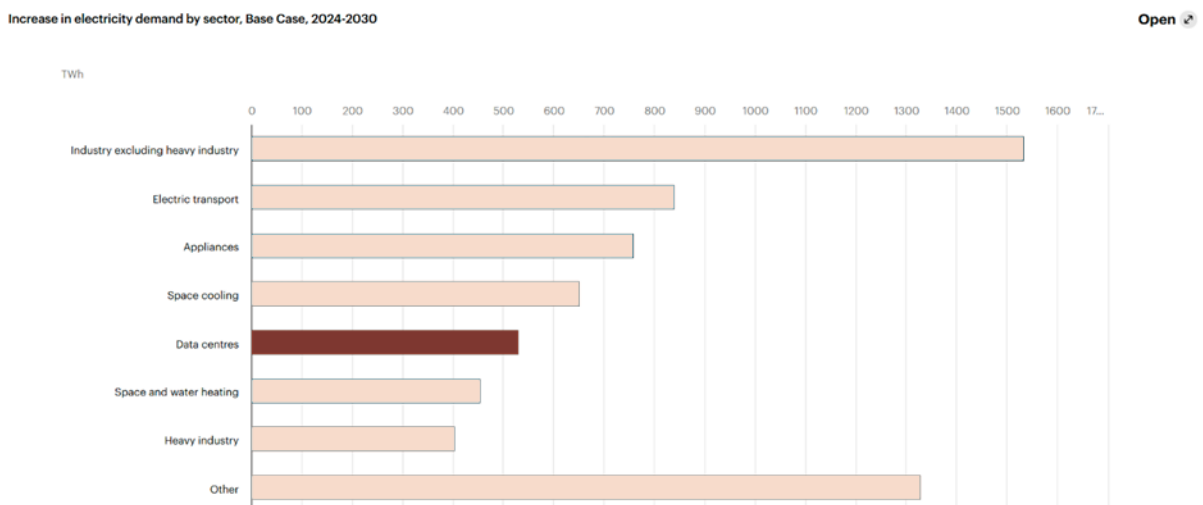


Figure 9: Increase in electricity demand by sector (Base Case 2024-2030)

in efficiency. As a result, two data centres with similar computational capacity can produce very different environmental loads. High water use becomes especially relevant in locations where freshwater resources are under pressure or where heatwaves become more frequent. This variation shows that the environmental cost of AI cannot be understood as a single universal value but depends on the local configuration of energy supply climate and infrastructure.

Taken together these developments show that AI systems rely on energy intensive and resource intensive infrastructures. The ecological footprint of these infrastructures forms an essential part of the theoretical boundary conditions for any study that positions AI within environmental decision making. The growth of data centres shown in Figure 2 and the rise in electricity demand shown in Figure 3 illustrate the scale of these boundary conditions and highlight why digital sustainability must be understood as an inherent part of the broader environmental context in which AI operates.

2.5 Key Takeaways chapter 2

Chapter 2 confirms that sustainable design choices in Dutch corporate real estate is shaped by three mutually influencing dimensions, with the Environmental dimension forming the primary analytical entry point in the early design phase. At this stage, Social and Governance considerations do not yet generate substantive tensions because their interpretation depends on how Environmental strategies take shape: user expectations only become meaningful once material and energy concepts are defined, and governance requirements only become operational when there is clarity about what must be verified, monitored or reported. The Environmental dimension is therefore the most determinative focus at this early moment, as it is here that sustainable design choices are first articulated and can still be aligned.

Although the Environmental domain spans themes such as water use, ecological value, climate adaptation and resource flows, the literature shows that the most consequential early-stage challenge arises from the relationship between operational energy demand and material-related (embodied) emissions. These two components dominate life-cycle impact and are among the first Environmental considerations to shape design direction. This study therefore narrows its analytical scope to this intersection, while recognising that the interpretive approach and mechanisms explored can be extended to other Environmental themes when required.

At the same time, the AI literature conceptualises LLM-based agents as interpretive systems capable of structuring fragmented information, aligning terminology and reconstructing decision making across distributed documents. When viewed alongside the ESG literature, a clear conceptual correspondence emerges: the points where Environmental decision making becomes fragmented correspond to the types of complexity that AI systems are designed to help structure and interpret. This does not imply that AI replaces human judgement, but

that it can function as a reflective lens, making assumptions, dependencies and points of misalignment more explicit in early-stage decision-making.

Taken together, these insights clarify why the interaction between Environmental decision making and AI forms the logical starting point for this research. Chapter 2 provides the conceptual justification for focusing on the Environmental dimension, for narrowing that focus to the relationship between operational and material-related emissions, and for examining how an AI agent can support early design conversations when fragmentation is high and design freedom is still at its peak. This theoretical foundation prepares the ground for the subsequent chapters, which explore how Environmental decision making unfolds in practice and how the agent can meaningfully support sustainable design choices.

Chapter 03

Methodology

- 3.1 Research Through Design method
- 3.2 Research Techniques
- 3.3 Round 1: Exploration (problem space)
- 3.4 Round 1: Exploration (Solution space)
- 3.5 Round 2: Prototype
- 3.6 Round 3: Validation
- 3.7 Research output
- 3.8 Trustworthiness and Rigour
- 3.9 Ethical Considerations
- 3.10 Key Take Aways Methodology

3. Methodology

The theory presented in Chapter 2 suggested that there is a relationship between environmental decision-making in the early design phase and the way in which an AI agent can function as a support system. To reveal these dynamics, the research follows an abductive Research Through Design approach. In the first phase, interviews are conducted to understand how environmental decision making comes about in practice and to gain insight into the possibilities and limitations of AI. These findings form the basis for the iterative prototyping of the agent. The prototype is then tested in a design workshop, which examines how the agent influences human decision-making. The technical functioning of the agent is also assessed and validated by AI experts.

3.1 Research Through Design method

The integration of environmental decision making and AI interpretation creates a research situation that cannot be captured by conventional analytical methods. Such methods are based on stable variables, linear causality and reproducible patterns, whereas this research focuses on understanding human decision-making. As (Stappers, 2017) highlight, positivist research models are too rigid to adequately understand such contextual, human-related processes. This calls for an approach that focuses on meaning-making and interaction rather than measurability.

Research Through Design (RTD) closely aligns with this need. RTD is a research approach in which knowledge is generated during the design process: by creating, experimenting, reflecting and observing how that artefact changes the thinking of those involved. The method is based on the idea that designers construct meaning while working under uncertainty. This is in line with the interpretivist basis of RTD, in which knowledge grows through systematic reflection on experiences in situations where little is yet certain. (NKK Gamage, 2025)The value of an artefact does not lie in whether it 'works', but in what it reveals about the underlying decision making.

RTD also fits within a pragmatic research orientation, which Gamage (1997) describes as focused on usability and solving problems in real practices. The introduction of an AI agent as a Digital Building Master is precisely such a practice-oriented intervention: its effect on environmental decision making can only be understood by testing the agent in realistic design contexts. RTD provides the right framework for this, because the artefact functions as a analysis.

In this study, the AI prototype was therefore deliberately designed not as a finished instrument, but as a research probe. The prototype reveals how environmental decision making shifts when the system organises documents, harmonises definitions or makes implicit connections explicit. This use of an artefact as a source of insight is typical of RTD: knowledge arises because the prototype

provokes interactions, not because it is a product (NKK Gamage, 2025) This makes RTD the most logical and scientifically sound method for this research.

The application of RTD in this research is divided into three rounds. RTD assumes that a complex issue cannot be understood in one go, but that understanding grows by first exploring, then developing an artefact, and then observing that artefact in use (Zimmerman et al., 2007). These research rounds and the associated techniques are explained in the following sections.

3.2 Research techniques

This choice of an RTD approach also means that the research will be structured abductively. An abductive logic, as described by (Dubois & Gadde, 2002a) fits seamlessly with the iterative nature of RTD, in which theoretical understanding and empirical insights constantly reinforce each other. Instead of testing pre-formulated hypotheses, knowledge develops step by step, as observations from practice feed back into the design of the artefact. (Zimmerman et al., 2007) show that abduction is a core principle within design-oriented research, because it offers scope to incorporate surprising insights and refine the design as it progresses.

This abductive RTD logic forms the basis for the research design. The research is structured in three rounds that influence each other iteratively: first, the problem domain is explored; then, a prototype is developed; and finally, it is investigated how this prototype changes environmental decision making. Table 3 shows how this logic has been translated into concrete research techniques, their sources and their purpose. The paragraphs below explain these rounds separately.

In the first round of research, the problem space is explored through semi-structured interviews with developers, architects, consultants and AI specialists. (Bryman, n.d.) These conversations reveal how environmental transactions are interpreted in practice, where knowledge gaps arise and what conditions an AI agent must fulfil to contribute meaningfully as a reflective system. This analysis also yields initial insights into how an agent should function. These functions are only specified after this empirical phase and form the basis for the design requirements.

In the second round, these empirical insights are translated into a prototype, developed within a specific PV case in which the trade-off between operational and material-related emissions is central. The numerical values and sustainability parameters used in this case are derived from an existing EDGE Technologies project, EDGE stadium, ensuring that the scenario reflects a realistic and professionally validated design context. This case, which recent literature describes as a representative tension zone in sustainability design (Tirelli & Besana, 2023) offers a manageable yet authentic setting for testing the agent's interpretative mechanisms. The prototype is based on LLM technology in combination with Retrieval-Augmented Generation (Guo et al., 2024), so that project-specific documents can be incorporated into the interaction. The prototype is then explored in a design workshop to reveal implicit assumptions,

and argumentation structures the workshop shows how the prototype influences the formulation of environmental decision making, which assumptions it makes explicit, and in what ways it supports shared interpretation.

In the third round of research, experts assess the transparency, usability and limitations of the system. This evaluation is informed by guidelines on the responsible use of AI, including the requirements of the EU AI Act (2024) and literature on human interpretation of model output (Gao et al., 2023). As a result, the agent is assessed as reflective support within environmentally oriented decision-making, not as an autonomous decision-making system. A summary of the research approach is shown in Table 3. Each round is explained in the following paragraphs.

Research	Sub question	Research techniques	Source	Purpose
<i>Theoretical background</i>	Backbone of the project	Literature review	Academic papers, regulations, ESG documents	Understanding ESG decision-making in corporate real estate and the technical application of AI.
<i>Round 1</i>	SQ1: problem space	Semi-structured interviews	Developers, architects, consultants and AI experts	Understanding how environmental decision making processes currently work, where tensions arise and what AI knowledge is lacking.
	SQ2: Solution space	Thematic analysis of interviews + translation to design requirements	Developers, architects, consultants and AI experts	Formulating design requirements for the agent; determining what input, forms of explanation and timing of intervention are needed.
<i>Round 2</i>	SQ3: Prototype	Research through design	Copilot Studio prototype + real project PV-embodied carbon trade-off case	Developing and refining the AI prototype as a reflective system in a concrete design situation.
		Case study	EDGE Technologies	This case study was used to create an existing design solution.
<i>Round 3</i>	SQ4: Validation	Expert panel	AI experts, developer, sustainability consultants	Evaluating the transparency, usability and legitimacy of the agent in decision-making.

Table 2: Research techniques (own work)

3.3 Round 1: Exploration (problem space)

As mentioned in Section 3.1, the first round is the exploration round. This round consists of a problem space and a solution space. In the problem space, a conceptual framework is developed to structure the understanding of the AI agent. In the solution space, this framework is translated into functional requirements for the prototype in the subsequent phase.

Development of conceptual framework

The three roles of the Digital Building Master emerged from the integration of Environmental decision-making literature and AI mechanisms. When the

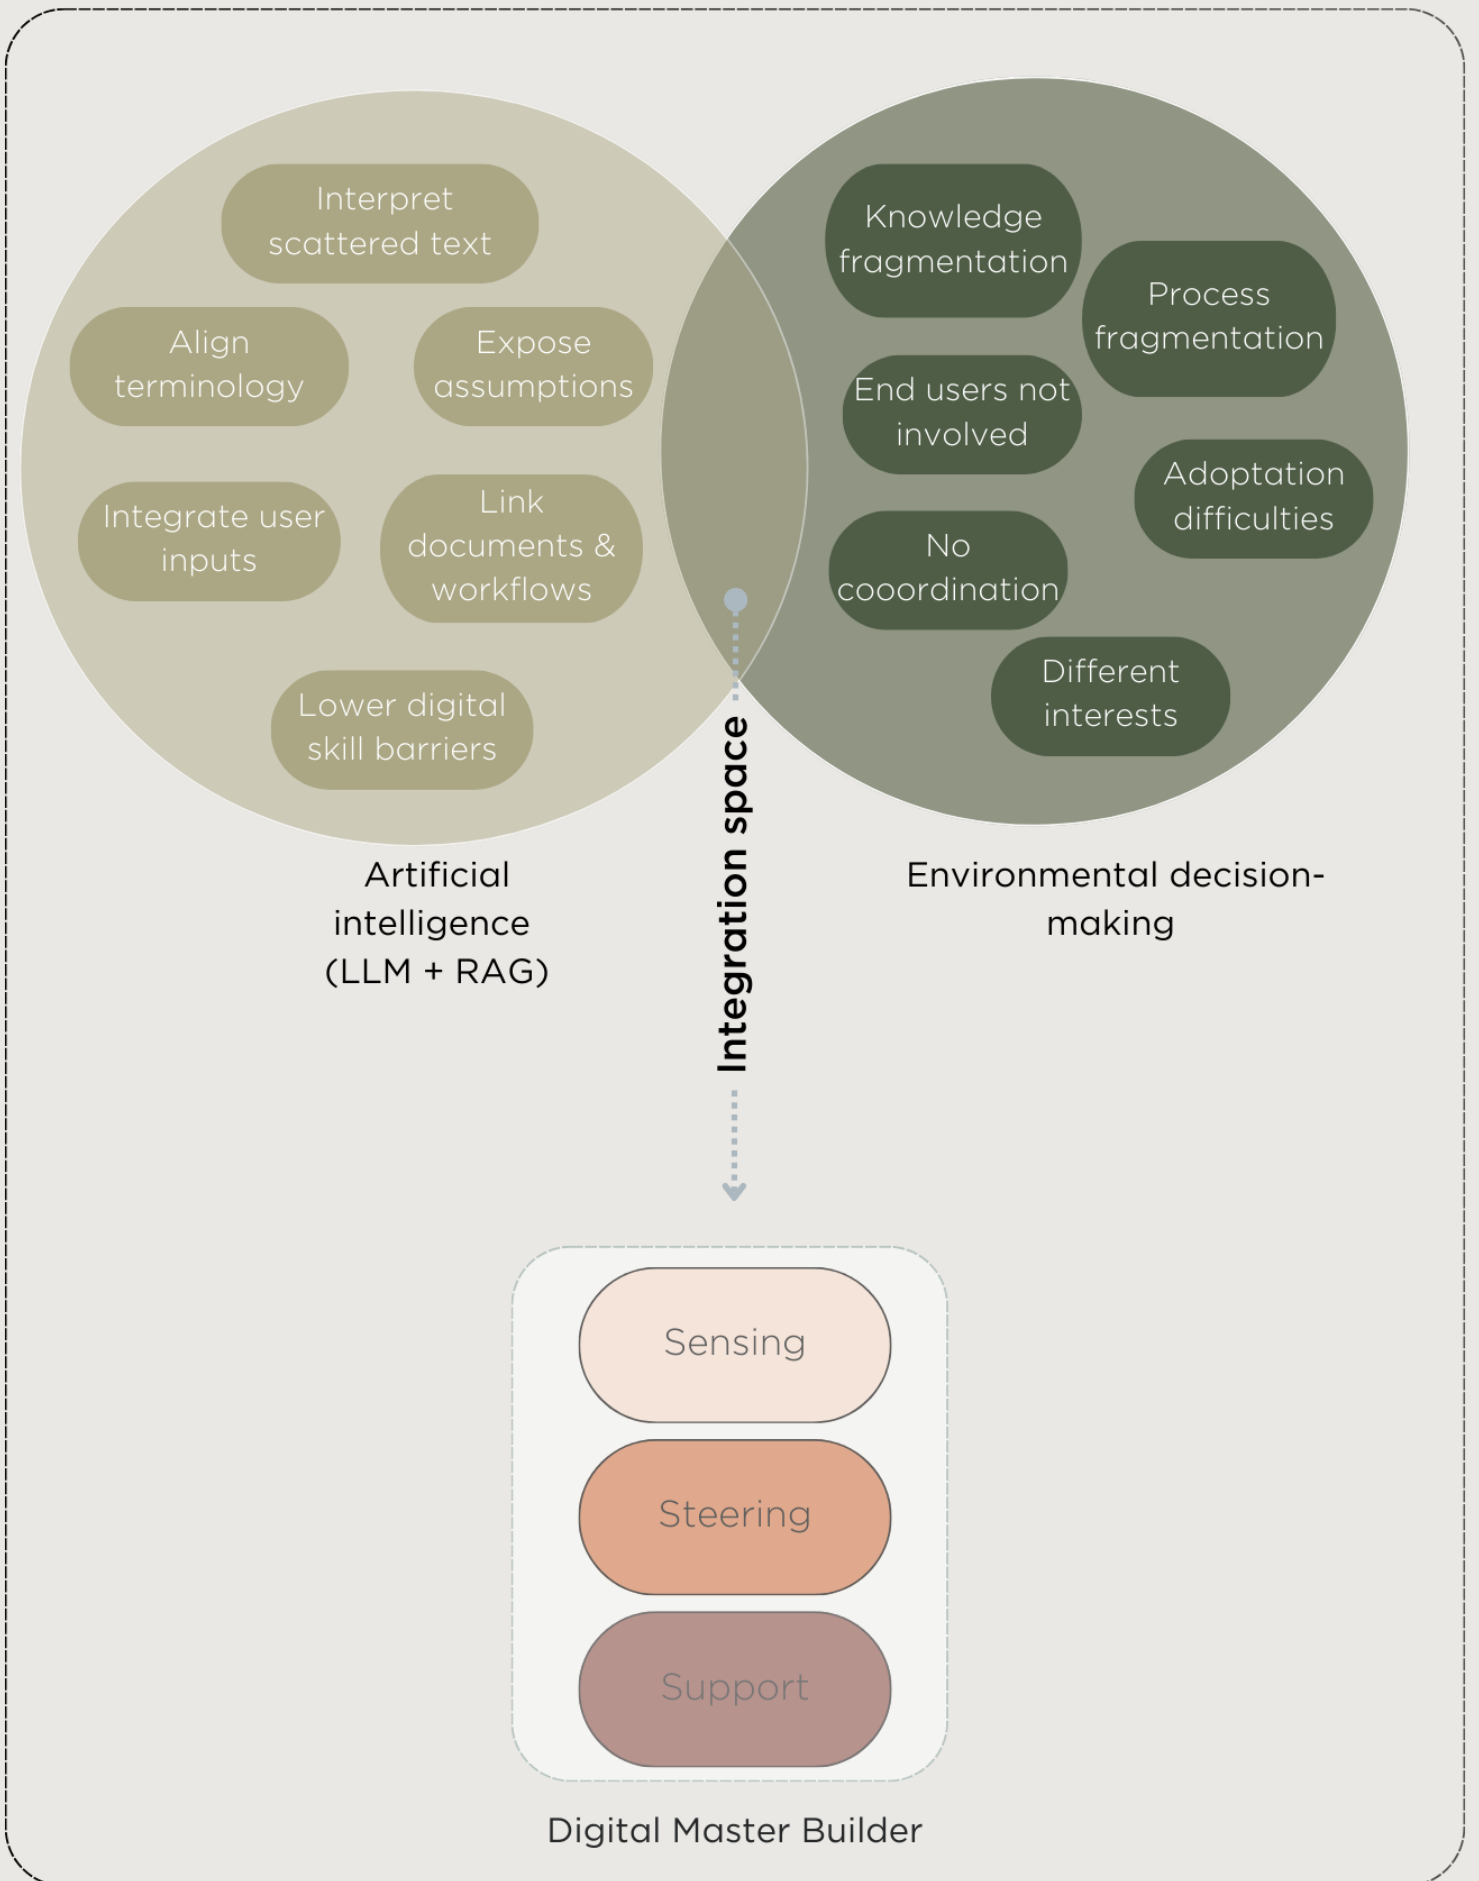
mechanisms identified in both domains are analysed side by side (as synthesised in Table 1 and the integration figure), it becomes clear that the AI agent is analytically relevant in situations where early design choices are constrained by fragmentation, uncertainty and dispersed information. Rather than addressing individual decisions, the agent's role is conceptualised through three functions, sensing, steering and support, which together describe how the system can contribute to structuring early-stage design conversations.

Sensing refers to the agent's capacity to detect patterns, assumptions and differences in interpretation that arise from fragmented knowledge, conflicting frames and limited coordination in Environmental decision-making. Drawing on AI capabilities such as interpreting dispersed textual information, identifying semantic inconsistencies and exposing implicit assumptions, the agent adds an additional layer of sensitivity to the design process. In this role, the agent does not evaluate or resolve tensions but makes them visible at a moment when they can still be discussed and interpreted.

Steering describes how the agent helps to structure decision making when process fragmentation and technical complexity make it difficult for teams to maintain overview. By connecting documents, aligning terminology and translating complex information into accessible natural language, the agent provides direction in a non-prescriptive manner. Steering does not involve selecting or optimising design options; instead, it supports the organisation of information and decision making so that design teams can better navigate incomplete or contradictory inputs.

Support captures the stabilising role of the agent in situations characterised by divergent stakeholder perspectives and limited user involvement. AI mechanisms that allow the integration of user-related observations, comparison of viewpoints and traceability of arguments enable the agent to function as a shared reference point. This supports alignment across disciplines by reducing reliance on individual interpretations and by providing a transparent basis for discussion and justification.

Together, sensing, steering and support form the conceptual foundation of the Digital Building Master. They do not represent sequential steps or a fixed decision-making workflow, but an analytical framework that describes how an AI agent can enhance coherence, transparency and reflection in early-stage Environmental decision making. The relevance of each role is context dependent. In situations where Environmental ambitions and assumptions are still implicit, sensing becomes most salient by revealing differences in interpretation. As information increases and design complexity grows, steering gains importance by helping teams' structure fragmented inputs without prematurely closing design options. Support becomes particularly relevant in multi-stakeholder settings, where shared reference points are needed to stabilise discussion and justification across disciplines. These roles therefore become prominent depending on the design phase, the problem context and the actors involved, and form the basis for the functional requirements developed in the subsequent phase of the research.



Empirical research: Semi-structured interviews

Semi-structured interviews were conducted to explore how the three preliminary roles of the Digital Building Master: sensing, steering and support, manifest in sustainable design choices. These roles were derived from the theoretical integration of ESG-related fragmentation and AI interpretive mechanisms (Figure 9), but remained, at this stage, conceptual. The interviews therefore served as an initial empirical exploration to assess whether, how and under which conditions these roles resonate with practitioners' descriptions of their decision-making processes.

Semi-structured interviewing is well suited to this task because it allows predefined themes (the three roles) to guide the conversation while still leaving space for unexpected insights to surface. (Bryman, 2016). This flexibility is important, since sustainable design choices is only partly explicit; much of it emerges through tacit heuristics, informal coordination and fragmented information flows. The method also aligns with the abductive logic of this study, in which empirical findings refine rather than validate the conceptual framework. (Dubois & Gadde, 2002a)

A targeted sampling strategy ensured variation in professional backgrounds, including environmental specialists, developers, architects and AI experts (Appendix A: Overview Interviewees) A total of fifteen professionals were interviewed. This diversity made it possible to reveal differences in language use and assumptions.

Each interview lasted approximately 45 minutes and was conducted in a personal setting whenever feasible. This setting encouraged participants to articulate not only formal decision making but also doubts, informal heuristics and implicit considerations that typically remain invisible in project documentation. All interviews were recorded, transcribed using Fireflies and imported into ATLAS.ti for systematic analysis. See Appendix B: Interview protocol for more insights in the questions and Appendix D overview code groups Atlas.TI.

The analysis in Atlas.TI began with a deductive coding structure based on the three preliminary roles of the Digital Building Master. In keeping with an abductive approach (Dubois & Gadde, 2002), the coding framework was intentionally kept open so that empirical patterns could reshape the theoretical structure. New codes were added when recurring insights emerged.

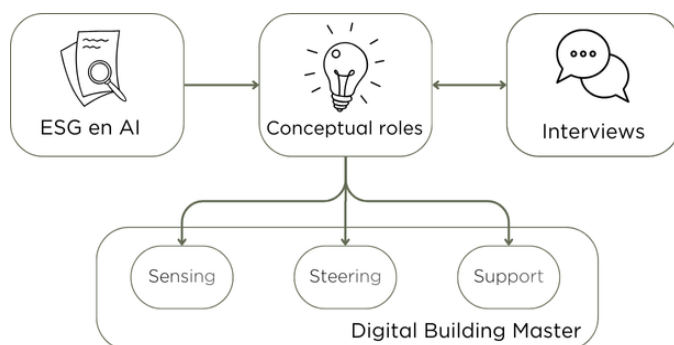


Figure 10: Development of the Digital Building Master roles (own work)

3.4 Round 1: Exploration (Solution space)

The formulation of the requirements translates the empirical insights from the interviews into the operational conditions for the AI agent. These interviews revealed that environmental decision making in the early design phase is shaped by unevenly distributed information, implicit trade-offs and fragmented responsibilities. This is precisely why there was a need for a system that does not make decisions but rather reveals the underlying decision making. Based on these insights, the agent is positioned as a Digital Building Master who identifies tensions, organises information and creates common references, the three conceptual roles of sensing, steering and support.

To translate these roles into an agent that can function in design practice, the requirements are linked to the technical layers of an AI system (Figure 11). These layers describe the steps by which the agent receives, processes, structures and returns information to the user. Positioning the requirements within this layered structure creates a direct link between the conceptual model and the concrete behaviour that the agent must exhibit. At the same time, it becomes clear which components are technically trainable and where human interpretation remains indispensable.

This approach creates a consistent framework in which the three conceptual roles are translated into an operational system. The layered structure forms the foundation on which the prototype is built and enables the agent to determine not the content of environmental decision making, but the way in which that decision making is formed, more explicitly, more traceably and more easily replicable. In the next phase, these requirements will be used to train the prototype.

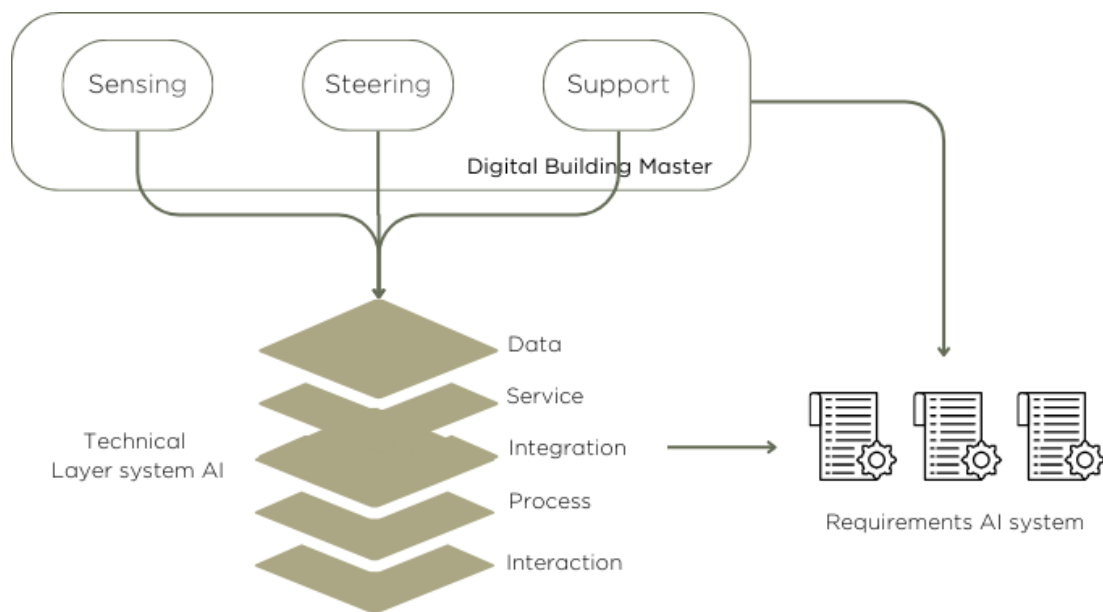


Figure 11: Translation of the Digital Master Builder's conceptual roles into the technical layers of the AI system (own work, 2025)

3.5 Round 2: Prototype

Within this Research Through Design approach, as mentioned earlier, the prototype functions as a medium to better understand the influence of AI on decision making. RTD regards prototypes as epistemic artefacts: a system that reveal how people reason, interpret information and recognise tensions in complex design situations (Stappers, 2017). From that perspective, the prototype in this study was designed around the three roles of the Digital Building Master, so that the influence on environmental decision making could be effectively investigated.

The prototype was developed in Microsoft Copilot Studio, a choice that makes sense both methodologically and practically. EDGE Technologies has a corporate licence, which means that all data remains within the secure Azure tenant and is not used for external model training (Security for Microsoft 365 Copilot, 2025). This creates a controlled, privacy-secure research environment in which project documents can be integrated without data loss or risks. Appendix H: Copilot Studio Agent shows the agent configuration.

To prevent the complexity of environmental decision making from becoming too broad, a defined but realistic use case was chosen. This is in line with guidelines for validation in design research, in which limited scenarios help to make shifts in decision making clearly observable (De Boeck & Vaes, 2024). The selected case, the trade-off between reducing operational CO₂ and increasing material-related CO₂ in the application of PV panels, was repeatedly mentioned during the interviews as a typical area of tension in early design phases.

First step: reconstruction of human decision making

In the first development step, the sustainable design choices underlying the PV-related trade-off between operational and material-related CO₂ were reconstructed using realistic project data from the completed EDGE Stadium project. This project provides a suitable basis because it represents a fully realised corporate real estate development in which energy performance, circularity and material-related emissions were explicitly weighed as part of early design decision-making. The PV case reflects a recurring type of sustainable design choice in corporate real estate portfolios, where comparable decisions must be made across assets under similar regulatory, financial and spatial conditions.

The photo in Figure 11 shows the roof of EDGE Stadium and visualises the physical design context from which the PV parameters were derived. The roof geometry, available surface area and positioning of installations are directly visible, allowing the sustainable design choice to be understood as a response to concrete spatial constraints rather than as an abstract optimisation exercise. This grounding clarifies how environmental ambitions are translated into design decisions under real project conditions.

The parameters for the PV calculation come directly from EDGE Stadium and have been validated by an installation specialist Appendix G: Human Workflow Solar

Panel. Because the project is completed and its technical data is available within EDGE, the case provides a transparent and verifiable basis for analysing sustainable design choices. This level of objectivity is essential in corporate real estate, where environmental decisions must be explainable, repeatable and defensible at portfolio level.

An additional advantage of selecting EDGE Stadium is that many of the developers, consultants and sustainability specialists involved in the original project remain active within the organisation. This enabled the reconstructed design choice to be discussed with practitioners who were involved in the original process, ensuring that assumptions, constraints and priorities accurately reflect corporate real estate practice. In this way, EDGE Stadium functions as a reference project through which sustainable design choices can be examined and contextualised.

Together, the completed project status, validated technical data and organisational continuity create a bounded yet realistic context for prototype development. This aligns with Research Through Design principles, which emphasise the use of concrete cases to study how design choices are shaped and structured in practice (De Boeck & Vaes, 2024) The PV case therefore offers a robust foundation for investigating how an AI agent, positioned as reflective support, can help structure and clarify sustainable design choices in early-stage corporate real estate development without determining their outcome.



Figure 12: Rooftop configuration of EDGE Stadium used as basis for the PV trade-off case. (Edge, 2025)

Second step: translation into a working workflow

In the second step, the reconstructed decision making was converted into a workflow in Copilot Studio, built up from layered prompts that correspond to the technical layers from the requirements package. The most important workflow components are shown in Appendix I: Workflow AI in Copilot. Appendix J: Example CoPilot Response illustrates how the agent converts this structure into concrete output.

The functionality is deliberately limited to the three senses of the Digital Building Master: This ensure that the agent provides precisely the cognitive support that was identified as lacking in the interviews and theoretical analysis.

The prototype is therefore deliberately partial but analytically powerful: operational enough to elicit decision making, but not so complete that it would take over the decision-making process. This is in line with recommendations from design-method research, in which prototypes are primarily intended to generate insight rather than form a final product (De Boeck & Vaes, 2024).

In this form, the prototype forms the basis for the validation phase, in which it is investigated how an AI agent influences environmental decision making in the early design phase. Figure 11 shows this methodological cycle in context.

3.6 Round 3: Validation

The validation round was designed as a design-oriented comparative study, in line with recommendations for evaluating early-stage design methods. (De Boeck & Vaes, 2024) Rather than assessing technical accuracy or decision quality, the aim of this phase was to examine how the presence of the AI agent influences the way sustainable design choices are interpreted, structured and discussed by designers in a realistic corporate real estate context. This aligns with the core premise of this research, which positions the AI agent not as a decision-maker, but as a reflective support tool within early-stage design decision making.

To capture variation in interpretation and use, the participant group was deliberately heterogeneous, in Appendix O: Overview of participants in the validation phase. It included professionals with backgrounds in development, sustainability consultancy, design and technical advisory roles. Design research shows that such diversity is essential for revealing implicit assumptions and differences in framing, as professionals from different practices bring distinct heuristics and priorities to the design process. (Designerly Ways of Knowing, n.d.) This heterogeneity created an “argumentation landscape” in which the relevance of sensing, steering and support could be observed across different actor positions rather than assumed to operate uniformly.

