CHOICE

Or the terrain of uncertainty, a study on how surveillance fragments memory

Master of Science in Architecture for Heritage

CHOICE

Or the terrain of uncertainty, a study on how surveillance fragments memory

Master's Thesis
Kosar Mohammadi

Supervisor
Prof. Tomà Berlanda

Student Num.
S329412

Politecnico di Torino | Feb. 2026

00|Abstract

This thesis investigates how surveillance fragments spatial perception and memory through agent-based computational modeling. Using women's navigation under surveillance in University of Tehran Faculty of Fine Arts as a case study, this research examines how systematic threats reshape the cognitive maps of individuals in controlled environments.

The methodology integrates spatial cognition theory [Tolman, Mallot, Ingold] with agent-based simulation to model navigation under stress and uncertainty. The simulation operationalizes three interacting systems: constrained perception (heading-dependent intake zones, occlusion), dynamic memory (dual-layer storage distinguishing spatial detail from learned threat, with differential decay rates), and stress-driven learning through which agents psychologically erase physically accessible paths. The model tests how individual risk tolerance thresholds determine breaking points where viable navigation collapses.

Results demonstrate that surveillance operates not through direct control but through the fragmentation of spatial knowledge itself. Even moderately cautious agents lose 60-76% of accessible space, while highly risk-averse agents experience paralysis in 33% of movement attempts. Critically, the simulations reveal an "inaccessible void", spaces rendered unusable not by direct threat but by network topology, as surrounding paths are psychologically erased. This research contributes both a methodological framework for quantifying experiential distortion and conceptual insights into how power operates through the fragmentation of spatial memory and the systematic erosion of the right to inhabit public space.

"Of course it is happening inside your head, Harry,
but why on earth should that mean that it is not real?"

Professor Albus Percival Wulfric Brian Dumbledore, Headmaster of Hogwarts, Order of Merlin (first class), Supreme Mugwump, Chief Warlock and Grand Sorcerer

" ..But we should resist the temptation to assume
that since stories are stories, they are, in some sense,
unreal or untrue, for this is to suppose that the only
real reality, or true truth, is one in which we, as living,
experiencing beings, can have no part at all."

The temporality of the landscape,
Tim Ingold

To all the brave, bright Iranian women and men,
and Baba.

\Table of Contents

01 | Introduction: 13
Following the rabbit: the terrain of uncertainty -
1.1 Research question and significance 14
1.2 Core contribution 15
1.3 Methodology: Agent-based Modeling 15
1.4 Case study : the Faculty of Fine Arts 16
1.5 Thesis structure 16

02 | Context: 19
Surveillance and Spatial Constraint in Iran -
2.1 The gendered construction of public space 19
2.2 Evolution of surveillance infrastructure 21
2.3 Resistance and adaptation 22
2.4 Educational spaces as controlled environments 26
2.5 Personal Narrative: Experience of Surveillance 31
2.6 Spatial Fragmentation as Lived Reality 36

03 | Theoretical Framework: 39
Spatial Cognition Under Threat -
3.1 Defining Space: Dichotomies and the Borderland 39
3.2 Space as an Agent: The Terrain of Uncertainty 41
3.3 Spatial Cognition: Perception, Memory, and Planning 43
3.4 Memory Systems and Cognitive Maps 44
3.5 Adaptation, Memory Dynamics, and Fragmentation 49
3.6 Operational Bridge: From Theory to Simulation Variables 52

04 | Methodology: 55
Agent-Based Modeling -
4.1 The Agents 56
4.2 The Mechanisms 58
4.3 The Parameters 61
4..4 The Model: Simulating the Faculty of Fine Arts 62

05 | Results: 67
The Fragmented Landscape -
5.1 First scenario: Control, the privileged explorer 68
5.2 Second scenario: mildly brave explorer 71
5.3 Third scenario: risk-avoider 76
5.4 Fourth scenario: The brave 81
5.5 Conclusion and Discussion 86

06 | Discussion: 95
Design implications -
6.1 The Designer's Ethical Crisis: Defending Optionality 96
6.1.1 Mutual Visibility 96
6.1.2 Efficient Zoning 97
6.1.3 Functional Ambiguity 98
6.1.4 Network Redundancy 100
6.1.5 Material Complexity 101
6.1.6 Revisiting the corners 104
6.2 Synthesis: Architecture as Accumulation of Friction 104

07 | Conclusion: 109
The Fragmented Narrative -
7.1 The Mechanisms of Fragmentation 110
7.2 From Simulation to Lived Reality 111
7.3 Limitations and Future Directions 112
7.4 The right to the city 113

0B| Bibliography 117
0G| Glossary 121
0A| Appendices 125

\List of Figures

Cover *a close-up of map of memory and navigation- control scenario [Author]*

02/Context

Figure 2.1 *Defaced graves of women [DW Persian, Accessed 27/09/2025, https://p.dw.com/p/4Eb6F]*

Figure 2.2 *Elementary school students in Iran and obligatory hijab [ISNA, Accessed 27/09/2025, isna.ir/xkpW9]*

Figure 2.3 *Photos of University of Tehran [CAOI, Author]*

Figure 2.4 *Map of main campus of University of Tehran [Municipality of Tehran, Author]*

Figure 2.5 *Plan of Faculty of Fine Arts and peer-mapped surveillance system [Author]*

Figure 2.6 *Map of individual safe spots for discretionary time [Author]*

Figure 2.7 *a-e. Parallel Sections of the Faculty [Author]*

Table 2.1 *The general materials of the Chastity and Hijab Bill [Zarbighalehhammami et al., 2025, Radio Farda, 2023]*

03/Theoretical Framework

Figure 3.1 *Patchy cognitive structure [Author]*

Figure 3.2 *The fragmentation of ToU and the options it offers [Author]*

Figure 3.3 *Isovist and Local chart [A. Mallot, 2024]*

Figure 3.4 *The remembered local charts and their egocentric orientation [Author]*

Figure 3.5 *Diagram of the cognitive map [A. Mallot, 2024]*

Figure 3.6 *The loop of decision-making in spatial behavior [A. Mallot, 2024]*

Table 3.1 *Operational bridge: translating theoretical framework into operational ABM concepts. [Author]*

04/Methodology

Figure 4.1 *The loop of main commands in the simulation [Author]*

Figure 4.2 *The prototype of the model [Author]*

Figure 4.3 *Translation and simplification of Faculty to environment patches in Netlogo [Author]*