The abductive logic of this validation phase reflects the conditions of early design in corporate real estate, where uncertainty, incomplete information and normative considerations dominate. In such settings, insight does not emerge through convergence or consensus, but through the confrontation of divergent

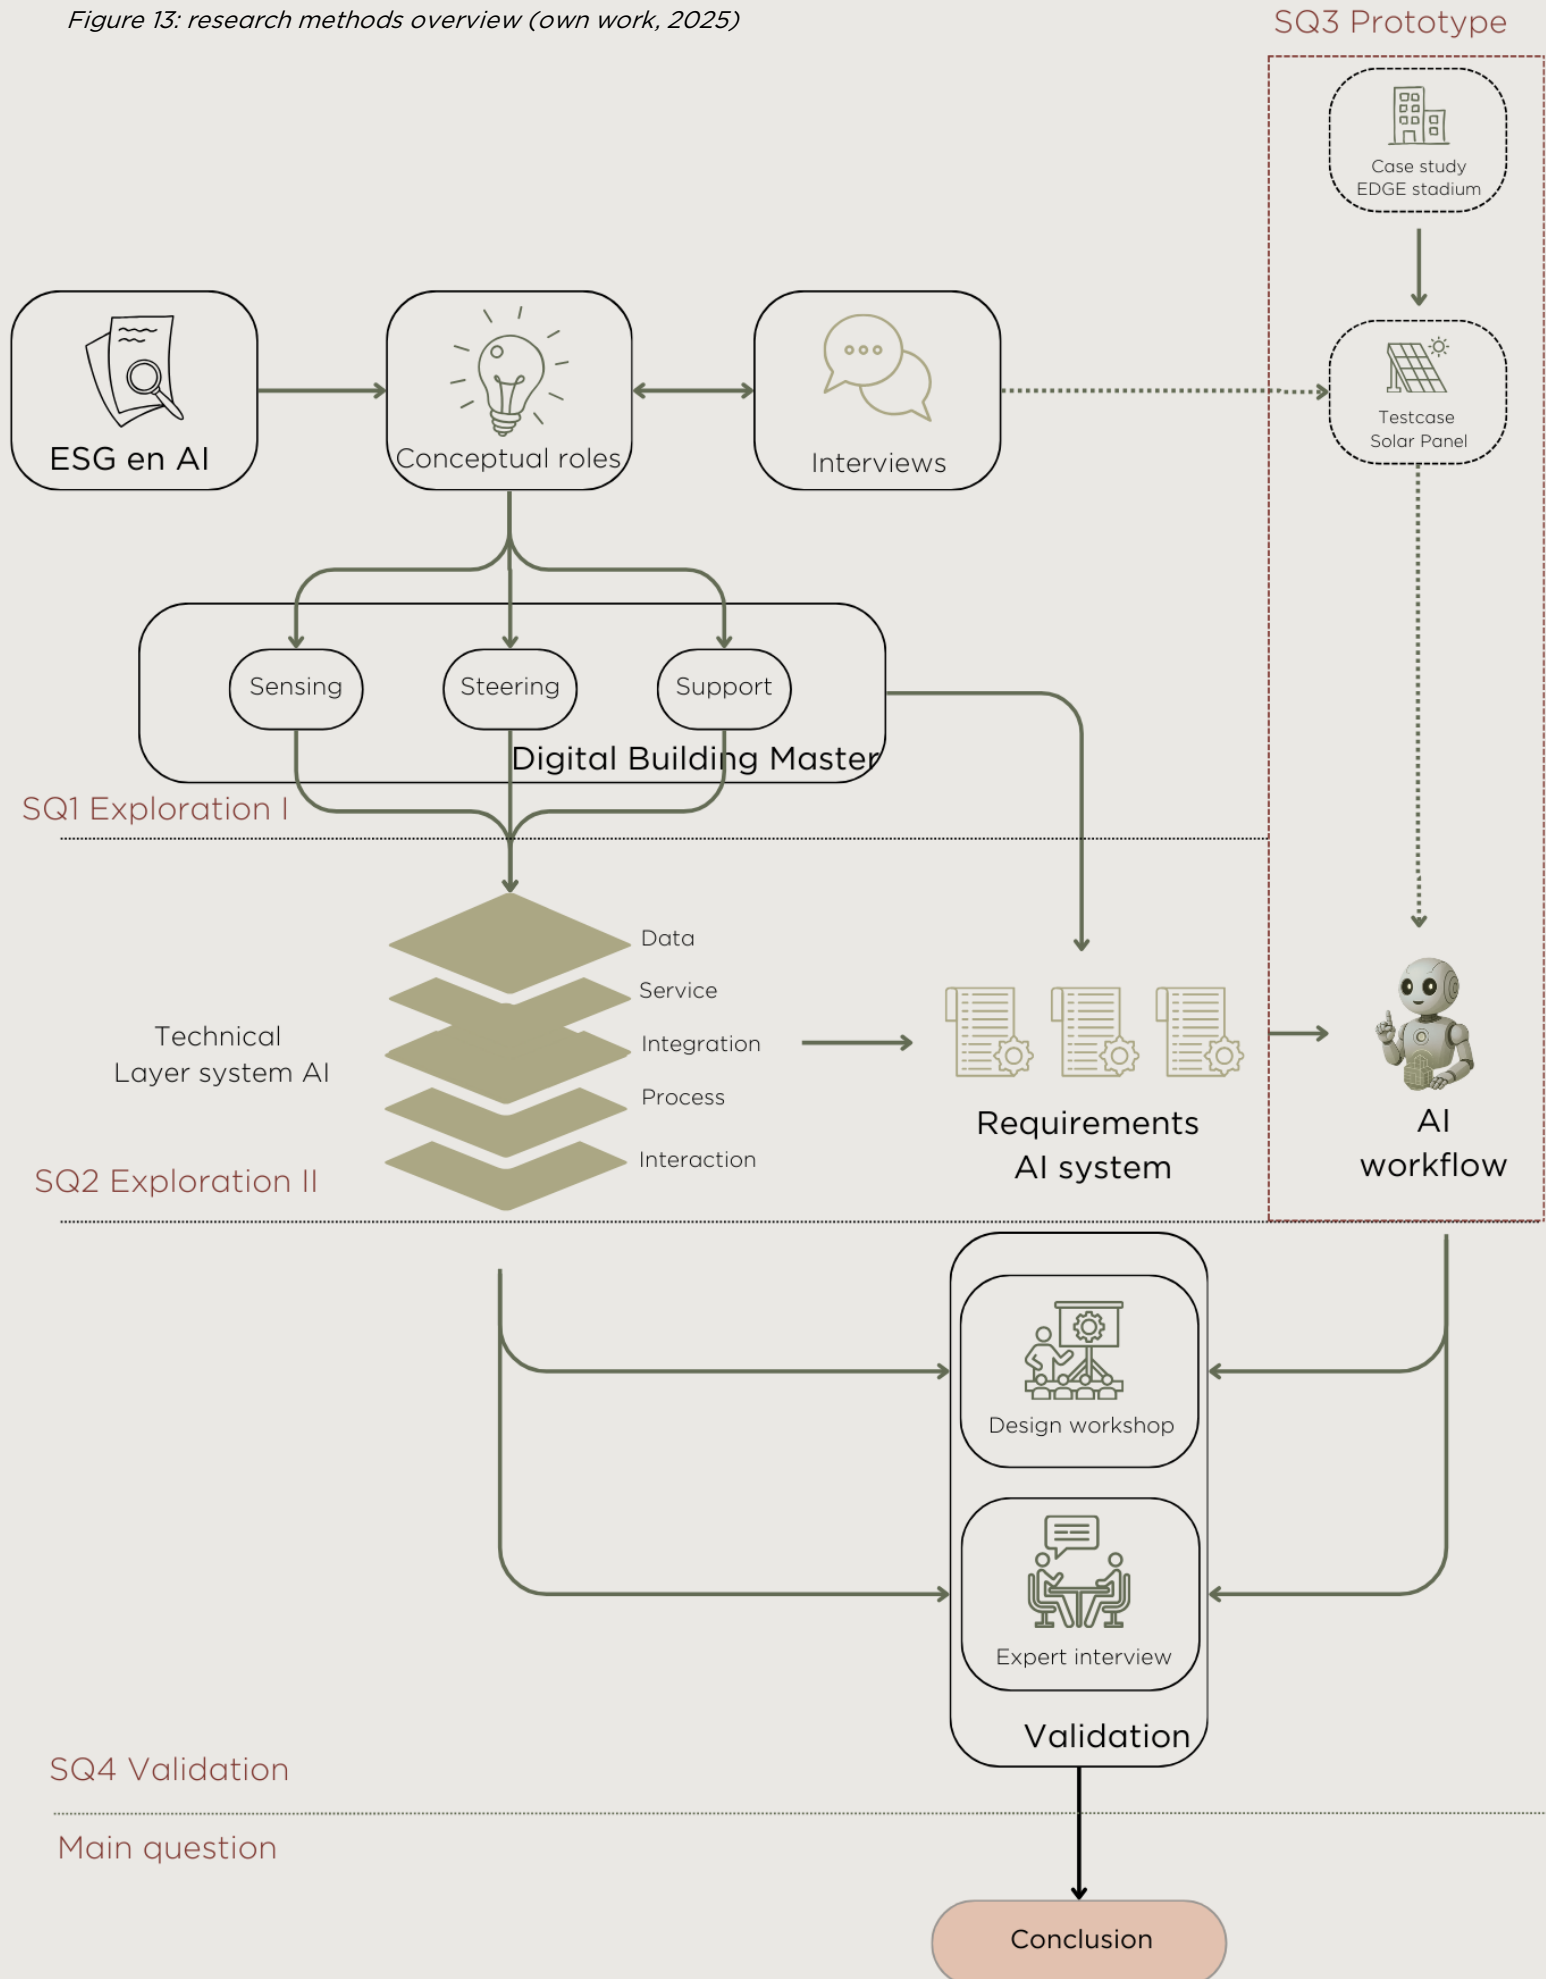
perspectives (Dorst, 2011). This made it possible to observe when specific AI roles became salient. Sensing proved most relevant when Environmental assumptions were still implicit or contested, steering gained importance as information volume and complexity increased, and support became critical when participants sought to justify or stabilise arguments across disciplinary boundaries. In this way, the validation explicitly addressed the feedback that these roles do not operate continuously, but become useful depending on context, phase and actor position.

The validation consisted of two consecutive parts, shown in Appendix P: Design Workshop material. First, participants worked on the PV case as a representative sustainable design choice within corporate real estate, without any digital support, using only a blank worksheet, sketches and their own experience. This case was deliberately selected because similar PV-related decisions recur across corporate real estate portfolios and involve the integration of Environmental ambitions, financial considerations and organisational constraints rather than isolated optimisation. In the second part, participants addressed the same case using the AI agent's structured output (Appendix J: Example Copilot Response), in which assumptions, tensions, uncertainties and source references were explicitly organised.

This comparative setup is a recognised validation approach in design research, as it makes shifts in decision making strategies more visible than standalone evaluations (Doellken et al., 2024). All sessions were recorded audio visually and complemented with sketches and notes, which served as external representations of the decision making process. These artefacts made it possible to analyse not only what choices were discussed, but how designers framed problems, integrated Environmental considerations and justified sustainable design choices under different conditions.

Following the workshop, additional interpretative expert interviews were conducted with an AI specialist and a sustainability consultant. These interviews served to contextualise the observed shifts in decision making, particularly in relation to the role of the designer. Rather than being displaced by the AI agent, designers remained responsible for interpretation, judgement and legitimisation. The agent influenced how information was structured and discussed, but not what decisions were ultimately made. Within Research Through Design, such interpretative layers are essential, as prototypes are evaluated not as finished systems but as interventions that can provoke new ways of seeing and decision making. (Zimmerman et al., 2007). Figure 13 (own work, 2025) summarises how the validation phase relates to the research questions and rounds, positioning the workshop not as a test of optimisation, but as an empirical exploration of how AI-mediated structuring can support the integration of sustainable design choices in early-stage corporate real estate development.

Figure 13: research methods overview (own work, 2025)



3.7 Research output

Deliverables and data set

The central deliverable of this research is a structured understanding how sustainable design choices in the early design phase of Dutch corporate real estate change when an AI agent is introduced as reflective support. This contribution materialises in the form of a structured workflow that operationalises the three roles of the Digital Building Master: sensing, steering and support, within a realistic early-stage design context.

The workflow structures how environmental considerations are surfaced, related and justified while design freedom is still high. By comparing human-only decision making with agent-supported decision making in the PV case, the research shows how digital mediation reshapes the way sustainable design choices are articulated, integrated and explained. In this sense, the workflow functions both as an analytical lens on early-stage decision-making and as a transferable intermediate artefact that can be explored further in corporate real estate practice, where similar design choices recur across portfolios.

The prototype underpinning this analysis is intentionally partial yet operational within a clearly defined scope. It reproduces the decision making dynamics around a recurring corporate real estate design choice, balancing operational CO₂ reduction and material-related emissions, while leaving other system components conceptual. This deliberate incompleteness creates methodological control while maintaining ecological validity. The agent does not generate solutions; it structures the interpretative space in which designers, developers and sustainability specialists negotiate sustainable design choices. This makes it possible to empirically compare how integration unfolds with and without AI mediation.

The dataset reflects the iterative Research Through Design logic of the study. Semi-structured interviews with developers, designers and sustainability consultants establish how sustainable design choices are currently approached in corporate real estate, including implicit assumptions, portfolio-driven constraints and fragmented responsibilities. Interviews with AI practitioners define the technical boundary conditions of the agent and clarify where human judgement remains indispensable. Document analysis and manually reconstructed PV calculations based on the EDGE Stadium project provide a verifiable and transparent empirical basis for the prototype. Finally, the validation workshop introduces the agent into a live design setting, producing comparative material that captures shifts in framing, justification and integration of environmental considerations. All empirical materials were anonymised and processed in accordance with TU Delft ethical and data management standards (Appendix K).

Dissemination and audiences

The research speaks directly to practitioners who navigate environmental trade-offs in the early phases of corporate real estate projects, where uncertainties, incomplete data and competing objectives shape the decision making behind design choices. Developers, sustainability consultants, designers and engineers increasingly operate in a landscape that demands transparent justification of environmental implications, particularly when CO₂-related decisions carry both regulatory and reputational consequences. The workflow developed in this study offers a structured way to approach such complexity: it reveals the underlying logic of trade-offs without reducing them to simplified metrics and provides a means to articulate tensions that usually remain implicit. In this sense, the findings are equally relevant for scholars investigating reflective design support systems and for AI developers working on domain-aware decision making systems that operate within environmentally driven design contexts.

Dissemination will occur through multiple channels to ensure accessibility across academic and professional communities. The full thesis will be made publicly available through the TU Delft Repository, enabling open access for researchers, educators and industry stakeholders. Within EDGE, the findings will be shared through internal presentations, supplemented by visual walkthroughs of the workflow and prototype outputs, reflecting the organisation's interest in transparent environmental decision making during early project stages. The visual components serve as communicative tools that allow practitioners to engage with the material in an intuitive and critically reflective manner.

A condensed and practice-oriented version of the work will be shared via LinkedIn to reach professionals active in sustainability, corporate real estate development and AI-assisted design. The workflow is intentionally presented as an intermediate deliverable: sufficiently concrete to demonstrate its analytical and communicative value, yet open-ended enough to accommodate further development by organisations such as EDGE, academic research groups and design teams exploring environmental decision making frameworks. The modularity of the workflow design makes it suitable for future integration into existing decision-support infrastructures or as a foundation for subsequent prototyping efforts.

Central to the dissemination strategy is the explicit positioning of the AI-agent as a reflective support system. The agent does not evaluate options or optimise outcomes; instead, it foregrounds environmental tensions, exposes underlying assumptions and enhances the traceability of design arguments. By doing so, it reinforces rather than replaces human judgement, acknowledging the contextual and value-driven character of environmental decisions in early design. This perspective will be emphasised across all dissemination formats to ensure that the technology is interpreted within the boundaries and intentions defined by the research.

3.8 Trustworthiness and Rigour

Trustworthiness in this research is grounded in a transparent methodological design, formal ethical approval and a responsible approach to the use of AI. This orientation reflects broader developments in design research and environmental decision-making, where scholars increasingly stress that AI systems must strengthen transparency, traceability and accountability rather than introduce new layers of opacity (Lee & Liang, 2024) . To ensure methodological rigour, the study draws on a triangulated empirical foundation consisting of semi-structured interviews, document analysis, manual PV scenario calculations and the validation workshop. Each of these sources captures different dimensions of early-stage environmental decision making, and their combination enables interpretative claims to be examined from multiple angles. This reduces the risk of single-perspective bias and provides a coherent analytical structure in which findings reinforce rather than contradict one another.

The analytical procedures were documented systematically. Progression from raw interview transcripts to thematic interpretations, from manual PV calculations to layered prompting structures, and from workshop observations to decision making analysis was recorded in detail to preserve the continuity and verifiability of the research process. Such documentation aligns with guidelines for responsible AI research, which emphasise that decision-support systems are only legitimate when their decision making paths remain traceable to human interpretation and data provenance (Khan et al., 2021). The validation workshop added an interpretative dimension: by observing how experts assessed, contested and elaborated on the agent's structured decision making, the study incorporated a form of peer scrutiny that strengthens interpretative reliability. This corresponds with international recommendations that position participatory reflection as a cornerstone of responsible AI use within design and planning processes (Ethics of Artificial Intelligence, 2023)

Taken together, the combination of methodological triangulation, explicit documentation and expert reflection provides a rigorous foundation for evaluating how a reflective AI-agent influences early-stage environmental decision making. It ensures that the study's conclusions rest not on a single method or perspective, but on a layered and well-substantiated account of how decision making emerges, shifts and becomes traceable within the context of environmental design decisions.

3.9 Ethical Considerations

Ethical reflection was integrated into the research design from the outset, acknowledging that the study engages directly with human decision making and with the influence of AI on interpretative processes. All activities were conducted under the approval and guidelines of the TU Delft Human Research Ethics Committee (HREC). Participants received clear and accessible information about the aims, procedures and potential implications of the study, and written informed consent was obtained with an explicitly stated right of withdrawal. Personal data were anonymised and stored securely within TU Delft systems in accordance with

the institutional Data Management Policy and the General Data Protection Regulation (GDPR). These procedures reflect a commitment to transparency, accountability and participant autonomy, in line with broader ethical standards for research involving AI-supported decision-making processes (Liang et al., 2024)

The role of the AI-agent was deliberately constrained to safeguard against ethical risks associated with automated influence on value-laden design choices. The agent did not prioritise options, generate recommendations or make normative judgements. Instead, its operation was limited to structuring information, surfacing environmental tensions and clarifying decision making paths. This aligns with literature cautioning against assigning evaluative authority to automated systems in domains where decisions carry contextual, normative and ecological implications (Khan et al., 2021) All outputs produced by the agent were manually reviewed for traceability, coherence and contextual appropriateness. The prompts, workflow configuration and interaction procedures were fully documented to enable audit, replication and critical assessment, consistent with best-practice expectations for transparency and reproducibility in AI research (Liang et al., 2024)

Beyond procedural safeguards, the research adopted a reflective stance toward the broader implications of introducing AI into early-stage environmental decision making. Ethical considerations were therefore treated not only as compliance requirements but as part of an ongoing interpretative responsibility held by the researcher. This resonates with UNESCO's ethical frameworks, which emphasise that responsible AI arises from a continuous reflection on the values embedded in both technological systems and research practices (UNESCO, 2023). The combination of methodological rigour, formal ethical safeguards and reflective engagement strengthens the credibility, integrity and societal relevance of the study's conclusions.

3.10 Key take aways Methodology

The methodology of this study is grounded in the premise that sustainable design choices in the early design phase cannot be adequately understood through conventional analytical approaches alone. Early-stage Environmental decision making is exploratory, value-laden and shaped by interaction rather than optimisation. For this reason, a Research Through Design (RTD) approach was adopted. RTD treats design artefacts as epistemic tools through which decision making becomes visible and can be examined in context. In this study, the AI agent functions as such an epistemic artefact: not as a system that makes decisions, but as a medium for exposing how decisions are formed, structured and justified.

An abductive research logic underpins this approach. Theory and empirical material are developed in parallel, allowing insights from interviews, observations and validation exercises to directly inform the design and refinement of the prototype. This iterative movement between theory and practice enables emerging tensions, assumptions and integration patterns to be captured and

translated into design adjustments, rather than being treated as deviations from a predefined model.

The methodology is structured into three consecutive research rounds. The first round, Exploration, uses semi-structured interviews to examine how sustainable design choices are currently approached in corporate real estate and to concretise the three roles of the Digital Building Master—sensing, steering and support. These roles are subsequently translated into technical and functional requirements for the agent. The second round focuses on Prototype development, in which these requirements are operationalised into a working workflow using a realistic and validated corporate real estate case. The third round, Validation, places the prototype in a live design setting to analyse how it influences the formation, articulation and justification of sustainable design choices in practice.

Together, this structure provides a methodologically coherent framework in which the roles of the Digital Building Master are not only conceptually defined, but also empirically examined within the dynamics of early-stage design. The methodology therefore supports the central aim of the research: to understand how AI can function as reflective support for integrating Environmental considerations into sustainable design choices, while preserving human judgement and design agency.

Chapter 04

Findings in the problem space

4.1 Structure of the problem space.

First coding round: understanding environmental decision-making

Second coding round: interpreting AI and the three roles of the Digital Building Master

4.2 Thematic analysis first coding round.

ESG-Environmental in corporate real estate.

Early Design Phase.

The ESG feedback flow

4.3 Thematic analysis: second coding round.

AI System in Practice

The Digital Building Master

4.4 Contextual activation of sensing, steering and support

4.5 Key take aways H4

4. Findings in the problem space

This chapter presents the results of the first phase: the problem space. As mentioned in Chapter 3, this was designed by conducting semi-structured interviews. The aim of the chapter is to gain a better understanding of environmental decision-making and the use of AI in the Dutch corporate real estate sector.

4.1 Structure of the problem space

First coding round: understanding environmental decision-making

In the first coding round, there is a reconstruction of the two core components of the main question: ESG–Environmental in corporate real estate and environmental decision making in the early design phase. These two domains serve as deductive starting points, sensitising concepts in the sense of Blumer, (1954): they provide direction without dictating the data. They are derived from the theoretical framework and the interview protocol but remain deliberately open so that unexpected patterns can emerge.

Within these two domains, the interviews were coded in an open manner (Strauss & Corbin, 1990), allowing themes that were not defined in advance to emerge spontaneously. The figure illustrates this process: the white boxes show the themes that emerged empirically during the analysis within each domain.

On the ESG–Environmental side, themes such as value creation through sustainability, certification strategy & credits, commercial strategy & market, and stakeholder profiling emerge. These categories show how respondents relate sustainability ambitions to market logic, regulations, and strategic positioning. On the early design phase side, themes such as design philosophy, decision-making & stakeholders, and the design brief & certification in combination with the EDGE-specific development method emerge. These themes reveal how environmental decision making takes shape in practice: through informal interpretations, professional routines and project logics that are rarely made explicit.



Figure 14: Open coding themes (Atlas.TI, 2025)

Methodologically, this approach is in line with (Braun & Clarke, 2006) thematic analysis, in which deductive structures are combined with inductively generated categories to ensure both theoretical coherence and empirical openness. The abductive research logic (Dubois & Gadde, 2002) becomes visible because theory and data constantly refer to each other: the deductive domains provide focus, while the inductive themes form the concrete content of the problem space.

This first round of coding shows how ESG ambitions and design decisions influence each other, but also where tensions arise, often not because of explicit choices, but because of fragmented responsibilities, differences in interpretation and implicit trade-offs that are not recorded anywhere. These vulnerabilities form the empirical basis on which the three roles of the Digital Building Master emerge in the second coding round. They represent the minimum functional package needed to address the structural weaknesses in the problem space.

Second coding round: interpreting AI and the three roles of the Digital Building Master

The second coding round builds directly on the methodological logic described in chapter 3. Because the three preliminary roles of the Digital Building Master, sensing, steering and support, are still conceptual at this stage, the purpose of this round is to examine how these roles can be translated into conditions for an AI-agent in practice. In line with an abductive approach, the coding structure is therefore partly deductive and partly inductive: the roles provide direction, but the data determine their substantive form.

To create analytical coherence, two overarching deductive domains are predefined: AI in practice and Digital Building Master(agent) explained in the three conceptual roles. These green domains in Figure 14: Coding round 2 (Atlas.TI, 2025)Figure 14 reflect the two perspectives necessary for shaping the agent: how practitioners currently experience AI, and how the three roles can meaningfully take shape within real environmental decision making processes according to the experts.

Within these deductive structures, inductive coding is used to identify recurring clusters of meaning. The white themes in the figure, such as attitude towards AI, AI's role in research, and the functional requirements for the agent, are not predefined categories but empirical groupings that emerge from patterns that practitioners repeatedly bring forward. This methodological combination follows (Braun & Clarke, 2006) approach to thematic analysis, in which an initial structure can be combined with inductive code development to retain both theoretical grounding and empirical openness.

The three roles of the Digital Building Master guide the interpretation throughout this round, but they do not constrain it. Instead, they operate as interpretive anchors that help identify when practitioners describe situations that correspond to sensing (surfacing tensions), steering (bringing structure to dispersed decision making) or support (providing a shared reference point). This abductive use of

conceptual roles aligns with Dubois & Gadde's (2002) view that concepts should evolve iteratively through interaction with the field material.

The coding round results in 311 coded segments, grouped into fourteen categories across the four overarching thematic domains 17Appendix D overzicht code groups Atlas.TI. Together, these domains provide insight into how practitioners articulate Environmental decision making, how they experience the opportunities and limitations of AI, and under which conditions they perceive an AI-agent as appropriate, helpful or legitimate in early-stage design work.

By combining deductive structure with inductive emergence, this second coding round consolidates the link between the reconstructed problem space and the agent's conceptual roles. The resulting thematic structure forms the analytical bridge to the solution space, where these roles are translated into technical and functional requirements for the prototype (Section 4.2).



Figure 15: Coding round 2 (Atlas.TI, 2025)

4.2 Thematic analysis first coding round

This section explains the thematic domains from Figure 13 and Figure 14 with representative quotes that show how participants express environmental decision making and AI in practice. The detailed coding and full analysis are included in Appendix E: Example sub codes by group

ESG-Environmental in corporate real estate

The literature in chapter 2 shows that environmental objectives in corporate real estate development are structured by tensions between ecological values, corporate feasibility and institutional frameworks. The interviews confirm this, but above all reveal how these tensions play out in practice. In particular, the role of certification systems emerges as a structural force that both supports and limits ambitions. The empirical themes discussed below are derived from recurrent patterns across interviews.

Certifications such as BREEAM and WELL are seen by the interviewees as necessary to make sustainability measurable and legitimate, but at the same time they guide the design. As one developer said:

"BREEAM is sometimes just the playing field. You know where the credits are and that guides the choices you make."

This is in line with the literature that describes certification systems as 'formative frameworks' that not only measure but also define what sustainability means. The interviews show that in practice this leads to prioritisation of scoreable elements.

The EN01 credit for local energy generation is often decisive in this regard. In the case of an urban project, the location proved so unfavourable that achieving a higher score was virtually impossible. One project manager put it this way:

"You want to build a super sustainable building. But there's not much you can do about the Energy part... and that means you can't achieve Excellent or Outstanding."

This makes it clear that environmental goals are not only ecological, but also methodological: what counts is what can be demonstrated. In addition, the interviews confirm an area of tension that is mentioned theoretically in the literature but becomes concrete and tangible here: operational CO₂ versus material-related CO₂. A design manager said about this:

"Producing the panel is, of course, very CO₂-intensive... and on the material side, we simply have a limit value."

The choice for solar panels can therefore be both good and bad, depending on the lens through which one views it. This shows that environmental goals do not have a single direction; they always require interpretation.

Finally, the interviews make the market logic behind environmental decision-making explicit. One developer stated: "With that extra investment in, for example, a wooden structure, we also need to be able to charge a higher rent. Otherwise, it simply won't happen."

Here, the tension between idealism and feasibility becomes concrete and recognisable: the question is not whether a solution is better, but whether it can be supported within the project model. The literature mentions this tension, but the interviews show where and when it becomes decisive.

Early Design Phase

Against this backdrop, it becomes clear why sustainability ambitions at the start of a project often appear both strongly, formulated and vulnerable. The interviews show that terms such as Paris Proof serve as a guideline, but without a common definition.

"One designer said: "Paris Proof has been the standard for us for some time. Only... what people mean by it exactly differs." The ambition is shared, but the meaning is negotiable."

The early design phase considered as the SO/VO phase appears to be the decision-making moment when this flexibility takes effect. This is the moment when choices about construction, building mass and energy concept become virtually final. As one developer summarised:

"You'd rather go for something that stands than something that is fantastic in theory but doesn't come to fruition."

The interviews directly confirm the theoretical literature: decisions made early on are difficult to reverse later and therefore determine the environmental profile of the building. Two recurring areas of tension consistently emerge in this phase:

- Operational vs. material-related CO₂
- Energy performance vs. quality of use

What the interviews add to the theory is the insight that these considerations are not made explicit in practice. They arise in design meetings, spreadsheets and informal conversations. One project manager put it this way:

"The knowledge is carried by the person. When someone leaves, you lose that context."

This leads to the central vulnerability: It is not the ambition that shifts, but the meaning of the ambition shifts along with the design decisions, without this being explicitly discussed. It is precisely this mechanism that forms the core of the ESG ambition feedback loop, which is visualised and explained in the following section.

The ESG feedback flow

Based on both theoretical insights and interview findings, the decision-making process has been reformulated into a deep understanding of environmental decision making. The ESG ambition feedback loop gives more detailed insights in the results from the environmental research. Figure 15 presents the ESG ambition feedback loop developed in this research. The model synthesises theoretical insights and interview findings and reconstructs how environmental ambitions shift during the early design phase. At the start of a project, ambitions are typically framed in broad, high-level terms, supported by regulatory incentives and financial frameworks. As the figure shows, however, these ambitions immediately come under pressure once design choices must be made, particularly in the well-documented tension between operational and embodied emissions.

This design pressure introduces costs, which in turn feed directly into feasibility assessments. Interviewees described how such considerations rarely trigger explicit renegotiation of ambition. Instead, ambition is gradually adjusted informally, influenced by feasibility thresholds, market expectations and certification logic. Certification systems such as BREEAM, which literature describes as formative frameworks, play a visible role here: what counts as 'environmental' becomes tied to what can be demonstrated through credits. As one interviewee put it:

"BREEAM is sometimes just the playing field... that guides the choices you make."

In the lower part of the loop, the figure visualises how this informal renegotiation contributes to knowledge loss across handovers. When decision making steps are not documented, later teams reinterpret earlier intentions through flexible terminology and weak traceability. The meaning of the original ambition becomes

ambiguous, eventually feeding back into the next formulation of ambition levels. Ambition thus changes without ever being formally lowered—its definition drifts.

The feedback loop therefore does not depict a failure of sustainability goals, but a failure of visibility. It shows how environmental intentions become unstable when costs, certification incentives and feasibility considerations shape decisions implicitly rather than explicitly. The implication is clear: strengthening environmental performance requires not only better targets, but greater transparency and traceability of the decision making that shapes those targets in the early design phase.

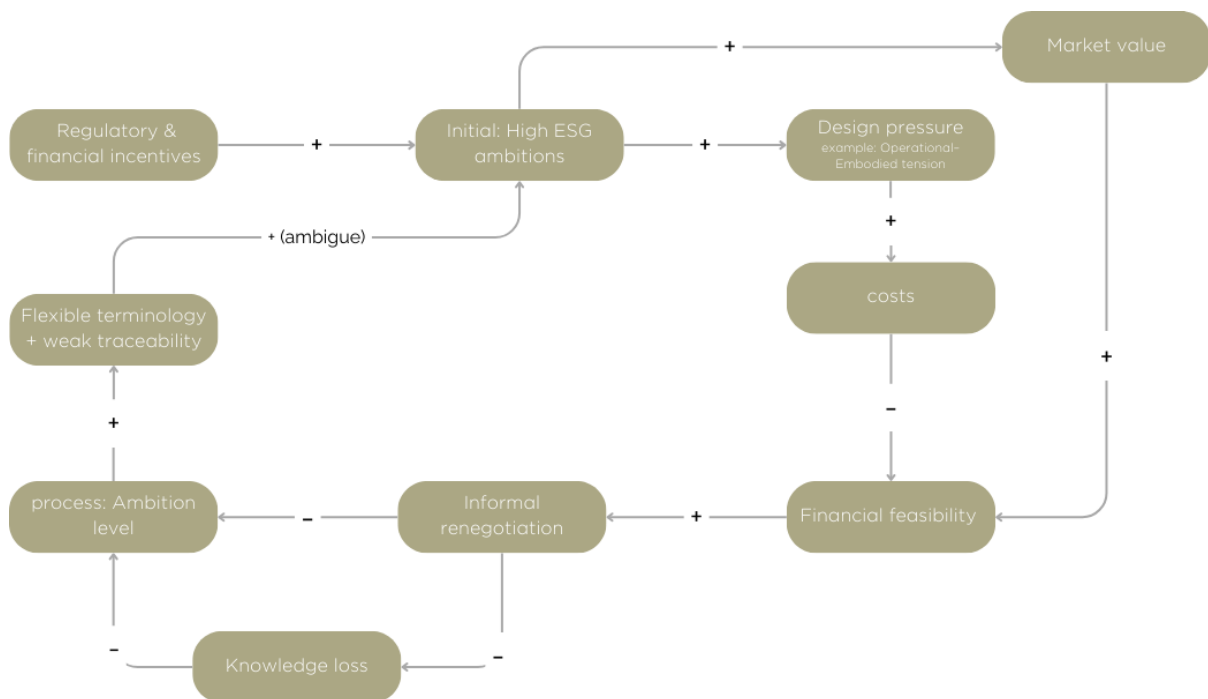


Figure 16: ESG feedback loop (own work, 2025)

4.3 Thematic analysis: second coding round

The second coding round focuses on how experts make sense of the three roles of the Digital Building Master when they speak about AI in the context of early design work. The first coding round revealed where environmental decision making becomes vulnerable; in this round, the focus shifts to how an AI agent should behave in that field of influence to be seen as appropriate and legitimate.

Therefore, this round does not look again for tensions in environmental decision making, but for interpretations, expectations and boundaries that professionals associate with AI. The deductive domains of AI in practice and the three roles of the Digital Building Master(agent) form the structure here, while inductive themes arise within AI that show how designers, developers and consultants talk about reliability, control, transparency, integration into existing workflows and the scope for human judgement.

The result is a thematic structure that clarifies the conditions under which a Digital Building Master is considered helpful. This coding round thus forms the analytical link between the problem space and the functional specifications of the prototype in the next chapter

AI System in Practice

The interviews show that practitioners in the early design phase do not view AI as a decision-maker but as a digital system that provides structure, clarity and analytical support. The attitude towards AI is cautiously positive: respondents welcome this form of assistance, as long as human judgment remains central.

*"I want insight. I want to make the choice myself."
Another said: "We have to make the decisions."*

This perspective mirrors insights from human-in-the-loop research, which demonstrates that AI is most effective when it assists rather than replaces human judgement. (Wang et al., 2024) In practice, environmental and design choices are treated as contextual and interpretative. Interviewees therefore stress that AI should not define outcomes but should help make decision making visible.

This expectation becomes concrete in how respondents describe useful AI behaviour. They imagine AI as a system that helps articulate questions, refine definitions and expose implicit assumptions. One interviewee described this explicitly as support in problem framing:

*"To generate research questions." Another pointed to definitional clarity:
"It would already help to get definitions of terms right."*

This aligns with recent work framing AI as a semantic mediator that connects documents, knowledge domains and stakeholder perspectives (Wang et al., 2024) Interviewees also emphasise process-related benefits, such as organising project

information, linking requirements to design implications and generating scenario-based overviews. Importantly, these scenarios are not expected to steer decisions, but to enable informed discussion.

As one respondent summarised: the agent may explore options, “but the final decision is made by humans.”

These expectations translate into concrete system requirements. Practitioners emphasise that AI-generated insights must be transparent and traceable. They want to understand what evidence underpins a scenario.

“I want to be able to see what scenarios are based on.”—and where the underlying data originate “You must be able to trace where the data comes from.”

This requirement aligns directly with Retrieval-Augmented Generation, which separates project documentation from model knowledge and thereby enhances traceability (Sanderson, 2024).

Integration with existing tools and workflows forms another recurrent theme. Respondents stress that the agent must function within established environments such as BIM models, spreadsheets and Programme of Requirements documentation.

“How do you bring all disciplines together in one system?” one interviewee asked. Another added a concrete operational expectation: “It must scan the design and produce the report five minutes later.”

These comments reflect broader adoption research, which shows that usability and workflow compatibility—rather than model sophistication—determine whether AI becomes part of daily practice (Bansal et al., 2021). The interviews add an important nuance to this: speed is not a convenience factor but a precondition for adoption. In early design work, decisions unfold rapidly and information needs shift from meeting to meeting; practitioners therefore emphasise that an AI-agent must operate at the same pace as the design process itself.

Taken together, these insights show that AI is viewed as most valuable when it acts as a reflective support system rather than an autonomous actor. Practitioners embrace AI when it organises information, makes assumptions explicit and preserves decision making, but reject it when it attempts to optimise or decide. This understanding directly informs the elaboration of the Digital Building Master in the roles of sensing, steering and support, which together translate this practitioner perspective into concrete functional expectations for the agent.

The Digital Building Master

The ESG feedback loop from Section 4.2 shows that environmental goals in the early design phase are not the problem, but rather the way in which they are handled. In the second coding round, these findings are concretised by investigating how an AI agent can support these vulnerable moments. The three roles introduced earlier: Sensing, Steering, and Support, form the interpretative lens for this. The interviews show how these roles take on meaning in practice and under what conditions they are perceived as legitimate and useful by designers, developers, and consultants.

Sensing

Sensing refers to the early recognition of tensions that, as becomes apparent in the feedback loop, often only surface at a late stage. Interviewees emphasise that such conflicts—such as the tension between operational and material-related CO₂—do exist but are discussed too late.

“It just means that there is a design challenge there.”

At this stage, the agent is particularly valued when he identifies early on where design choices clash, so that alternatives are still open. As one participant put it:

“Would particularly like to understand where the points of attention lie. What should I keep in mind?”

Sensing shifts awareness forward and makes it possible to discuss tensions before they become “ingrained” in cost estimates or technical choices.

Steering

Steering ties in with the insight that many design decisions arise in informal conversations where arguments quickly fade away (Section 4.1). The interviewees recognise themselves in a role for the agent in which the agent structures the conversation without giving direction:

“So, the scenarios are worked out by the agent, but the final decision is made by humans.”

The agent helps by making assumptions visible, juxtaposing alternatives and asking questions at the right moment. This prevents decision making from disappearing in interactions between disciplines:

“So, what emerges is not the only way; it's about the conversation that follows.”

Steering thus ensures a shared and stable line of thinking, precisely at moments when the feedback loop shows that ambitions are becoming diffuse.

Support

Support focuses on the problem that is one of the greatest risks in the feedback loop: knowledge loss when teams change or phases overlap. As one interviewee summarised:

| *"It travels with people."*

The agent is seen as a system that retains decision making and definitions, so that environmental goals not only remain formally the same but also remain consistent in terms of content when new team members join.

| *"It helps if we have the same definition in the first place."*

Support thus safeguards interpretation and information, which strengthens the continuity of environmental decision making.

The interviews make it clear that the three roles of the Digital Building Master fit in with recognisable moments in the design process when decision-making comes under pressure. Practitioners not only describe that these moments exist, but also what kind of contribution an AI agent must make at such moments to be perceived as supportive and legitimate.

This means that sensing, steering and support should not be understood as continuous or uniform functions, but as roles that become relevant under specific conditions. The following section therefore considers when and for whom each role becomes dominant, and how this variation relates to different moments of Environmental integration in the early design process.

4.4 Contextual activation of sensing, steering and support

Figure 17 visualises how the relevance of sensing, steering and support shifts across the project phases of corporate real estate development. The graph does not suggest that these roles operate sequentially or replace one another. Instead, it shows how their relative impact varies depending on the phase, the dominant design pressure and the actors involved. The interviews confirm that practitioners recognise these shifts and implicitly adapt their way of working accordingly.

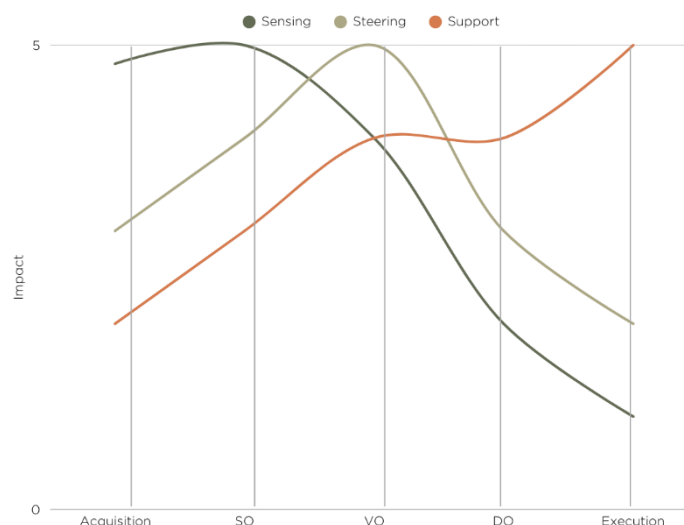


Figure 17: Impact - project phase each role (own work)

Figure 17 visualises how the relevance of sensing, steering and support develops across successive project phases, set against a scale of impact (0-5). The horizontal axis follows the project phases from acquisition and preliminary design (SO), through provisional design (VO), to final design and implementation (DO). The vertical axis does not represent the intensity of use, but the extent to which a role contributes meaningfully to dealing with uncertainty and accountability in that phase. The graph shows that the three roles never disappear, but that their relative importance shifts as uncertainty decreases and formal legitimacy increases.

In the acquisition and early SO phase, the project is located on the left side of the graph, where uncertainty is at its highest and formal accountability is at its lowest. This explains why sensing reaches its highest value here. The graph shows a clear peak in sensing in this phase, because this is the moment when environmental assumptions are still implicit and decisions about building mass, construction and materialisation cause a strong lock-in. The interviews confirm this pattern. A sustainability advisor points out that less experienced developers find it difficult to oversee the implications of combined sustainability ambitions:

"For a novice development manager, this can be an assistant who says: pay attention, ask these questions now." In terms of the graph, this means that sensing functions here as a mechanism for making uncertainty explicit before choices are finalised.

When the project moves on to the preliminary design phase, the graph shows that sensing decreases slightly, while steering increases significantly. This crossover is analytically relevant. Uncertainty decreases as scenarios, calculations and designs become more concrete, but it does not disappear. At the same time, there is a growing need to align different perspectives and disciplines. The graph therefore shows that steering has its greatest impact in this phase. Interviews with designers and consultants make it clear that this is related to semantic frictions and the renegotiation of ambitions. A consultant describes the need to gain insight into which 'buttons' there are and what their effects are overall. Steering here does not support decision-making, but rather the structuring of the conversation in which uncertainty is converted into provisional certainty.

What is striking in the graph is that support and sensing intersect in the VO. This marks a moment of transition. While sensing remains necessary to reveal remaining tensions, support begins to increase in importance because choices now must be explicitly recorded. This explains why support never drops to zero: even in phases of high uncertainty, there is already a latent need to retain decision making, even if it is not yet dominant.

In the DO and implementation phases, the graph shifts further to the right. Uncertainty is relatively low here, but the need for formal accountability is high. This explains why support reaches its highest value in this phase. The graph shows support as the dominant role because sustainable design choices must now stand

up to new team members, implementing parties and external stakeholders. A sustainability advisor recognises the importance of this:

"That during the construction phase, you can reason back to what was discussed in the SO phase."

This applies even more strongly to financiers. A real estate developer indicates that sustainable ambitions are directly linked to finance ability, making traceability and evidence crucial. In graphic terms, the function of the AI agent shifts here from dealing with uncertainty to ensuring certainty.

The fact that no role is ever at zero in the graph underlines an important empirical insight: certainty in corporate real estate is never absolute. New risks continue to arise in later phases, for example due to cost increases or implementation constraints, which means that sensing remains relevant. Similarly, steering remains necessary to discuss changes without losing coherence. The graph therefore does not visualise a linear process, but a dynamic field of tension in which the AI agent places a different emphasis each time.

Taken together, the graph shows that the digital building master does not perform a constant function but is activated situationally based on the relationship between uncertainty and legitimacy. By explicitly linking sensing, steering and support to the project phase and stakeholder position, it becomes clear how the AI agent can contribute to the integration of sustainable design choices without reducing them to fixed decision rules or optimisation.

4.5 Key take aways H4

The analysis in this chapter provides an answer to sub-question 1: How are Environmental (E) considerations currently understood and handled by design teams and where do gaps arise that could be relevant for AI-agent support? The findings show that Environmental ambitions do not disappear during the early design phase. Instead, their meaning gradually shifts as projects progress. Targets such as Paris Proof or BREEAM Excellent formally remain in place, but their interpretation changes under the influence of what is perceived as technically feasible, financially acceptable, or certifiable at a given moment. This shift rarely results from an explicit decision. Rather, it emerges through a series of small design choices, technical assumptions, and cost-related adjustments. Tensions are experienced but seldom articulated. As a result, Environmental ambition subtly adapts without being consciously reconsidered, leaving the team largely unaware of the change.

The revised ESG feedback loop makes this dynamic visible and discussable. By introducing three intervention streams, the loop does not alter the ambition itself but changes the conditions under which that ambition is interpreted and integrated into the project. Each stream intervenes at a different moment in the design process, responding to distinct forms of uncertainty and coordination pressure.

Sensing

In the original process, tensions within the Environmental dimension tend to surface only once the design has largely stabilised. At that point, design commitments, cost pressures, and financial feasibility reinforce one another. Environmental frictions therefore emerge late, when alternatives are politically, technically, or financially difficult to reopen. This timing allows ambitions to drift towards what appears feasible, without that shift being recognised as a deliberate choice.

The introduction of sensing alters this timing. The Digital Building Master makes Environmental tensions traceable as soon as they arise, by linking design intentions to their potential side effects at an early stage. When a design direction has simultaneous structural, material, and energy consequences, the agent signals this during the SO or early VO phase, before decisions are locked in. In this way, the feedback loop is opened at the front end of the process. Tension no longer returns only through downstream cost or feasibility constraints but becomes visible at the same moment as the ambition is first articulated.

Sensing does not prioritise one Environmental parameter over another. Its function is to shift the moment of consideration. Trade-offs become visible while the design space is still open and uncertainty is high. This prevents ambition from eroding through informal renegotiation or knowledge loss, as observed in the original loop. Instead of an implicit drift, the process allows for explicit exploration. Environmental decision making becomes investigative rather than reactive, weakening the reinforcing effect of late-stage pressure.

In the revised ESG feedback loop, sensing therefore operates as an early interruption mechanism. It counters silent ambition erosion by making Environmental tensions tangible before the project narrows into a fixed trajectory.

Steering

Where sensing makes Environmental tensions visible, steering focuses on how these tensions are *integrated* into sustainable design choices. In the original feedback loop, integration often fails because Environmental ambitions are discussed in fragmented and informal ways. Terms such as Paris Proof, CO₂-neutral or low-impact materials circulate in meetings without shared meaning, while relevant information is spread across documents, disciplines and moments in time. As a result, Environmental considerations are not actively integrated with spatial, technical and financial choices, but gradually adjusted to what seems feasible in the moment. Sustainable design choices emerge implicitly, shaped by conversational dynamics rather than deliberate alignment.

The steering role of the Digital Building Master intervenes precisely at this point. Instead of resolving trade-offs, the agent structures how Environmental considerations are connected to other design dimensions. By clarifying definitions, surfacing assumptions and presenting comparable scenarios, the agent helps

teams understand how different interpretations of Environmental ambition lead to different design consequences. This supports integration by making explicit how Environmental goals relate to material strategies, energy concepts and spatial decisions, rather than allowing them to drift apart during discussion.

Within the ESG feedback loop, steering acts as an integrative force. Where the original loop allows Environmental ambition to fragment under the pressure of feasibility and time constraints, the agent creates a shared interpretative space in which sustainable design choices can be weighed coherently. Financial, technical and Environmental arguments are not merged into a single metric but positioned alongside each other in a structured way. This enables teams to negotiate trade-offs consciously, instead of absorbing them implicitly into the design.

Steering therefore does not reduce complexity or uncertainty. Instead, it supports teams in holding multiple considerations together long enough to make integration possible. Sustainable design choices remain the outcome of human judgement, but the conditions under which these choices are formed become more transparent and less dependent on individual framing. In this way, steering stabilises the feedback loop without closing it, allowing Environmental considerations to remain actively connected to design decisions as uncertainty gradually decreases and the project moves toward definition.

Support

In the original process, the coherence of Environmental objectives does not disappear because sustainable design choices are deliberately reversed, but because the logic that integrated those choices gradually dissolves. Early decisions are often well argued now they are made, yet their rationale remains embedded in conversations, individual expertise or temporary project constellations. As the project progresses through phases, teams change, documentation fragments and attention shifts toward delivery. What remains visible is *what* was decided, while *why* certain Environmental options were selected over others becomes increasingly opaque. As a result, Environmental ambition erodes not through formal redefinition, but through loss of interpretive continuity.

The support role of the Digital Building Master intervenes at this point by preserving the integrative logic behind sustainable design choices. Rather than storing outcomes as static decisions, the agent links Environmental considerations to their design context, alternatives and justifications. By documenting why specific material strategies, energy concepts or spatial solutions were preferred, and under which constraints other options were set aside, the agent maintains the connective tissue between ambition and execution. This transforms Environmental integration from a momentary alignment into a traceable design narrative.

Within the ESG feedback loop, support functions as a stabilising memory structure. Where earlier phases are characterised by uncertainty and exploration, later phases demand accountability, transferability and legitimacy. The agent enables this transition by making Environmental decision making accessible across time,

roles and organisational boundaries. New team members, external reviewers or financiers can reconstruct how sustainable design choices emerged without reopening resolved discussions or relying on informal recollection.

Support therefore does not lock decisions in place, nor does it prescribe future action. Instead, it safeguards the coherence of Environmental ambition by ensuring that integration remains intelligible beyond the moment of choice. In doing so, it allows sustainable design choices to persist as meaningful commitments throughout the project lifecycle, rather than dissolving into isolated technical outcomes.

Structure

Figure 17 shows Environmental decision-making as a feedback structure rather than a linear sequence. Environmental ambitions remain formally present throughout the process, but their meaning shifts as design pressure, informal renegotiation and financial feasibility interact. Integration therefore does not fail through explicit trade-offs, but through gradual reinterpretation of what ambitions are understood to require.

Within this structure, sensing, steering and support operate as overlapping integrative roles, not as sequential steps. Sensing links initial Environmental ambitions directly to emerging design pressure, such as the operational-embodied carbon tension, at moments of high uncertainty. By making frictions visible early, it prevents ambitions from being adjusted only through late-stage cost feedback.

Steering intervenes where interpretation becomes unstable. In the central loop of flexible terminology and informal renegotiation, Environmental ambitions drift because concepts lack shared meaning. Steering stabilises this loop by organising interpretation, aligning assumptions and structuring discussion, so that ambitions remain integrated within design conversations rather than diluted through ad hoc compromise.

Support acts where certainty increases but coherence is at risk. As projects move towards feasibility, earlier decision making tends to disappear through phase transitions and team changes. Support preserves the logic behind Environmental choices, linking decisions to their original context and maintaining traceability across phases and stakeholders.

Chapter 05

Findings in the solution space

5.1 Functional Architecture of the AI agent

Perception Layer

Reasoning / Planning Layer

Tool Layer

Specialized Functions Layer

Memory / Provenance Layer

5.2 Operational requirements

5.3 Key take aways H5

5. Findings in the solution space

This chapter translates the conceptual roles of the Digital Building Master into a functional architecture that can operate within early-stage Environmental design practice. While the previous chapters clarified when sensing, steering and support become relevant in the design process, this chapter explains how these roles are technically and functionally enabled within an AI-agent configuration.

The architecture should not be read as a sequential system design. Instead, it provides a functional framework that remains constant, while the relative activation of its components shifts depending on project phase, uncertainty and the need for accountability. In this way, the architecture reflects the situational logic identified in Chapter 4: sensing, steering and support are never fully absent, but their contribution varies as Environmental integration moves from exploration toward formalisation.

5.1 Functional Architecture of the AI agent

Figure 17 presents the functional architecture of the AI agent, structured into five interconnected layers: Perception, Decision making/Planning, Tool, Specialised Functions and Memory/Provenance. This structure was developed inductively, grounded in empirical insights into how AI systems reason, retrieve information and support design discussions in practice. Together, the five layers operationalise the three roles of the Digital Building Master without hardwiring them to specific project phases.

The architecture translates the agent cycle introduced in Chapter 2, perceiving, interpreting and acting (Huang & Ai, 2023), into a configuration that can adapt to shifting design conditions. As shown in Chapter 4, early phases are characterised by high uncertainty and implicit assumptions, while later phases demand traceability and legitimacy. Rather than changing the system itself, these shifting conditions activate different layers more prominently within the same architecture.

The architecture is therefore designed to strengthen Environmental integration by structuring fragmentation and making assumptions explicit, not by optimising outcomes or automating decisions. It supports the design team's reflective capacity by enabling early detection of tensions, structured interpretation of trade-offs and preservation of decision making over time. In this sense, the five layers also serve as the analytical framework for the functional requirements imposed on the agent, which are elaborated in the remainder of this chapter and tested in the PV case in Chapter 6.

Perception Layer

The Perception Layer forms the sensory input of the system and operationalises the first step of the agent model (Huang & Ai, 2023). In Figure 17, this layer corresponds to the light pink zone, where the role of sensing is most prominent. Its

primary function is not data processing as such, but the early recognition of Environmental relevance within fragmented and incomplete design information.

This layer becomes particularly active in phases characterised by high uncertainty, such as acquisition and early SO, as identified in Chapter 4. At this stage, Environmental assumptions are often implicit, and design intentions are formulated before their consequences are fully understood. The Perception Layer translates raw inputs into an enriched design context in which Environmental domains are explicitly activated. This allows early tensions to surface before design choices become fixed.

Technically, the agent integrates heterogeneous inputs into a shared semantic space, in line with principles of input modality management and distributed semantic representation (Mikolov et al., 2013) Through Natural Language Understanding and attention mechanisms, relevant entities, criteria and potential frictions are identified, including missing assumptions or unarticulated dependencies. (Vaswani et al., 2017) Rather than resolving these tensions, the layer signals where Environmental complexity emerges.

In this way, the Perception Layer implements sensing as an early warning function. It shifts Environmental consideration forward in time, creating the cognitive basis for structured decision making, retrieval and scenario exploration in later layers, as demonstrated in the PV case in Chapter 6.

Decision making / Planning Layer

The Decision making/Planning Layer forms the analytical core of the AI agent. Information surfaced in the Perception Layer is structured here into explicit lines of decision making that can be shared and discussed within the design team. In Figure 17, this layer corresponds to the orange zone, where steering becomes dominant: not by directing outcomes, but by organising how Environmental considerations are interpreted and connected to design choices.

This layer is most relevant in phases where uncertainty starts to decrease but interpretative pressure increases, particularly during the SO-VO transition identified in Chapter 4. At this stage, multiple disciplines interact and Environmental ambitions risk fragmenting through informal negotiation. The Decision making/Planning Layer intervenes by stabilising interpretation: it structures arguments, aligns assumptions and makes differences between alternatives explicit without prescribing a preferred solution.

Technically, this aligns with the description of the agent as a meaning-construction system based on pattern recognition in LLMs (Guo et al., 2024). The layer operates as a Cognitive or Planning Engine (Krishnan, 2025), translating semantic input into coherent decision making structures. Through Knowledge Representation and Decision making (KRR), recognised patterns are converted into comparable design implications. This becomes visible in Chapter 6 when the agent separates

compliance-driven logic (e.g. ENE01) from broader Environmental effects such as embodied CO₂.

Explicit decision making strategies, such as Chain-of-Thought and Tree-of-Thought, are used to make interpretative steps traceable rather than implicit. These mechanisms allow sustainable design choices to be explored as structured scenarios instead of collapsing into ad hoc compromise. The Planning Engine sequences these decision making steps while remaining responsive to new input from the design team, supporting iterative dialogue rather than linear optimisation.

In this way, the Decision making/Planning Layer operationalises steering as an integrative function. It does not resolve Environmental tensions but creates the conditions under which sustainable design choices can be consciously weighed, compared and connected across disciplines before they harden into fixed decisions.

Tool Layer

The Tool Layer is the place in the architecture where the agent leaves the system boundaries and gains access to external sources. In line with Chapter 2, this is necessary: an LLM can interpret patterns but has static knowledge and can hallucinate when crucial information is missing (Gao et al., 2023). For sustainable design choices to remain credible, Environmental decision making must therefore be grounded in verifiable external sources.

In Figure 17, this layer aligns primarily with the role of support, which becomes increasingly relevant as projects move from uncertainty toward accountability. The agent functions here as a mediator between interpretative decision making and formal Environmental reference systems, such as standards, emission databases and project-specific technical data. Rather than interpreting or prioritising these sources, the agent links them explicitly to the design questions under discussion.

In the literature, this mechanism is described as Tool Invocation (Shen et al., 2007). The agent can request data via APIs; when APIs are missing, UI understanding techniques can be applied. The core of this layer is Retrieval-Augmented Generation (RAG), as explained in Chapter 2 (Sanderson, 2023). RAG semantically searches external documentation, selects verified fragments and prevents hallucination (Sun et al., 2025, pp. 5-6).

This layer becomes particularly dominant in later project phases, as visualised in Figure 16, when sustainable design choices must be substantiated toward external stakeholders. In Chapter 6, this is operationalised by linking ENE01, NTA 8800, PVGIS and NMD data directly to the PV case through controlled chunking, retrieval and augmentation. The Tool Layer thus ensures that Environmental claims remain auditable and transferable without constraining interpretation.

Specialized Functions Layer

The Specialised Functions Layer contains the internal thinking functions that enable the agent to perform specific subtasks. In Figure 17, this layer corresponds to the green zone. The layer mainly supports the role of steering: it organises alternatives, makes assumptions visible and unfolds decision making without giving direction. In the literature, such configurations are described as ensembles of specialised modules (Sun et al., 2025, pp. 6–7). Here, they function as internal routines within a single central agent. Using internal selection logic (Zhou et al., 2024, p. 10), the agent determines which module is activated when and how information is brought together.

A key mechanism is task decomposition: complex questions are broken down into smaller cognitive sub-steps. This mechanism is essential for Chapter 6, where the PV question is automatically broken down into an energy component (ENE01/NTA 8800) and a material component (NMD emission factors).

This layer thus operationalises the theoretical building blocks from Chapter 2: transformer architectures provide semantic structure (Vaswani et al., 2017), vector representations ensure semantic consistency (Mikolov et al., 2013) and cognitive decomposition techniques enable transparent sub-analyses.

Memory / Provenance Layer

The Memory/Provenance Layer forms the agent's memory and supports both sensing and support. Whereas other layers focus primarily on perception, interpretation and action, this layer ensures that information is not lost and that decisions can be traced later. This is important because LLMs can generate meaning, but do not automatically have a stable or controllable memory (Gao et al., 2023). This layer therefore stores three types of knowledge: interaction history, general knowledge patterns and semantic representations that enable relevant information to be retrieved quickly (Mikolov et al., 2013). In addition, the layer contains functions that ensure that documents, assumptions and source rules remain linked. This creates a consistent trail of information, which is exactly what is needed in Chapter 6 when the agent shows what its scenarios are based on.

In Figure 17, this layer supports both sensing and support. Early in the process, it enables sensing by preserving emerging uncertainties and flagged tensions. Later, it becomes central to support, as accountability increases and sustainable design choices must be explained to new team members, contractors or financiers.

By linking documents, assumptions and source references, the Memory/Provenance Layer creates a continuous line of meaning across project phases. This makes it possible to reconstruct why certain Environmental options were chosen and under which conditions alternatives were set aside. As demonstrated in Chapter 6, this traceability allows the agent to function as a reflective support system rather than a decision authority.

Functional architecture of the AI agent across roles

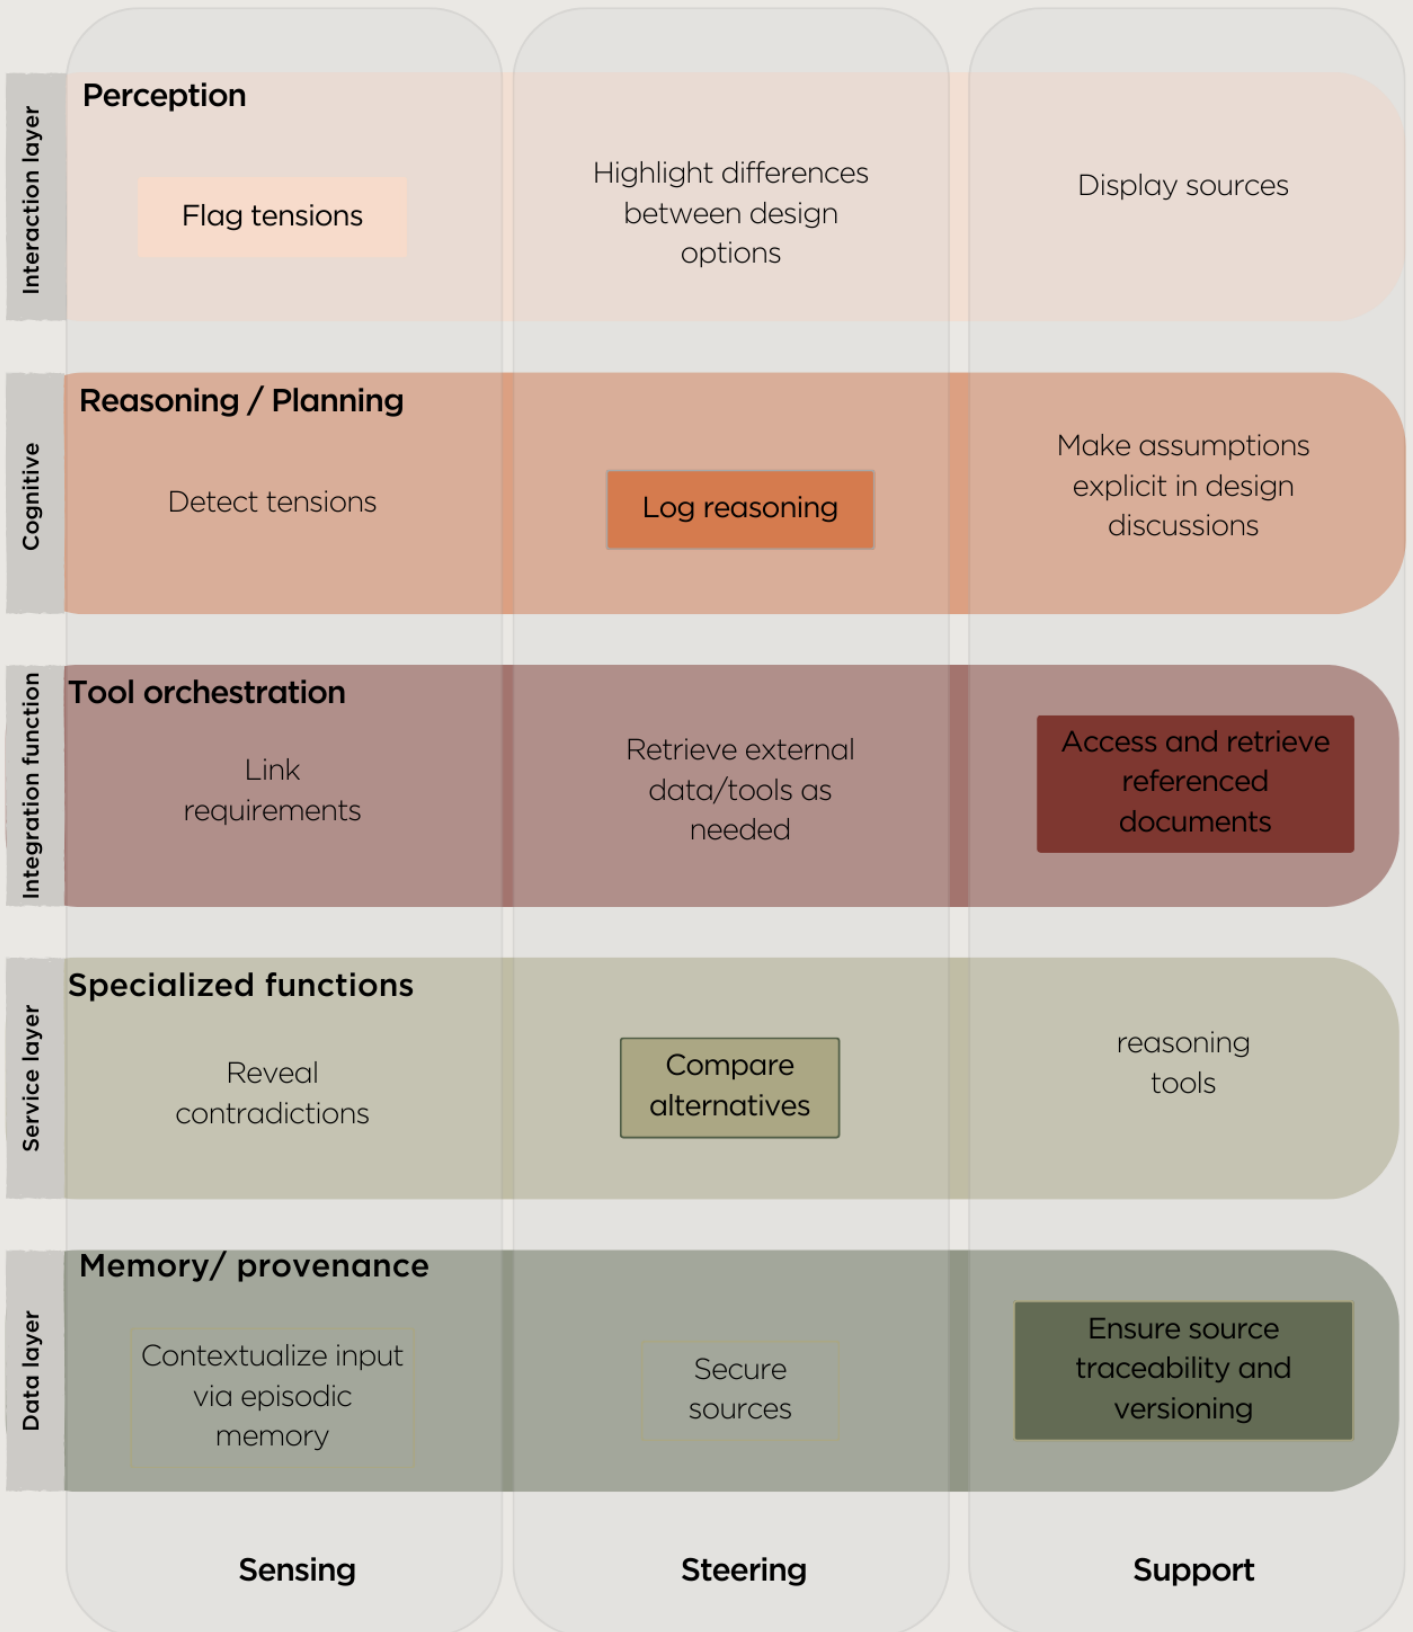


Figure 19 Functional Architecture of the AI agent across roles (own work , 2025)

5.2 Operational requirements

The requirements in Table 4 translate the three roles of the Digital Building Master into concrete behavioural conditions across the five layers of the functional architecture. Rather than optimising or automating sustainable design choices, they specify how the agent should act to make Environmental decision making more open to discussion in early-phase corporate real estate development.

Table 3: Requirement framework outlining the AI-agent's Environmental roles and functions. (own work, 2025)

Layer	ID	Requirement
PERCEPTION	SE-P01	The agent should surface environmental tensions and missing assumptions early, without prioritising one interpretation over another.
	ST-P02	The agent should present design options side-by-side in comparable formats to support discussion rather than convergence.
	ST-P03	The agent should prompt questions that make implicit assumptions explicit.
	SU-P04	The agent should display the source and justification of information directly within the interface
DECISION MAKING	SE-R01	The agent should detect when environmental objectives are affected by early-phase design decisions.
	SE-R02	The agent should surface potential environmental risks before design lock-in.
	ST-R03	The agent should log argumentative steps and decision rationales in a structured, traceable format.
	ST-R04	The agent should highlight when key assumptions change between iterations.
	SU-R05	The agent should retain decision making across project phases to prevent knowledge loss.
TOOL ORCHESTRATION	SE-T01	The agent should compare the environmental requirements to regulatory and contextual constraints to reveal inconsistencies.
	SE-T02	The agent should map environmental objectives to specific design parameters to support evaluability.
	ST-T03	The agent should link each decision to evidence (calculations, benchmarks, LCA/EPG references) without interpreting the evidence.
	SU-T04	The agent should maintain persistent access to referenced documents throughout the project timeline.
FUNCTIONS	SE-S01	The agent should detect when key environmental terms (e.g., Paris Proof, biobased, net-zero) are used with divergent meanings.
	SE-S02	The agent should flag conceptual contradictions between sources or interpretations.
	ST-S03	The agent should structure disagreements by making interpretive differences transparent rather than resolving them
	ST-S04	The agent should generate reflection prompts that maintain design plurality rather than converge on a single solution.
MEMORY	SU-M01	The agent should provide stable, shared definitions and benchmarks agreed upon by the design team.
	SU-M02	The agent should surface domain knowledge proportionally to the expertise level of the interacting user.
	SE-M03	The agent should signal uncertainty where retrieved or stored data is incomplete, ambiguous, or outdated
	SU-M04	The agent should preserve data lineage, so that referenced evidence is always traceable back to its source
	SU-M05	The agent should operate within representational constraints to avoid duplication of sensitive project information.

5.3 Key take aways H5

SQ2 asks which functional qualities an AI agent must possess to meaningfully support Environmental (E) decision making in early-phase design processes. The findings show that effective support does not depend on a single capability, but on a coherent configuration of five interrelated qualities, each embedded in one layer of the functional architecture and activated in relation to sensing, steering and support.

First, the agent must be capable of early Environmental perception. The Perception Layer enables the agent to interpret fragmented inputs and translate them into a coherent Environmental context. By surfacing implicit assumptions and early tensions between design intentions and Environmental domains, the agent supports sensing at moments of high uncertainty, before design lock-in occurs. Without this early interpretative capacity, Environmental considerations enter the process too late to remain meaningfully integrated. Second, the agent must support structured interpretation rather than evaluation. The Decision making/Planning Layer shows that Environmental trade-offs only become discussable when assumptions, argumentation paths and consequences are made explicit. This layer enables steering by organising decision making before positions harden, allowing sustainable design choices to be explored while uncertainty is still negotiable.

Third, the agent must ground interpretation in verifiable Environmental evidence. The Tool Layer demonstrates that internal decision making alone is insufficient. By retrieving standards, emission factors and regulatory constraints through controlled RAG pipelines, the agent connects design decision making to external references without prioritising or optimising outcomes. This ensures legitimacy and prevents speculative Environmental claims.

Fourth, the agent requires modular reflective capacities. The Specialized Functions Layer enables the decomposition of complex Environmental questions into energy-related, material-related and contextual components. This supports steering by preserving plurality and keeping alternative interpretations open, rather than converging prematurely on a single solution.

Finally, the agent must sustain Environmental coherence over time. The Memory/Provenance Layer preserves definitions, assumptions and data lineage across phases and team transitions. This layer supports both sensing and support by maintaining traceability where uncertainty decreases and accountability increases, enabling sustainable design choices to remain intelligible beyond the moment of decision.

Taken together, these qualities describe an AI agent that adapts its emphasis as projects move from uncertainty to certainty. In doing so, the agent strengthens the reflective capacity of design teams and stabilises the integration of Environmental considerations within sustainable design choices across the early project phases.

Chapter 06

Findings in the prototype

6.1 Testcase

6.2 Human Workflow

6.3 AI workflow

Baseline

Retrieval

Augmentation

Generation

6.4 How the Three Roles Become Visible in the Generated Scenario

6.5 Key takeaways H6

6. Findings of the prototype

In this phase, the research moves from methodological preparation to practical application. The prototype is used to observe how an AI agent can sustainably design choices in the early SO/VO phase. Rather than evaluating design outcomes, the prototype focuses on how environmental considerations are structured, articulated and related to one another. One recurring environmental tension identified across all interviews is translated into a working workflow. This allows the functioning of the Digital Building Master to be examined as a mechanism for structuring decision making, not for determining design decisions. The prototype is therefore intentionally limited. It clarifies how decision making unfolds under environmental complexity, while leaving responsibility for design choices with the project team.

6.1 Testcase

During the prototype phase, the focus was on the tension surrounding the use of PV panels on the roof in the early SO/VO design stage. This tension appeared in every interview as a recurring, decisive moment within the environmental dimension of office development. PV improves operational energy performance and helps meet targets such as BENG-2 and ENE01 (Sartori, n.d.). yet simultaneously increases the material-related CO₂ emissions captured in LCA and MPG assessments.(Mohan & Vijayalaxmi, 2024) . The result is a predictable trade-off between energy generation and material efficiency.

This trade-off is analytically significant because it emerges early, while its consequences only surface later. In practice, the rise in the MPG score typically becomes visible once materialisation and architectural principles are already fixed (Gustavsson et al., 2010) often leading to redesign loops, budget pressure or the loss of alternatives. The PV issue therefore functions as an archetype of early design decisions that shape the final environmental performance more strongly than initially assumed. Different disciplines also weigh the issue differently. For some, PV is a straightforward route to stronger operational performance or certification scores; for others, the added material burden demands compensation elsewhere. Because assumptions are dispersed across documents, models and expert judgement, the decision making behind these choices remains only partially visible (Sartori, n.d.) The test case is therefore well suited to showing how divergent interpretations and priorities arise in the early design phase.

This prototype reconstructs this assessment using project data from EDGE Stadium. This project offers a fully documented and internally accessible design process, with data on roof geometry, available surfaces and installation locations. This made it possible to reconstruct the PV assessment based on real conditions rather than hypothetical values. Figure 11: Rooftop configuration of EDGE Stadium used as basis for the PV trade-off case. (EDGE, 2025) shows the roof configuration on which the PV parameters are based. All PV data has been validated by an installation specialist, which reinforces the reliability of the case study.

6.2 Human Workflow

The complete human workflow for the PV case is included in Appendix G. It is constructed around three calculation frameworks: the BENG method I *NTA8800: 2024*, the ENE01 reference values and improvement thresholds from BREEAM-NL Guidance Note 41, and the material-related emission factors for PV systems from the NMD Category 3 report (620 kg CO₂/kWp). The case study follows these frameworks closely and is deliberately based on a shell calculation: user installations, consumption profiles and tenant-specific adjustments are excluded so that only the building envelope and fixed systems remain. This mirrors the development phase of EDGE Stadium, where ENE01 was assessed alongside broader design considerations.

The workflow diagram in Appendix G shows how the calculation formulates ambitions, interprets standards, determines usable roof area and computes both operational yield and material-related emissions. These two lines of decision making, operational (Steps 2–2d) and material (Steps 3–3d), run entirely in parallel and intersect only in Step 4, where the combined CO₂ balance is derived.

A notable feature of the workflow is that ENE01 depends on far more than PV alone: energy requirements (BENG1), building installations, heating systems, ventilation requirements, internal heat gains and usage profiles are all considered within NTA 8800. In the case study, these parameters are intentionally fixed to isolate the PV tension. This creates analytical clarity, but also fragility: in real projects, changes in any of these inputs immediately shift ENE01 performance, yet this dependency remains invisible in the workflow's early steps.

This is where the most significant blind spots appear. The workflow reveals the tension between operational CO₂ reduction and material-related emissions only after both calculations have been completed separately. Although the consultant works rigorously within NTA 8800, Guidance Note 41 and the NMD factors, these frameworks do not communicate with each other. As a result, changes in roof layout, emission factors or standard interpretations can alter outcomes without being detected until late in the SO/VO phase.

Appendix G also shows that much of the underlying decision making is implicit. Assumptions are often agreed informally; interpretations vary between specialists and choices are rarely documented systematically. The model itself is robust, but its use is highly personal, making the workflow sensitive to variation and limiting the traceability of key considerations. Understanding this human decision making is essential before introducing AI support. Zaimoglu et al., (2023) underline that data-driven decision-making always relies on interpretative judgement, while (Chen et al., 2022) show that AI becomes meaningful only when it aligns with existing cognitive frameworks. In this workflow, that alignment requires making parallel calculation routes visible earlier, surfacing implicit assumptions and connecting standards that remain strictly separated in current practice. It is precisely on these points that an AI agent can provide analytical value later in this research.

6.3 AI workflow

The human workflow shows how consultants reconstruct a PV scenario by navigating standards, assumptions and design constraints. The AI workflow mirrors this logic but does so within the behavioural requirements defined in Chapter 5. These requirements determine how the agent must interpret environmental questions, surface uncertainties and structure alternatives without steering decisions.

RAG provides the technical mechanism through which the agent retrieves information, yet it does not prescribe the decision making. That decision making is shaped by the sensing-, steering- and support-requirements. Retrieval, augmentation and generation therefore form the layered translation of these requirements: retrieval activates the correct environmental domains, augmentation connects retrieved rules to project-specific parameters, and generation synthesises these into transparent scenario structures.

Baseline

The workflow is grounded in a knowledge base configured in Copilot Studio, structured according to the architecture in Appendix H and Appendix N. This ensures that the agent consistently enacts the behavioural requirements by drawing directly on the documents that underpin environmental decision making in Dutch practice. For operational performance, the agent relies on ENEO1, the BENG-2 method and NTA 8800, which define system yields, conversion factors and minimum thresholds.

Geometric project information is linked to PVGIS, allowing the agent to compute irradiation based on location, tilt and orientation. This supports the sensing-requirements by exposing context-dependent variation rather than relying on generic assumptions. Material-related CO₂ impacts are based on emission factors from the National Environmental Database. In Step 4, operational and material values are juxtaposed in line with the steering-requirements, which require transparency without preference for any outcome.

Accurate and traceable project data is essential. The rooftop of EDGE Stadium provides this: its geometry, usable zones and installation layout can still be observed directly, and key team members remain available to verify assumptions. This reduces interpretive uncertainty and strengthens both the internal validity of the augmentation phase and the plausibility of generated scenarios.

This setup also follows design-oriented methodological guidance: a test case should be complex enough to reveal model behaviour while remaining interpretable. EDGE Stadium meets this condition; its documentation is cohesive enough to reconstruct operational and material consequences yet detailed enough for realistic application of the environmental rules.

All documents in the knowledge base are segmented into small, semantically coherent fragments, enabling the agent to reason at the same granularity as

human consultants. Topic workflows restrict which fragments can be accessed. The simple “Goodbye” example in Figure 21 Figure 23 illustrates this routing logic; Appendix I show the equivalent PV-specific workflows.