Figure 4.4 *Map of individual safe spots for discretionary time [Author]*

Table 4.1 *The variables of environment patches [Author]*

Table 4.2 *The variables of threat agents [Author]*

Table 4.3 *The variables of protagonist [Author]*

Table 4.4 *Parameters of the simulation [Author]*

Table 4.5 *Translation of Faculty of Fine Arts into simulation patches [Author]*

05/Results

Figure 5.1 *Four scenarios of fragmentation based on 4 individual characters. [Author]*

Figure 5.2 *The map of memory and navigation in Control Scenario. [Author]*

Figure 5.3 *The map of detail memory in Control Scenario. [Author]*

Figure 5.4 *The map of memory and navigation in Mildly Brave Scenario. [Author]*
Figure 5.5 *The map of detail memory in Mildly Brave Scenario. [Author]*
Figure 5.6 *The map of avoidance memory in Mildly Brave Scenario. [Author]*
Figure 5.7 *The map of memory and navigation in Risk-avoider Scenario. [Author]*
Figure 5.8 *The map of detail memory in Risk-avoider Scenario. [Author]*
Figure 5.9 *The map of avoidance memory in Risk-avoider Scenario. [Author]*
Figure 5.10 *The map of memory and navigation in Brave Scenario. [Author]*
Figure 5.11 *The map of detail memory in Brave Scenario. [Author]*
Figure 5.12 *The map of avoidance memory in Brave Scenario. [Author]*
Figure 5.13 *The fragmentation of memory at the end of the simulation for all of the scenarios [Author]*
Graph 5.1 *Exploration rate in control scenario [Author]*
Graph 5.2 *Graph Coverage in control scenario [Author]*
Graph 5.3 *Exploration rate in Mildly Brave scenario [Author]*
Graph 5.4 *Graph Coverage in Mildly Brave scenario [Author]*
Graph 5.5 *Percentage of avoided, active and never-visited areas over time in Mildly Brave scenario [Author]*
Graph 5.6 *Exploration rate in Risk-avoider scenario [Author]*
Graph 5.7 *Graph Coverage in Risk-avoider scenario [Author]*
Graph 5.8 *Percentage of avoided, active and never-visited areas over time in Risk-avoider scenario [Author]*
Graph 5.9 *Exploration rate in Brave scenario [Author]*
Graph 5.10 *Graph Coverage in Brave scenario [Author]*
Graph 5.11 *Percentage of avoided, active and never-visited areas over time in Brave scenario [Author]*
Table 5.1 *Control Scenario - Graph structure and navigation statistics over the simulation. [Author]*
Table 5.2 *Mildly Brave Scenario - Graph structure and navigation statistics over the simulation. [Author]*
Table 5.3. *Risk-Avoider Scenario - Graph structure and navigation statistics over the simulation. [Author]*
Table 5.4 *Brave Scenario - Graph structure and navigation statistics over the simulation. [Author]*
Table 5.5 *Comparative Table of all the scenarios. [Author]*

06/Discussion

Figure 6.1. *The decreased intensity of propagated stress by introducing public gaze to the design, or inducing stress for the threat agents [Author]*
Figure 6.2. *Maintaining security of assets while minimizing control [Author]*
Figure 6.3. *Sample of increasing friction through architectural design [Author]*
Figure 6.4. *Network redundancy as a tool to increase friction and optionality [Author]*
Figure 6.5 *The remembered landscape in varying scales based on intensity of detail memory [Author]*
Figure 6.6 *Detailed Plan of the Faculty of Fine Arts and identification of some of the proposed strategies used in its design [Author]*

07/Conclusion

Figure 7.1 *A young man breaking the security cameras used to identify and eventual arrest of protesters, during the national protests of Jan. 2026, Image enhanced slightly using Google Gemini to increase the resolution by the author [VahidOnline, Jan. 2026]*

01 | INTRO

Following the Rabbit
The terrain of uncertainty

Unconventionally, we are going to start by following a rabbit, foraging in a forest or over a hill at dawn. The rabbit is constantly aware of its body, its landscape. Eyes racing, ears alert, muscles ready to be strained for a run.

It is rare among humans. We have artificial environments, the cities, and the architecture, and all it means when one uses the word environment, the state, the structure, that gives us the privilege of wandering, rather unaware of the setting. This is spatial freedom, when every crossing presents equally probable future trajectories, and it is just the whim of the agent, the main character of her life, to choose the destination and action.

However, this primal state of spatial hyper-awareness is not an anthropological relic; it is the daily reality for subjects whose bodies are politicized, whose navigation is a constant negotiation of risk. The rabbit can be seen as an allegory for the subject of this thesis, the navigation of women under surveillance in Iran, where conformity is obliged and each navigational choice is a risk. From the rabbit to the constrained campus, the environment reveals itself as a terrain of uncertainty, a complex system of possibilities and risks.

The agent, the protagonist, is simultaneously creating this network and being affected by it. As time passes, she will learn which ways are easier on the eye or under the sole, and begin to inhabit. Adaptation is the very essence of the relationship between environments and agents. And, unlike what might sound, adaptation isn't solely a reaction/response, but is also an act, a design. The environment gains a sort of responsiveness and adaptation as it is the constant infrastructure of life, acting as an agent itself. Each place holds a value to the eye of the observer, associated with what has been experienced there, and what it has to offer.

To search for the foci of uncertainty and the representation of the landscape they can build is a site-specific search, which, for this thesis, focuses on the surveillance system. Under threat and extreme conditions, the adaptation is much more visible, as it should happen fast and is enabled by specific, identifiable threats. The internalized space becomes a distorted representation of the omniscient, Euclidean structure. A new internal map is constructed not just from physical layouts but from risk anticipation, cultural coding, and accumulated experience.

1.1 Research Question and Significance

In such threatening environments, familiar public spaces are no longer accessible in ordinary ways. They become fragmented, stripped of their meanings, and ultimately stolen from those who once laid claim to them through daily use. This erosion targets not just liberty, but memory. When control and surveillance distort these spaces, the disruption is not merely physical. It ruptures the psychological foundation of belonging—the right to inhabit, remember, and shape one's environment. Consequently, the core question guiding this research emerges:

How does surveillance fragment spatial perception and memory?

1.2 Core Contribution

This thesis argues that surveillance operates not merely through direct physical control but through the cognitive corruption of spatial memory, leading to the psychological erasure of public space. By integrating spatial cognition theory with computational modeling, this work offers two primary contributions:

1. Empirical Insight: Identifying the "breaking points" where persistent, threat-induced avoidance behavior collapses an agent's functional cognitive map, leading to measurable spatial constraint.
2. Methodological Framework: Providing a repeatable Agent-Based Model (ABM) framework for operationalizing abstract spatial cognition theories (e.g., dual memory, latent learning) to quantify and visualize subjective experiential distortion.