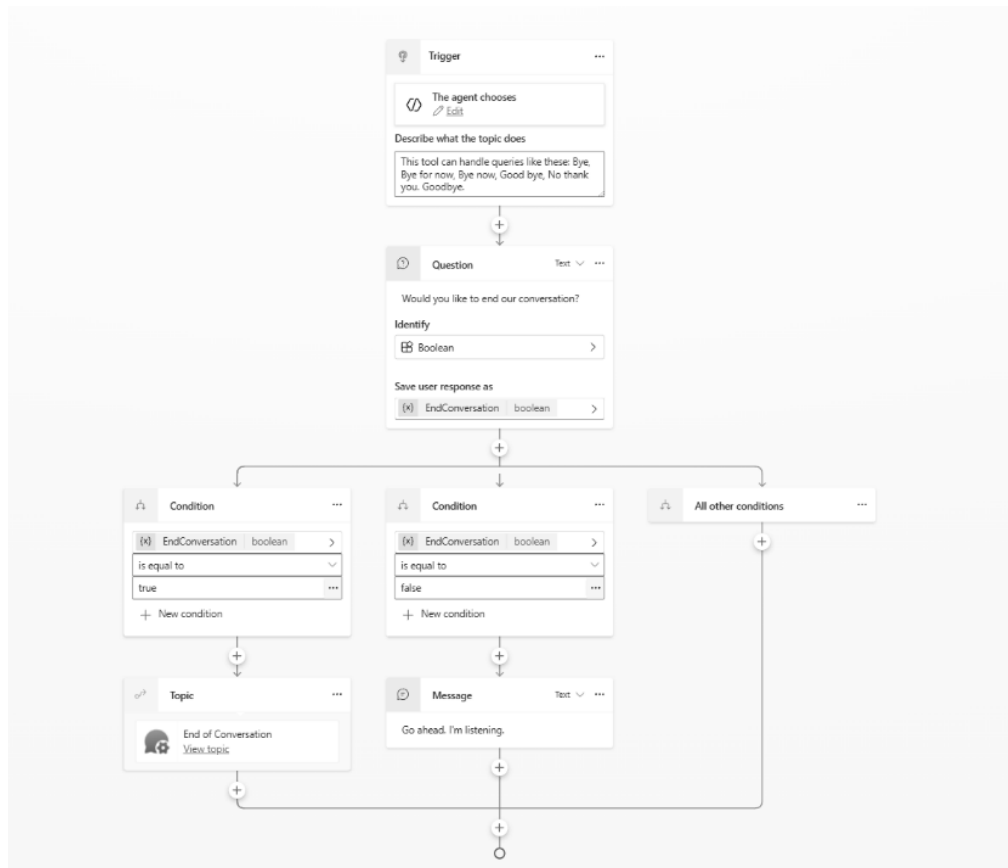


Figure 21: Conversation flow for ending or continuing a dialogue in an AI agent (Copilot, 2025)

Retrieval: activating the correct environmental domains

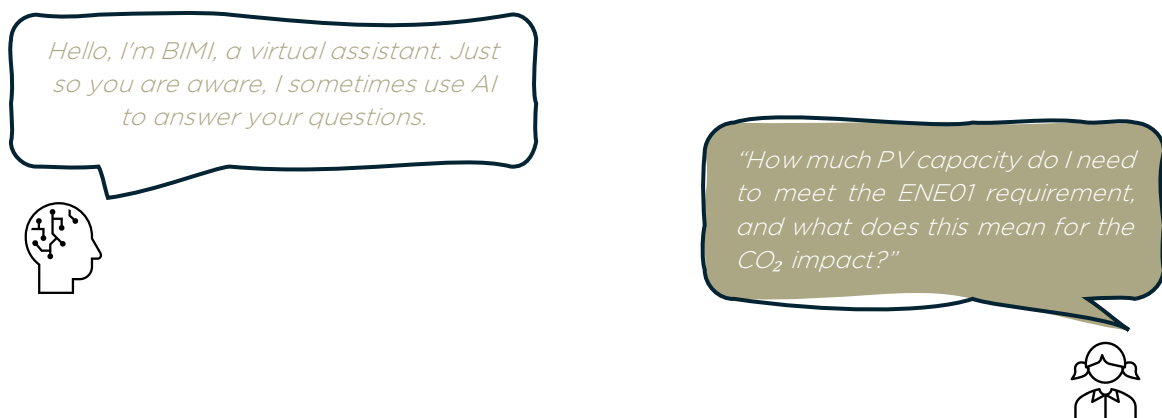


Figure 22: Dialogue between AI agent and human (own work, 2025)

To understand how retrieval works, it is instructive to look at the central question posed in the PV test case. Shown in the figure above. For clarity within the workshop environment, the agent was given the practical nickname BIMl, reflecting its integration into the Copilot Studio interface. This name is functional

rather than conceptual: it helps participants recognise the agent as a consistent interaction partner during the test scenario. An experienced sustainability consultant immediately recognises that this question contains two distinct lines of decision making: meeting the operational energy requirement and understanding the material-related CO₂ consequences. Through its sensing-requirements, the agent must make this distinction explicit.

The chunk example in Table 4 demonstrates this mechanism. The fragment “*meet the ENE01 requirement*” activates the operational energy domain and is linked to ENE01 and NTA 8800. From these sources, the agent retrieves the BENG-2 reference and target values, establishing the first building block of the information package.

Chunk	Citation	Activated domain	Sources	Value
ENE01 OBJECTIVE	'Meet the ENE01 requirement'	Operational energy performance	ENE01, NTA8800	ENE01 requires ≥10% reference for BENG-2

Table 4: Chunk example (own work, 2025)

A parallel retrieval occurs for the fragment “*what does this mean for the CO₂ impact?*”, which activates the material domain and links it to the emission factors in the National Environmental Database. This establishes the second building block. The full elaboration of all chunks is provided in Appendix M: Overview Chunks.

Retrieval therefore does more than gather information: it activates the sensing-requirements by breaking the query into its underlying components and linking each part to the appropriate standard. This prevents the question from being processed as a single undifferentiated request and establishes the structured foundation on which the subsequent workflow steps depend.

Augmentation: connecting norms to the project environment

Where retrieval clarifies what the question refers to, augmentation clarifies how the retrieved rules apply to the specific project context. Retrieval surfaces normative conditions; augmentation shows how they manifest within the building’s physical and material boundaries. Here, the support-requirements become dominant: preserving definitional consistency, tracing numerical values to their origin and making implicit assumptions explicit.

Chunk	Citation	Activated domain	Sources	Value	Information package
ENE01 OBJECTIVE	'Meet the ENE01 requirement'	Operational energy performance	ENE01, NTA8800	ENE01 requires ≥10% reference for BENG-2	Reference energy: Rref = 816,200 kWh/yr (20,405 m ² × 40 kWh/m ² ·yr)

Table 5: Chunk example (own work, 2025)



This “information package” column expands the retrieval table. Normative values are now linked to project geometry, gross floor area, roof layout and validated assumptions. A textual rule becomes a calculable design constraint.

In the human workflow, this contextual information is scattered across documents and informal discussions. Augmentation consolidates these fragments into a coherent analytical unit. Once this enriched structure is in place, the agent no longer operates on disconnected pieces but on a meaningful, context-aware foundation from which realistic scenarios can be generated.

Generation: forming scenarios through sensing, steering and support

With retrieval and augmentation complete, the agent moves from interpretation to decision making. At this point, it no longer processes fragments individually but works with the complete information packages in which normative rules and project parameters have been combined. Appendix I: Workflow AI in Copilot shows the endpoint of this preparation: each retrieved requirement has been anchored in the geometric and material conditions of the EDGE Stadium case.

Generation starts beyond this table. Here, the agent combines multiple information packages into coherent scenario patterns. The ENE01 objective becomes one anchor among several, interacting with the PVGIS irradiation profile, the available roof area, the PV density, the installation layout and the material emission factors. Together, these variables define the corridor of feasible design options within which meaningful alternatives can emerge.

In this phase, the three requirement clusters become visible in the agent’s behaviour. Sensing appears when the agent highlights uncertainties, missing data or intersecting logics, for example, when it notes that yield depends on shading or that the calculation results form a range rather than a single value. Steering becomes visible in the way the agent structures scenario variants neutrally, showing how options differ in capacity, yield or embodied carbon without suggesting which path should be chosen. Support emerges in the consistent tracing of definitions, references and calculation logic back to their sources, ensuring that each value remains transparent and verifiable.

The LLM functions here as a processing engine: it turns the combined information packages into scenario alternatives that are realistic, context-specific and analytically transparent. The agent does not optimise or decide; it exposes the structure of the design space. By making the interaction between operational performance and material-related emissions explicit, the agent enables practitioners to recognise environmental tensions far earlier than is possible in human workflows. The following paragraph analyses this behaviour through the three roles, using direct citations from the CoPilot responses.

6.4 How the Three Roles Become Visible in the Generated Scenario

Figure 16 in Chapter 4 conceptualised sensing, steering and support as roles whose relative relevance shifts across project phases in corporate real estate development. The early SO/VO phase was identified as a moment of overlap, characterised by high uncertainty, emerging formalisation and increasing coordination pressure. In that phase, sensing remains critical, steering reaches its peak, and support begins to gain importance without yet becoming dominant.

The PV scenario analysed in this chapter corresponds precisely to this zone of overlap. The generated response does not activate the three roles sequentially, nor does it privilege one role over the others. Instead, sensing, steering and support become simultaneously visible within a single interaction. This empirical observation confirms the dynamic pattern suggested by Figure 16 and demonstrates that the roles can coexist in practice rather than replacing one another.

Sensing

In its sensing role, the agent foregrounds technical uncertainty rather than resolving it. It explicitly states that PV yield varies with roof layout and assumed yield ranges (SE-P01), and that variation arises from tilt, azimuth and shading as defined within NTA 8800 (SE-PM03). The agent also consistently identifies the origin of reference values used in the assessment, such as the BENG-2 benchmark (SE-P04). In doing so, it makes visible which parameters remain variable and which assumptions are still provisional in the early SO/VO phase.

At the same time, the absence of a citation for SE-R02 reveals a clear limitation. While uncertainty within the calculation is exposed, potential risks and lock-in effects that extend beyond the immediate parameter space are not identified. Sensing therefore operates at the level of technical variability, not yet at the level of strategic constraint.

This means that the generated scenario only partially reflects the sensing role as conceptualised in Chapter 4. It aligns with the early-phase function of making assumptions explicit while design freedom is still high, but it does not yet perform the anticipatory sensing required to signal when early choices begin to narrow future design corridors. The prototype thus confirms the relevance of sensing in the early SO phase, while simultaneously revealing where this role remains underdeveloped.

ID	Copilot-citation	Interpretation
SE-P01	"Range depends on roof layout/orientation and assumed yield"	Uncertainties are actively identified; missing assumptions are made visible.
SE-PM03	"Yield varies with tilt/azimuth/shading as handled in NTA 8800"	Variation due to design choices is highlighted, without preference for any alternative.
SE-P04	"BENG-2 reference is 40 kWh/m ² -yr (GN41)."	The source is explicitly stated; origin remains transparent.
SE-R02	-	Agent reports no risks or lock-in effects.

Table 6: Sense, more in appendix K

Steering

The steering role becomes visible where the agent structures the design space without directing the outcome. It presents multiple PV capacity scenarios side by side as equivalent planning variants (ST-P02), unfolds the underlying decision-making step by step (ST-P03), and explicitly clarifies that these scenarios are exploratory rather than prescriptive (ST-S03). This provides the user with an organised and transparent set of alternatives, without implicit preference or normative guidance.

At the same time, the absence of a citation for ST-S04 marks a clear boundary of the current prototype. While the agent structures scenarios, it does not introduce reflective prompts that question targets, balances or strategic intent. Steering is therefore activated as organisational structuring, not as interpretative interrogation.

This pattern corresponds to an early form of steering as described in Chapter 4. It aligns with the moment in the SO-VO transition where uncertainty has begun to narrow, but shared interpretation is still unstable. At this stage, steering does not resolve tensions or negotiate ambition but organises how different perspectives and calculations are related. The generated scenario reflects this function accurately, while also showing that more advanced steering behaviour remains undeveloped.

ID	Copilot-citation	Interpretation
ST-P02	"You should plan roughly 170-180 kWp... 300-320 kWp... 430-455 kWp"	Equivalent variants side by side.
ST-P03	"Reference primary energy Rref = ..." → "If the modelled design still needs X%..."	Step-by-step decision making visible and neutral.
ST-S03	"These are planning scenarios, not certification requirements."	Emphasises that variants remain alternatives, not final recommendations.
ST-S04	No quote found	Agent does not actively reflect; only provides scenarios.

Table 7: Steering, see more in appendix K

Support

The support role is the most fully operationalised in the current prototype. Throughout the generated scenario, definitions remain stable and numerical values are consistently anchored to their source documents. As shown in Table 8, the BENG-2 reference value is held constant and explicitly cited (SU-M04/PO4). Subsequent steps remain linked back to NTA 8800 and PVGIS (SU-M05), while material-related impacts are anchored in NMD Category 3 data once the product list is fixed (SU-T04).

In this role, the agent functions as a memory structure rather than a decision aid. Decision making remains traceable across steps. Documents remain connected. Assumptions do not drift unnoticed as the scenario unfolds. This directly addresses the pattern identified in Chapter 4, where Environmental ambition often erodes because the rationale behind earlier sustainable design choices becomes opaque during phase transitions.

The presence of this support behaviour in an early SO/VO scenario anticipates the growing importance of traceability and accountability shown on the right-hand side of Figure 16. Even before formal decisions are fixed, the agent already preserves the coherence of Environmental decision making in a way that allows choices to remain intelligible across later phases and stakeholder changes.

Id	Copilot-citation	Interpretation
SU-M04/PO4	"BENG-2 reference is 40 kWh/m ² ·yr (GN41)."	Definitions and sources are displayed immediately.
SU-M05	No quotation needed	No infringement: all data is public project material.
SU-M04	"The anchoring of NMD categories... makes the source lines verifiable"	Bronsporen blijven consistent door de workflow heen.
SU-T04	"We will use NMD Category 3 data... once product list is fixed"	Shows link to documents, although long-term storage is limited.

Table 8: Support, see more in appendix K

Although the prototype already displays recognisable behaviour across sensing, steering and support, the comparison with the requirement framework makes clear where the PV case remains incomplete. The analysis is based on a single Copilot response. While additional prompts could reveal further nuance, the current output already exposes the main boundaries of the system.

Several forms of decision making that are central to PV-related sustainable design choices are not yet activated. The agent does not articulate conceptual tensions between ENE01, NTA 8800 and the NMD emission factors, even though these frameworks rely on different assumptions and system boundaries. Nor does it make interpretative differences explicit, even though such distinctions often shape responsibility, pacing and momentum in early design phases.

Signals of emerging lock-in are likewise absent. Requirements such as SE-R02 and ST-R03 anticipate that the agent would warn when early design choices begin to restrict later options, for example when roof geometry limits viable orientations or when performance gaps narrow the range of feasible PV capacities. While the prototype calculates accurately, it does not yet recognise when calculations start to constrain design freedom.

Reflective prompting is also missing. Although the agent structures scenarios clearly, it does not question why a particular improvement target is applied, whether alternative balances between operational and material considerations might be strategically relevant, or how uncertainty should influence early interpretation. As a result, part of the interpretative work that characterises real Environmental decision making remains outside the generated response.

These gaps do not undermine the concept. They delineate the current limits of this iteration and clarify where a Digital Building Master could add the most value in future development: by linking frameworks that are currently applied in isolation, signalling narrowing design corridors at an earlier stage, and sustaining multiple interpretations long enough to support deliberate reflection. Taken together, the findings show both what the prototype already contributes and where further refinement is required.

6.5 Key takeaways H6

The three qualities defined in the requirement framework—sensing, steering and support—are operationalised in the prototype through their explicit embedding in the agent’s workflow. Rather than acting as abstract roles, they function as behavioural constraints across a layered RAG architecture. Each quality is primarily enacted within one processing layer, while remaining analytically connected across the workflow.

Sensing is enacted in the retrieval layer. The agent decomposes the user query into discrete Environmental components and links each component to its corresponding normative source. In the PV testcase, this activates ENE01, NTA 8800 and NMD emission factors as parallel yet non-merged logics. This separation is intentional. It prevents premature harmonisation and exposes tensions between operational energy performance and embodied emissions. Through semantic segmentation and topic-based routing, the agent surfaces parameter variability, dependencies and missing inputs at calculation level. Sensing is thus operationalised as the explicit articulation of technical uncertainty, rather than as early optimisation.

Steering becomes visible in the generation layer. Retrieved and augmented information is assembled into structured scenario representations. Variants are presented side by side, with their underlying decision making unfolded step by step. Outcomes are framed as exploratory comparisons, not as recommendations. In the PV testcase, this reveals how photovoltaic panels perform differently depending on whether operational energy targets, embodied carbon thresholds or portfolio-level consistency are prioritised. Steering therefore supports interpretation and trade-off awareness rather than decision-making, aligning with the SO-VO transition described in Chapter 4.

Support is enacted in the augmentation layer. Normative sources are anchored to project-specific parameters such as roof geometry, reference energy use, applicable emission factors and spatial constraints. Using the rooftop of EDGE Stadium as a controlled reference environment ensures definitional stability, traceability of assumptions and verifiable source lines. At the same time, the testcase is explicitly framed within a Corporate Real Estate context. The PV intervention is not treated as a universally valid solution, but as one asset-level option whose implications must be understood in relation to portfolio strategies, refurbishment cycles and long-term Environmental performance targets.

Taken together, the layered implementation of sensing, steering and support renders the prototype suitable for empirical evaluation. The workflow mirrors how Environmental design choices are explored in early-phase CRE development, while making that decision making explicit enough to be inspected, compared and critiqued. The PV testcase demonstrates that the prototype already enacts core structuring behaviours of a Digital Building Master, while also revealing where interpretative depth, portfolio-level decision making and reflective prompting require further development.

Chapter 07

Findings of the validation space

- 7.1 The structure of the design workshop
- 7.2 Human Workflow
- 7.3 AI workflow
- 7.4 Comparison AI and human workflow
- 7.5 Differences in behaviour and decision-making
- 7.6 Stakeholder roles and redistribution of agency
- 7.7 Technical and Role-Related Limitations
- 7.8 Synthesis
- 7.9 Key Take aways H7

7. Findings of the validation space

This chapter evaluates the Digital Building Master through a structured design workshop. The validation does not assess technical correctness or optimisation outcomes. Instead, it examines how the AI agent reshapes the design process itself. The focus lies on how the presence of the agent affects decision making, interaction and the handling of Environmental considerations in the early SO/VO phase. By comparing design decisions made with and without the AI agent, the workshop reveals how quickly teams converge, which assumptions remain implicit, and where Environmental tensions are either absorbed or made explicit. Attention is given to how time is distributed across the process. Rather than accelerating decisions, the agent influences where effort is invested: slowing down interpretation and articulation in the early phase to stabilise outcomes later. The validation further explores how different stakeholders engage with the agent in distinct ways. The AI does not function as a single-user tool, nor as a decision-maker. Instead, it operates within a shared design space, where developers, consultants and designers activate different roles of the agent depending on their responsibilities. Behavioural effects observed during the workshop are discussed with a sustainability advisor, while the technical and procedural assumptions of the agent are reviewed with an AI expert.

7.1 The structure of the design workshop

The validation workshop was set up as a controlled contrast test to assess the extent to which the AI agent behaves as theoretically and empirically assumed in the previous chapters. Central to this design was not the evaluation of accuracy, but the investigation of how the agent influences the design decision.

The workshop consisted of two rounds with the same design challenge. In the first round, participants worked without AI. The PV case was based on a realistic design assignment with ambitious sustainability goals such as BREEAM outstanding and Net Zero Carbon. The aim was to make a choice regarding the amount of PV and the material for the PV. The entire assignment is presented in Appendix x. Participants filled in an A3 canvas, making their assumptions and arguments visible. Afterwards, time was set aside for reflection on the human workflow.

In the second round, AI agent BIMl was integrated via six predefined prompts. BIMl's answers were displayed on the board. The same A3 canvas layout was used in this round. The concluding reflective questions were formulated identically to the set from the first round, allowing for a systematic comparison of differences in collaboration, decision making and attention to environmental aspects.

By maintaining the same assignment, time pressure and canvas structure, it became clear where the agent adds value. The workshop demonstrated the extent to which BIMl fulfils the roles of sensing, steering and support in early real estate development and how it contributes to providing insight into environmental decision-making.

7.2 Human workflow

The human-only design round was characterised by rapid convergence driven by experiential knowledge and pragmatic constraint handling. From the outset, spatial and energy parameters were translated almost immediately into decisive assumptions. The relationship between a roof area of approximately 1,000 m² and a programme of 20,000 m² led directly to the conclusion that the building would require roughly twenty floors. This inference was not debated, but accepted as self-evident, illustrating how empirical knowledge functions as a shortcut under time pressure.

A similar pattern emerged in technical decisions. The size of the air handling unit was set at 70 m² without discussion. The atrium requirement was resolved pragmatically by redefining it as a compact sky lobby, thereby avoiding a full-height atrium shaft. This adjustment freed up roof space for PV, which was then estimated at 825 m², with an assumed usability of eighty percent. None of these assumptions were calculated explicitly; they functioned as working hypotheses that were sufficiently plausible to allow the process to move forward.

Material choices followed the same logic. Aluminium-glass foil PV panels were selected almost automatically, primarily because they were familiar and perceived as lightweight and efficient. Alternatives such as glass-glass panels were dismissed early, not based on quantified comparisons, but due to assumed cost and structural implications. These judgements were grounded in prior experience rather than in explicit evaluation.

Energy consumption was similarly fixed at 60 kWh/m², based on what participants considered a realistic benchmark from previous projects. Installation space and system sizing were also determined intuitively. Collectively, these assumptions formed a coherent but coarse decision framework: internally consistent, yet weakly validated. They enabled speed, but limited reflection.

A notable shift occurred when the ENEO1 credit was discussed. Once it became clear that this credit is mandatory for BREEAM Outstanding, the group's attention narrowed sharply. The requirement to generate at least fifty percent of energy consumption on-site became the dominant concern. By reversing the calculation, the mismatch between achievable generation (approximately 170,000 kWh) and required demand (approximately 600,000 kWh) was quickly identified. This numerical gap legitimised an immediate expansion of the solution space to include façade PV, landscape-based generation and even wind energy.

At this point, Environmental decision making became almost exclusively operational. Embodied carbon was mentioned, mainly in relation to the footprint of PV panels, but remained secondary. Other Environmental considerations disappeared from the discussion altogether. This illustrates how, under time pressure, attention shifts toward parameters that are directly calculable and verifiable, while less tangible or less immediate impacts are temporarily sidelined.

Decision-making dynamics were also unevenly distributed. One participant clearly took the lead in formulating assumptions and performing rapid calculations. The remaining participants largely affirmed these steps, reinforcing a hierarchical but efficient workflow. Despite the complexity of the assignment, a design proposal was produced within fifteen minutes.

From an analytical perspective, this round exemplifies fast convergence under uncertainty. Decisions are taken quickly, assumptions remain largely implicit, and integration across Environmental dimensions is limited. This aligns with the left-hand trajectory in the time-impact figure: early decisions are reached rapidly, but their robustness and traceability remain fragile. The human workflow therefore provides a clear baseline against which the effects of the AI-supported round can be evaluated.

7.3 AI workflow

The AI-supported round differed immediately from the human-only workflow, both in tempo and in interaction pattern. Where the human round was characterised by rapid convergence, the introduction of BIMi slowed the process and produced a fragmented rhythm. Response delays, prompt adjustments and occasional system failures interrupted the flow of discussion. Participants alternated between waiting for output, reformulating prompts and temporarily improvising without the agent. These technical frictions were inherent to the prototypical setup and added an additional layer of work: alongside designing, participants now also had to manage the AI interaction itself.

While this initially led to moments of silence and mild frustration, the slowdown had a clear analytical effect. Decision-making no longer progressed through uninterrupted heuristic jumps. Instead, decision making was broken into explicit steps. This confirms the dynamic illustrated in the time-impact figure: the AI does not accelerate early decision-making but redistributes effort by slowing down interpretation in the SO/VO phase.

Despite the reduced pace, the use of predefined prompts introduced more structure into the discussion than in the human-only round. Participants again began with roof limitations, but in contrast to the first round, the tension between operational and embodied carbon surfaced immediately. This was not because participants had changed priorities, but because BIMi's prompts framed this tension as an explicit starting point. In this moment, sensing and steering were activated simultaneously: uncertainties were surfaced while the space of possible interpretations was organised.

When BIMi did respond, its output primarily functioned as validation rather than invention. The estimated usable PV surface area of 550–650 m² closely matched the earlier human quick scan. This alignment positioned the agent not as an alternative designer, but as a control mechanism that formalised existing assumptions. At the same time, BIMi articulated spatial dependencies that had remained implicit in the human round, such as the relationship between roof layout,

shading and yield. This added formal coherence to the emerging design logic and strengthened its external justifiability.

A second recurring pattern was continuous calibration of the agent against internal expert knowledge. Participants explicitly questioned whether BIMi should be “smarter” than an experienced colleague and tested its outputs against established rules of thumb. This interaction shifted the role of the technical expert. Rather than being the sole source of authority, the expert acted as a validator of AI-generated scenarios. At the same time, the exercise exposed expectations toward a Digital Building Master. Participants assumed that such an actor would have immediate access to embodied carbon data for different PV types. The absence of this information made visible where human intuition still dominated and where the agent’s knowledge base remained incomplete.

Throughout the round, BIMi functioned less as an active designer and more as a slow but structuring advisor. It did not propose solutions, nor did it resolve tensions. Instead, it forced participants to articulate assumptions, dependencies and Environmental implications more explicitly than in the human workflow. Roles were not sequential. At different moments, the agent activated sensing through uncertainty exposure, steering through scenario structuring, and support through source anchoring, often within the same interaction.

From an analytical perspective, the AI-supported workflow demonstrates how a Digital Building Master reshapes early-stage decision-making without taking over authority. The process slows down, but integration increases. Parallel Environmental logics remain visible longer, assumptions become traceable, and convergence is delayed until tensions are explicitly acknowledged. Even in its technically limited form, the agent formalises Environmental decision making in a way that reduces silent prioritisation and prepares the ground for more stable decisions in later phases.

7.4 Comparison AI and Human workflow

This paragraph links AI and the human workflow. First, the results are assessed based on the findings from H2.

The literature reference

Chapter 2 summarises the literature on ESG and AI in a table. This section links the same findings to the results of the prototype test. This shows where AI and human workflow add value.

Environmental aspect	Human workflow	Ai workflow
KNOWLEDGE FRAGMENTATION	“We’ll just assume 60 kWh/m ² , shall we? That’s what we always use.” This confirms that knowledge is not consistent across the board.	“Oh, indeed, BREEAM ENE01 stands for...” This demonstrates that BIMi provides knowledge to the team, enabling them to act in a consistent manner.
PROCESS FRAGMENTATION	One participant performs all the calculations (“Let me do the math, otherwise it will take too long”), while the rest follow.	Prompts force you to think about the following: “First run the roof capacity check”. This ensures that the

END USERS NOT INVOLVED	-	conversation follows a structured line, regardless of who is speaking.
DIFFICULT ADOPTION OF TECHNOLOGY	-	The technology crashes, prompting participants to immediately seek the expert's knowledge. This emphasises that not everyone is aware of IT developments.
NO COORDINATION IN EARLY PHASES	The informal hierarchy guides the design process. The discussion is clearly uncoordinated.	The pre-set prompts encourage participants to think about structure. What is useful to ask first?
DIFFERENT STAKEHOLDER INTERESTS	Conflicts of interest remain implicit: "Aluminium is simply light, so let's go with that."	BIMI explicitly identifies contradictions: "Weight vs. efficiency trade-off identified", thereby revealing unspoken preferences.

Table 9: Differences between human and AI-driven environmental decision making in the early design phase (own work, 2025)

What is remarkable about the table is that some aspects cannot be measured using the test case. The end users were not relevant enough in the test case to make a statement about this. In addition, the adoption of technology is not relevant in the human workflow. An important conclusion from the table is that the human workflow relies heavily on intuition and rules of thumb. The AI workflow ensures that humans focus entirely on the agent until it responds. In addition, the AI has its own technical limitations.

7.5 Differences in behaviour and decision-making

Figure 19 visualises a fundamental shift in how decision-making unfolds when an AI agent is introduced into the early design process. The figure does not compare speed in isolation but illustrates how cognitive and interpretative effort is redistributed over time between divergence, interpretation and convergence.

In the human-only workflow, decision-making is characterised by rapid convergence. Divergence and convergence overlap almost immediately: options are generated and narrowed in quick succession, guided by experiential knowledge and informal hierarchy. This produces momentum and allows teams to reach a preliminary decision quickly. At the same time, this compressed process limits integration. Assumptions remain largely implicit, Environmental considerations are selectively prioritised, and tensions between operational and material impacts are absorbed into heuristics rather than explicitly discussed.

The AI-supported workflow follows a different temporal logic. As shown in the upper part of Figure 19, divergence expands and convergence is deliberately delayed. The design phase becomes longer and more fragmented, not because participants hesitate, but because assumptions, dependencies and Environmental tensions are surfaced and kept active. BIMI maintains parallel lines of decision making and prevents premature closure by structuring options without resolving them. This creates a broader and flatter "AI diamond" in which interpretation precedes decision-making.

The effect of this shift becomes visible at the level of outcomes. Although decisions take longer in the SO/VO phase, the resulting design direction stabilises earlier in the overall project trajectory. Integration work that would otherwise reappear

later as redesign, correction or reinterpretation is moved forward in time. The “gain” indicated in Figure 19 therefore does not lie in faster early decisions, but in earlier and more explicit integration of Environmental considerations while design freedom is still high.

Importantly, the validation does not demonstrate that the AI-supported workflow produces objectively better design outcomes in a technical or optimised sense. What it does show is a qualitative shift in how outcomes are formed. With BIMi, design results are less dependent on individual heuristics and informal hierarchy, and more explicitly grounded in articulated assumptions and documented decision making. Environmental considerations are not optimised but integrated more coherently.

The value of the Digital Building Master therefore lies in improving the conditions under which decisions are made rather than in improving the decisions themselves. Time is not saved by accelerating choice, but by reducing the need for reinterpretation, correction and renegotiation later in the process. In this sense, BIMi functions as a reflective support tool that reshapes decision-making behaviour without replacing human judgement.

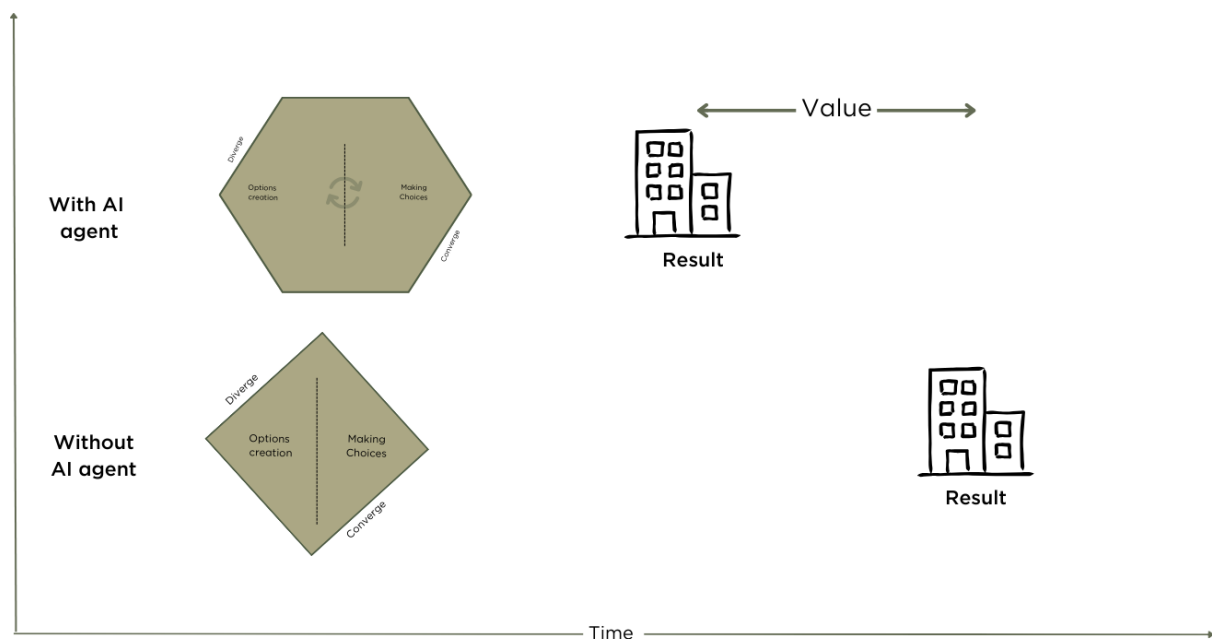


Figure 23: The effect of an AI agent on decision making (own work, 2025)

7.6 Stakeholder roles and redistribution of agency

The introduction of the AI agent reshapes how agency is distributed among stakeholders during early design decision-making. Rather than transferring control to the system, BIMi functions as an epistemic catalyst that reorganises interaction patterns, suspends informal hierarchies, and redistributes responsibility for interpretation and validation. This section explicitly contrasts the human-only workflow with the AI-supported workflow to analyse shifts in agency rather than performance.

In the human workflow, authority concentrates early around experiential expertise. Decisions are driven by heuristics and shared intuition rather than explicit calculation. This is visible in statements such as:

“We’re usually around 60 kWh per square metre, I’d say” and “One hundred and fifty square metres for the air handling unit... that’s already very large.”

These remarks illustrate how plausibility is established through familiarity and precedent. Such heuristics create momentum, but they also compress uncertainty and limit integration across Environmental dimensions.

With BIMi present, this dynamic changes. Technical specialists no longer function as the sole source of truth but shift into a validating role. Participants explicitly tested the agent’s output against expert rules of thumb, captured in the remark: *“It has to be smarter than the installation advisor” Expertise is not replaced, but repositioned.*

The agent generates structured scenarios; the specialist assesses whether these scenarios align with professional judgement. This introduces a second reference frame and weakens the automatic authority of experience.

Architects engage the agent at the interface between spatial design and Environmental ambition. In the human workflow, the designer primarily operates as a feasibility translator; with BIMi, this role shifts toward mediation between spatial quality and Environmental ambition. In the human round, spatial compromises were made pragmatically to maintain feasibility, for example:

“It becomes more of a sky lobby than an atrium... you’re not pulling that atrium through all twenty floors.” With BIMi, spatial decisions are immediately linked to Environmental consequences. The agent explicitly connects roof layout and shading to PV yield, confronting architectural freedom with requirements such as ENE01.

This activates sensing at a moment when design flexibility is still high and prevents spatial decisions from functioning as neutral carriers of ambition.

Developers engage most strongly through the steering role. In the human workflow, the developer tends to narrow options quickly to protect feasibility and momentum, expressed in statements like:

“That’s why our buildings can never be BREEAM Outstanding unless we start making logistics halls.” With BIMi, this reflex is interrupted.

The agent opens what participants described as an “AI diamond”: options expand rapidly, but convergence slows down. This forces explicit negotiation about trade-offs between cost, feasibility and Environmental performance, rather than absorbing these tensions implicitly. Importantly, decision authority remains human.

As one participant stated: “The agent opens things up... but in the end, it’s still us who have to choose.”

A distinct role emerges for the AI expert during the workshop. This role does not contribute substantive design content, but safeguards process integrity. The transcripts show that participants were acutely aware that the quality of the agent's output depended on prompt formulation, captured in the statement:

“The art of using AI is describing it as precisely as possible.” The AI expert ensures that BIMi does not overstep missing data or generate unsupported claims, reinforcing the position of the agent as reflective support rather than autonomous actor.

These roles are not sequential and do not map linearly onto project phases. Instead, different stakeholders activate different roles of the agent depending on the moment and the question at hand. When ambitions are set, human actors dominate. When exploration begins, sensing and validation become central. When scenarios are evaluated, steering structures the discussion. When convergence is required, human hierarchy is reasserted. Agency therefore remains human throughout but is temporarily redistributed to allow Environmental integration before decisions solidify.

Taken together, the validation workshop shows that BIMi does not function as a designer or decision-maker. It operates as a reflective support tool that slows down premature convergence, exposes implicit assumptions and keeps multiple interpretations active long enough for negotiation. Its contribution lies in structuring reflection rather than delivering superior performance. Integration emerges not through automation, but through a reconfiguration of how stakeholders interact with complexity.

7.7 Technical and Role-Related Limitations

The technical validation builds directly on the requirement-based evaluation presented in Section 6.4 and is complemented by the observations from the validation workshop. Together, these analyses clarify not only *what* the prototype does and does not support, but why these behaviours emerge in practice. The validation was conducted with two AI experts using the same requirement framework included in Appendix P, allowing for a direct link between observed social behaviour, technical causes and architectural constraints.

Limitations in sensing and perception

Section 6.4 already demonstrated that the prototype detects Environmental tensions and uncertainties only when these are explicitly articulated in the input. The technical validation confirms that this is a structural limitation of the perception layer. Requirements such as SE-P01 (early variable detection), SE-P02 (risk recognition before design lock-in) and SE-P03 (making assumptions explicit) are only partially supported. The agent does not autonomously scan for latent tensions or implicit risks. As a result, sensing remains reactive rather than proactive.

This limitation is directly reflected in the workshop behaviour. As shown in Table 9, moments of silence and loss of initiative occurred when BIMl was invoked. Because the agent does not initiate questions or signal emerging concerns by itself, the design team entered a temporary waiting mode. This explains why requirements such as SE-R01 and SE-R02 were largely absent in the observed interaction. The agent can respond to uncertainty when prompted, but it does not perceive or foreground it independently.

Steering without interpretation

Within the steering role, the prototype performs largely as designed. Requirements ST-P02 (side-by-side scenarios) and ST-P03 (structured decision making) are fully supported. This is visible in both the CoPilot responses analysed in Section 6.4 and the workshop behaviour, where BIMl consistently broadened the design space by presenting multiple PV scenarios.

However, this structuring function comes with a clear limitation. Reflective steering, as defined in ST-S04, is not supported. The agent does not generate interpretative prompts, nor does it question why certain targets, assumptions or framings are used. As a result, the design space expands but does not automatically converge. Table 9 captures this asymmetry clearly: many options are generated, but little direction emerges.

This explains the characteristic “flat AI diamond” observed in the workshop. BIMl organises information but does not interpret it. It structures without prioritising and clarifies without evaluating. Steering therefore remains procedural rather than reflective.

Support, memory and continuity

The most significant technical limitations appear in the support role. While source accountability (SU-P04), stable definitions (SU-M01, SU-M04) and traceable provenance (SU-D04) function reliably within a single interaction, the prototype lacks persistent multi-turn memory. Requirements such as SU-R05 are therefore completely absent.

This limitation manifests socially as a fragmented interaction rhythm. As shown in Table 9, participants had to repeatedly reconstruct context and assumptions, leading to a “choppy” conversational pattern. Each AI response was treated as an isolated event rather than part of a cumulative decision-making process. Although the agent maintains internal consistency per response, it cannot safeguard continuity across iterations.

Conceptual interpretation is also limited. Requirements related to detecting multiple meanings or conceptual disagreement (SU-S01, SU-S02) are not supported. Terms are used exactly as provided, without recognising alternative interpretations or contested definitions. This explains why spontaneous hierarchy

disappeared during the workshop: the agent does not recognise expertise, role differences or interpretative authority, treating all input as equivalent.

Dependency, framing and strategic use

Beyond these technical limitations, the validation workshop revealed important role-related and governance risks. The agent's decision making is entirely dependent on the documents included in the RAG configuration and the prompts formulated by users. This creates a strong dependency on input selection and framing.

If certain standards, emission factors or perspectives are excluded, the agent cannot surface them. Conversely, selective document inclusion can legitimise interpretations of Environmental ambition while marginalising others. Similarly, prompt formulation can steer attention toward operational performance while downplaying material impacts, or vice versa. The dominance of operational energy observed in the workshop is partly explained by the absence of numerical embodied carbon data in the knowledge base (SU-D05), but also by how the question was framed.

This dependency introduces risks of bias and strategic manipulation. Because BIMi does not decide but structures option spaces, users can repeatedly prompt the agent until a preferred framing emerges. In such cases, the agent reflects user intent rather than challenging it. Without procedural safeguards or shared governance over inputs, the Digital Building Master risks reinforcing existing power asymmetries instead of mitigating them.

7.8 Synthesis

Table 9 explicitly links observed behaviour to unmet or partially met requirements and their technical causes. Together with the requirement validation in Appendix P, this shows that the prototype's influence on the design process is not incidental but structurally grounded in its architecture. The agent excels at consistency, traceability and structured scenario generation, but lacks autonomous perception, interpretative steering and long-term memory.

These limitations do not undermine the concept of the Digital Building Master. Instead, they clarify its current boundaries. The prototype functions as a reflective support tool that formalises Environmental decision making, but it remains dependent on human initiative, governance and interpretation. Advancing the prototype toward a fully operational Digital Building Master therefore requires not only technical extensions, but also careful consideration of who controls inputs, how prompts are formulated and how AI-supported decision making is embedded within professional decision-making practices.

Observed behaviour	Requirement-related limitation	Appendix reference	Interpretation
Temporary silences and loss of initiative when bimi is invoked	Absence of autonomous perception and signalling (se-p01, se-p02) and lack of reflective prompts (st-s04)	Se-p01: uncertainties recognised only when explicitly stated (✓); st-s04: no question generation (-)	The architecture does not support proactive signalling or question initiation. As a result, the agent remains reactive, which places responsibility for continuation back on the human participants and explains the observed waiting behaviour.
Fragmented interaction rhythm due to repeated waiting for output	No persistent multi-turn memory (su-r05); each interaction treated independently	Su-r05: no persistent memory (-), explicitly mapped	The lack of conversational memory prevents accumulation of assumptions and intermediate conclusions. This forces participants to repeatedly reconstruct context, resulting in a discontinuous and “choppy” interaction pattern.
Many options generated with limited directional convergence (“broad and flat ai diamond”)	Structuring without prioritisation (st-p02, st-p03) and absence of conceptual tension detection (su-s02)	St-p02 / st-p03: full support (✓); su-s02: tensions not recognised (-)	The system expands the design space by organising scenarios side by side but lacks mechanisms for interpretation or prioritisation. This asymmetry explains why divergence increases while convergence does not follow automatically.
Agent perceived primarily as validating rather than designing	No interpretative logic (st-t03) and no access to numerical material or emission data (su-d05)	St-t03: links to standards without interpretation (✓); su-d05: numerical environmental data unavailable (-/-)	The agent can confirm compliance with normative references but cannot reason about design consequences beyond validation. This positions bimi as a verifier rather than an active design contributor.
Spontaneous hierarchy weakens; experts are slower to assume leadership	No recognition of roles, expertise or interpretative authority (su-s01, st-s03)	Su-s01: no detection of multiple meanings (-); st-s03: differences mentioned but not indicated (±)	Because all input is treated equivalently, the system suppresses implicit expertise hierarchies. Leadership patterns dissolve temporarily, shifting coordination back to collective negotiation.
Operational energy dominates discussion; embodied impacts fade into the background	Absence of lca or material data in rag configuration (su-d05)	Su-d05: exclusively textual, normative sources (✓)	The dominance of operational energy follows directly from data availability. Without quantitative embodied impact data, the system cannot sustain material-related decision making, reinforcing an operational focus.
Ai-supported workflow perceived as more structured and formal than human workflow	Strong source consistency and stable definitions (st-d02, su-m01)	St-d02: correct provenance (✓); su-m01: stable definitions (✓)	Consistent sourcing and definition management produce a cleaner and more formal decision making structure. This increases perceived rigour, even though interpretative depth remains limited.
Direction of discussion follows prompt formulation and document scope	No counter-framing, intent recognition or bias detection (se-r03, su-s01)	Mapping confirms absence of intent awareness (-)	User-defined prompts and document selection determine which environmental dimensions become salient. This enables framing effects and strategic steering without being detected or corrected by the system.

Table 9: Key take aways H7 do not indicate system failure but demonstrate how specific architectural and configurational choices directly shape interaction patterns, role dynamics and Environmental decision making.

7.9 Key take aways H7

The validation workshop demonstrates that the added value of the Digital Building Master does not lie in producing technically superior or optimised design outcomes, but in reshaping how Environmental decisions are formed in the early SO/VO phase. The presence of the AI agent does not accelerate decision-making. Instead, it redistributes effort over time by slowing down interpretation and articulation upfront, while stabilising outcomes later in the project trajectory. With BIMi, assumptions remain visible for longer, parallel Environmental logics stay active, and premature convergence is delayed. This does not result in objectively “better” designs in a quantitative sense, but in decisions that are more coherent, traceable and defensible across phases.

The workshop further shows that agency remains fully human yet is temporarily redistributed. BIMi does not act as a designer or decision-maker, but as a reflective support tool that reorganises interaction patterns. Informal hierarchies weaken when the agent is introduced, expertise shifts from authoritative rule-setting to validation, and decision authority ultimately returns to the developer during convergence. Different stakeholders activate different roles of the agent depending on the moment and the question at hand. Sensing, steering and support are therefore not sequential roles, but situational capacities that emerge within a shared design space.

Importantly, the validation makes clear that objectivity is not guaranteed by the system itself. BIMi is entirely dependent on user input, document selection and prompt formulation. This creates sensitivity to framing, bias and strategic steering. Rather than eliminating these dynamics, the Digital Building Master makes them visible by exposing how Environmental decision making is structured and constrained. Integration therefore arises not from automation, but from transparency.

Taken together, the findings of this chapter confirm that the Digital Building Master contributes to Environmental decision-making by improving the conditions under which decisions are made, not by replacing human judgement. Its value lies in early integration, explicit negotiation and reduced reinterpretation later in the process. Time is not saved through faster choices, but by preventing silent prioritisation and deferred conflict.

Chapter 08

Conclusion

- 8.1 Answer to sub-question 1
- 8.2 Answer to sub-question 2
- 8.3 Answer to sub-question 3
- 8.4 Answer to main question

8. Conclusion

This chapter answers the main research question by returning to the four sub-questions. Together, these answers demonstrate how an AI agent can structure the integration of sustainable design choices in the early SO/VO phase of corporate real estate development.

8.1 Answer to sub-question 1

How are Environmental (E) considerations currently interpreted and handled by design teams, and where does fragmentation create space for the AI agent to act as a Digital Building Master?

The findings show that design teams enter the early SO/VO phase with strong Environmental ambitions. Labels such as Paris Proof, Net Zero Carbon and BREEAM Outstanding function as shared reference points. However, these ambitions are rarely translated into explicit and consistently shared decision-making structures. Environmental considerations are interpreted through discipline-specific perspectives, experiential knowledge and informal heuristics. Assumptions remain implicit, meanings vary between actors and Environmental decision making becomes distributed across documents, calculations and conversations.

This results in structural fragmentation. While a shared vocabulary exists, the underlying Environmental logic does not fully align. Operational energy performance, material-related impacts and certification requirements evolve in parallel strands that only partially intersect. Tensions between these strands typically surface late, when design choices have already narrowed the solution space. Environmental ambition therefore does not disappear, but gradually shifts in response to feasibility, cost and compliance pressures.

This fragmentation is particularly consequential in the early design phase. Decisions taken during SO/VO have a disproportionate impact on Environmental performance, while formalisation and accountability remain low. It is precisely within this condition of high impact and low structure that space emerges for an AI agent to act as a Digital Building Master.

The agent does not intervene as a decision-maker. Instead, it strengthens Environmental decision making through three complementary roles. As sensing, it exposes implicit assumptions and reveals Environmental tensions at an early stage. As steering, it structures information and presents Environmental alternatives in a comparable and neutral manner. As support, it stabilises definitions and anchors decision making to traceable sources. Together, these roles transform fragmented Environmental decision making into a more explicit, coherent and manageable interpretative process.

Sub-question 1 thus shows that the added value of the AI agent arises directly from existing fragmentation. The agent does not resolve this fragmentation through

optimisation, but by making it visible at a moment when design flexibility is still high.

8.2 Answer to sub-question 2

What functional qualities should an AI agent possess to support the structuring of Environmental (E) considerations in early design discussions?

The findings show that an AI agent should not be designed as an optimisation engine or decision-making authority. Its primary function is to structure Environmental decision making under conditions of uncertainty and complexity. To fulfil this role, the agent requires three interrelated functional qualities: sensing, steering and support.

Sensing refers to the ability to make implicit Environmental assumptions explicit. This does not involve autonomous risk detection, but structured decomposition of design questions into Environmental components that can be linked to relevant standards and datasets. In doing so, uncertainties, dependencies and tensions become visible early in the process.

Steering operates without prescription. Environmental integration benefits from comparison rather than prioritisation. Steering therefore consists of organising Environmental information into transparent and comparable scenarios. This delays premature convergence while preserving design freedom and disciplinary autonomy.

Support provides continuity and traceability. Environmental decision making in early design phases is vulnerable to drift as actors change, documents evolve and assumptions are overwritten. The agent stabilises definitions, anchors calculations to their sources and maintains coherence across project phases without fixing outcomes.

Together, these qualities define the functional profile of the Digital Building Master. Rather than improving Environmental outcomes directly, the agent improves the conditions under which sustainable design choices are discussed, negotiated and legitimised.

8.3 Answer to sub-question 3

How can these qualities be translated into a prototype AI agent that can be explored and evaluated in practice?

The functional qualities identified in this research were translated into a prototype designed explicitly as a research artefact. The objective was not to create a deployable system, but to make the behavioural effects of a Digital Building Master observable within a realistic early-design context.

The prototype, BIMi, was developed through a Research Through Design approach. This allowed sensing, steering and support to be embedded directly into the structure of the agent's workflow. BIMi was evaluated not on technical

performance or optimisation outcomes, but on its capacity to structure Environmental decision making during the SO/VO phase.

To ensure relevance to practice, the prototype was implemented in Microsoft Copilot Studio within a controlled Azure environment. The scope was deliberately limited to a single but characteristic Environmental tension: the relationship between operational energy reduction and material-related emissions associated with rooftop PV systems. Interviews confirmed that this tension is representative of early-stage Environmental decision making in corporate office development.

The prototype was grounded in a reconstructed human Environmental workflow based on validated project data from EDGE Stadium. This ensured that evaluation focused on changes in structure and interpretation rather than computational novelty.

Sensing was operationalised through decomposition of user queries and linkage to distinct Environmental domains. Support was operationalised through source anchoring and parameter stability. Steering was operationalised through neutral side-by-side scenario generation. Crucially, the prototype remained deliberately incomplete. It structured decision making but did not resolve tensions or make decisions. This incompleteness was essential to observe how the agent influences human Environmental decision making rather than replacing it.

Validation showed earlier exposure of Environmental tensions, delayed convergence and increased traceability of assumptions. The prototype therefore provides a concrete and evaluable translation of the Digital Building Master concept.

8.4 Answer to sub-question 4

How does the presence of the AI agent influence how design teams articulate, compare and reflect on Environmental (E) considerations during decision-making?

Validation demonstrates that the AI agent primarily influences how Environmental decision making is structured rather than which decisions are made. The agent does not change Environmental ambitions or introduce new technical knowledge. Instead, it reshapes the cognitive and social structure through which Environmental considerations are interpreted and integrated.

In human-only workflows, Environmental decision making is characterised by rapid convergence driven by experience, informal hierarchy and pragmatic constraint handling. Divergence and convergence occur almost simultaneously, compressing Environmental decision making into heuristics.

With the AI agent present, this structure shifts. Divergence expands and convergence is delayed. Assumptions, dependencies and Environmental tensions are surfaced and kept active. The agent structures Environmental decision making by decomposing questions, organising information into parallel scenarios and anchoring interpretations to explicit sources. This prevents premature closure and redistributes interpretative effort to an earlier stage.

As a result, Environmental considerations that would normally surface late become part of early design discussions. Informal hierarchies temporarily weaken. Expertise shifts from assumption-making to validation. Decision authority remains human, but Environmental decision making becomes more explicit, collective and traceable.

8.5 Answer to the main question

How can an AI agent, operating as a Digital Building Master, structure human Environmental (E) decision making processes in the early design phase of Dutch corporate real estate development?

This research shows that an AI agent can support Environmental (E) decision making in the early SO/VO phase of corporate real estate development by structuring the way sustainable design choices are articulated, compared and integrated. The agent does not optimise outcomes, automate decisions or replace professional judgement. Its contribution lies in reshaping the interpretative conditions under which Environmental decision making takes place in contexts characterised by uncertainty, time pressure and accountability.

By making assumptions explicit, keeping parallel Environmental logics visible and anchoring reasoning in traceable sources, the Digital Building Master alters how decisions are formed rather than what decisions are made. Environmental tensions surface earlier, interpretations become comparable, and choices remain intelligible across actors and project phases. This strengthens integration without accelerating convergence or enforcing trade-offs.

Sustainable design choices therefore remain human decisions, but they are made within a more transparent, coherent and accountable decision-making structure. In this way, the Digital Building Master supports not the content of decisions, but the process through which Environmental considerations are integrated into sustainable design choices over time.

Chapter **09**
Discussion

9. Discussion

This research shows that the value of an AI agent in the early SO VO phase does not lie in better calculations or technical optimisation. The added value lies in how environmental decision-making is shaped at a time when design choices are still open. The Digital Building Master does not change what is decided, but how decisions are made. By making assumptions explicit and comparing different environmental arguments, a shared framework is created. Choices remain traceable across project phases and involved parties. As a result, environmental integration shifts from individual experience to joint and transparent decision-making. Considerations become less dependent on who is at the table and easier to follow over time.

Integration as a cognitive and organisational process

The results of the study show that integrating sustainable design choices in the early design phase is primarily a cognitive and organisational issue. Previous research in corporate real estate shows that environmental ambitions are often mentioned but rarely explicitly developed jointly (Olteanu & Ionaşcu, 2023). In practice, they remain intertwined with routines, experiential knowledge and informal coordination within design teams.

This picture also emerges from the interviews and workshops. Design teams use similar terms such as Paris Proof or BREEAM Outstanding but assign different meanings to them. As a result, environmental ambition gradually shifts through pragmatic choices rather than explicit renegotiation. Ifediora et al. (2024) describe this as a negotiation process in which sustainability is not fixed but is constantly being interpreted.

In this context, integration fails not because of a lack of data or tools, but because of fragmented decision-making across disciplines and moments in time. Digital tools often support analysis but leave interpretation to individuals (Motalebi et al., 2025a). The Digital Building Master intervenes precisely here. By making assumptions and dependencies visible, a shared interpretative framework is created in which environmental decision-making becomes discussable and reproducible.

Within corporate real estate development, this fragmentation is exacerbated by complex stakeholder structures and long project durations. Early choices must remain defensible later to financiers, certification bodies and portfolio managers (Riratanaphong & Pewklieng, 2025). Integration here is less about optimisation and more about robustness: the ability to keep environmental decision making consistent over time and across stakeholders.

Integration without optimisation or enforced trade-offs

An important insight from this research is that environmental integration does not automatically mean that design choices must be optimised or weighed against each other. In the PV case, two environmental perspectives coexist: the

contribution of solar panels to operational energy performance and the material-related impact of the system itself. These perspectives do not disappear into a single outcome, nor are they played off against each other.

The AI agent does not force these logics together into a single score or conclusion. It does not choose which is more important. Instead, it keeps both perspectives visible and shows how they influence each other. This makes it clear where tensions lie, but also where there is room for further design. This proves particularly valuable in an early design phase, when much is still uncertain and there is considerable design freedom.

Solar panels function as a connecting design element in this process. They simultaneously touch on spatial design, material choices and energy ambitions. This is precisely why they lend themselves to integration without the need to force immediate decisions. Instead of closing the discussion, they keep it open. This is at odds with optimisation-oriented sustainability approaches, in which integration is often understood as combining, ranking or weighing indicators to arrive at a single best solution.

The findings show that such optimisations at an early stage can narrow the design process. Early decisions eliminate alternatives and leave assumptions undiscussed. This research shows that integration becomes stronger when multiple environmental perspectives are allowed to coexist. Sustainable design choices then arise not as technical outcomes, but as joint and substantiated choices.

Integration here does not mean deciding early but carefully postponing. By keeping environmental dimensions visible and connected while the design is still open, space is created for better considerations later in the process. This makes integration less a question of calculation and more a matter of timing, dialogue and shared understanding.

Temporal shifts and the quality of integration

The comparison between the human design process and the AI-supported process reveals a clear difference in pace. In the human process, exploration and decision-making are almost immediately intertwined. Design teams quickly arrive at choices based on experience, established routines and informal hierarchies. This ensures progress, but it also reduces the scope for truly comparing different environmental perspectives. Assumptions often remain unspoken and tensions disappear into rules of thumb or previous project experiences.

When the Digital Building Master is deployed, this pace changes deliberately. In the SO VO phase, the interpretation process is slowed down. Exploration takes longer and decisions are made less quickly. Different Environmental arguments continue to coexist before being brought together. This does not lead to technically better solutions in terms of performance or optimisation. The added value lies elsewhere.

By slowing down the pace, the quality of integration shifts. Sustainable design choices become less dependent on individual intuition and more strongly linked to explicit assumptions, applicable standards and available evidence. Design choices thus become easier to explain and easier for others to follow.

It is important that this work does not disappear but is shifted in time. Integration that normally only becomes visible later in redesign, reinterpretation or negotiation now takes place earlier. From an academic perspective, this requires a different view of efficiency in sustainable design. Efficiency here does not mean faster decision-making, but more stable and consistently integrated decision-making across project phases. This is particularly relevant in corporate real estate development, where early choices have long-term financial and portfolio effects.

Redistribution of expertise and interpretative authority

This research shows that the Digital Building Master influences who within the design team directs environmental decision-making. In the human process, this role often lies with a limited number of experts. Their experience serves as an implicit reference point and guides choices without assumptions always being explicitly stated.

With the AI agent present, this shifts. Assumptions, calculations and reasoning become visible to the entire team. Expertise does not disappear but takes on a different position. Instead of being the sole source of environmental knowledge, experts take on a more assessing and interpreting role. They assess and interpret the information structured by the agent instead of just providing it themselves.

This shift temporarily flattens decision-making. Hierarchies play a less automatic role and fixed leadership patterns are disrupted. At the same time, new forms of distinction emerge. The ability to properly guide the agent, ask critical questions and place outcomes in the right context becomes a new professional skill. Designers and consultants who combine environmental expertise with an understanding of AI appear to be better positioned in this regard.

The role of sustainability consultants is changing as a result. They are acting less as primary knowledge carriers and more as validators of AI-structured decision-making. Professional judgement remains essential but is used within a more explicit and transparent framework. Integration remains a human responsibility but takes place within an environment in which assumptions and choices are more visible.

Within corporate real estate development, this shift is in line with broader changes in how sustainability expertise is organised. Knowledge is no longer concentrated in a single role but distributed across internal teams, external advisors and digital tools. The Digital Building Master reinforces this development by making environmental decision-making more accessible and collaborative without diminishing the importance of human expertise.

Objectivity, framing and integration risks

During validation, the same tension kept coming up. The output of the Digital Building Master is quickly seen as objective and reliable. The information is clear. The reasoning is consistent. That inspires confidence. At the same time, this conceals an important dependency. The agent can only work with what it is given. What becomes visible in terms of environmental considerations is determined earlier in the process. The selection of documents, the composition of the knowledge base and the formulation of prompts determine which assumptions are made explicit and which remain out of sight. The agent does not distort information on its own. It reinforces and stabilises the framework that is inherent in the input. Precisely because this framework is presented clearly and systematically, it becomes more difficult to question it. This creates a specific risk. Not of deliberate manipulation, but of selective exclusion. Environmental aspects that are not included in the input do not gradually disappear from the conversation. They simply do not come into view. Integration thus appears complete, when it is incomplete.

AI-supported integration does not take power away from environmental decision-making. It shifts it. Influence lies with those who determine what data is used, how questions are asked and where analytical boundaries are drawn. In corporate real estate development, where environmental arguments converge with financial interests and certification logic, this dependency can be used strategically. This underlines that the Digital Building Master should not be seen as a neutral actor. Its epistemic authority requires a counterbalance. This means joint agreements on input. Awareness of what is excluded. And ongoing human reflection on what is and is not made visible. Integration therefore remains a human responsibility. The agent can reinforce this process, but only when its dependencies are recognised and actively managed.

Methodological and theoretical reflection on integration

The findings of this study should be interpreted within the limitations of a qualitative single-case study in the SO VO phase of corporate real estate development. The contribution of the Digital Building Master does not lie in general applicability or demonstrable improvements in environmental performance or optimisation. Its value lies in the insight it provides into how environmental decision-making is structured, stabilised and legitimised at an early stage under conditions of uncertainty and open interpretation.

By explicitly positioning the AI agent as reflective support rather than a decision-maker, this research does not align with optimisation-oriented sustainability literature. Instead, it aligns with scientific discussions on integration, sensemaking and organisational decision-making. The agent does not resolve trade-offs or produce new environmental knowledge but influences the conditions under which sustainable design choices are discussed and recorded before choices become final.

At the same time, the research shows that AI-supported integration introduces new dependencies. The appearance of objectivity masks a strong dependence on previous choices in document selection, knowledge base and prompt formulation. Interpretative authority does not disappear but shifts to those who determine these inputs and analytical boundaries. As a result, integration does not appear as a neutral or purely technical achievement, but as a process that must be actively steered and monitored. Coherence and transparency remain dependent on human judgement and continuous reflection.

Chapter **10**
Limitations

10. Limitations

This chapter reflects on the limitations of the study in direct relation to the research design and the four sub-questions. Each limitation follows from deliberate methodological and conceptual choices made to investigate how an AI agent can support Environmental (E) decision making and the integration of sustainable design choices in the early SO/VO phase of corporate real estate development. The limitations constrain interpretation and generalisability, while simultaneously clarifying the conditions under which the findings are valid.

Limitations related to sub-question 1 and qualitative interviews

This sub-question was investigated through semi-structured interviews with professionals involved in the early stages of corporate real estate development. A first limitation relates to the interpretative nature of this method. Interviews provide insight into how decision-making is articulated and reflected upon, but not into actual decision outcomes. Implicit assumptions, informal negotiations and power relations may therefore remain partly hidden, despite targeted questioning techniques.

In addition, the interview group is limited in size and situated within an organisational context. However, people from other organisations were also interviewed, but they are only familiar with the working methods of the main organisation. This allows for in-depth analysis and context sensitivity, but limits statements about how environmental fragmentation manifests itself in other organisations or project types. The patterns identified should therefore be read as context-specific configurations of environmental decision-making, not as universal characteristics of design teams. Finally, the analysis focuses primarily on environmental considerations that have already been partially formalised through standards, certifications and calculation practices. More informal or implicit forms of environmental decision-making are therefore less prominent. This influences the extent to which fragmentation can be interpreted outside highly structured environmental domains.

Limitations related to sub-question 2 and conceptual modelling

The functional qualities of sensing, steering and support were developed through abductive decision-making, combining insights from interviews with conceptual reflection. Together, these qualities form an analytical model and do not constitute an exhaustive or definitive classification of possible AI roles. Other ways in which AI can support environmental decision-making are conceivable but fall outside the scope of this study. The model assumes that environmental decision-making benefits from greater explicitness, comparability and traceability. This assumption is supported by the empirical material, but at the same time reflects a normative choice about how sustainable design choices should ideally be made. In design contexts where speed, intuition or informal coordination are central, these functional qualities may be experienced differently or even seen as obstructive.

Therefore, the proposed functional profile of the Digital Building Master should be read as context sensitive. It offers insight into how AI Environmental can structure decision-making within this specific research setting but does not claim universal applicability.

Limitations related to sub-question 3 and Research Through Design

The prototype was developed using a Research Through Design approach and deliberately positioned as a research artefact, not as a system ready for practical application. As a result, no statements can be made about technical robustness, scalability or long-term operational use. The value of the prototype lies in exploration and understanding, not in technical maturity.

In addition, the prototype focuses on one specific environmental tension: the relationship between operational energy performance and material-related emissions from rooftop solar panels. This tension is analytically sharp and recognisable within corporate real estate development but represents only one type of environmental integration issue. The findings therefore show how AI can structure decision-making within clearly defined environmental dilemmas, not how it functions across the full spectrum of environmental considerations.

Furthermore, the prototype operates within a limited document environment. Environmental aspects that are not included in the knowledge base cannot be visible in the decision-making process. The selection and delimitation of sources thus function as epistemic filters. This increases the transparency of existing decision-making, but at the same time introduces a new form of framing that requires active attention and governance.

Finally, the research does not assess the technical quality or performance outcomes of sustainable design choices. The focus is on how choices are articulated, compared and stabilised under uncertainty. This limits statements about environmental optimality but strengthens insight into the integration process itself.

Limitations related to sub-question 4 and validation design

The influence of the AI agent was investigated through a single validation workshop with a limited number of domain experts. This setup allowed for detailed observation of interaction patterns, decision-making behaviour, and role shifts within the team. At the same time, this setting does not allow for statistical generalisation. The observed effects should therefore be interpreted as trends within this specific context and not as sector-wide patterns.

In addition, differences in AI skills among participants were found to influence their interaction with the prototype. Participants who had difficulty with the structure or pace of the agent temporarily lost influence in the discussion, regardless of their substantive expertise. This points to a risk of epistemic inequality, whereby technological skill temporarily outweighs professional knowledge.

The structured and consistent presentation of AI-generated reasoning also introduces a risk of premature consensus. Under time pressure, teams may be too quick to rally behind AI-supported scenarios without sufficiently questioning the underlying assumptions. This underlines that human judgement and explicit responsibility for sustainable design choices remain necessary, even in AI-supported decision-making.

Finally, the observed role shifts are rooted in a single organisational and project context. Design teams with different organisational cultures, contract types or levels of digital maturity may experience these effects differently. The findings provide insight into possible dynamics, but do not claim to represent the universal effect of AI in environmental decision-making.

Cross-cutting limitations related to scope and context

Across all sub-questions, this study deliberately focuses on the Environmental (E) dimension to maintain analytical clarity. Social and governance considerations were excluded, not because they are less relevant, but because their inclusion would have obscured the specific mechanisms through which Environmental decision making is structured in early design phases. The findings should therefore not be read as a contribution to integrated ESG decision making, particularly in domains that rely more heavily on political negotiation or informal governance.

The research is situated within corporate real estate development, a context characterised by certification-driven ambition, long investment horizons and formal accountability structures. Other real estate segments operate under different governance regimes and temporal dynamics. As a result, the observed mechanisms of Environmental integration cannot be directly extrapolated beyond corporate real estate without additional empirical investigation.

Finally, the Digital Building Master, including the roles of sensing, steering and support, is positioned as an analytical lens rather than a comprehensive theory of Environmental decision making. Alternative theoretical perspectives may foreground different dynamics or yield different interpretations of the same empirical material.

Implications of the limitations

Taken together, these limitations clarify the conditions under which the Digital Building Master can be understood as a reflective support tool for Environmental decision making. The agent does not resolve uncertainty, optimise outcomes or replace professional judgement. Its contribution lies in reshaping how Environmental considerations are articulated, compared and legitimised within the bounded context of early-stage corporate real estate development.

Rather than weakening the findings, these limitations specify the type of insight this research offers: not prescriptive solutions, but an analytical understanding of how Environmental integration can be made more explicit, traceable and discussable when design choices are still open.

Chapter **11**
Recommendations

11. Recommendations

The findings of this research show that the primary value of an AI agent in the early SO/VO phase of corporate real estate development lies in its capacity to structure Environmental (E) decision making and support the integration of sustainable design choices. The recommendations therefore do not concern technological optimisation or system performance, but focus on organisational, professional and governance conditions that enable such an agent to function as a reflective support tool rather than as a decision authority.

11.1 Recommendations for practice

Invest in shared AI literacy as a collective condition

The validation shows that differences in familiarity with AI have a direct impact on who can actively participate in environmental decision-making. Participants who understand the structure and working methods of the agent retain influence in the discussion. Others temporarily lose that influence, regardless of their substantive expertise.

This points to a risk of epistemic inequality that does not stem from professional knowledge, but from technological skill. To prevent the agent from reinforcing existing inequalities, a shared basis of AI literacy is needed within design teams. This is not about technical depth, but about a common understanding of what the agent does, what it does not do, and how its output should be interpreted. Only then can the agent function as collective reflective support rather than as a mechanism of exclusion.

Reposition the role of sustainability advisors

The findings show that the role of sustainability advisors in early environmental decision-making is shifting. Because the agent provides direct access to standards, reference values and structured scenarios, the work is shifting from providing information to interpreting, validating and monitoring consistency.

This shift requires explicit recognition in roles, work processes and collaboration agreements. The value of the sustainability advisor lies less and less in providing figures and increasingly in interpreting assumptions, dealing with uncertainty and monitoring consistency between disciplines. When the agent is seen as a calculation tool and the advisor as an operator, both remain underutilised. Their strength lies precisely in the combination of structured AI output and professional judgement.

Position the agent explicitly as reflective support

The validation shows that the agent contributes most when it is used to reveal assumptions, keep environmental tensions open and structure alternatives without immediately resolving them. When teams approach the agent as an authoritative source or decision-making logic, there is a risk of excessive dependence and premature consensus.

It is therefore important to explicitly position the agent as reflective support. This means that AI output is always provisional, remains interpretative and must be subject to professional consideration. Sustainable design choices remain the result of consultation, negotiation and responsibility, even when decision-making is structured by AI.

Curate and govern the Environmental knowledge base

The agent's output is entirely dependent on the quality, selection and structure of its knowledge base. Incomplete, selectively curated or strategically framed inputs directly shape which Environmental considerations become visible and comparable.

Organisations should therefore invest in a central, curated Environmental knowledge library that includes stable normative references, emission factors, internal policies and project-specific data. This library should function not only as technical input for the agent, but also as institutional memory that preserves the traceability of Environmental decision making across project phases and team changes.

Importantly, decisions about inclusion, exclusion and updating of sources should be treated as governance issues rather than technical tasks. Curating the knowledge base is itself an act of Environmental framing and requires transparency and accountability.

Establish explicit governance and accountability arrangements

Finally, clear governance arrangements are essential. While the agent influences how Environmental decision making is structured and when assumptions surface, it does not take responsibility for decisions. Organisations must explicitly define who formulates prompts, who validates AI-supported decision making, how conflicting interpretations are handled and where final accountability lies.

Without such arrangements, the agent risks reinforcing framing effects or enabling strategic use of AI output to legitimise predetermined choices. As AI increases transparency in Environmental decision making, it simultaneously increases the need for explicit procedural and ethical boundaries. Governance therefore becomes more important, not less.

11.2 Recommendations for future research

The academic contribution of this research lies in explaining how Environmental (E) decision making is structured in early-stage corporate real estate development and in conceptualising the role of AI as a reflective support mechanism rather than as a decision-maker. Building on this contribution, several directions for future research emerge.

Extending process-oriented analysis beyond Environmental decision making

This study demonstrates the value of analysing sustainability not as an outcome, but as an integration process shaped by interpretation, comparison and

stabilisation of decision making. Future research could extend this process-oriented perspective to other domains of sustainable decision-making in the built environment. In particular, studies could examine whether similar structuring effects occur in domains that are less formalised than Environmental decision making, such as social value creation or governance-related trade-offs. Such research would test the transferability of the Digital Building Master concept beyond analytically structured domains.

Comparative studies across Environmental domains and project types

The empirical contribution of this study is grounded in a single Environmental tension within corporate office development. Future research could investigate how AI-supported structuring operates across different Environmental themes, such as circular material use, climate adaptation or biodiversity integration. Comparative studies across multiple project types and organisational contexts would allow researchers to assess whether the observed mechanisms of delayed convergence, increased traceability and redistributed interpretative effort remain consistent or vary with context.

Longitudinal research on AI-supported decision making over project lifecycles

This research focuses on the early SO/VO phase, where uncertainty is high and commitments are provisional. Future research could adopt a longitudinal approach to examine how AI-supported Environmental decision making evolves across subsequent project phases. Such studies could investigate whether early structuring effects persist, how assumptions are reinterpreted over time, and whether traceability supports accountability during later decision moments such as investment approval or certification.

Methodological development of Research Through Design for AI studies

By using a deliberately incomplete prototype, this study demonstrates how Research Through Design can be employed to study cognitive and organisational effects of AI rather than technical performance. Future research could further develop this methodological approach by comparing different prototype fidelities, interaction designs or degrees of AI autonomy. This would contribute to methodological debates on how to empirically study AI-human interaction in design contexts where knowledge is provisional and distributed.

Professional role development and AI literacy in design teams

The findings reveal shifts in professional roles and emerging forms of expertise, such as prompting and interpretation of AI-supported decision making. Future research could examine how these roles develop over time, how AI literacy becomes institutionalised within organisations, and how responsibility and authority are renegotiated in repeated project settings. Such studies would deepen understanding of the socio-professional implications of AI as reflective support.

Governance, ethics and accountability of AI-supported decision making

Finally, this research highlights that AI-supported structuring increases the need for explicit governance rather than reducing it. Future research could explore governance models that address framing effects, responsibility allocation and ethical boundaries in AI-supported sustainability decision making. This includes studying how organisations prevent over-reliance on AI output, manage conflicts between human judgement and AI-supported scenarios, and ensure accountability for sustainable design choices.

Chapter **12**
Implementation strategy

12. Implementation strategy

This chapter translates the findings of the research into an implementation strategy for corporate real estate organisations that consider introducing AI to support Environmental decision making in the early SO/VO phase. The strategy is not framed as a technological roadmap, but as a set of organisational, professional and procedural conditions that must be in place for the Digital Building Master to function as intended. Each element of the strategy is directly grounded in the empirical findings, the prototype evaluation and the abductive Research Through Design approach adopted in this thesis.

Build a shared Environmental knowledge base

The interviews and validation show that environmental considerations are spread across various documents, disciplines and individual routines. This fragmentation makes it difficult to trace assumptions and track choices across project phases. The effectiveness of the Digital Building Master proved to be entirely dependent on the availability of a shared and consistent knowledge base.

Therefore, a curated environmental knowledge environment is not a technical prerequisite, but an organisational core requirement. The knowledge base functions as an institutional memory in which assumptions, standards and references remain explicit, even when teams or project phases change.

Develop shared AI literacy

During validation, it became apparent that differences in AI literacy had a direct impact on participation and interpretative authority in the conversation. Participants who understood the structure and logic of the agent retained influence; others temporarily lost their position, regardless of their substantive expertise.

This underlines that AI literacy must be a collective professional skill. It is not technical depth that is needed, but a shared understanding of how it works, its limitations and the interpretation of its output, to prevent technological skill from outweighing domain knowledge.

Remain critically human

The validation showed that AI-structured output is quickly perceived as objective and authoritative. However, this apparent neutrality masks the dependence on input, source selection and framing. The agent makes assumptions visible but does not generate truth.

Therefore, human alertness must be explicitly organised. The value of the agent lies not in certainty, but in evoking reflection. Human judgement remains necessary to recognise blind spots, missing perspectives and convincing but partial reasoning.

Support team dialogue and dynamics

One important effect of the agent was that convergence in the SO/VO phase was delayed. Whereas teams normally make quick decisions based on experience and hierarchy, the agent kept multiple environmental arguments visible in parallel. This did not improve the technical outcome, but it did improve the quality of integration.

Implementation must support this dynamic rather than correct it. Teams must learn that delayed convergence is not inefficiency, but a prerequisite for explicit and shared Environmental decision-making.

Redefine the role of the sustainability advisor

The research shows that the agent redistributes interpretative authority. Environmental expertise shifts from implicit individual reference to explicit team reflection. This changes the role of sustainability advisors.

Their value no longer lies primarily in providing information, but in interpreting, validating and monitoring consistency in AI-structured decision-making. This shift in role must be explicitly recognised in order to prevent resistance and mispositioning.

Treat the early design phase as decisive

The findings show that environmental integration is particularly vulnerable in the early stages, when assumptions are still implicit but later have major consequences. The agent proved to be valuable precisely there, because it makes assumptions explicit before they become entrenched in routines or checklist logic.

This confirms that implementation should focus on the SO/VO phase, where there is still considerable room for interpretation and environmental decision-making can still be shaped rather than justified.

Aim for robust, not faster, decision-making

The research shows that AI support does not lead to faster or better optimisation. In the PV case, parallel Environmental logics deliberately continued to coexist. The value lay in traceability and comparability, not in reduction to a single metric.

Therefore, implementation should focus on robust decision-making that stands the test of time, phases and stakeholders, rather than on speed or efficiency now.