1.3 Case Study: The Faculty of Fine Arts

It is often not possible to study the mentioned environments directly, especially in cases of high surveillance or danger. One such case exists in Iran, where women's presence in public space is persistently shaped by systematic threats.

The site selected for this thesis is the Faculty of Fine Arts of the University of Tehran, chosen primarily for methodological reasons. It serves as a controlled urban laboratory that reduces the complexity of the study:

1. Contained Complexity: It represents an isolated urban block, allowing us to eliminate external factors like public street traffic.
2. Modellable System: The contained environment and known administrative structure allow for a precise definition of the surveillance system and its associated stress propagation, which is essential for the model's parameters.
3. Auto-Ethnographic Focus: The model parameters and observed patterns are derived from the single agent's auto-ethnographic account, transforming the challenge of external validation into a study of experiential congruence.

Using this case, this thesis develops a model of the fluid terrain of uncertainty under top-down spatial control, producing visualizations of fragmentation of an individual's spatial memory.

1.4 Methodology: Agent-Based Modeling

The medium of this exploration is an agent-based simulation—a tool that allows us to observe how space is shaped by stress-informed navigation. This simulation is not a literal reconstruction, but a research-informed abstraction of space-making and movement under surveillance and uncertainty.

Building on the theoretical framework, the simulation models decisions made by the agent navigating space under conditions of:

1. Limited Perception: The agent only receives cues from a local Peripersonal/ Perception Zone.
2. Dual Memory: An internal structure that stores both spatial detail and a secondary layer of avoidance memory (based on reinforcement and latent learning).
3. Stress Propagation: Environmental risk is operationalized as a decay function tied to surveillance sources.

In the following sections, the theoretical foundations are discussed in greater depth, leading to the construction and application of this rigorous, subjective space-making platform.

1.5 Thesis Structure

This thesis proceeds through an integrated investigation that moves from contextual framing and theoretical development to computational modeling and back to critical interpretation.

Chapter 2: Context: Surveillance and Spatial Constraint in Iran

Establishes the political and social reality of surveillance targeting women in public spaces and justifies the selection of the Faculty of Fine Arts as a methodologically optimal, contained urban proxy, detailing its physical characteristics and the observed subjective avoidance patterns.

Chapter 3: Theoretical Framework: Spatial Cognition Under Threat

Lays the intellectual foundation by synthesizing philosophical concepts of space, landscape theory (Ingold, Gibson), and cognitive processes (Tolman, Mallot). This chapter defines how threat necessitates the modeling of dual memory and adaptive, risk-averse navigation.

Chapter 4: Methodology: Agent-Based Modeling

Justifies the choice of ABM for subjective research and details the rigorous operationalization of the theory: defining the agent's inner state (avoidance threshold), the environmental inputs (Stress Propagation), and the Avoidance

Threshold rule that governs fragmentation. Furthermore, details the translation of the physical site into the simulation, including the lockdown of all subjective parameters and scenario design.

Chapter 5: Results: The Fragmented Landscape

Presents the simulation outcomes, providing a quantitative analysis of route fragmentation and the identification of the "breaking point" where the agent's viable options collapse due to threat accumulation.

Chapter 6: Discussion: Design Implications

Critically interprets the results, focusing on the cognitive erasure of public space, the mechanisms of spatial oppression revealed by the model, and offering design implications for architecture in response to surveillance.

Chapter 7: Conclusion

Summarizes the key findings and contributions, reiterating the value of the methodological framework and suggesting future research directions.

02 | CONTEXT

Context
Surveillance and Spatial Constraint in Iran

Following the 1979 Islamic Revolution and the establishment of theocratic rule, the Iranian state developed a comprehensive framework of systematic oppression targeting women. This framework operates through constitution, local regulations, and reinforcing social norms, transforming enforcement from a purely state function into a societal mechanism that recruits not only state officers but also ordinary citizens and family members as enforcers.

At the foundation of this system lies mandatory hijab. Far from being merely a dress code, hijab serves as the architectural element enabling a broader discriminatory structure. By making women "distinguishable" and "marked," it replaces individual identity with gender category, applying universally from seven-year-old schoolgirls required to veil for school attendance to images of women on tombstones, which must be covered even in death. [Figures 2.1, 2.2]

2.1 The Gendered Construction of Public Space

The Iranian religio-patriarchal system constructs public space itself as fundamentally threatening to women, justifying their exclusion unless they submit to male-dictated boundaries. As Tony El-Hage observes, this construction operates through "quasi-religious discourses concerning moral and sexual risk and purity" that render public space taboo for women. Transgression through movement outside "traditional" domains means women risk not only physical danger but social ruin. [T. El-Hage, 2002]

This framing positions women simultaneously as vulnerable (requiring protection through seclusion) and dangerous (potential sources of moral corruption). The contradiction reveals the system's actual function: not protection of women, but control through constructed vulnerability.

The foundation of this structure is in the definition of women simultaneously as threatening and threatened. This construction perpetuates an aggressive environment as "natural" and "unquestionable" while labeling the female body as a source and symptom of deviance. Unregulated female sexuality threatens patriarchal social order precisely because it represents autonomy—the capacity for self-determination that undermines hierarchical control.

The veil operates as a physical manifestation not only of religious belief but of ideological subjugation. It serves as "a constant reminder to individual women and to society at large of prohibited places and spaces," simultaneously marking women as different and signaling their restricted access to public life.

Docile Bodies and Disciplinary Power

Foucault's concept of "docile bodies" sheds light on how this system functions. Women's bodies become objects invested with power—marked, trained, disciplined to "emit signs" of compliance. The woman must discipline her own body to avoid punishment, creating what Foucault terms "a carceral texture" where surveillance "assures both the real capture of the body and its perpetual observation." [M. Foucault, 1977]

Like Bentham's Panopticon, the surveillance system induces a state of conscious and permanent visibility that assures the automatic enforcement of power. The perfection of power renders its actual exercise increasingly unnecessary—women become self-policing agents, suspects, and offenders simultaneously. The uncertainty of observation (Am I being watched now? Will this camera identify me? Will that shopkeeper report me?) creates constant

Figure 2.1 *Defaced graves of women, or "rectification of women's graves with unveiled pictures on them" [DW Persian]*

Figure 2.2 *Elementary school students in Iran and obligatory hijab [ISNA]*

vigilance that normalizes restriction as self-protection.