Establish clear governance and accountability

Because the agent is completely dependent on its input environment, power shifts to those who select sources, formulate prompts and set boundaries. This creates new governance issues. Clear agreements on responsibility, validation and decision-making are necessary to prevent AI output from being used as a means of legitimising predetermined choices. Transparency increases the need for governance, not the other way around.

Chapter 13

Reflection

- 13.1 Product, process, planning
- 13.2 Method choice
- 13.3 Looking ahead
- 13.4 Master track reflective
- 13.5 Ethical reflection
- 13.6 Transferability
- 13.7 Extra

13. Reflection

This chapter reflects on the research design and the development of the prototype. The focus is not on repeating results, but on weighing up choices, evaluating their effects and clarifying the limits within which the findings are meaningful. Reflection is not a final step in this research, but an integral part of the process.

AI is a relatively new and complex subject within property development. This made it necessary to constantly review choices, clarify assumptions and actively limit the scope of the research during the study. The value of the results is therefore directly linked to the way in which the research developed iteratively. In this chapter, I make my role as a researcher explicit and position the results within the choices that were made along the way.

13.1 Product, process, planning

The research product developed gradually. The starting point was a broad question about the relationship between AI and environmental ambitions in property development. In this early phase, the focus was on ESG in a broad sense and on the idea that a digital tool could structure fragmented sustainability ambitions. How AI would function in practice within the design team was still unclear at that point.

Literature research and interviews brought clarity to this. As my understanding of AI and environmental decision-making grew, it became clear that the original ambition to develop a multi-agent system was too broad for the purpose of this research. That breadth made it difficult to visualise mechanisms and led to conceptual ambiguity. The decision to zoom in on a single central AI agent was therefore not a simplification, but a necessary demarcation to achieve depth.

This choice brought about a clear shift. The focus shifted to a defined and recognisable environmental tension: the relationship between operational emission reduction and material-related emissions in PV systems. By explicitly limiting this domain, it became possible not only to describe environmental decision-making, but also to understand it analytically. At the same time, this choice made the limitations of the research explicit. The findings say something about decision-making within sharply defined environmental dilemmas, not about sustainability in a broad sense.

This shift also changed the nature of the research product. The AI agent was no longer seen as the goal, but as an analytical tool. The prototype functioned as a lens to reveal differences between human and AI-structured decision-making. Thus, the design became not a solution, but a means to enable reflection. This is directly in line with the Research Through Design approach, in which research and design develop in parallel and mutually reinforce each other.

The planning of the research supported this iterative approach. The theoretical basis was laid early on, which allowed room for experimentation and adjustment in the subsequent phases. The alternation between progress and reconsideration was sometimes confrontational but proved essential to achieving a stable and coherent research product. It was precisely by taking steps back that assumptions could be tested and choices could be more clearly substantiated. This dynamic is characteristic of Research Through Design and contributed to the quality and consistency of the research.

13.2 Method choice

The methodological approach of this study was deliberately adjusted during the process. The original approach was to analyse the digital role within the design team from an actor-oriented perspective. In practice, this approach proved insufficiently aligned with the objective of the research. The focus became too strongly on theoretical positioning, while the core question revolved around how a digital tool could strengthen environmental decision-making in practice. This gap between method and research objective made it necessary to change the approach.

The choice of Research Through Design offered a more suitable framework. This approach made it possible to directly link research and prototyping. Insights from literature and interviews could be immediately translated into the prototype, while design decisions in turn raised new analytical questions. This was essential because my knowledge of AI was limited at the start of the research and developed through designing and testing. The prototype thus served not only as a research object, but also as a learning tool.

The concept of the Digital Building Master, elaborated in the three functions of sensing, steering and support, provided guidance within this iterative working method. These functions are not intended as fixed or simultaneously active roles, but as context-dependent forms of support. Depending on the project phase, type of uncertainty and actors involved, the relevant function shifts. This distinction helped to organise complexity without reducing it and prevented the agent from being positioned as an all-encompassing solution.

At the same time, this method presented challenges. The iterative nature of Research Through Design made it difficult to keep the scope sharp from the outset. The initial breadth of the sustainability domain created uncertainty about what did and did not belong in the research. Explicitly focusing on the environmental dimension was not an obvious choice, but the result of repeated reflection and feedback moments. By revisiting earlier steps, the boundaries became clearer and the method gained more focus and stability.

Ultimately, this approach confirmed the suitability of Research Through Design for this study. The prototype did not develop into a final product, but into an analytical tool that provided insight into how environmental decision-making is structured. It was precisely in the interaction between research and design that assumptions

became visible, tensions became explicit and the boundaries of the system became clear. In this way, the chosen method was consistent with the aim of the research: not to provide solutions, but to deepen understanding.

13.3 Looking ahead

The final phase of the research was devoted to refinement and clarification. Feedback from the P4 presentation and subsequent discussions was used to clarify the structure of the thesis and position its core contribution more explicitly. In this phase, the focus was deliberately narrowed further to environmental decision-making, concepts were used more consistently, and the line between problem definition, method, results and discussion was sharpened.

In addition, the evaluation of the AI agent was deepened. By looking back at the prototype interactions and the validation workshop, what the agent does and does not do was articulated more clearly. The emphasis has shifted from technical operation to the effects on interpretation, comparison and reflection within design teams. This has led to a more explicit positioning of the Digital Building Master as reflective support rather than as a decision-making system.

The reflective chapters were also further elaborated during this period. Limitations are explicitly linked to methodological choices and sub-questions, and the social and academic implications are more clearly separated. This has sharpened the scope of the findings and prevented the research from claiming more than it can substantiate.

In addition to the academic process, the research was also discussed outside the university during this phase. A presentation at EDGE provided an opportunity to test the clarity and recognisability of the concept in a professional context. This interaction did not yield any new data, but it did contribute to reflection on the applicability and communicative power of the research.

This concluding phase thus marks both a conclusion and a starting point. The research has been completed in its current form, but the questions surrounding AI and environmental decision-making in real estate development remain in flux. The insights gained and the framework developed form a basis for further exploration, both in research and in practice.

13.4 Master track reflective

The subject of this research ties in directly with the core of the Master Track Management in the Built Environment, which focuses on decision-making, design strategies and process organisation. Although the research has a clear technological component, its contribution lies not in technology itself, but in understanding how digital support influences decision-making and collaboration in the built environment. It is precisely this focus that makes the research relevant within MBE.

This master's track focuses on the question of how complexity is managed in projects involving multiple actors, divergent interests and changing definitions.

This research shows that environmental decision-making in the early design phase is highly dependent on consultation, interpretation and repeated coordination. The Digital Building Master makes these processes more explicit and easier to structure, without replacing human judgement. In this way, the research contributes to an understanding of how digital tools can support management processes rather than take them over.

At the same time, this research required a different form of analysis than is usual within MBE. Instead of focusing on formal decision-making structures or quantitative performance indicators, the emphasis was on interpretation, meaning-making and process dynamics. This broadening of focus is in line with the increasing role of digital systems in property development and shows that management issues are increasingly taking place at the intersection of technology, design and organisation.

The research is also relevant within the broader context of AUBS. Digital technology influences all disciplines, but the quality of design and decision-making remains dependent on human expertise. In architecture and design, craftsmanship remains essential, precisely because not everything can be automated. This research underlines that AI is not a substitute for design skills but can be a tool to make decision-making more explicit and better substantiated. The tension between digital support and human expertise is therefore a core reflection of this research and is in line with both the substantive and didactic objectives of MBE and the MSc AUBS.

13.5 Ethical reflection

The academic contribution of this research lies in revealing how AI Environmental can structure decision-making in the early design phase. The research shows that AI is not merely a technical tool but fulfils an epistemic role. By making assumptions explicit and structuring reasoning, the agent influences what is considered relevant, comparable and defensible. In doing so, AI directly addresses questions about knowledge formation and responsibility in design and decision-making processes.

The social significance of this shift lies in the way it changes design and development processes in real estate practice. By structuring interpretation and making differences explicit, Environmental reasoning becomes less dependent on individual routines and informal positions of power. At the same time, this redistributes the way influence is exercised within teams. Ethics plays a role here not only at the level of outcomes, but especially at the level of process and power: who determines what becomes visible, which assumptions apply, and when a discussion is considered closed.

A first explicit ethical tension concerns the environmental impact of AI itself. As elaborated in Appendix M, AI has both embodied emissions in hardware and data centres and operational emissions during use. This creates a paradox. The technology is used to analyse the environmental performance of buildings, but it

itself has an ecological footprint. This research cannot resolve this tension, but it does make it visible. The question of whether the process-related gains in decision-making outweigh the environmental costs of AI remains open and requires further empirical substantiation, particularly in the case of large-scale application in design practices.

A second ethical dimension concerns dependency and apparent objectivity. The Digital Building Master makes assumptions explicit but does not produce a neutral truth. Output remains dependent on input, source selection and system demarcation. Precisely because AI presents reasoning in a structured and convincing manner, there is a risk that outcomes will be seen as self-evident or objective. This can reduce critical reflection rather than enhance it.

The research shows that the ethical value of AI lies not in autonomy, but in reflection. The agent supports decision-making by raising questions, revealing tensions and enabling comparison. The responsibility for interpretation and consideration remains with humans. This responsibility cannot be delegated to technology. On the contrary, as AI structures decision-making, the need for explicit human judgement, transparent governance and conscious limitation increases.

Ethics in this research is therefore not about avoiding technology, but about positioning it carefully. The Digital Building Master can only be used responsibly when its epistemic influence is recognised and actively managed. This requires awareness, reflection and continued human involvement in the application of digital resources in the built environment.

13.6 Transferability

The transferability of the research lies in the structural design of the digital building master and not in the exact form of the prototype. The technical implementation is project-dependent, but the underlying logic is generic. The three roles — Sensing, Steering and Support, provide a stable structure for organising environmental decision making. This structure can be applied in other projects and other environmental themes, if the input data is adapted to the design issue.

The prototype can be scaled up in organisations that work with environmental goals in the early design phase. The reflective function of the agent makes assumptions visible and supports consistent decision-making. This is in line with the need for greater transparency in development processes. Organisations can integrate the agent into internal design meetings or use it to supplement existing analysis tools.

The design is also useful in education. Within the Master MBE programme, the digital architect can serve as a didactic tool to teach students how environmental decision making develops in multidisciplinary teams. The agent's roles help to make fragmentation and assumptions discussable. The agent thus functions as a training and reflection tool.

There are still limitations. The application requires access to project data and users who understand how to control the agent. The quality of the outcome remains dependent on human expertise. Nevertheless, the reflective effect of the agent is reproducible. The research shows that the way in which the agent makes assumptions visible and keeps design teams in an open design stage for longer is also valuable outside the case study.

The transferability therefore lies in the possibility of scaling up the prototype and deploying it in organisations and education. The core lies in the way the agent structures environmental decision making and not in the specific technical implementation.

13.7 Extra

How does the value of human intuition change when AI analyses faster and more consistently than humans?

During this research, I became increasingly aware of the persuasive clarity of the AI agent. The agent responds quickly, organises information and presents assumptions in a structured and coherent way. This creates a sense of efficiency and reliability. However, this clarity should not be confused with correctness. The agent does not know what is right; it only structures what is available.

Discussions with designers highlighted that human intuition operates in a fundamentally different way. Intuition does not compete with analytical consistency but complements it. It reads context, senses direction and recognises tensions that do not surface in datasets or formalised inputs. This became particularly evident during the prototyping phase. While the agent revealed patterns and relationships, decisions about relevance, priority and meaning remained human. This experience reshaped my understanding of intuition. Rather than being replaced by AI, intuition gains space when analytical noise is reduced. The relationship between human judgement and machine analysis shifts from competition to complementarity.

A related concern is the risk of AI gaining disproportionate influence because of the confidence with which it communicates. Throughout testing, the agent consistently presented its output in a structured and assertive manner. This can easily lead to over trust. At several moments, I found myself momentarily overlooking the fact that every output was dependent on my own inputs, assumptions and framing choices. The agent appears neutral, but it never is.

This observation underscores the need for safeguards. Effective use of AI requires deliberate moments of pause, in which users reflect on what the agent does not know, which data are absent and which assumptions are implicitly introduced through prompts and source selection. The agent excels at making certain aspects visible, but it cannot account for what remains outside its frame. Responsibility for recognising these blind spots remains firmly with humans.

Within design teams, this implies that AI should function as a conversation partner rather than a decision-maker. Clear agreements are needed about how AI output is interpreted, challenged and contextualised. These safeguards are not primarily technical. They are cultural and cognitive. They depend on awareness, critical engagement and reflective practice.

The central question is therefore not whether AI will gain influence in design processes, but whether teams remain sufficiently attentive to shape and limit that influence appropriately. This research suggests that such balance is achievable. However, it requires sustained human involvement and a continuous recognition that AI can be persuasive, but never omniscient.

What safeguards are needed to prevent AI from gaining too much influence because it sounds convincing?

During testing, one characteristic of the AI agent became particularly apparent: it consistently sounds certain. Its output is well structured, coherent and confident. This presentation creates trust and, at times, too much trust. I noticed moments in which I briefly overlooked that every output was shaped by my own inputs, assumptions and framing choices. The agent appears neutral, but it never is.

This observation makes the need for safeguards clear. Effective use of AI requires deliberate pauses to reflect on what the agent does not know, where the limits of the available data lie and which assumptions are implicitly introduced by the user. The agent is effective at revealing patterns and making reasoning explicit, but it cannot identify what remains outside its frame. Recognising these blind spots remains a human responsibility.

Within design teams, this means that AI should function as a conversation partner rather than a decision-maker. Clear agreements are needed about how AI output is interpreted, questioned and contextualised. These safeguards are not primarily technical. They concern mindset and practice: maintaining awareness, fostering critical engagement and creating space for reflection.

The central issue is therefore not whether AI will gain influence in decision-making processes, but whether teams remain sufficiently conscious to shape and limit that influence appropriately. This research suggests that such balance is achievable. However, it depends on continued human involvement and an explicit recognition that AI can be persuasive but is never omniscient.

Bibliography

- Allard Jolles, Nicoline Kok, Saskia Naafs, & Iris Thewessen. (2021). Bouwmeesters-2021-PDF. *College van Rijksadviseurs*.
- Blumer, H. (1954). What is Wrong with Social Theory? *American Sociological Review*, 19(1), 3. <https://doi.org/10.2307/2088165>
- Böttcher, T. P., Weber, M., Weking, J., Hein, A., & Krcmar, H. (2020). Value Drivers of Artificial Intelligence Value Drivers of Artificial Intelligence Completed Research. In *Americas Conference on Information Systems*.
- Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77-101. <https://doi.org/10.1191/1478088706qp063oa>
- Bryman_StructuredInterviewing_197-219 (1)*. (n.d.).
- Cambridge Dictionary*. (2025). <https://dictionary.cambridge.org/>
- Chen, C., Feng, S., Sharma, A., & Tan, C. (2022). Machine Explanations and Human Understanding. *Transactions on Machine Learning Research*, 2023. <https://doi.org/10.1145/3593013.3593970>
- Chopra, S. S., Senadheera, S. S., Dissanayake, P. D., Withana, P. A., Chib, R., Rhee, J. H., & Ok, Y. S. (2024). Navigating the Challenges of Environmental, Social, and Governance (ESG) Reporting: The Path to Broader Sustainable Development. In *Sustainability (Switzerland)* (Vol. 16, Issue 2). Multidisciplinary Digital Publishing Institute (MDPI). <https://doi.org/10.3390/su16020606>
- Data, Privacy, and Security for Microsoft 365 Copilot | Microsoft Learn*. (n.d.). Retrieved November 9, 2025, from https://learn.microsoft.com/en-us/copilot/microsoft-365/microsoft-365-copilot-privacy?utm_source=chatgpt.com
- Davis, M. C., Leach, D. J., & Clegg, C. W. (2020). Breaking Out of Open-Plan: Extending Social Interference Theory Through an Evaluation of Contemporary Offices. *Environment and Behavior*, 52(9), 945-978. <https://doi.org/10.1177/0013916519878211>
- De Boeck, M., & Vaes, K. (2024). Validation of a Design Method for Human Augmentation. *Augmented Human Research*, 9(1). <https://doi.org/10.1007/s41133-024-00072-w>
- Designerly Ways of Knowing*. (n.d.).
- Doellken, M., Nelius, T., & Matthiesen, S. (2024). Evaluation of the ACAP analysis method for process-based validation of textual and graphical design methods. *Journal of Engineering Design*, 35(4), 460-482. <https://doi.org/10.1080/09544828.2024.2320018>

- Dorst, K. (2011). The core of “design thinking” and its application. *Design Studies*, 32(6), 521–532. <https://doi.org/10.1016/j.destud.2011.07.006>
- Dubois, A., & Gadde, L. E. (2002a). Systematic combining: an abductive approach to case research. *Journal of Business Research*, 55(7), 553–560. [https://doi.org/10.1016/S0148-2963\(00\)00195-8](https://doi.org/10.1016/S0148-2963(00)00195-8)
- Dubois, A., & Gadde, L.-E. (2002b). Dubois). *Journal of Business Research*, 55, 553–560.
- Dutch Data Center. (2025). *Sustainability and Energy - Dutch Data Center Association*. <https://www.dutchdatacenters.nl/en/data-centers/sustaina/>
- Ethics of Artificial Intelligence / UNESCO*. (n.d.). Retrieved October 22, 2025, from https://www.unesco.org/en/artificial-intelligence/recommendation-ethics?utm_source=chatgpt.com
- Felgueiras, F., Mourão, Z., Moreira, A., & Gabriel, M. F. (2023). Indoor environmental quality in offices and risk of health and productivity complaints at work: A literature review. *Journal of Hazardous Materials Advances*, 10. <https://doi.org/10.1016/j.hazadv.2023.100314>
- Gao, Y., Xiong, Y., Gao, X., Jia, K., Pan, J., Bi, Y., Dai, Y., Sun, J., Wang, M., & Wang, H. (2023). *Retrieval-Augmented Generation for Large Language Models: A Survey*. <https://github.com/Tongji-KGLLM/>
- Grant Sanderson. (2023). *Attention and transformers*. [Video recording].
- Grzegorzewska, M., & Kirschke, P. (2021a). The impact of certification systems for architectural solutions in green office buildings in the perspective of occupant well-being. *Buildings*, 11(12). <https://doi.org/10.3390/buildings11120659>
- Grzegorzewska, M., & Kirschke, P. (2021b). The impact of certification systems for architectural solutions in green office buildings in the perspective of occupant well-being. *Buildings*, 11(12). <https://doi.org/10.3390/buildings11120659>
- Guo, T., Chen, X., Wang, Y., Chang, R., Pei, S., Chawla, N. V., Wiest, O., & Zhang, X. (2024). *Large Language Model based Multi-Agents: A Survey of Progress and Challenges*. <http://arxiv.org/abs/2402.01680>
- Gustavsson, L., Joelsson, A., & Sathre, R. (2010). Life cycle primary energy use and carbon emission of an eight-storey wood-framed apartment building. *Energy and Buildings*, 42(2), 230–242. <https://doi.org/10.1016/J.ENBUILD.2009.08.018>
- Huang, Y., & Ai, R. (2023). *Levels of AI Agents: from Rules to Large Language Models*.

IEA. (2025). *World Energy Outlook Special Report Energy and AI* IEA. CC BY 4.0. www.iea.org/terms

Ifediora, C. O., Idhlolo, I. O., & Halim, H. C. (2024). Environmental, Social and Governance (ESG) Factors as Key Factors in Real Estate Investment Decision: Mixed Use Properties. *International Journal of Development and Economic Sustainability*, 12(4), 16–28. <https://doi.org/10.37745/ijdes.13/vol12n41628>

Jebril, I., Al-Zaqeba, M. A. A., Al-Khawaja, H. A., Al Obaidy, A. L. A., & Marashdah, O. S. (2024). Enhancing estate governance using blockchain technology through risk management in estate governance of business sustainability. *International Journal of Data and Network Science*, 8(3), 1649–1658. <https://doi.org/10.5267/j.ijdns.2024.3.002>

Khan, A. A., Badshah, S., Liang, P., Waseem, M., Khan, B., Ahmad, A., Fahmideh, M., Niazi, M., & Akbar, M. A. (2021). Ethics of AI: A Systematic Literature Review of Principles and Challenges. *ACM International Conference Proceeding Series*, 383–392. <https://doi.org/10.1145/3530019.3531329>

Kolltveit, B. J., & Grønhaug, K. (2004). The importance of the early phase: The case of construction and building projects. *International Journal of Project Management*, 22(7), 545–551. <https://doi.org/10.1016/j.ijproman.2004.03.002>

Krizmanić, V., & Milovanović, A. (2025). From Space to Satisfaction: Investigating Architectural Interior Determinants of Quality Work Environments. *Buildings*, 15(13). <https://doi.org/10.3390/buildings15132256>

Kumar, D., Maurya, K. K., Mandal, S. K., Mir, B. A., Nurdiawati, A., & Al-Ghamdi, S. G. (2025). Life Cycle Assessment in the Early Design Phase of Buildings: Strategiess, and Future Directions. *Buildings*, 15(10), 1612. <https://doi.org/10.3390/BUILDINGS15101612/S1>

Larisa Olteanu, A., & Elena Ionaşcu, A. (2023). *The Importance of ESG in the Real Estate Investments*.

Lee, C. L., & Liang, J. (2024). The effect of carbon regulation initiatives on corporate ESG performance in real estate sector: International evidence. *Journal of Cleaner Production*, 453. <https://doi.org/10.1016/j.jclepro.2024.142188>

Liang, C. J., Le, T. H., Ham, Y., Mantha, B. R. K., Cheng, M. H., & Lin, J. J. (2024). Ethics of artificial intelligence and robotics in the architecture, engineering, and construction industry. In *Automation in Construction* (Vol. 162). Elsevier B.V. <https://doi.org/10.1016/j.autcon.2024.105369>

Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient Estimation of Word Representations in Vector Space. *1st International Conference on Learning Representations, ICLR 2013 - Workshop Track Proceedings*. <https://arxiv.org/pdf/1301.3781>

Miller, N., & Pogue, D. (2018). Sustainable real estate and corporate responsibility. *Routledge Handbook of Sustainable Real Estate*, 19–36. <https://doi.org/10.1201/9781315622750-2/SUSTAINABLE-REAL-ESTATE-CORPORATE-RESPONSIBILITY-NORM-MILLER-DAVE-POGUE>

Mohan, S., & Vijayalaxmi, J. (2024). Embodied and Operational Carbon in Buildings: Strategies to Decarbonize. *SpringerBriefs in Architectural Design and Technology, Part F4015*, 1–135. <https://doi.org/10.1007/978-981-97-7187-5/COVER>

Motalebi, M., Heffernan, E., McCarthy, T., Marzban, S., & Rashidi, A. (2025a). Sustainability and Stakeholder Engagement in Building Information Modelling-Enabled Construction: A Review of Critical Success Factors in Design and Planning Phases. *Sustainability 2025, Vol. 17, Page 1086, 17(3)*, 1086. <https://doi.org/10.3390/SU17031086>

Motalebi, M., Heffernan, E., McCarthy, T., Marzban, S., & Rashidi, A. (2025b). Sustainability and Stakeholder Engagement in Building Information Modelling-Enabled Construction: A Review of Critical Success Factors in Design and Planning Phases. In *Sustainability (Switzerland)* (Vol. 17, Issue 3). Multidisciplinary Digital Publishing Institute (MDPI). <https://doi.org/10.3390/su17031086>

NKK Gamage, A. (2025). Research Design, Philosophy, and Quantitative Approaches in Scientific Research Methodology. *Scholars Journal of Engineering and Technology Abbreviated Key Title: Sch J Eng Tech*, 13(2), 91–103. <https://doi.org/10.36347/sjet.2025.v13i02.004>

NTA 8800_2024 nl (2) (1). (n.d.).

Palomar-Torres, A., Rey-Hernández, J. M., Rey-Hernández, A., & Rey-Martínez, F. J. (2025). Decarbonizing Near-Zero-Energy Buildings to Zero-Emission Buildings: A Holistic Life Cycle Approach to Minimize Embodied and Operational Emissions Through Circular Economy Strategies. *Applied Sciences (Switzerland)*, 15(5). <https://doi.org/10.3390/app15052670>

Panesar, A. (2020). Machine learning and AI for healthcare: Big data for improved health outcomes. In *Machine Learning and AI for Healthcare: Big Data for Improved Health Outcomes*. Apress Media LLC. <https://doi.org/10.1007/978-1-4842-6537-6>

Purcell, A. T., & Gero, J. S. (1998). Drawings and the design process: A review of protocol studies in design and other disciplines and related research in cognitive psychology. *Design Studies*, 19(4), 389–430. [https://doi.org/10.1016/S0142-694X\(98\)00015-5](https://doi.org/10.1016/S0142-694X(98)00015-5)

Riratanaphong, C., & Pewklieng, N. (2025). ESG as a Sustainability Approach in Corporate Real Estate Management. *Journal of Architectural/Planning Research and Studies*, 22(1). <https://doi.org/10.56261/jars.v22.270055>

Sartori, I. (n.d.). *Net zero energy buildings: A consistent definition framework*. Retrieved November 11, 2025, from

Stappers, P. J. , & G. E. (2017). *Research through Design | The Encyclopedia of Human-Computer Interaction, 2nd Ed.* <https://www.interaction-design.org/literature/book/the-encyclopedia-of-human-computer-interaction-2nd-ed/research-through-design>

Strauss, A., & Corbin, J. (1990). Introduction. *Basics of Qualitative Research 2nd Edition*, 3-14. <https://doi.org/10.4135/9781452230153>

The MacLeamy curve - Daniel Davis. (n.d.). Retrieved December 1, 2025, from <https://www.danieldavis.com/macleamy/>

Tirelli, D., & Besana, D. (2023). Moving toward Net Zero Carbon Buildings to Face Global Warming: A Narrative Review. In *Buildings* (Vol. 13, Issue 3). MDPI. <https://doi.org/10.3390/buildings13030684>

Vaswani, A., Brain, G., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). *Attention Is All You Need*. 1. <https://arxiv.org/pdf/1706.03762>

Wang, Z., Zhang, T., Ren, X., & Shi, Y. (2024). AI adoption rate and corporate green innovation efficiency: Evidence from Chinese energy companies. *Energy Economics*, 132. <https://doi.org/10.1016/j.eneco.2024.107499>

Xi, Z., Chen, W., Guo, X., He, W., Ding, Y., Hong, B., Zhang, M., Wang, J., Jin, S., Zhou, E., Zheng, R., Fan, X., Wang, X., Xiong, L., Zhou, Y., Wang, W., Jiang, C., Zou, Y., Liu, X., ... Gui, T. (2023). *The Rise and Potential of Large Language Model Based Agents: A Survey*. <http://arxiv.org/abs/2309.07864>

Zaimoglu, A. K., Pratt, L., & Fisher, B. (2023). Epistemological role of human decision making in data-informed decision-making. *Frontiers in Communication*, 8, 1250301. <https://doi.org/10.3389/FCOMM.2023.1250301/BIBTEX>

Zimmerman, J., Forlizzi, J., & Evenson, S. (2007). Research through design as a method for interaction design research in HCI. *Conference on Human Factors in Computing Systems - Proceedings*, 493-502. https://doi.org/10.1145/1240624.1240704/SUPPL_FILE/1240704.MP4

14. Appendix A: Overview Interviewees

This appendix provides an overview of everyone who was interviewed or participated in the expert panel.

Interview	Role	Company
1	Senior project developer Building Physics	DGMR
2	Expert sustainable buildings and Building Physics	DGMR
3	Advisor in Sustainability and Building Physics	LBPSight
4	Architect Director / BREEAM expert in-use	BroekBakema
5	Director projects	EDGE Technologies
6	Senior Development Manager	EDGE Technologies
7	Development manager	EDGE Technologies
8	Senior Development Manager	EDGE Technologies
9	Development Manager	EDGE Technologies
10	Director Projects	EDGE Technologies
11	Director Concepts & Acquisitions	EDGE Technologies
12	WELL AP + Indoor climate specialist	BBA Binnenmilieu
15	Installation Advisory	EDGE Technologies
13	AI expert	NextSense
14	AI expert	Struck

15. Appendix B: Interview protocol

The interview was conducted in Dutch, as the interviewees are all Dutch speakers. The interview questions were divided according to the schedule in the previous appendix. This means that some questions were only asked to 1 to 12 and some only to the AI experts. The questions were shared in advance so that the interviewee could prepare.

Group 1: Professionals in corporate real estate development (interviews 1 to 13)

Achtergrond & context (3-4 min)

- Wat is je rol en in welk type projecten werk je?
- Met welke duurzaamheids- en rapportagekaders werken jullie, specifiek gericht op de Environmental (E) component? (Bijv. EU Taxonomy, CSRD, BREEAM, WELL,)

Vroege ontwerpfase: informatie & besluitvorming (8-10 min)

- Welke SO/VO-beslissingen hebben volgens jou de grootste invloed op de milieuprestaties?
- Welke documenten zijn in deze fase leidend (PvE, normteksten, berekeningen, rekenrapporten, SharePoint, BIM-modellen)?
- Wat wil je bereiken met deze documenten?
- Hoe worden deze documenten beheerd en gedeeld?

De Bouwmeester als sensing rol: spanningen & misalignment (10-12 min)

- Kun je een recent spanningsveld beschrijven waarin milieu-afwegingen centraal stonden?
- Wat wilden jullie bereiken (wenselijke situatie)?
- Hoe is het uiteindelijk gegaan (werkelijke situatie)?
- Wat zie jij als de grootste kloof tussen wens en werkelijkheid?
- Welke informatie of afstemming miste je in dat proces?
- Hoe is de keuze vastgelegd (bijv. notulen, PvE-wijziging, bronverwijzing)?
- Waar ontstaan volgens jou de meeste afstemmingsproblemen: tussen ontwikkelaar, adviseurs, gemeente, belegger?

De Bouwmeester als steering rol: informatie ordenen (8-10 min)

- Welke tools en workflows gebruik je nu (BIM, Revit, LCA, Excel, dashboards)?
- Wat zie jij als een betrouwbare bron (norm-ID, paragraaf, versie/datum, link)?
- Wat maakt een bron overtuigend in een besluit?
- Hoe gaat versiebeheer en traceerbaarheid?
- Als twee bronnen conflicteren, wie bepaalt welke leidend is?

De Bouwmeester als support rol: transparantie (5-6 min)

- Wat zou voor jou waardevolle output zijn in een verrijkt PvE (PvE+)?
- (Bijv. spanningsveldentabel, PvE-diff, risicokaart, aannamesregister.)
- Welke velden moeten daarin altijd voorkomen?
- (Beschrijving, bewijs/bron, aannames, opties, eigenaar, deadline, link.)

Interviewprotocol - Groep 2: AI-experts (interview 14 + 15)

1) Achtergrond & rol (kort, 2-3 min)

- Wat is je ervaring met AI-modellen, LLM-architecturen, of agent workflows in ontwerp- of datadomeinen?

2) Technische werking van LLM-agents (8-10 min)

- Hoe interpreteert een LLM informatie (embedding spaces, attentie, contextvensters)?
- Wat zijn bekende valkuilen: hallucinations, probabilistische fouten, contextverlies, bias?
- Hoe beïnvloeden prompt-design en workflow-structuur de betrouwbaarheid?

3) RAG, databronnen en technische vereisten (8-10 min)

- Wat is essentieel voor een goede RAG-implementatie in bouwprojecten?
- Welke type documenten zijn goed/niet goed te verwerken?
- Hoe ga je om met versiebeheer, documentstructuur, inconsistenties en metadata?
- Welke infrastructuur is minimaal nodig om project-specifieke analyses betrouwbaar te draaien?

4) Rollen van de Digital Building Master(8-10 min)

- In hoeverre zijn de drie conceptuele rollen (sensing, steering, support) technisch realiseerbaar?
- Welke functies zijn risicovol of moeilijk zonder aanvullende controlelagen?
- Wat moet altijd door mensen worden beoordeeld en wat kan een agent veilig structureren?
- Onder welke voorwaarden is een “reflective support” AI-agent verantwoord inzetbaar?

5) Grenzen, risico's en governance (5-6 min)

- Welke risico's zie jij voor AI-gebruik in milieubesluitvorming (traceerbaarheid, modeldrift, interpretatieproblemen)?
- Hoe verhoudt dit zich tot regelgeving zoals de EU AI Act?
- Wat zijn randvoorwaarden om een agent veilig in te zetten binnen een ontwerpsteam?

16. Appendix C: Key Insights interviewees

This appendix summarises the key insights from the interviews. The table shows how respondents interpret the role of AI in early design practice and what tensions they experience in this regard. The insights help to recognise patterns in the way professionals deal with environmental considerations during the design process.

	Key insight of Notable findings	Unexpected insight
1	The interviewee feels that many tensions stem from the lack of unambiguous definitions of sustainability terms, such as Paris Proof. I find it interesting that the AI is seen here as a kind of 'dictionary' that can help streamline discussions and avoid misunderstandings.	The interviewee outlined a clear gap in adoption: experienced professionals with established routines are reluctant, while the younger generation will naturally incorporate AI into their way of working. I found this a striking observation, as it shows that the adoption of AI is not only technically but also generatively determined.
2	The interviewee believes that tensions in projects are inevitable and should be resolved by the design team. I find it striking that he sees AI's role as limited to generating research questions and scenarios, especially in support of novice development managers	The interviewee expresses concerns about the sustainability (energy and water consumption of data centres) and reliability of AI, especially for juniors. I found it interesting that he simultaneously acknowledges that AI is already saving a lot of time in administrative tasks.
3	The interviewee sees the AI agent as a critical sparring partner who helps inexperienced professionals understand sustainability requirements and visualise areas of concern in the design phase. I find this interesting because it gives the AI not the role of a decision-maker, but rather that of a facilitator in the learning process.	The interviewee expresses a pronounced fear that AI will lead to standardisation, turning buildings into 'uniformity sausages'. I find this striking, as it shows that tension and over-standardisation are seen as necessary for true sustainability, and that AI should instead stimulate creativity rather than restrict it.
4	The interviewee sees the AI as a kind of "extra conscience" that helps make choices from the early design phase traceable and hold ambitions. I find this an apt image, as it shifts the role of AI from advisor to guardian of consistency and legitimacy.	The interviewee also sees AI as a possible "definition list" to overcome language tensions around sustainability, while adoption varies greatly by generation. I find this striking, as it highlights that the value of AI is not only in content output, but also in removing misunderstandings and in the extent to which users are willing to trust the system.
5	The interviewee sees AI as a means of providing corporate justification: it should show how sustainable choices lead to higher costs but also to potential returns, for example through a choice menu or sliding scale. I find this interesting because AI thereby provides not only technical but also financial legitimacy.	The interviewee stresses that AI should also conduct market research on the ESG priorities of tenants and investors. I find this striking because it shifts AI's role from internal analysis to external market orientation, which can be crucial for rentability and returns.
6	The interviewee emphasises that AI should mainly save time by testing designs against standards and certifications at lightning speed, with instant visual output and a review of missing inputs. I find this interesting because AI is almost positioned here as an accelerator of compliance, which can greatly reduce uncertainty in the early phase.	The interviewee sees circularity mainly as a logistical problem: the availability, timing and price of materials from donor buildings determine whether reuse is feasible. I find this striking, as it shows that circular ambitions are constrained less by design choices than by external chain processes.
7	The interviewee sees AI mainly as a decision map that visually juxtaposes scenarios and their consequences. I find this striking because the emphasis is on control: the user wants quicker insight but ultimately decides for himself.	The interviewee notes that sustainability ambitions are often discounted due to financial and technical constraints, partly due to the conservative attitude of investors. I find this interesting because it shows how strongly investors determine the playing field for innovation in sustainability.

8	The interviewee highlights that linguistic tension around terms like Paris Proof shows the need for recognition and certification. I find this interesting because AI is positioned here as an overarching knowledge system that can organise and define standards.	The interviewee expresses distrust towards AI, fuelled by personal negative experiences and ethical sensitivities, right down to the naming of the agent. I find this striking because it shows that trust in AI is not only technical, but also emotionally and culturally determined.
9	The interviewee sees the composition of the design team as decisive for sustainability; personal ambition, he says, outweighs actual design choices. I find this striking because it places the role of human commitment above technical interventions.	The interviewee points to great knowledge loss in team changes and therefore emphasises the need for AI with learning capabilities. I find this interesting because he simultaneously expresses scepticism about the neutrality of AI advice, which illustrates well the ambivalent attitude towards technology.
10	The interviewee emphasises that the biggest tension is in finding a contractor willing to take risk with a still incomplete design. I find it interesting that he sees AI's role mainly in providing insight into scenarios but clearly draws the line at making investment decisions himself.	The interviewee argues that financial feasibility determines choices more strongly than expected, emphasising the "kilowatt reduction per euro" balance. I find it striking that he simultaneously emphasises the contradiction between Paris Proof via solar panels and the high embodied carbon this causes. His desire for an AI that automatically fills HBAs and visually displays results shows the importance of speed and understandability in decision-making Risk sharing, financial feasibility, embodied vs operational carbon, visual output.
11	The interviewee emphasises that feasibility is the all-important factor. I find it interesting that he sees AI mainly as a system to automatically process data in the HBA and visually show the balance between costs and returns but clearly sets the boundary at making investment decisions himself.	The interviewee points to tensions between embodied and operational carbon, for example due to solar panels or the height of buildings. I find this striking because it makes sustainability not only more complex but also more expensive, which further sharpens the discussion on priorities.
12	The interviewee feels that tensions belong primarily to the design team. I find it interesting that he limits the role of AI to supporting juniors with research questions and scenario thinking, so more as a learning system than as a decision-maker.	The interviewee expresses concerns about sustainability and reliability of AI, but at the same time acknowledges that deployment in administrative tasks saves time. I find this striking because it shows the ambiguity of AI well: useful in practice, but controversial in trust and impact.
13	The installation consultant demonstrated that early-stage PV calculations rely on a sequence of tacit assumptions that are rarely documented but strongly influence final decision making. It helped with the calculation of the solar panels.	Despite strong technical expertise, the consultant noted that "the biggest mistake is thinking we're calculating; we're actually estimating." This distinction revealed that even experts operate within substantial uncertainty in the early phase, reinforcing that the agent should not aim for precision but for clarity about uncertainty.
14	The AI expert sees that LLM decision making is path-dependent: small differences in prompt structure or document retrieval radically shift the logic the model follows. This confirmed the need for layered prompt workflows rather than single-shot answers and validated the decision to align the agent with the five-layer system (interaction, process, integration, service, data).	He highlighted that "models hallucinate most when humans hide their own uncertainty." This was unexpected: it suggests that poor documentation in ESG-related design phases does not only hinder humans, but also actively induces hallucinations in AI. The agent therefore needs explicit uncertainty markers. .
15	The central insight was that a digital agent becomes useful only when it mirrors the structure of the human workflow: a chain of steps, each with its own checks and contextual grounding. His explanation of agent workflows validated the choice to treat the AI agent as a decision making partner rather than a calculator.	The human teams often "lose the decision making chain long before they lose the data," meaning that misalignment arises not from missing information but from missing structure. This shifted the focus: the agent's value lies more in preserving the decision making path than in providing additional content, an insight that strengthened the conceptualisation of the roles sensing, steering and support.

17. Appendix D: overview of code groups Atlas.TI

This appendix provides an overview of the code groups used in Atlas.ti to systematically analyse the interview data. The table shows how the themes are constructed from both inductive categories from the fieldwork and deductive categories from the literature. This overview provides insight into how the analysis is structured and how the themes surrounding the E-dimension, early design practice, and the role of the AI agent were developed. The diagram enhances the transparency of the research process and supports the interpretation of the findings in the main text.

Main theme	Code group	Provenance	Brief explanation
ESG- ENVIRONMENTAL IN CORPORATE REAL ESTATE	Waardecreatie door duurzaamheid (18)	Literatuur + data	Sustainability as a system for value development and market positioning
	Certificering strategie & credits (20)	Deductive	Role of certifications as a steering or legitimising system
		Deductive	The sustainability criteria are translated into a Programme of Requirements.
	Design brief & certificering (29)		
	Commerciële strategie & markt (10)	Inductive	Area of tension between sustainability and corporate feasibility
EARLY DESIGN PHASE	Marktintelligentie & stakeholder profilering (5)	Inductive	Mapping ESG priorities of stakeholders
	Ontwerpfilosofie (30)	Inductive	Differences in design approaches
	Edge development werkwijze (20)	Inductive	EDGE's specific process approach to integrating ESG ambitions
AI IN PRACTICE	Besluitvorming & stakeholders (48)	Literatuur & data	Fragmentation in decision-making and role of different actors in early design phase
	Houding t.o.v. AI (18)	Inductive	Expectations and ethical concerns of AI
	AI-rol in onderzoek (12)	Inductive	Interviewees' reflections on how AI research and decision-making can influence
THE BUILDING DIGITAL MASTER (AGENT)	Sensing (61)	Deductive	Role of AI in identifying tensions within E goals
	Steering (29)	Deductive	AI as a coordinator and support for collaboration in design and decision-making processes
	Support (21)	Deductive	AI as a source of transparency and accountability towards external parties
	AI-system requirements (34)	Inductive	Practical conditions for deployment of AI (transparency, reliability, integration into workflow)

18. Appendix E: Example sub codes by group

This appendix shows how the main themes from the analysis have been elaborated into subcodes. The excerpts from the interviews illustrate how respondents speak concretely about environmentally oriented considerations in design practice. The examples show the variety of arguments and demonstrate how values related to sustainability are expressed in daily decision-making. This level of detail supports the interpretation of the main findings in the thesis and makes the analytical process transparent.

E1 ESG - Environmental in Dutch corporate real estate

Value creation through sustainability

Subcodes	Illustrative quote from interview	Interview
Role of frontrunners in supra-legal sustainability (1)	<i>"Ik wil het eerste gebouw maken dat dat doet, omdat ik duurzaamheid erg belangrijk vind. Dus als je voorop wilt lopen of een duurzamer gebouw wilt maken dan anderen, kom je dingen tegen."</i>	3
Material selection & quality (4)	<i>'Denken wij dat we met die extra investering in bijvoorbeeld een houtconstructie, we ook een hogere huur kunnen vragen bij potentiële huurders.'</i>	5
Socio-spatial value of sustainability (2)	<i>"De mens staat centraal in het gebouw. Het gaat niet om het gebouw zelf, maar het gaat erom dat mensen daar op een gezonde manier een plek krijgen waarin ze in de productiviteit worden aangemoedigd, zich kunnen ontwikkelen. Dus dat vraagt om gezonde gebouwen. Maar ook zijn het plekken voor ontmoetingen, vormen van communities. En daar zit een sociaal aspect aan"</i>	2
Daylight (9)	<i>Toegang tot daglicht hebt, veel daglicht, maar ook dat je die houten vloeren weer ziet. Dus het is een afweging tussen je CO2 footprint. Maar daarnaast ook kijken van ja, wat doet het met de beleving van je toekomstige werknemers of gebruikers die in dat gebouw gaan</i>	5
Value creation beyond certification: healthy buildings (1)	<i>k denk dat we moeten oppassen ook, dat we niet alleen kijken naar de harde feiten van vierkante meters en labels en die kwaliteiten, maar dat je ook kijkt naar de Bijvoorbeeld, dat zie je bij ABN. En dat is dan misschien wel heel veel breder.</i>	4

Certification strategy & credits

Subcodes	Illustrative quote from interview	Interview
BREEAM guidelines codes (17)	<i>"Dus bijvoorbeeld die energierecredit, dat is ENE01 bijvoorbeeld. Die is cruciaal om te behalen, want dat heeft te maken met je bankberekening. Als je die niet haalt, dat is een verplichte credit, als je daar een bepaald niveau niet haalt, kan je überhaupt niet naar bijvoorbeeld de BM Outstanding of BM Excellent."</i>	5
Embodied carbon overweging	<i>'J we willen een gebouw maken met een lage embodied carbon footprint. We willen het liefst rond of onder de 250 kilogram per vierkante meter zitten. In het begin hebben we niet zo sterk die targets</i>	5

gezet. Het moet 205 zijn, maar het moet een lage CO2 footprint hebben.”

Paris proof unclear (5)	<i>‘Mensen weten vaak niet zo goed als ze paris proof zeggen wat ze nou precies willen en of dat haalbaar is voor hun gebouw. Dus wat wij ook bemerken is, het is vaak ook soms gebrek aan kennis’</i>	3
Selective use of certification credits (6)	<i>‘Je kijkt naar specifieke credits, dus je weet vanuit ervaring en adviseurs weten gewoon, oké, bepaalde credits zijn tricky om te halen.’</i>	5

Commercial strategy & market (10) + Market intelligence & stakeholder profiling (5)

Subcodes	Illustrative quote from interview	Interview
Business case as a driver (7)	<i>‘Er moet wel een haalbare business case hebben. Anders heb je überhaupt geen project. Dus dan is de vraag van, oké, ga je dan je ambitieniveau verlagen?’</i>	6
Commercial interest in design choices (5)	<i>‘Wij zijn op zich best nog wel dat we innovatief zijn en nieuwe technieken willen toepassen. Maar een belegger is oer-conservatief. Dus daar in die innovatie zit ook vaak nog wel een spanning. We kunnen het niet, want een belegger vindt het allemaal veel te spannend en kan niet mee in deze innovatie’</i>	11
External regulations & contextual factors (6)	<i>‘Door de jaren heen verandert er in de markt heel veel en verandert misschien ook wel de ESG doelen en doelstellingen, waardoor je project achter begint te lopen.’</i>	12

E2 Early design phase

Design philosophy

Subcodes	Illustrative quote from interview	Interview
Architectural ambition (4)	<i>‘Het moet een lage carbon footprint hebben. Het moet uiteindelijk een energiezuinig gebouw worden of Paris-proof gebouw. En het moet een gebouw worden wat ook een soort van architectuur een soort van, ik noem het een statement is, maar in ieder geval Ja, dat we een bijzonder gebouw maken.’</i>	6
Human-centred design principles (12)	<i>‘Maar daarnaast ook kijken van ja, wat doet het met de beleving van je toekomstige werknemers of gebruikers die in dat gebouw gaan.’</i>	5
Design choice wood (7)	<i>‘En dan zit die duurzaamheidsambitie daar tussen. Dus in hoeverre gaat een hogere ambitie zorgen voor hogere inkomsten. En compenseer je daarmee de hogere kosten. En dat begint eigenlijk helemaal in het begin al met de eerste keuze van materiaal’</i>	7
Constructive material choices (6)	<i>‘En het belangrijkste daarbij is, denk ik, gewoon welke bouwmethodiek je gaat gebruiken en hoe je hoofdopzet van je gebouw gaat zijn, zeg maar. Dus ofwel ga je met hout bouwen of ga je wat traditioneler met beton en staal bouwen, zeg maar’</i>	7
Installations as a determining factor in design (5)	<i>‘Natuurlijk kun je een vorm van een gebouw meegeven. Dan kunnen ze er al wat uithouden, dat er voldoende dak beschikbaar is ofzo, weet ik veel. Maar het hangt ook vaak samen met het installatieconcept, of je iets wel of niet haalt.’</i>	6

Decision-making & stakeholders

Subcodes	Illustrative quote from interview	Interview
Dominance of budget and feasibility in decision-making (9)	<i>Ja, dat is maar net wie je stakeholders zijn. Wie je al aan boord hebt. Bij Royale werk ik samen met de eigenaar. Dus eigenlijk bepaalt hij het aan zijn portemonnee</i>	6
Guiding role of client/customer (8)	<i>a, en dan is het echt zoeken van wat wil de opdrachtgever. Soms wil de opdrachtgever dat zelf ook niet. Ja. Wat willen ze nou precies voor duurzaamheid? Maar wat zijn de eisen? Waar moeten we op sturen? En dat zijn eigenlijk best vermoeiende projecten, want dan is het heel lang zoeken om er überhaupt achter te komen wat ze willen</i>	3
Stakeholder role and cooperation (9)	<i>Maar in een ander project heb ik gezeten met een duurzaamheidsexpert die ook heel goed is, maar van zichzelf veel ingetogener is en nooit het gesprek zou overnemen van hé, we moeten ook hieraan denken. Dus dan wordt het automatisch al wat minder de boventoon.</i>	5
Complexity, uncertainty and conflicting interests (16)	<i>Ja kijk, uiteindelijk ons werk is heel erg mensenwerk. Het is allemaal best technisch en inhoudelijk. Zeker commerciële gasten zien dat en denken dat wij enorme techniek nerds zijn. Maar het gaat heel veel over mensen en over mensen meekrijgen in je ideeën en over hoe je bepaalde dingen met elkaar wil regenen.</i>	7

Design brief & certification (29) + EDGE development methodology

Subcodes	Illustrative quote from interview	Interview
Edge design brief (17)	<i>Dus ja, dat is een allesomvattend document, maar daarnaast is het ook een heel breed document. Ja, we geven onze design brief altijd mee aan... De architect en alle technische adviseurs die werken aan het project. Het is wel van belang om dat document zo snel mogelijk project specifiek te maken.</i>	1
Certification and sustainability as a design system (12)	<i>'Dat heeft heel veel effect. Het heeft ook prijskonsequenties. Het heeft ook layouttechnische consequenties. Het heeft ook te maken met de hoogte die je ingaat. Dus die is best wel maatgevend. En ook wel je installatieconcept kun je bijvoorbeeld wel of geen WKO toepassen. Het heeft een enorme impact op wat voor score je bijvoorbeeld ook kunt halen op Brium.</i>	12
Collaboration within design teams (7)	<i>'Ja. Kijk, wij kunnen een heel hoge duurzaamheidsambitie hebben en bijvoorbeeld met hout willen bouwen. Maar als uiteindelijk een aannemer dat niet kan, of er is maar één leverancier die bijvoorbeeld dat type hout kan leveren, en die bij dat project wil meewerken of dat niet kan, ja, dan heb je wel een clash, zeg maar. Dan heb je heel hoog doelstelling, maar je kan het niet waarmaken</i>	5

E3 AI in practice & Attitude towards AI & AI's role in research

Subcodes	Illustrative quote from interview	Interview
Doubts about AI (13)	<i>Dat is dus de vraag. Hoe de hel gaat hij dat doen?</i>	4
Potential of AI (13)	<i>Alleen maar zeggen dat daar dus een ontwerpuitdaging ligt. Het wil niet zeggen dat het per definitie waar is dat Als je deze slider naar die kant zet, dat die slider dan naar die kant gaat.</i>	12
Pain and tensions in the design process (4)	<i>'Ik krijg gewoon heel goede antwoorden, maar het wordt allemaal heel erg hetzelfde.</i>	3

AI system requirements

Subcodes	Illustrative quote from interview	Interview
Carbon trade offs	<i>Ik denk dat dat de AI-system überhaupt ook heel erg zou kunnen helpen in het duidelijk maken van als je deze dingen wilt voor je gebouw, wat betekent dat dan? Wat gaat het ook betekenen tijdens het proces?"</i>	8
Costs vs sustainability	<i>"Tijdwinst en betrouwbaarheid. Ik denk dat tijdwinst als dat model op een gegeven moment goed in elkaar zit. Dat je daar ook inderdaad een bibliotheek in hebt van bouwkosten, materialen, MPG."</i>	5
Stakeholder tension	<i>"Je wil wel altijd weten wat de vraag is achter de vraag. [...] Dus dat de scenario's worden uitgewerkt door de agent, maar dat het eindbesluit wordt gemaakt door de mens."</i>	4
Design process	<i>"Ik denk dat je wel echt moet afwegen wanneer het nuttig zeg maar is. En waar het kan helpen en ja zo een beetje de de middenweg vinden."</i>	3

E4 The Digital Building Master(agent)

Sensing

Subcodes	Illustrative quote from interview	Interview
BREEAM guidelines codes &	<i>"Dus bijvoorbeeld die energiecridit, dat is ENE01 bijvoorbeeld. Die is cruciaal om te behalen, want dat heeft te maken met je bankberekening. Als je die niet haalt, dat is een verplichte credit, als je daar een bepaald niveau niet haalt, kan je überhaupt niet naar bijvoorbeeld de BM Outstanding of BM Excellent."</i>	7
Embodied carbon consideration	<i>we willen een gebouw maken met een lage embodied carbon footprint. We willen het liefst rond of onder de 250 kilogram per vierkante meter zitten. In het begin hebben we niet zo sterk die targets gezet. Het moet 205 zijn, maar het moet een lage CO2 footprint hebben."</i>	10
Paris proof unclears	<i>"Mensen weten vaak niet zo goed als ze paris proof zeggen wat ze nou precies willen en of dat haalbaar is voor hun gebouw. Dus wat wij ook bemerken is, het is vaak ook soms gebrek aan kennis"</i>	8
Business case as a driver (costs vs sustainability)	<i>"Er moet wel een haalbare business case hebben. Anders heb je überhaupt geen project. Dus dan is de vraag van, oké, ga je dan je ambitieniveau verlagen?"</i>	

Steering

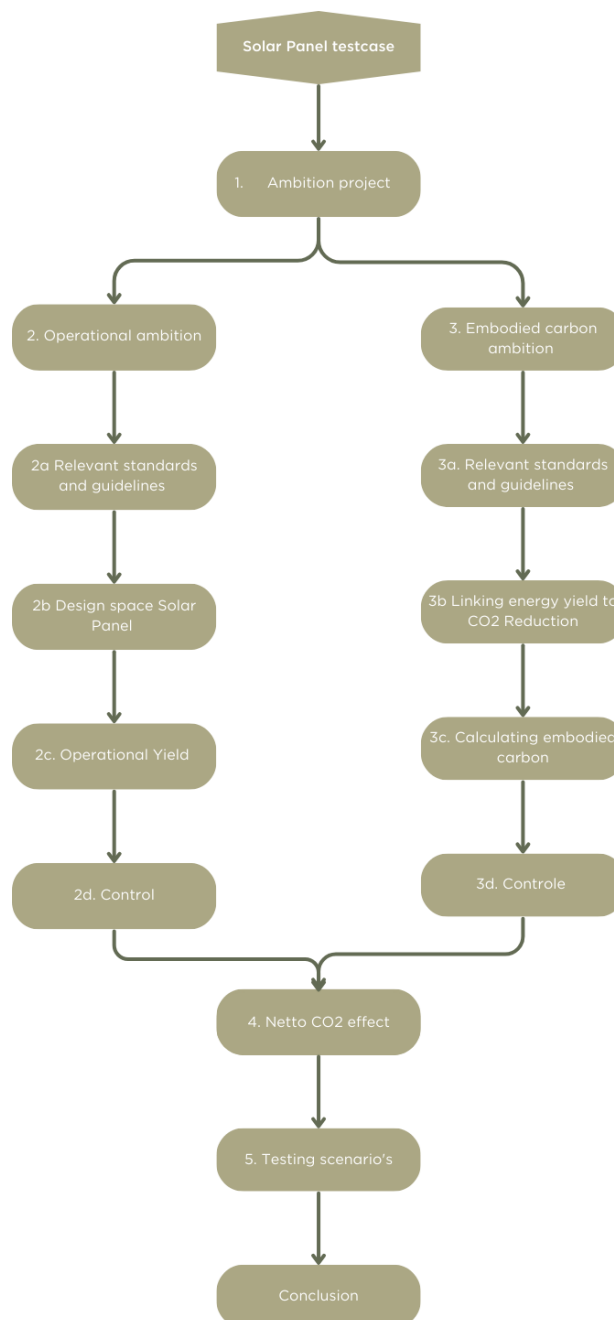
Subcodes	Illustrative quote from interview	Interview
Carbon trade-offs (Materials vs Energy)	<i>"Omdat vaak om een gebouw energiezuiniger te maken, heb je eigenlijk meer materiaal nodig en dan is het meest concrete voorbeeld is PV panelen."</i>	2
Generating research questions	<i>"AI kan de gebruiker ondersteunen in het formuleren van onderzoeksvragen, bijvoorbeeld 'Vraag je ontwerpteam hoe [houtbouw en hoge Well ambitie] verwezenlijkt kunnen worden'."</i>	5
Costs vs sustainability (Additional material for certificate)	<i>"Dus en dan dan gebruik je dus wel extra materiaal om dus die [PV] toe te voegen. Ja. [om Brium Excellent te halen]"</i>	4
Stakeholder tension (Regulation vs ambition)	<i>Maar ja, dat dat dat dat vooral vanuit ontwikkelaars, maar ook vanuit gemeentes die die die schrijven zo'n eis op. En die zeggen, dit mag de C. O twee uitstoot zijn."</i>	7

Support

Subcodes	Illustrative quote from interview	Interview
Support for unclear PVE's	<i>"AI kan de gebruiker ondersteunen door de ingevoerde PVE's te vergelijken met lokale regelgeving en context zodat een verrijkt PVE ontstaat."</i>	2
Traceability data origin	<i>"We willen wel weten waar die informatie vandaan komt, wat daar zit in die AI, dat is belangrijk voor de traceerbaarheid."</i>	7
Knowledge retention & Consistency	<i>De AI agent kan functioneren als 'extra geweten' of 'bewaker van consistentie' en helpen ambities vast te houden, vooral bij veranderingen in het projectteam."</i>	5
Design process tension	<i>"Een logische outcom [is] een besliskaart. Weet je of een of een of hoe hoe zou ik zou eigenlijk willen dat je gewoon gevisualiseerd krijgt van oké, dit is het topic en je hebt deze twee kampen."</i>	2

19. Appendix F: Human Workflow Solar Panel

This appendix shows how the solar panel case was processed step by step in the human workflow. The visualisation shows how energy performance and material effects are weighed within the frameworks of BREEAM-NL and the calculation method of NTA 8800. The workflow provides insight into how operational energy and CO₂ reduction are calculated and how these values guide design choices. This example forms the reference point for comparison with the analysis by the AI agent in the main text. The workflow is based on the following sources: BREEAM-NL. (2020). BREEAM-NL Nieuwbouw en Renovatie 2020: beoordelingsrichtlijn. Dutch Green Building Council. NEN. (2020). NTA 8800:2020 Energieprestatie van gebouwen - Bepalingsmethode. Nederlands Normalisatie-instituut.



Workflow Solar panel calculation

Step-by-step plan solar panels testcase

Step 1. Ambition project

- BREEAM Outstanding
- Net Zero Carbon
- Building function: Office

Step 2. Operational ambition

2a. Relevant standards and guidelines

The operational carbon ambition of the project is defined by the requirement to achieve BREEAM-NL Outstanding, in which the ENE 01 credit is mandatory. ENE 01 evaluates the reduction of primary fossil energy use (BENG 2). To obtain the minimum of 10 ENE 01 points required for an Outstanding score, the project must demonstrate a clear performance improvement compared to the reference building defined in Guidance Note 41.

Energy performance is calculated in accordance with NTA 8800, which provides the methodology for determining both energy demand and primary fossil energy use.

Parameter	Reference value $\leq 1,8$ Als/Ag
Energy requirement (BENG 1)	90 kWh/m ² ·jr
Primary fossil energy consumption (BENG 2)	40 kWh/m ² ·jr

To secure 10 ENE 01 points, the BENG 2 value must be at least 10% lower than the reference threshold:

- Reference BENG 2: 40 kWh/m²·yr
- Target BENG 2: 36 kWh/m²·yr (-10%)

For the project (for example total usable floor area $A_g = 27,151$ m²):

$$E_{\text{ref}} = 40 \times 27,151 = 1,086,040 \text{ kWh/yr}$$

$$E_{\text{target}} = 36 \times 27,151 = 977,436 \text{ kWh/yr}$$

Thus, the minimum operational energy reduction required from PV is:

$$\Delta E_{\text{min}} = 108,604 \text{ kWh/yr} \approx 108.6 \text{ MWh/yr}$$

This target forms the baseline for determining the minimum PV contribution.

2b. Design space for PV integration

Only the roof is considered for PV placement. The roof surface is divided into four usable zones, accounting for obstructions, technical installations, access, and safety clearance. These percentages already incorporate the reduction factor typically required for PV spacing.

Roof	Surface	Area %	Pv area
1	460 m2	50%	230
2	1310 m2	40%	524
3	590m2	50%	295m2
4	1310	30%	393 m2
TOTAL			



To translate usable area to PV capacity, a standard density of 200 Wp/m² is applied:

$$P_{DC,max} = 1,442 \times 200/1000 = 288.4 \text{ kWp}$$

Thus, the maximum technically and spatially feasible PV capacity is 288.4 kWp.

2c. Operational yield (scenario analysis)

Three PV scenarios are evaluated to allow staged ambition levels:

Scenario	Capacity	Number of panels	Net roof area
ENE 1 MINIMUM	105	233	520
INTERMEDIATE VARIANT	200	444	993
MAXIMUM ROOF CAPACITY	288.4	641	1433

Annual PV generation is estimated using PVGIS performance data for Amsterdam (Y_f = 1,037.5 kWh/kWp·yr):

$$E_{PV} = P_{DC} \times Y_f$$

Scenario	Pv capacity	Annual generation E _{PV} (kwh/yr)	Co2 reduction
ENE 1 MINIMUM	105	108,937	40.3
INTERMEDIATE VARIANT	200	207,500	76.8
MAXIMUM ROOF CAPACITY	288.4	299,146	110.7

*Using Dutch grid emission factor: 0.37 kg CO₂/kWh (RVO, 2024).

All scenarios meet spatial feasibility. The minimum scenario satisfies the ENE 01 reduction requirement, while the intermediate and maximum scenarios increase long-term operational CO₂ savings.

The ENE 01 minimum scenario is sufficient to meet the required operational reduction of 108.6 MWh/yr established in Step 2a. The intermediate and maximum scenarios provide increased CO₂ reduction potential but require proportionally more material input, which will be evaluated in Step 3.

This step concludes the operational branch and establishes the quantitative reference used to assess embodied carbon implications.

2d. Control

This step confirms whether:

1. The projected PV system meets the minimum required energy reduction (108.6 MWh/yr).
2. Assumptions used in the design phase match performance outcomes during commissioning and operation.

This validation step closes the operational branch and provides the input to Step 3 (Embodied carbon ambition), in which the same three PV scenarios are evaluated for material-related emissions.

Step 3. Embodied carbon ambition

3a. Relevant standards and system boundaries

The embodied carbon assessment follows the National Environmental Database (NMD), Category 3. Since December 2023, PV systems are assigned a generic emission factor of:

$$EF_{PV} = 620 \text{ kg CO}_2\text{-eq/kWp}$$

This factor includes:

- Module A (Production & installation): raw materials, manufacturing, transport (A1-A5)
- Module C (End-of-life): dismantling, waste processing, transport (C2-C3)
- Module D (Recovery benefits): recycling and reuse already accounted for

System boundary: Cradle-to-Gate with options, consistent with Dutch embodied carbon assessment practice.

The functional lifetime of the PV system is set to 30 years, in alignment with the long-term Net Zero Carbon ambition of the project. This lifetime is carried consistently through Step 4 to ensure comparability.

3b. Link to operational scenarios

The embodied carbon calculation directly corresponds to the three design scenarios established in Step 2c-2d, ensuring that operational energy benefits and material impacts are evaluated on the same system boundaries.

Net roof area	Capacity	Number of panels	
ENE 1 MINIMUM	105	233	520
INTERMEDIATE VARIANT	200	444	993
MAXIMUM ROOF CAPACITY	288.4	641	1433

These configurations serve as the reference frame for material impact comparison.

3c. Embodied carbon calculation

Total embodied carbon is calculated using:

$$E_{emb} = P_{DC} \times EF_{PV}$$

Scenario	Capacity (kwp)	Emissiefactor (kg co ₂ /kwp)	Embodied carbon (kg co ₂)	Tco ₂ embodied carbon
ENE 01 MINIMUM	105	620	65.100	65,1
INTERMEDIATE VARIANT	200	620	124000	124
MAXIMUM ROOF CAPACITY	288.4	620	178808	178.8

This step quantifies the upfront climate impact of implementing each PV scenario.

3d. Control / verification

To ensure methodological consistency:

- Emission factors are sourced from NMD Category 3, compliant with Dutch LCA regulations.
- System configuration (modules, mounting, inverter) falls within the validated scope of the emission factor; therefore, no additional component modelling is required.
- The lifetime assumption of 30 years is explicitly documented for alignment with Net Zero Carbon project ambitions and subsequent payback calculation in Step 4.

Step 4. Net CO₂ effect — integration of operational and embodied performance

This step integrates the operational CO₂ reduction potential (Step 2d) with the embodied carbon impacts of the PV system (Step 3c) in order to determine the net climate effect over the system lifetime. This comparison makes the underlying

trade-off explicit: increasing PV capacity yields higher operational gains but also requires additional material input.

The net CO₂ effect is calculated as:

$$\text{Net CO}_2 \text{ effect (tCO}_2\text{)} = (\text{Operational CO}_2 \text{ savings per year} \times 30 \text{ years}) - \text{Embodied CO}_2$$

Scenario	Pc capacity (kwp)	Operational co ₂ savings (tco ₂ /yr)	Embodied co ₂ (tco ₂)	Co ₂ payback (years)	Lifetime net co ₂ savings (tco ₂)
ENE MINIMUM	105	40.3	65.1	1.6	1.144
INTERMEDIATE	200	76.8	124.0	1.6	2.180
MAXIMUM ROOF CAPACITY	288.4	110.7	178.8	1.6	3.141

Interpretation

All scenarios produce a net positive environmental effect within the 30-year operational horizon. The CO₂ payback time is consistent across scenarios (≈ 1.6 years), meaning that the embodied emissions of the PV system are offset early in the building's operational phase.

- The intermediate scenario presents the most balanced trade-off:
- Substantial lifetime CO₂ reduction,
- Moderate material requirements,
- Stable payback period.

The maximum feasible scenario further increases total CO₂ savings but shows diminishing marginal returns: the embodied carbon grows faster than the operational gain. This scenario may still be strategically valuable where surplus generation, energy autonomy or offsetting capability are prioritised.

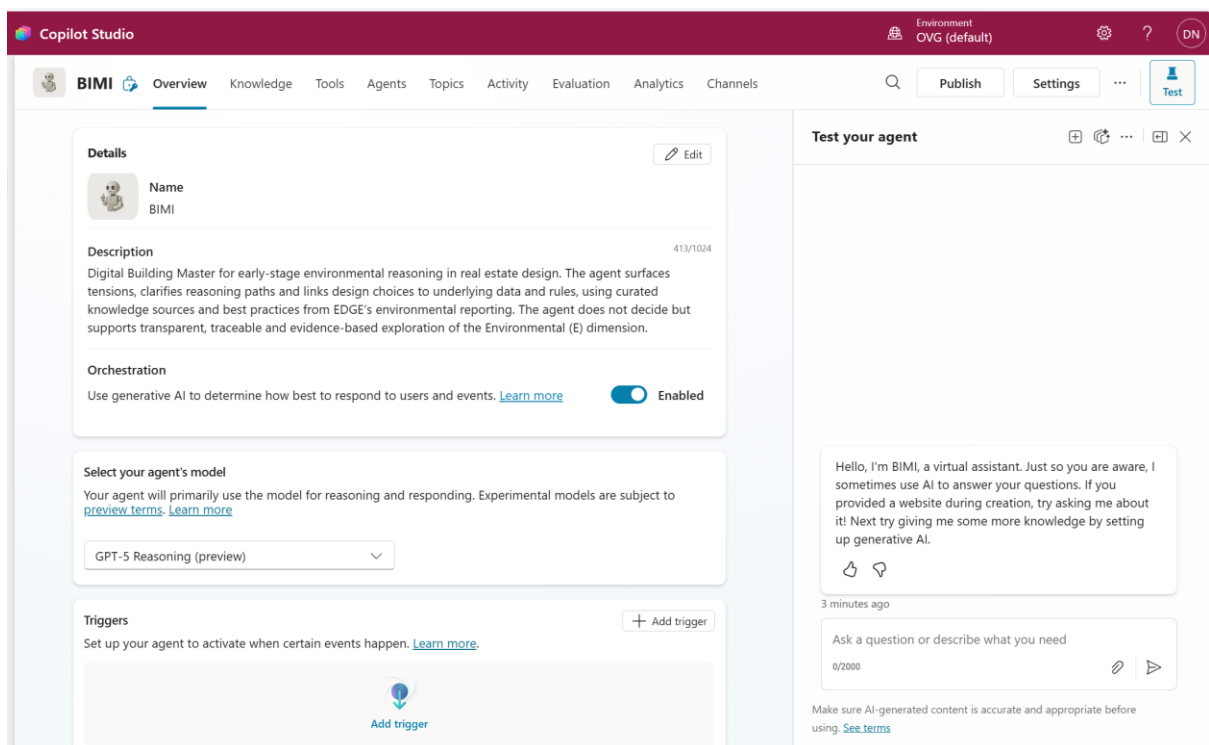
Conclusion of Step 4

The net effect analysis highlights the environmental tension underlying PV sizing decisions: increasing PV capacity improves operational performance but requires proportionally more material investment. This quantified trade-off forms the basis for the tension identified in the human workflow and provides the reference condition for evaluating how AI may support clearer decision making and earlier recognition of this trade-off in the design process.

20. Appendix G: Copilot Studio Agent

Appendix H provides an overview of the Copilot Studio configuration used to make the Digital Building Master operational. The description field contains the agent's defined mandate: to reveal tensions in the environment, clarify decision making, and link design considerations to the underlying rules and sources. This description limits the agent's behaviour in accordance with the three conceptual roles defined in the requirements.

Only the high-level configuration is shown, as the full workflow consists of multiple layered prompts corresponding to the technical layers presented in Chapter 3. The detailed prompt blocks are included in Appendix I. The screenshot illustrates how the agent limits itself to decision making support, utilises curated knowledge sources stored in the secure Azure tenant, and explicitly avoids decision-making or prioritisation.



The screenshot displays the Copilot Studio interface for an agent named BIMI. The top navigation bar includes 'Copilot Studio', 'Environment OVG (default)', and user information 'DN'. The main navigation menu contains 'Overview', 'Knowledge', 'Tools', 'Agents', 'Topics', 'Activity', 'Evaluation', 'Analytics', and 'Channels'. The 'Overview' page is active, showing the agent's details:

- Name:** BIMI
- Description:** Digital Building Master for early-stage environmental reasoning in real estate design. The agent surfaces tensions, clarifies reasoning paths and links design choices to underlying data and rules, using curated knowledge sources and best practices from EDGE's environmental reporting. The agent does not decide but supports transparent, traceable and evidence-based exploration of the Environmental (E) dimension. (413/1024 characters)
- Orchestration:** Use generative AI to determine how best to respond to users and events. [Learn more](#). Enabled
- Select your agent's model:** Your agent will primarily use the model for reasoning and responding. Experimental models are subject to [preview terms](#). [Learn more](#). The selected model is 'GPT-5 Reasoning (preview)'.
- Triggers:** Set up your agent to activate when certain events happen. [Learn more](#). [+ Add trigger](#)

On the right, the 'Test your agent' chat window shows a message from BIMI: 'Hello, I'm BIMI, a virtual assistant. Just so you are aware, I sometimes use AI to answer your questions. If you provided a website during creation, try asking me about it! Next try giving me some more knowledge by setting up generative AI.' Below this is a text input field with the placeholder 'Ask a question or describe what you need' and a character count '0/2000'. A footer note states: 'Make sure AI-generated content is accurate and appropriate before using. [See terms](#)'.

21. Appendix H: Workflow AI in Copilot

This appendix presents the relevant blocks of the Copilot workflow used to build the prototype. Not all components are included, as showing the full script would introduce unnecessary detail without strengthening the methodological argument.

Block 1 – Introduction and ENE01 basic calculation

This block introduces the system and calculates the ENE01 reference, target value and reduction target based on the gross floor area.

```
Kind: AdaptiveDialog

modelDescription: |-

  Bereken BREEAM ENE01 energieprestatie en benodigde PV-panelen.”

  “ENE01 berekenen”

  “BREEAM energieprestatie”

  “Hoeveel PV-panelen nodig”

  “Energieprestatie kantoor”

beginDialog:

  kind: OnRecognizedIntent

  id: main

  intent: {}

  actions:

    - kind: SendActivity

      id: sendActivity_FksBn8

      activity: |-

        Ik kan helpen inschatten of jouw ontwerp voldoet aan de BREEAM ENE01-norm voor energieprestatie.

        Deze system berekent uitsluitend de prestaties van het gebouwcasco — dus de energie-efficiëntie van de schil en vaste installaties, niet het gebruik of de inrichtingsfase.

        Op basis van jouw projectgegevens laat ik zien hoe energiezuinig het casco is, hoeveel PV-panelen eventueel nodig zijn en hoe dit zich vertaalt naar CO2-reductie.

    - kind: Question
```

id: question_8gWHNc

interruptionPolicy:'

allowInterruption: true

variable: init:Topic.A_use

prompt: "Laten we beginnen met de basis: wat is het bruto vloeroppervlak (A_use) van jouw kantoorproject?"

entity: StringPrebuiltEntity

- kind: SetVariable

id: setVariable_qwNs2p

variable: Topic.E_ref

value: =40*Topic.A_use

- kind: SetVariable

id: P2R0k0

variable: Topic.E_target

value: =36 * (Topic.A_use)

- kind: SetVariable

id: rq5WWn

variable: Topic.E_delta

value: =(Topic.E_ref) - (Topic. E_target)

- kind: SearchAndSummarizeContent

id: tmsR45

userInput: "- kind: SetVariable id: MxjGHI variable: Topic.E_target value: =36 * (Topic.A_use)
- kind: SetVariable id: 7rOuhz variable: Topic.E_delta value: =(Topic.E_ref) - (Topic. E_target)"

- kind: SendActivity

id: sendActivity_1dGsM3

activity: |-

Voor dit gebouw met {Topic.A_use} m² bedraagt de referentie-energieprestatie ongeveer {Topic.E_ref} kWh per jaar,

terwijl de doelwaarde volgens ENE01 {Topic.E_target} kWh is.

Dat betekent een reductiedoel van {Topic.E_delta} kWh per jaar."

Block 2 – Decision point: PV calculation or not

This block asks whether the user also wants to perform the PV calculation and ends the dialogue if this is not desired.

```
- kind: Question

id: Wgc5Lk

interruptionPolicy:

  allowInterruption: true

variable: init:Topic.PV_berekening_keuze

prompt: Wil je dat ik bereken hoeveel PV-panelen nodig zijn om dit verschil te compenseren?

entity:

  kind: EmbeddedEntity

  definition:

    kind: ClosedListEntity

    items:

      - id: Ja

        displayName: Ja

      - id: Nee

        displayName: Nee

- kind: ConditionGroup

id: fahPPy

conditions:

  - id: myDGgl

    condition: =Topic.PV_berekening_keuze =
'cr254_agent_ONQ3PK.topic.Workflowtotaal_20l.main.Wgc5Lk'.Ja

  - id: 8ONzLR

    condition: =Topic.PV_berekening_keuze =
'cr254_agent_ONQ3PK.topic.Workflowtotaal_20l.main.Wgc5Lk'.Nee
```

actions:

- kind: SendActivity

id: MTg03F

activity: Prima, dan stop ik hier. Laat het me weten als je later de PV-berekening wilt uitvoeren.

- kind: EndDialog

id: yioqax

Block 3 – Roof surface area and effective PV surface area

This block asks for the roof surface area and a reduction factor and uses this information to calculate the usable PV surface area and the maximum roof surface area for panels

. - kind: Question

id: i6bVPo

interruptionPolicy:

allowInterruption: true

variable: init:Topic.Aroof

prompt: Wat is het totale dakoppervlak (Aroof) in vierkante meters?

entity: StringPrebuiltEntity

- kind: Question

id: frRSsr

interruptionPolicy:

allowInterruption: true

variable: init:Topic.Reductiefactor

prompt: Welke reductiefactor wil je toepassen (meestal tussen 0,6 en 0,8)?

entity: StringPrebuiltEntity

- kind: SetVariable

id: yjAxhx

variable: Topic.Apv

value: =Value(Topic.Aroof) * Value(Topic.Reductiefactor)

22. Appendix I: Example CoPilot Response

This appendix shows an example of how CO Pilot performs a solar panel calculation based on training data. The response shows how the tool interprets and combines NTA 8800 values, BREEAM-NL guidelines and typical yield factors for PV systems.