2.2 Evolution of Surveillance Infrastructure

The enforcement of hijab began in the 1980s with the introduction of morality police patrolling streets with the authority to arrest women violating the dress code. This initial system relied on visible, mobile surveillance—patrols in vehicles or on foot, concentrated at key navigation points: main squares, public transport stations, pedestrian thoroughfares, parks, and shopping centers. [E.Mehrabi, 2022]

However, persistent women's rights movements, declining public support for state-imposed restrictions, and a widening chasm between society and state have forced the regime to intensify and modernize its surveillance apparatus. Recent decades have seen an expansion of control mechanisms:

1. Advanced Technological Surveillance: Security and traffic cameras equipped with facial recognition algorithms are being deployed increasingly to monitor public spaces, workplaces, and educational institutions, according to authorities. The state has attempted to incorporate private security cameras from residential complexes into this network, creating comprehensive visual coverage.
2. Distributed Enforcement: The surveillance burden has been redistributed beyond state agents. Commercial establishments face closure for serving unveiled women. Banks, government offices, hospitals, and even taxi services are conscripted as enforcement sites, with financial security positioned as a privilege granted only to compliant businesses. The general public becomes deputized enforcers, their own safety contingent on enforcing restrictions on others.

3. Punitive Escalation: Identification through surveillance now triggers consequences beyond immediate arrest: fines, property confiscation, bank account closure, trial, imprisonment. It is worth mentioning that the constant monitoring inside workplaces enforced by Harasat (the security unit of the universities) and Gozinesh (The admission unit, which controls the compliance of the student with state standards to admit them) units established throughout the massive government or semi-governmental part of the Iranian economy will certainly cause extra danger to women's already not-so-secure employment and livelihoods. In some cases, even their family members' jobs may be affected too. The threat operates not just at the moment of encounter but extends into the administrative and financial infrastructure of daily life.

The Chastity and Hijab Law: Panoptic Formalization

The 2023 "Bill to Support the Family through Promoting the Culture of Chastity and Hijab" [Table 2.1] represents the systematic formalization of panoptic control. The law operationalizes surveillance through multiple mechanisms:

1. Permanent Observation: Security forces function as the central tower in Foucault's panopticon—their presence signals constant potential surveillance even when direct observation is absent.
2. Technological Ubiquity: The "Hijab Enforcer" project deploys cameras and facial recognition as invisible observers, with citizens constantly under surveillance without knowing the exact time or place of monitoring.
3. Induced Self-Regulation: Awareness of potential surveillance compels automatic compliance. Women internalize dress codes not through conviction but through fear of punishment and social consequences—the essence of panoptic power.
4. Distributed Punishment: Penalties extend beyond individual women to business owners, service providers, and public spaces where unveiled women appear, creating social pressure for mutual surveillance and enforcement. [Zarbighalehhammami et al., 2025, Radio Farda 2023]

This represents evolution in surveillance tools while maintaining basic panoptic principles: continuous, unobtrusive observation generating fear of punishment as the mechanism of control.

2.3 Resistance and Adaptation

Despite—or because of—intensifying control, civil disobedience persists and grows. Women find ways to survive by adapting to threatening environments, but this survival exacts costs. Vast portions of public space become avoided:

Penalties of Improper Hijab in Public Spaces				
Women	Wearing revealing clothing or garments that expose other than face, feet (under the ankles), or hands (below the wrist).		First offence: Monetary fine/prison sentence Subsequent offenses: increased monetary fine/ prison sentence	
Men	Wearing revealing clothing or garments that expose under the chest, above the knees, or the shoulders.			
Penalties for Celebrities				
Prison sentence	Confiscation of 10% of total assets	Disqualification for work from 5 to 15 years	Ban on social media activities from 6 months to 2 years	
Penalties for Improper Hijab at Work				
Prison sentence	Payment of 2 to 4 months of income		Ban on Leaving the country from 6 months to two years	
Penalties on Promoting Improper Hijab and Freedom of Dress in Online and Virtual Platforms				
Prison sentence	Payment of 3 to 6 months of income	Ban on social media activities from 6 months to 2 years	Ban on Leaving the country from 6 months to 5 years	Removal of content/ confiscation of virtual account

Table 2.1. The general materials of the Chastity and Hijab Bill, as a sample of the legal framework of oppression against women in Iran [Zarbighalehhammami et al., 2025, Radio Farda, 2023]

certain subway stations, main squares, parks, and commercial centers. The process of adaptation leads to what this thesis terms a "broken, distorted version of the city for women"—not physically inaccessible but psychologically fragmented through accumulated experience of threat.

But this exchange is fundamentally coercive. It violates the most basic rights—to body, to place, to expression. The discrimination does not end with hijab but begins there, creating cascading restrictions that fragment society and enable constant, belittling control. Moreover, the line of "acceptable" compliance continuously shifts, for example, women considered properly veiled in the 1990s faced arrest in the 1980s. There is no stable point of sufficient submission; compliance generates further restriction.

Individualism as a Threat

Religious fundamentalism perceives women's individualism as dangerous precisely because it threatens communal identity. Enforced dress codes attempt to void individual personhood, presenting monotonous appearances that constantly remind subjects of their relationship to the Islamic state.

Yet this very erasure of difference generates resistance. Dressing alternatively becomes not merely personal expression but political challenge—an assertion of independent existence against a system designed to render women interchangeable, regulated bodies. The regime's attempt to eliminate individual identity paradoxically makes every assertion of individuality a political act. This dynamic of control through visibility extends beyond public streets into institutional spaces, where surveillance becomes systematized and architectural. [Abdmolae, 2014]

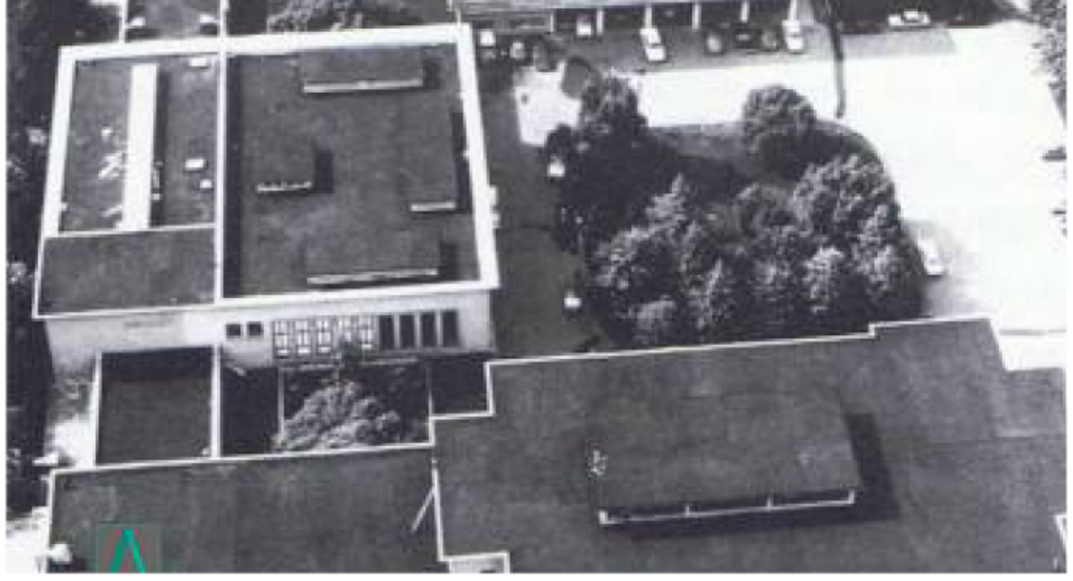

Figure 2.3 *Photos of University of Tehran [CAOI, Author]*
Figure 2.4 *Map of main campus of University of Tehran, entrance gates, and location of Faculty of Fine Arts [Municipality of Tehran, Author]*

University Confines
University Entrance Checks
Faculty of Fine Arts
Street Confines
Fences
Buildings
Water Features
Vegetation
M Subway Station

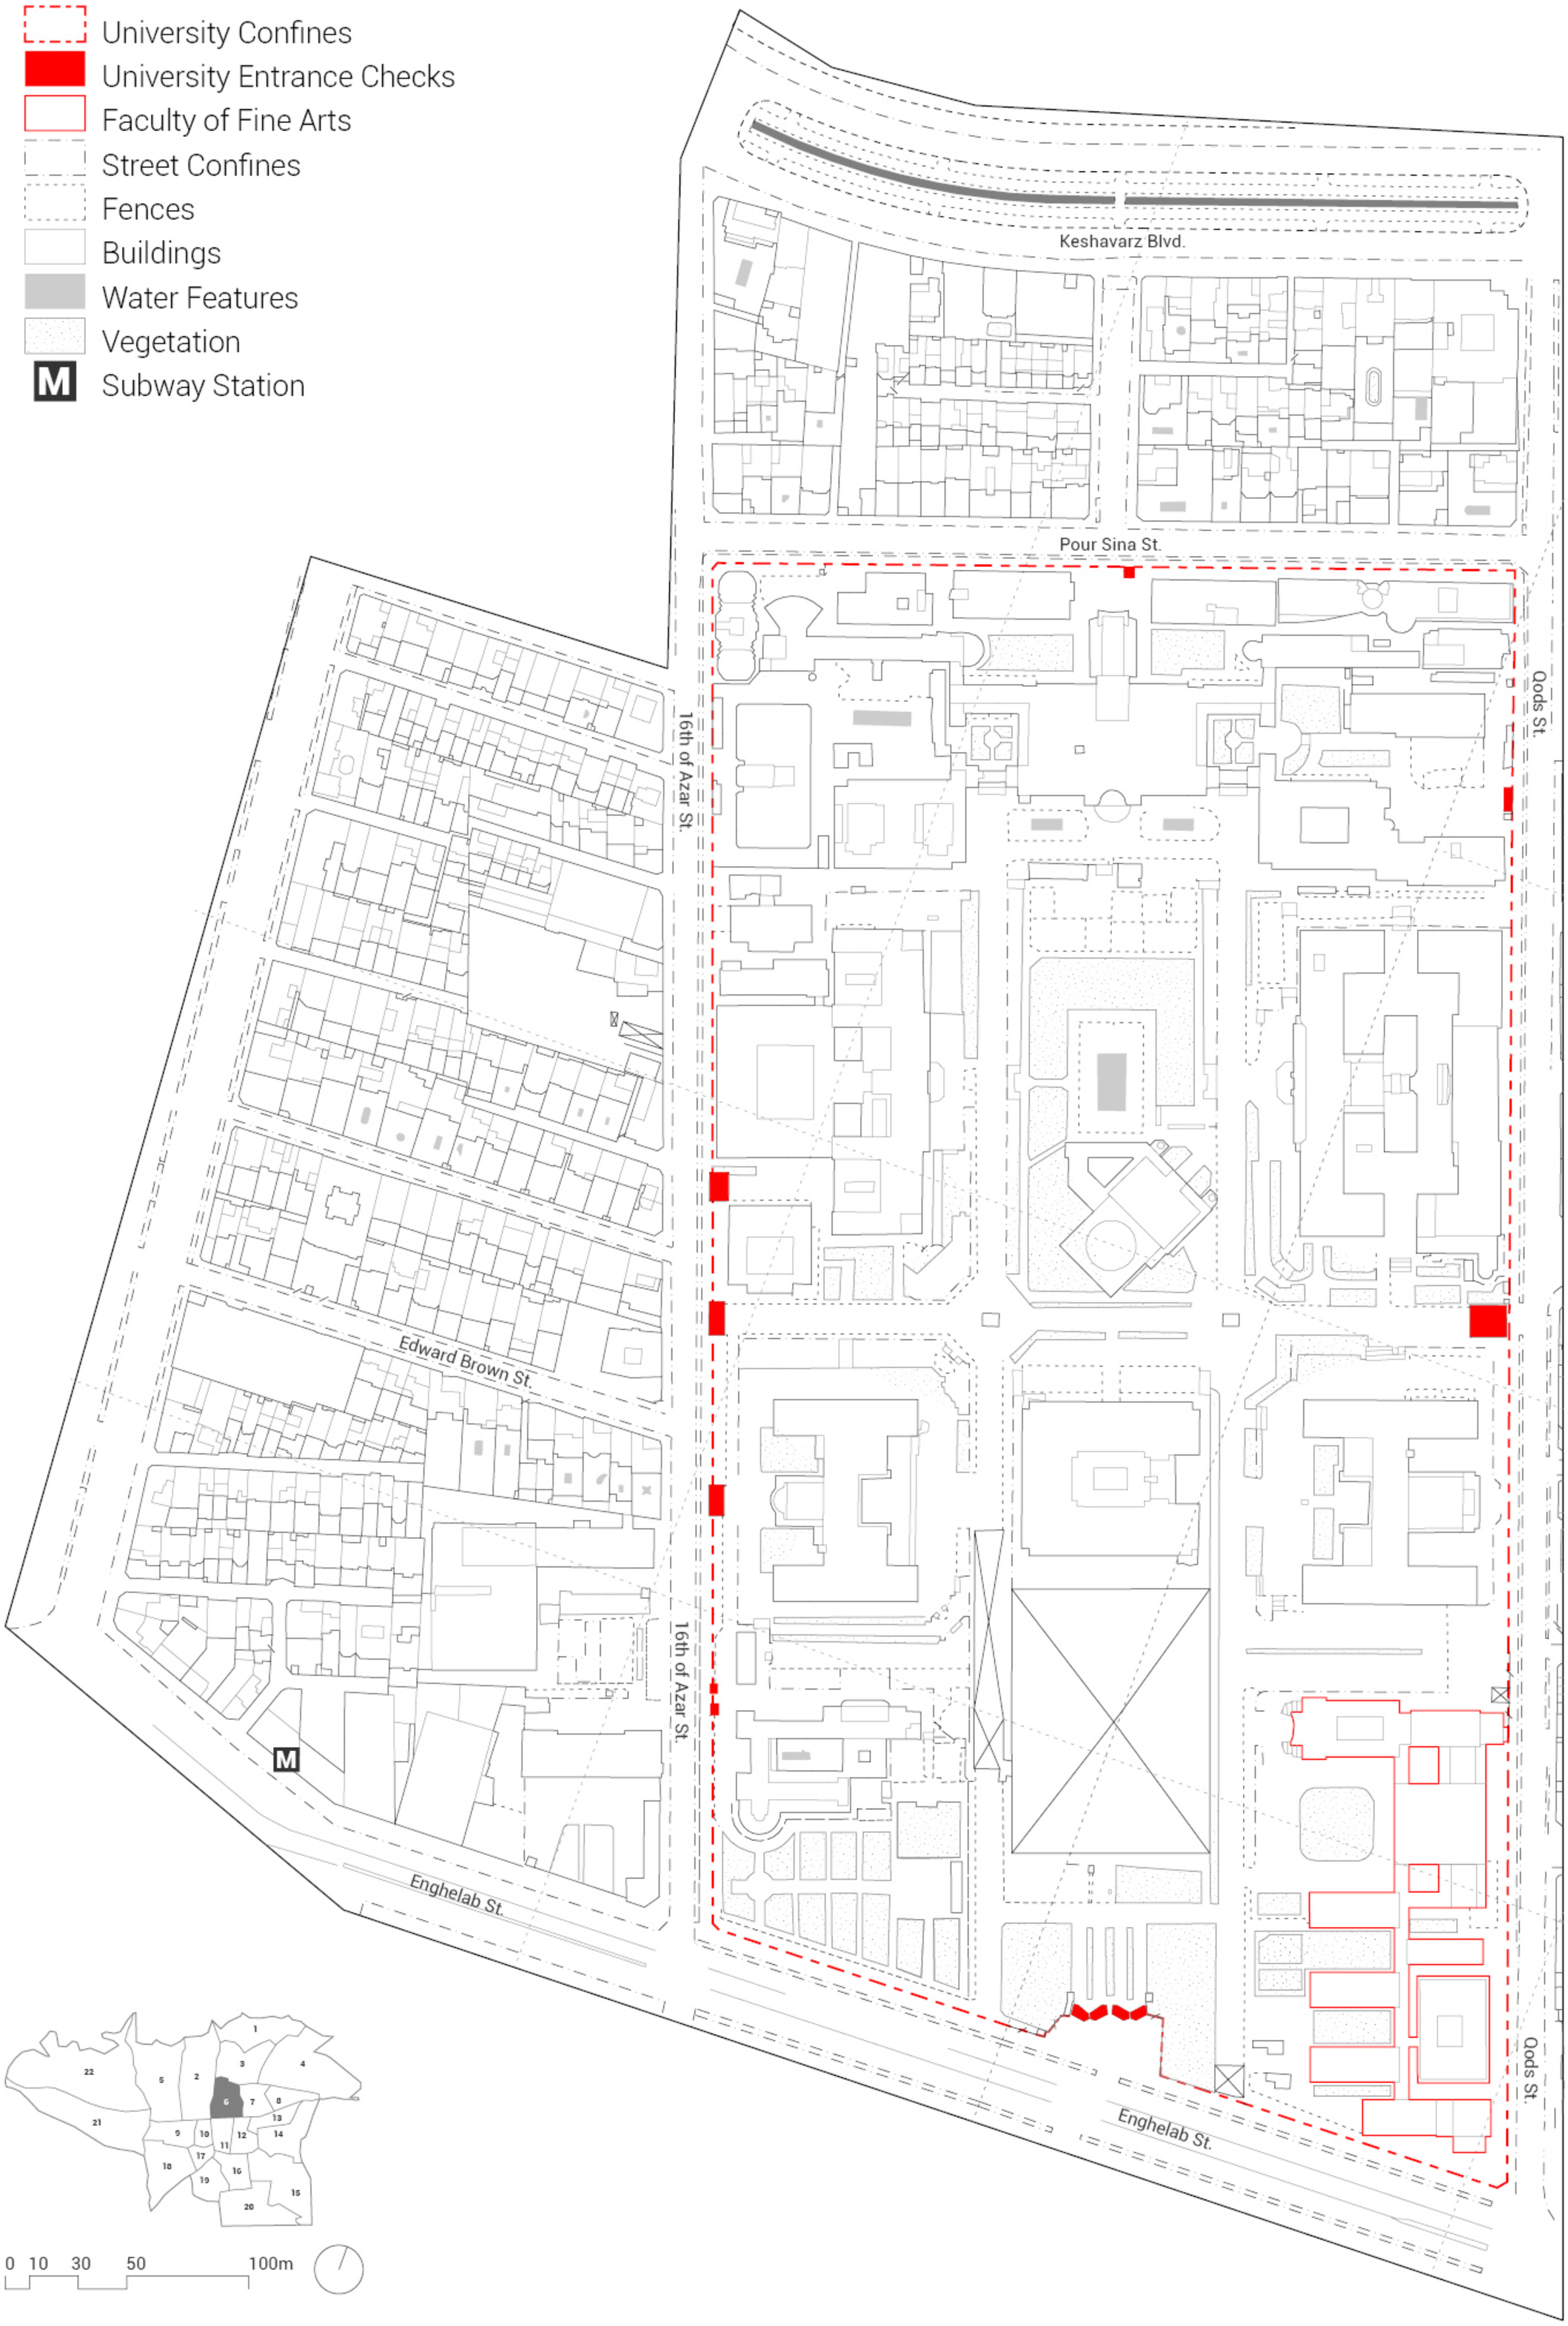

0 10 30 50 100m

2.4 Educational spaces as controlled environments

Historically, Iranian universities have played a crucial role in bolstering the state's authority and legitimacy. These institutions have fostered national identity, educated political leaders, and instilled deference for the state and authority among students. As formal institutions, universities in Iran, much like those elsewhere, have functioned as 'missionaries,' tasked with upholding and justifying the existing status quo through the dissemination of 'neutral' knowledge and scientific discourse. [P. Rivetti, 2020]