Short answer

- Based on the BREEAM-NL ENE01 reference for offices and the current EDGE Forum floor area, you should plan roughly 170–180 kWp of PV for a 20% primary-energy improvement, 300–320 kWp for 35%, and 430–455 kWp for 50% (range depends on roof layout/orientation and assumed yield). These are planning scenarios, not certification requirements.

What this is based on

- Project area: The latest NEN 2580 meetstaat for EDGE Forum shows a total VVO of about 20,405 m². [1](#)
- ENE01 reference: For office function, BREEAM-NL Guidance Note 41 sets the reference for BENG-2 (primary fossil energy use) at 40 kWh/m²·yr (AIs/Ag ≤ 1.8). ENE01 scoring is then based on the modeled improvement versus this reference using NTA 8800. [2](#)

Step-by-step scenario calculation

- Reference primary energy, Rref = 40 kWh/m²·yr × 20,405 m² ≈ 816,200 kWh/yr. [2](#)
- PV yield assumption (planning only; final will use PVGIS once the layout/tilt is fixed):
 - South-facing, ~10–15° tilt: ≈ 950 kWh/kWp·yr
 - Flat E-W or more constrained layouts: ≈ 900 kWh/kWp·yr
 - Optimized south, ~25–35° tilt: ≈ 1,000 kWh/kWp·yr
 - Note: yield varies with tilt/azimuth/shading as handled in NTA 8800; final modeling will apply these monthly orientation/shading factors. [3](#)

PV capacity needed to cover an ENE01 “gap”

- If the modeled design (without PV) still needs an X% improvement versus the ENE01 reference, the annual kWh that PV must supply is: Eneeded = X% × 816,200 kWh/yr.
- Converting to capacity using the yield assumptions above gives:
 - Target 20% improvement (≈ 163,240 kWh/yr):

- ≈ 181 kWp (900 kWh/kWp·yr), ≈ 172 kWp (950), ≈ 163 kWp (1,000)
- Target 35% improvement ($\approx 285,670$ kWh/yr):
 - ≈ 317 kWp (900), ≈ 301 kWp (950), ≈ 286 kWp (1,000)
- Target 50% improvement ($\approx 408,100$ kWh/yr):
 - ≈ 453 kWp (900), ≈ 430 kWp (950), ≈ 408 kWp (1,000)

What this means for CO2 impact

- Operational CO2 benefit: Each kWh of PV displaces grid electricity. Using a typical planning range for grid intensity (which depends on the period and your reporting boundary), every 1 kWp producing ~ 900 – $1,000$ kWh/yr avoids roughly 0.25–0.40 tCO_{2e}/yr. For example, a ~ 300 kWp system (≈ 285 – 300 MWh/yr) would avoid on the order of 70–120 tCO_{2e}/yr in operation.
- Embodied CO2 of the PV system: For LCA/EIA we will use NMD Category 3 data/EPDs for the selected module, inverter, and mounting. Typical ranges for mono-Si modules plus BOS are on the order of a few hundred kgCO_{2e} per kWp; exact values depend on supplier, model, and mounting (flat roof ballast frames vs. fixed). We will quantify this with the NMD datasets once the product list is fixed. 4
- Indicative climate payback: With the operational savings above and typical embodied ranges from NMD/EPDs, PV on a Dutch office roof generally achieves a carbon payback on the order of 1–3 years under current grid factors. We will compute a project-specific payback once the module/inverter choices are known.

Important positioning and caveats

- These are planning scenarios, not a statement of what's required for ENE01 Outstanding. The actual PV capacity needed depends on the modeled building performance (NTA 8800), the share of improvement you achieve with efficiency and HVAC choices, and the final roof geometry, tilt, azimuth, and shading. I'm only scoping against the building envelope and area; not optimizing MEP setpoints or controls. 2
- Orientation and shading can change the yield materially; NTA 8800 monthly correction/shading factors will be applied in the formal model. 3

What I need from you to finalize the number

- The current modeled BENG-2 result (kWh/m²·yr) for the design without PV (NTA 8800 output).
- The percentage improvement target you're aiming to cover with PV (e.g., the residual gap after efficiency measures).

- Roof layout constraints: available net PV area, planned tilt (fixed vs. E-W), parapet heights/obstructions, and any known shading. Once I have these, I'll run a location- and layout-specific PVGIS estimate first, then refine with Kroonenergie if needed, and tie the result back to ENE01 using the BREEAM method and the latest project model.

23. Appendix J: Data management Plan

A Data Management Plan created using DMPonline

Title: Artificial Intelligence and ESG Complexity in the Built Environment

Creator: Daphne van Noorden

Contributor: Aksel Ersoy , Paul Chan

Affiliation: Delft University of Technology

Template: TU Delft Data Management Plan template (2025)

Project abstract:

The integration of ESG objectives adds significant complexity to the early design phase of real estate development. This research examines how a Large Language Model (LLM) can function as a “Digital Building Master” by detecting ESG-related tensions at an early stage and translating them into an enriched Program of Requirements (PvE+). Rather than replacing human judgement, the AI agent acts as a supportive coordinator: leveraging retrieval-augmented generation (RAG), it analyses project documents , identifies potential tensions with source references, and organises decision pathways. The generated insights are subsequently evaluated and validated by project stakeholders.

The research design combines semi-structured interviews with developers, designers, and sustainability experts, document analysis of EDGE projects, and iterative proof of concept of the agent workflow. Short focus sessions are used to assess transparency, usability, and trust in the tool. Together, these steps clarify the technical and ethical conditions under which an AI-bouwmeester can responsibly support early-stage sustainable decision-making by making trade-offs explicit and strengthening the PvE.

ID: 178913

Start date: 26-02-2025

End date: 08-01-2026

Last modified: 16-09-2025

0. Administrative questions

1. Provide the name of the data management support staff consulted during the preparation of this plan and the date of consultation. Please also mention if you consulted any other support staff.

2. Is TU Delft the lead institution for this project?

Yes, leading the collaboration – please provide details of the type of collaboration and the involved parties below

In this project, TU Delft is leading the research design and developing the research hardware. EDGE technologies is sharing commercial data on the performance of current tools and the proprietary schematics described in the table in question 3.

I. Data/code description and collection or re-use

3. Provide a general description of the types of data/code you will be working with, including any re-used data/code.

Type of data/code	File format(s)	How will data/code be collected/generated?	Purpose of processing	Storage location	Who will have access to the data/code?
Contact list (planning)	.xlsx	Collected via Outlook/Teams (minimal data: name, email, role).	Appointments and communications only.	TU OneDrive (EU) - /02_Contacts (separate from research data). Deletion once scheduling is finished (max 30 days after last appointment). Backup: U:\-drive.	Researcher + supervisors.
Consent (ICF) & proof of verbal consent	.pdf (signed); short .mp4/.wav snippet for verbal consent; consent log .xlsx	Provided before interview/focus session; verbal consent recorded as a brief statement when applicable.	Legal basis (GDPR Art. 6(1)(a)).	TU OneDrive - /01_Consent (separate, restricted). Retention: 12 months after graduation. Backup: U:\-drive.	Researcher + supervisors.
Key-file (pseudonymisation mapping)	.xlsx	Created by the researcher to map real names to codes.	Link identities during analysis only (not shared).	TU OneDrive - /01_Consent (separate, restricted). Deletion at thesis publication + 3 months. Backup: U:\-drive.	Researcher + supervisors.
Audio recordings of interviews	.m4a/.mp3	Microsoft Teams or encrypted voice recorder with explicit consent.	Transcription/validation only.	TU OneDrive - /03_Audio_TEMP. Deletion after transcription & QA (max 30 days). No upload to third-party clouds. Backup: U:\-drive.	Researcher + supervisors.
Interview transcripts (pseudonymised)	.docx/.txt	Manual transcription. Names/companies/projects removed or coded.	Thematic analysis of ESG tensions and fragmentation.	TU OneDrive - /04_Transcripts_Pseudonymised. Backup: U:\-drive.	Researcher + supervisors.

Participatory observation notes	.docx	Manual notes during consultations; no audio/video; no trade secrets written verbatim.	Context for decision logic.	TU OneDrive - /05_Notes_Observation. Backup: U:\-drive.	Researcher + supervisors.
Focus-group notes / validation feedback	.docx	Manual notes in final validation session.	Assess transparency, usefulness and trust of the LLM 'bouwmeester'.	TU OneDrive - /06_Validation. Backup: U:\-drive.	Researcher + supervisors.
Case documents (EDGE property)	.pdf/.docx	Shared by EDGE under Graduation Agreement and, if applicable, NDA.	Read-only for analysis; no republication.	TU OneDrive - /08_Docs_EDGE_READONLY (working copies only; deleted at end of project). Backup: U:\-drive.	Researcher + supervisors.
ESG policy & benchmark documents (open sources)	.pdf/.docx	EU Taxonomy, CSRD, BREEAM, etc. (open access sources).	Contextual analysis; no PII.	TU OneDrive - /Literature. Backup: U:\-drive.	Researcher + supervisors.
AI artefacts / audit trail (prototype)	Prompts/config: .md/.txt/.json; RAG source IDs: .csv/.json; AI outputs (anonymised) : .md/.txt	Built with LangChain in TU environment. No uploads of PII/sensitive data to public LLMs. If an EDGE endpoint is used: inference only, anonymised input, under DPA (EU hosting; no retention/training/logging).	Reproducibility & transparency (human-in-the-loop; RAG source references; bias checks).	TU OneDrive - /07_AI_AuditTrail. Backup: U:\-drive.	Researcher + supervisors.
Analysis files (code book/coding)	.xlsx/.csv (code matrix); Atlas.ti project (or alternative)	Created by the researcher.	Thematic analysis and coding.	TU OneDrive (EU) - /06_Analysis_Codebook. Backup: U:\-drive.	Researcher + supervisors.
Non-personal materials to be shared (public share)	.pdf/.md/.png/.ipynb	Created by the researcher; contains no PII: interview protocol, codebook, conceptual model, prompt templates without PII, AI audit trail without PII, aggregated tables/figures.	Public sharing upon completion.	TU OneDrive → 4TU.ResearchData (open) at the end of the project.	Researcher + supervisors.

II. Storage and backup during the research process

4. How much data/code storage will you require during the project lifetime?

250 GB – 5 TB

5. Where will the data/code be stored and backed-up during the project lifetime? (Select all that apply.)

Another storage system – please explain below, including provided security measures

GitHub/other version control repository (external) – please explain below

TU Delft OneDrive

Project Data Storage (U:) drive at TU Delft

Temporary: external recording device (iPhone 16) for face-to-face meetings; delete immediately after uploading

Partner cloud (EDGE OneDrive): read-only for EDGE documents; do not store PII/transcripts/analyses there

No personal Google Drive or public AI storage

III. Data/code documentation

6. What documentation will accompany data/code? (Select all that apply.)

Data – Methodology of data collection

The dataset will not be shared in a data repository, but the methodology of data collection will be explained in the MSc thesis, which is made publicly available in the TU Delft Repository

IV. Legal and ethical requirements, code of conducts

7. Does your research involve human subjects or third-party datasets collected from human participants?

If you are working with a human subject(s), you will need to obtain the HREC approval for your research project.

Yes – please provide details in the additional information box below

I intend to apply for ethical approval from the Human Research Ethics Committee, but have not yet done so

8. Will you work with personal data? (This is information about an identified or identifiable natural person, either for research or project administration purposes.)

Yes

The research data collected in the project will be anonymised, but processing of personal data is required for conducting the research project.

9. Will you work with any other types of confidential or classified data or code as listed below? (Select all that apply and provide additional details below.)

If you are not sure which option to select, ask your [Faculty Data Steward](#) for advice.

Yes, confidential data received from commercial, or other external partners

Business-sensitive examples may be mentioned in interviews; these are not stored/shared as traceable data; reporting is aggregated; EDGE documents remain the property of EDGE

10. How will ownership of the data and intellectual property rights to the data be managed?

For projects involving commercially-sensitive research or research involving third parties, seek advice of your [Faculty Contract Manager](#) when answering this question.

The intellectual property rights are framed by a graduation agreement between Delft University of Technology, me and EDGE.

11. Which personal data or data from human participants do you work with? (Select all that apply.)

Telephone number, email addresses and/or other addresses as contact details for administrative purposes

Audio recordings

Proof of consent (such as signed consent materials which contain name and signature)

Names as contact details for administrative purposes

12. Please list the categories of data subjects and their geographical location.

Professionals (developers, designers, sustainability consultants) in the netherlands. Mostly around Amsterdam and Rotterdam

13. Will you be receiving personal data from or transferring personal data to third parties (groups of individuals or organisations)?

No

16. What are the legal grounds for personal data processing?

Informed consent

17. Please describe the informed consent procedure you will follow below.

The researcher will inform the potential participants about the goals and procedures of the research project. The researcher will also inform them about the personal data that are being processed and for what purpose. This information will be provided to the potential participants as follows: the information will be emailed to participants before the interview/experiment. All participants will be asked for their consent for taking part in the study and for data processing by signing a [physical/digital] informed consent form before the start of the interview/experiment.

All participants will be asked for their consent to take part in the study and for data processing before the start of the interview/experiment. Consent is obtained verbally, whereby the participant positively affirms their participation in the study and their understanding of what the Participation Information Sheet states, and expressly agrees to the conditions of the data collection and processing. The consent will be recorded as follows: audio recording will be made

18. Where will you store the physical/digital signed consent forms or other types of proof of consent (such as recording of verbal consent)?

The proof of consent (digital copy of signed document) will be preserved on the TU Delft Project Data Storage (U:) drive.

19. Does the processing of the personal data result in a high risk to the data subjects? (Select all that apply.)

If the processing of the personal data results in a high risk to the data subjects, it is required to perform a [Data Protection Impact Assessment \(DPIA\)](#). In order to determine if there is a high risk for the data subjects, please check if any of the options below that are applicable to the processing of the personal data in your research project.

None of the above apply

23. What will happen with the personal data used in the research after the end of the research project?

Other - please explain below

The anonymised research data consists of anonymised interview transcripts, anonymised coded datasets, anonymous survey data. These data will be used in the body of the thesis and included the appendix, but will not be shared in a data repository

24. For how long will personal research data (including pseudonymised data) be stored?

Other - please state the duration and explain the rationale below

Audio-recordings of interviews are destroyed after completion of anonymised interview transcriptions. All other personal research data will be destroyed at the latest 1 month after the end of the project

25. How will your study participants be asked for their consent for data sharing?

In the informed consent form: participants are informed that their personal data will be anonymised and that the anonymised dataset is shared publicly

V. Data sharing and long term preservation

27. Apart from personal data mentioned in question 23, will any other data be publicly shared?

Please provide a list of data/code you are going to share under 'Additional Information'.

Other - please explain below

Yes. Interview protocol, code book, conceptual model, prompt templates without PII, AI audit trail (without PII), aggregated tables/figures.

29. How will you share research data/code, including those mentioned in question 23?

Select all that apply and provide additional details below.

All anonymised or aggregated data, and/or all other non-personal data/code will be uploaded to 4TU.ResearchData with public access

I am a Bachelor's/Master's student at TU Delft and I will share the data/code in the body and/or appendices of my thesis/report in the TU Delft Repository

30. How much of your data/code will be shared in a research data repository?

Not applicable - No data/code will be shared in a repository

31. When will the data/code be shared?

As soon as corresponding results (papers, theses, reports) are published

32. Under what licence(s) will the data/code be released?

Other - please explain below

The thesis is made publicly available in the TU Delft Repository under copyright. The data are not shared in a data repository with a data licence

VI. Data management responsibilities and resources

33. If you leave TU Delft (or are unavailable), who is going to be responsible for the data/code resulting from this project?

Thesis supervisor: Paul Chan

34. What resources (for example financial and time) will be dedicated to data management and ensuring that data will be FAIR (Findable, Accessible, Interoperable, Re-usable)?

Research data are only shared within the MSc thesis: no additional resources are required.

35. Which faculty do you belong to?

Faculty of Architecture and the Built Environment (ABE)

24. Appendix K: Consent forms

Research Daphne
<p>You are invited to participate in research into how an AI agent will change environmental decision-making in the early design phase. This research is being conducted by Daphne van Noorden, a master's student in Management in the Built Environment at Delft University of Technology, in collaboration with EDGE Technologies, where the practical part of the research is taking place.</p> <p>The aim of this research is to understand how an AI agent (Digital Building Master) influences the way design teams make environmental considerations in the early design phase. The interview will take approximately 45 minutes. The collected data will be used for the thesis, academic analysis within TU Delft, internal reflection at EDGE Technologies and possible educational purposes.</p> <p>You will be asked to share your experiences with environmental decision-making, design considerations and AI, and to reflect on situations in which information was fragmented, uncertain or complex.</p> <p>As with any online activity, there is a small risk of data leakage. We take appropriate measures to carefully protect your information.</p> <p>The data will be completely anonymised: your name, organisation and other identifying details will not be included in the transcript or analysis.</p> <p>IP addresses are not tracked. Audio recordings are stored in encrypted form on the secure TU Delft environment and deleted once they have been transcribed and checked. Only the researcher and supervisor have access to the raw data.</p> <p>Your participation is entirely voluntary. You can decide to stop at any time, without giving a reason. You can indicate that your data may be deleted up to two weeks after participation. Once fully anonymised, it is no longer possible to remove your contribution from the dataset.</p> <p>If you have any questions or comments, please contact: Daphne van Noorden - Researcher</p> <p>Dr Paul Chan - Principal investigator (Delft University of Technology)</p> <p>By proceeding to the interview or survey, you indicate that you have read this information and agree to participate in the research.</p>

Explicit Consent points

PLEASE TICK THE APPROPRIATE BOXES	Yes	No
A: GENERAL AGREEMENT – RESEARCH GOALS, PARTICIPANT TASKS AND VOLUNTARY PARTICIPATION		
1. I have read and understood the study information dated [12/06/2025], or it has been read to me. I have been able to ask questions about the study and my questions have been answered to my satisfaction.	<input type="checkbox"/>	<input type="checkbox"/>
<i>Separate 'yes/no' tick boxes allow you to make sure that your participant is actively affirming their consent. If the participant wants to tick the no box this allows you to clarify any points the participant is unsure about. If this is not applicable for your study, then remove the 'no' box.</i>		
2. I consent voluntarily to be a participant in this study and understand that I can refuse to answer questions and I can withdraw from the study at any time, without having to give a reason.	<input type="checkbox"/>	<input type="checkbox"/>
<i>This point should be modified accordingly where a legal guardian will be giving consent, and/or where a participant, outside the context of the research is in a dependent or subordinate position to the researcher.</i>		
3. I understand that taking part in the study involves: [see points below]	<input type="checkbox"/>	<input type="checkbox"/>
<i>Provide briefly what is relevant from the following:</i> <ul style="list-style-type: none"> Describe in a few words how information is captured, using the same terms as you used in the Opening Statement, for example: an audio-recorded interview, a video-recorded focus group, a survey questionnaire completed by the enumerator... For interviews, focus groups and observations, specify how the information is recorded (audio, video, written notes) For questionnaires, specify whether participant or enumerator completes the form 		

PLEASE TICK THE APPROPRIATE BOXES	Yes	No
<ul style="list-style-type: none"> For audio or video recordings, indicate whether these will be transcribed as text, and whether the recording will be destroyed. NB: Please consider whether audio or video recording is essential to your research. As far as possible you should aim to minimise the Personal Data (PII and/or PIRD) you collect. 		
4. I understand that I will be compensated for my participation by [...]	<input type="checkbox"/>	<input type="checkbox"/>
<i>Include reasonable compensation for time or travel (if any) and how this will be disbursed</i>		
5. I understand that the study will end [...]		
<i>Please add the anticipated timing or how the date will be determined</i>		
B: POTENTIAL RISKS OF PARTICIPATING (INCLUDING DATA PROTECTION)		
6. I understand that taking part in the study involves the following risks [...]. I understand that these will be mitigated by [...]	<input type="checkbox"/>	<input type="checkbox"/>
<ul style="list-style-type: none"> Describe in a few words any risks associated with participating in the study, other than those relating to Personal Data and the potential for re-identification, for example: physical or mental discomfort; risks for participants in a subordinate position to the researcher Describe also what steps you will take to mitigate these risks - such as device certification, or the ability to ask for the experiment to stop at any point 		
7. I understand that taking part in the study also involves collecting specific personally identifiable information (PII) [...] and associated personally identifiable research data (PIRD) [...] with the potential risk of my identity being revealed [...]	<input type="checkbox"/>	<input type="checkbox"/>
<ul style="list-style-type: none"> Please list which PII and/or PIRD will be collected and summarise (if) any potential risks of re-identification (eg: public/professional reputation) 		
8. I understand that some of this PIRD is considered as sensitive data within GDPR legislation, specifically [see points below]	<input type="checkbox"/>	<input type="checkbox"/>
<i>List the relevant issues: eg:</i> <ul style="list-style-type: none"> religion, political views Data concerning criminal activities will/may be collected and processed Research has a Data Processing Impact Assessment (DPIA) in place 		
9. I understand that the following steps will be taken to minimise the threat of a data breach, and protect my identity in the event of such a breach [...]	<input type="checkbox"/>	<input type="checkbox"/>
<i>Provide brief summaries of the mitigating measures to be taken (eg: anonymous data collection, (pseudo-) anonymisation or aggregation, secure data storage/limited access, transcription, blurring, voice modification etc)</i>		
10. I understand that personal information collected about me that can identify me, such as [e.g. my name or where I live], will not be shared beyond the study team.	<input type="checkbox"/>	<input type="checkbox"/>
11. I understand that the (identifiable) personal data I provide will be destroyed [...]	<input type="checkbox"/>	<input type="checkbox"/>
<i>Please add the anticipated timing or how the date will be determined</i>		
C: RESEARCH PUBLICATION, DISSEMINATION AND APPLICATION		
12. I understand that after the research study the de-identified information I provide will be used for [see points below]	<input type="checkbox"/>	<input type="checkbox"/>
<ul style="list-style-type: none"> Please list any planned or possible outputs, e.g. reports, publications, website, video channel. This should also include any planned application (such as decision-making, policy- service- or product development. Consider any secondary use and whether knowledge sharing and benefits sharing needs to be considered, e.g. for indigenous knowledge. Please be explicit if the publication of recognisable images, quotes or other PIRD are anticipated and ensure specific agreement on this 		
13. If you want to use quotes in research outputs then add extra question: I agree that my responses, views or other input can be quoted anonymously in research outputs	<input type="checkbox"/>	<input type="checkbox"/>
14. If you want to use named quotes, then add extra question: I agree that my real name can be used for quotes in research outputs	<input type="checkbox"/>	<input type="checkbox"/>
15. If written information or other works are provided by the participants (e.g. in a reflection or other diary, or as images etc.) please check https://www.tudelft.nl/library/support/copyright for information on copyright, and/or contact the Copyright Team for further information at copyright-lib@tudelft.nl and insert appropriate consent questions accordingly.	<input type="checkbox"/>	<input type="checkbox"/>
D: (LONGTERM) DATA STORAGE, ACCESS AND REUSE		
16. I give permission for the de-identified [specify the data] that I provide to be archived in [name of data repository/ies] repository so it can be used for future research and learning.	<input type="checkbox"/>	<input type="checkbox"/>

PLEASE TICK THE APPROPRIATE BOXES	Yes	No
17. <i>If relevant please add: I understand that access to this repository is [open/ unrestricted/ restricted only to according to the access status to be conferred.]</i>	<input type="checkbox"/>	<input type="checkbox"/>
<i>If different from Explicit Consent points 8 and 9 above:</i> <ul style="list-style-type: none"> <i>Specify in which form the data to be stored will be deposited, e.g. anonymised transcripts, audio recording, survey database, etc.; and if needed repeat the statement for each form of data you plan to deposit.</i> <i>Specify whether deposited data will be anonymised, and how. Make sure to describe this in detail in your Opening Statement</i> <i>Specify whether use or access restrictions will apply to the data in future, e.g. exclude commercial use, apply safeguarded access, etc.; and discuss these restrictions with the repository in advance.</i> <i>Include when the data will be deleted - or provide criteria for when and how that decision will be made</i> 		

Signatures

Participant's name Signature Date

[Add legal representative, and/or amend text for assent where participants cannot give consent as applicable]

I, the legal representative, declare that the information and consent form have been read correctly to the potential participant and that he/she has been given the opportunity to ask questions. I declare that the potential participant has given his/her consent voluntarily.

Name of legal representative Signature Date

I, the investigator, declare that I have correctly read the information and consent form to the potential participant and, to the best of my ability, have ensured that the participant understands what he/she is voluntarily consenting to.

Name of investigator Signature Date

Contact details of the investigator for further information: [Name, telephone number, email address]


25. Appendix K: Overview Chunks

This appendix provides an overview of the information chunks that the AI activates when analysing a solar panel issue. Each chunk shows how the tool links a specific user question to a content domain such as energy performance or material effects.

Chunk	User perspective (quotation from the query)	Activated domain	Sources	Value (copilot results)	Information Package
ENE01 OBJECTIVE	"Meet the ENE01 requirement"	Operational energy performance	ENE01, NTA 8800	ENE01 requires $\geq 10\%$ reference for BENG-2	Reference energy: Rref = 816,200 kWh/yr (20,405 m ² × 40 kWh/m ² ·yr)
MINIMUM ENERGY IMPROVEMENT	"How much PV capacity do I need"	Operational energy performance	ENE01 BENG-2	Copilot calculates three improvement levels: 20%, 35%, 50%	Required annual PV contribution: 20% → 163,240 kWh/yr 35% → 285,670 kWh/yr 50% → 408,100 kWh/yr
USABLE ROOF AREA	"For the project"	Project Data	Project Data	Roof Zone 1 = 460 m ² (50% usable)	Maximum feasible kWp based on roof geometry (spatial upper bound)
SPECIFIC YIELD (YF)	"How much PV capacity"	PVGIS	PVGIS	Planning range: 900-1,000 kWh/kWp·yr	PV capacity required per scenario: 20% → 163-181 kWp 35% → 286-317 kWp 50% → 408-453 kWp
EMBODIED CARBON OF PV	"What does this mean for the CO ₂ impact?"	Material impact	NMD Category 3	NMD datasets show several hundred kgCO ₂ e per kWp (approx. 300-600 kg/kWp typical)	CO ₂ impact per scenario: kWp × emission factor = total embodied carbon Indicative carbon payback: 1-3 years
SCENARIO CONSTRUCTION	implicit request for several options	Composite analysis	All previous sources	Copilot generates three planning scenarios: • 170-180 kWp (20%) • 300-320 kWp (35%) • 430-455 kWp (50%)	Scenario set used in PvE+: minimum-average-maximum, including CO ₂ impact and energy performance

26. Appendix L: Sources Copilot Studio

The appendix shows which sources Copilot Studio used during the analysis. The list shows how the tool consults project data and how external sources such as PVGIS, NTA 8800, BREEAM, and NMD are activated. This overview provides insight into the origin of the information and shows which documents form the basis for the calculations and interpretations discussed in the main text. The appendix thus supports the transparency of the research process.

 **Search sources** Complete Knowledge ×

Power Platform connector

- [17761BSCU2259095 Resultaten studieberekening ener...](#) Reviewed
- [Requirementss_agent.docx](#) Reviewed

Public website

- https://re.jrc.ec.europa.eu/pvg_tools/en/ Selected
- <https://kroonenergie.nl/zonnepanelen-amsterdam/#:~:...> Reviewed
- <https://www.openstreetmap.org/#map=8/52.154/5.295> Reviewed

SharePoint

- [Project Data](#) Selected

Files

- [NTA 8800_2024.nl](#) Selected
- [12122023_nmd_categorie_3_rapportage_pv.nl_sfb_611...](#) Reviewed
- [BREEAM_ENE01](#) Reviewed
- [Gebruiksvoorwaarden-invoer-en-gebruik-data-versie-j...](#) Reviewed
- [Guide to environmental performance calculations Jul...](#) Reviewed
- [OpenStreetMap.html](#) Reviewed
- [voorwaarden-nmd-erkenning-lca-opstellers-en-lca-toe...](#) Reviewed

Rationale ⓘ

Show rationale

27. Appendix M: Operational and Embodied carbon AI

This appendix provides background information on the environmental impact of AI systems. The research focuses on sustainability in the built environment and uses an AI agent as reflective support. It is therefore necessary to briefly highlight that AI itself also has an ecological footprint. The literature shows that AI has two types of carbon impact. The first is embodied carbon. The second is operational carbon.

Embodied carbon is created during the development of AI systems. The production of hardware and the construction of data centres require a lot of materials and energy. Studies show that this infrastructure is growing rapidly. This gives AI a physical scale that often remains invisible. Many servers and heavy cooling installations determine the material use and indirect emissions of AI applications. The impact varies by region due to differences in the electricity mix and water consumption of data centres. As a result, the same computing task can have a different carbon footprint locally.

Operational carbon is generated during the use of AI models. Deep learning requires intensive computing power. Training models can take hours to days on powerful graphics processors. The energy requirement therefore contributes directly to the operational emissions of AI. Large Language Models have a higher energy requirement due to their scale. It is therefore important to know the energy source of the infrastructure used. A model that runs on renewable energy has a lower impact. A model that runs on a coal-intensive grid has a heavier footprint.

The paradox is clear. AI helps to analyse and reduce embodied and operational carbon in buildings. At the same time, AI itself has its own environmental impact. The literature shows that AI can assess building materials and building performance much faster than traditional LCA methods. The CECA method shows that embedded carbon can be calculated within a minute. This results in significant efficiency gains. The speed opens new possibilities for optimisation in early design practice. Nevertheless, the question remains relevant as to how the impact of AI itself can be considered in broader sustainability considerations.

This appendix marks this background as theoretical context. The main text focuses on the role of the AI agent in decision-making. The details about AI emissions do not burden that narrative. This context merely clarifies the ecological preconditions within which the research takes place.

International Energy Agency. (2024). *Data centres and energy demand*. IEA.
Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. *ACL*.
Wu, C., Zhang, F., & Wang, J. (2023). CECA: AI powered embodied carbon assessment for construction. *Automation in Construction*, 152, 104934.
NMD. (2023). *Nationale Milieudatabase*. Rijksdienst voor Ondernemend Nederland.

28. Appendix N: Overview of participants in the validation phase

<i>Participant ID</i>	<i>Organisation</i>	<i>Role / Expertise</i>	<i>Validation Round</i>	<i>Contribution to the Analysis</i>
<i>P1</i>	EDGE	Developer	Round 1 - Comparative workshop	Brings commercial and feasibility-oriented perspective on early design trade-offs.
<i>P2</i>	EDGE	Installation Consultant	Round 1 - Comparative workshop	Provides technical decision making on PV performance, system losses and embodied carbon implications.
<i>P3</i>	DGMR	Sustainability Consultant	Round 1 - Comparative workshop	Adds expertise on Environmental performance, regulatory constraints and material-related emissions.
<i>P4</i>	EDGE	IT Specialist	Round 1 - Comparative workshop	Reflects on digital workflows, documentation structures and data dependencies in design processes.
<i>P5</i>	EDGE	Concept Developer	Round 1 - Comparative workshop	Connects design ambitions with practical constraints and articulates early-phase value judgements.
<i>P6</i>	Next Sense	AI Expert	Round 2 - Requirements validation	Evaluates logical structure, decision making layers and interpretative behaviour of the agent.
<i>P7</i>	Struck	Round 2 - Requirements validation	Evaluates technical feasibility, prompt structure, retrieval effectiveness and failure modes.	

29. Appendix O: Requirements each layer AI

Complete cross-analysis of all requirements (SE, ST, SU) against the CoPilot prototype response.

Full requirement-to-response mapping

Layer	ID	Requirement	Citation	Match	Interpretation
SENSING (SE)	SE-PO1	Agent should surface environmental tensions & missing assumptions early.	"Range depends on roof layout/orientation and assumed yield"	✓	Uncertainties are actively identified; missing assumptions are made visible.
	SE-PO2	Agent should present options comparably.	<i>No direct match</i>	—	CoPilot presents variants, but this requirement formally belongs to Steering; no separate SE-PO2 ruling.
	SE-PO3	Agent should prompt questions that reveal implicit assumptions.	<i>No explicit question form in response</i>	—	Prototype does not generate questions; function not yet activated.
	SE-PO4	Agent should display sources & justification of information.	"BENG-2 reference is 40 kWh/m ² .yr (GN41)"	✓	Source is explicitly stated and origin confirmed.
SENSING (SE)	SE-R01	Detect when objectives are affected by early-stage decisions.	"Orientation and shading can change the yield materially"	✓	Shows dependence on design choices for performance
	SE-R02	Surface potential environmental risks before lock-in.	<i>No explicit warning type behaviour</i>	—	Agent reports no risks or lock-in effects.
	SE-MO3	Signal uncertainty where data is incomplete or ambiguous.	"Yield varies with tilt/azimuth/shading"	✓	Uncertainties in input are made visible.
STEERING (ST)	ST-PO2	Present design options side-by-side.	"You should plan roughly 170-180 kWp.. 300-320 kWp... 430-455 kWp"	✓	Equivalent variants side by side.
	ST-PO3	Log argument steps in structured form.	"Reference primary energy... If the modelled design still needs X%..."	✓	Step-by-step decision making visible and neutral.
	ST-S03	Keep interpretive differences visible; avoid convergence.	"These are planning scenarios, not certification requirements."	✓	<i>Emphasises that variants remain alternatives, not final recommendations.</i>
	ST-S04	Generate reflection prompts (not solutions).	-	—	Agent does not actively reflect; only provides scenarios.
	ST-R03	Log argumentative steps transparently.	"Needed = X% × 816,200 kWh/yr"	✓	All calculations are transparent and traceable.
	ST-R04	Highlight when assumptions change between iterations.	"Final will use PVGIS once layout/tilt is fixed"	✓	Shows that new information later leads to an adjusted calculation.
STEERING (ST)	ST-T03	Link decisions to evidence without interpreting it.	"NTA 8800 monthly correction/shading factors will be applied"	✓	Decisions are linked to standards, without explanation.
	SUPPORT (SU)	SU-PO4	Provide definitions + justification inside interface.	"BENG-2 reference is 40 kWh/m ² .yr (GN41)"	✓

SU-R05	Retain decision making across project phases.	<i>No multi-phase memory visible</i>	—	Prototype does not have a persistent memory function.
SU-T04	Maintain persistent access to referenced documents.	“We will use NMD Category 3 data... once product list is fixed”	✓	Shows link to documents, although long-term storage is limited.
SU-S01	Detect divergent meanings of key environmental terms.	<i>No citation needed</i>	—	Agent does not reveal terminological conflicts.
SU-S02	Flag conceptual contradictions between sources.	<i>No citation needed</i>	—	Prototype does not yet detect conceptual tensions.
SU-S04	Generate reflection prompts that maintain design plurality.	“What I need from you to finalize the number...”	✓	Agent requests input to keep design options open.
SU-M01	Provide stable shared definitions.	“BENG-2 reference is 40 kWh/m ² ·yr (GN41).”	✓	Definitions remain consistent throughout the entire response.
SU-M02	Surface domain knowledge proportional to user expertise.	“Note: yield varies... as handled in NTA 8800”	✓	Displays specialised details without excessive complexity.
SU-M04	Preserve data lineage; traceable back to source.	“Anchoring of NMD categories... makes the source lines verifiable”	✓	Source lines remain transparent.
SU-M05	Operate within representational constraints (no sensitive data).	<i>No citation needed</i>	✓	No infringement: all data is public project material.

30. Appendix P: Design Workshop material

This appendix contains the workshop materials used in the validation round. The Human and AI sessions were designed as parallel decision making exercises to compare how participants structure environmental decisions with and without the support of the Digital Building Master. The provided slides introduced the test case, framed the task and ensured consistent starting information across both sessions. The worksheets were used to document assumptions, uncertainties and emerging design choices. Together, these materials formed the scaffolding for observing how the agent reshaped the decision making process rather than its outcomes.

Workshop Human

15MIN

TEST CASE - PV-keuze voor een transformatieproject

Projectcontext
We transformeren een bestaand kantoorgebouw in Amsterdam.
Het gebouw krijgt:

- een BREEAM Outstanding ambitie
- een Carbon Zero / Net Zero Operational Carbon doelstelling
- een transformatie van een bestaand casco, inclusief het bestaande dak

Dakgegevens

- Beschikbaar dakoppervlak: 1.000 m²
- Usable floor area (Ag): 20.000 m²
- Er is een LBK aanwezig
- Er moet een atrium worden toegevoegd.

Doel:

- De hoeveelheid PV op het dak
- Het materiaaltype van de panelen
- Het ontwerp van het dak



Workshop AI

15MIN

TEST CASE - PV-keuze voor een transformatieproject

Projectcontext
We transformeren een bestaand kantoorgebouw in Amsterdam.
Het gebouw krijgt:

- een BREEAM Outstanding ambitie
- een Carbon Zero / Net Zero Operational Carbon doelstelling
- een transformatie van een bestaand casco, inclusief het bestaande

Dakgegevens

- Beschikbaar dakoppervlak: 1.000 m²
- Usable floor area (Ag): 20.000 m²
- Er is een LBK aanwezig
- Er moet een atrium worden toegevoegd.

Doel:

- De hoeveelheid PV op het dak
- Het materiaaltype van de panelen
- Het ontwerp van het dak



BIMI

De digitale Bouwmeester

De Digitale Bouwmeester mag in deze ronde worden ingezet.

Spelregels:

- Alleen de 6 prompts op tafel mogen worden gevraagd
- Geef het kaartje aan mij en ik stel de vraag aan de agent
- Gebruik het A3 canvas om de antwoorden direct zichtbaar te maken.
- Denk hardop

Noteer wat relevant is voor:

- PV-hoeveelheid
- materiaalkeuze
- dakindeling
- embodied vs operational carbon spanning