The universities in Iran are direct enforcers of the regime's ideology. They are part of the state structure, as they are founded, financed, and controlled by the government. Following the Islamic Revolution in 1979, under Khomeini's speech and orders, the control and purge of universities began, known as the Cultural Revolution. It began with the attack of Islamist militant forces on universities, the arrest of students, and the takeover of the premises. By order of the supreme leader, the High Council of the so-called Cultural Revolution was formed, and after two years of closure of higher education institutions, they reopened, fully integrated with the Islamic state. However, although this effort to redefine campuses as a matter of national security and neutralize

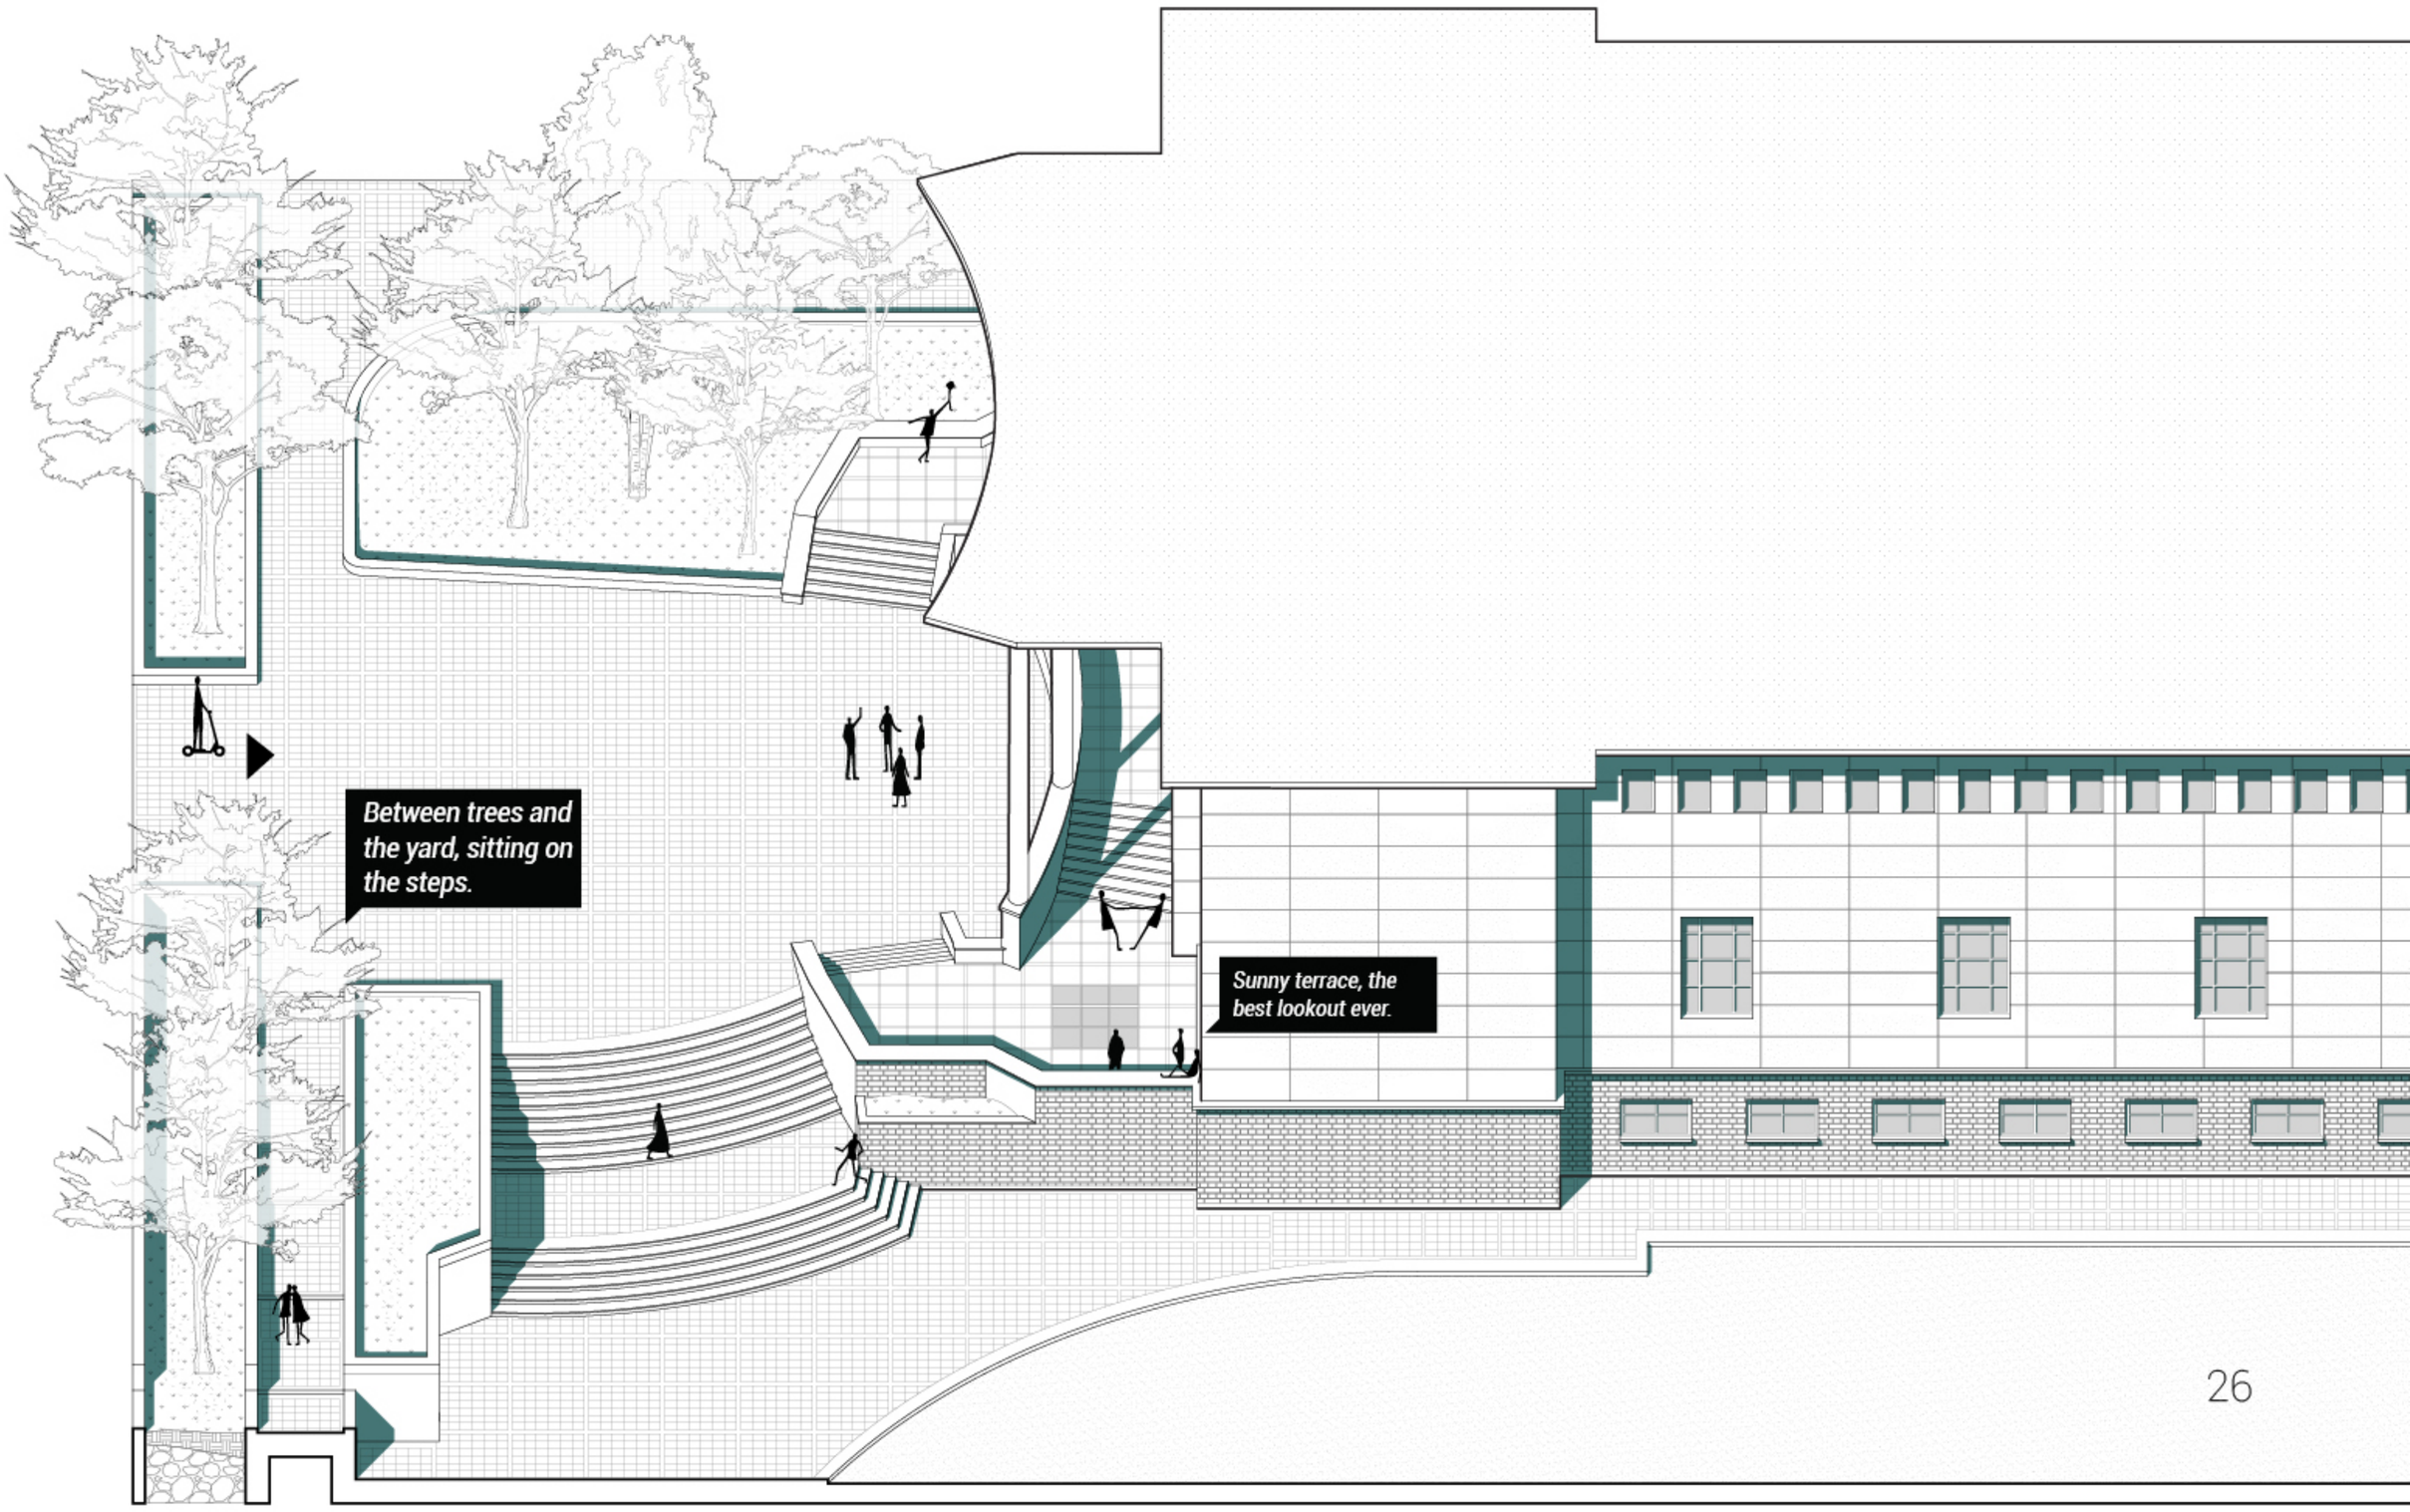

any sociopolitical unrest in universities was momentarily successful, over the following decades, the struggle for reclaiming campuses was revived. [S. Mojab, 2004]

As of now, regulating universities and student bodies happens in several ways, from admission regulations, obligatory religious and ideological courses being included in the curricula, penalties such as depreciation of the right to study because of political activities, ethnicity, or religion, surveillance and spatial regulation of campuses, and so on. As for the subject of this thesis, the same framework of controlling urban public spaces is applied to university campuses through the institutional surveillance system, entrance gates, university patrols, and security cameras. To enter the university, the student should present the student card to officers at the gender-segregated entrance gates, where, if the attire of the student is approved- also by the personal taste of the officer-, she can enter the university, where the surveillance continues.

A few additional regulations are also applied to the female students in addition to the general ones. For example, smoking cigarettes in university for women is forbidden, and mixed gender gatherings are frowned upon. However, universities usually have a simpler surveillance system imposed

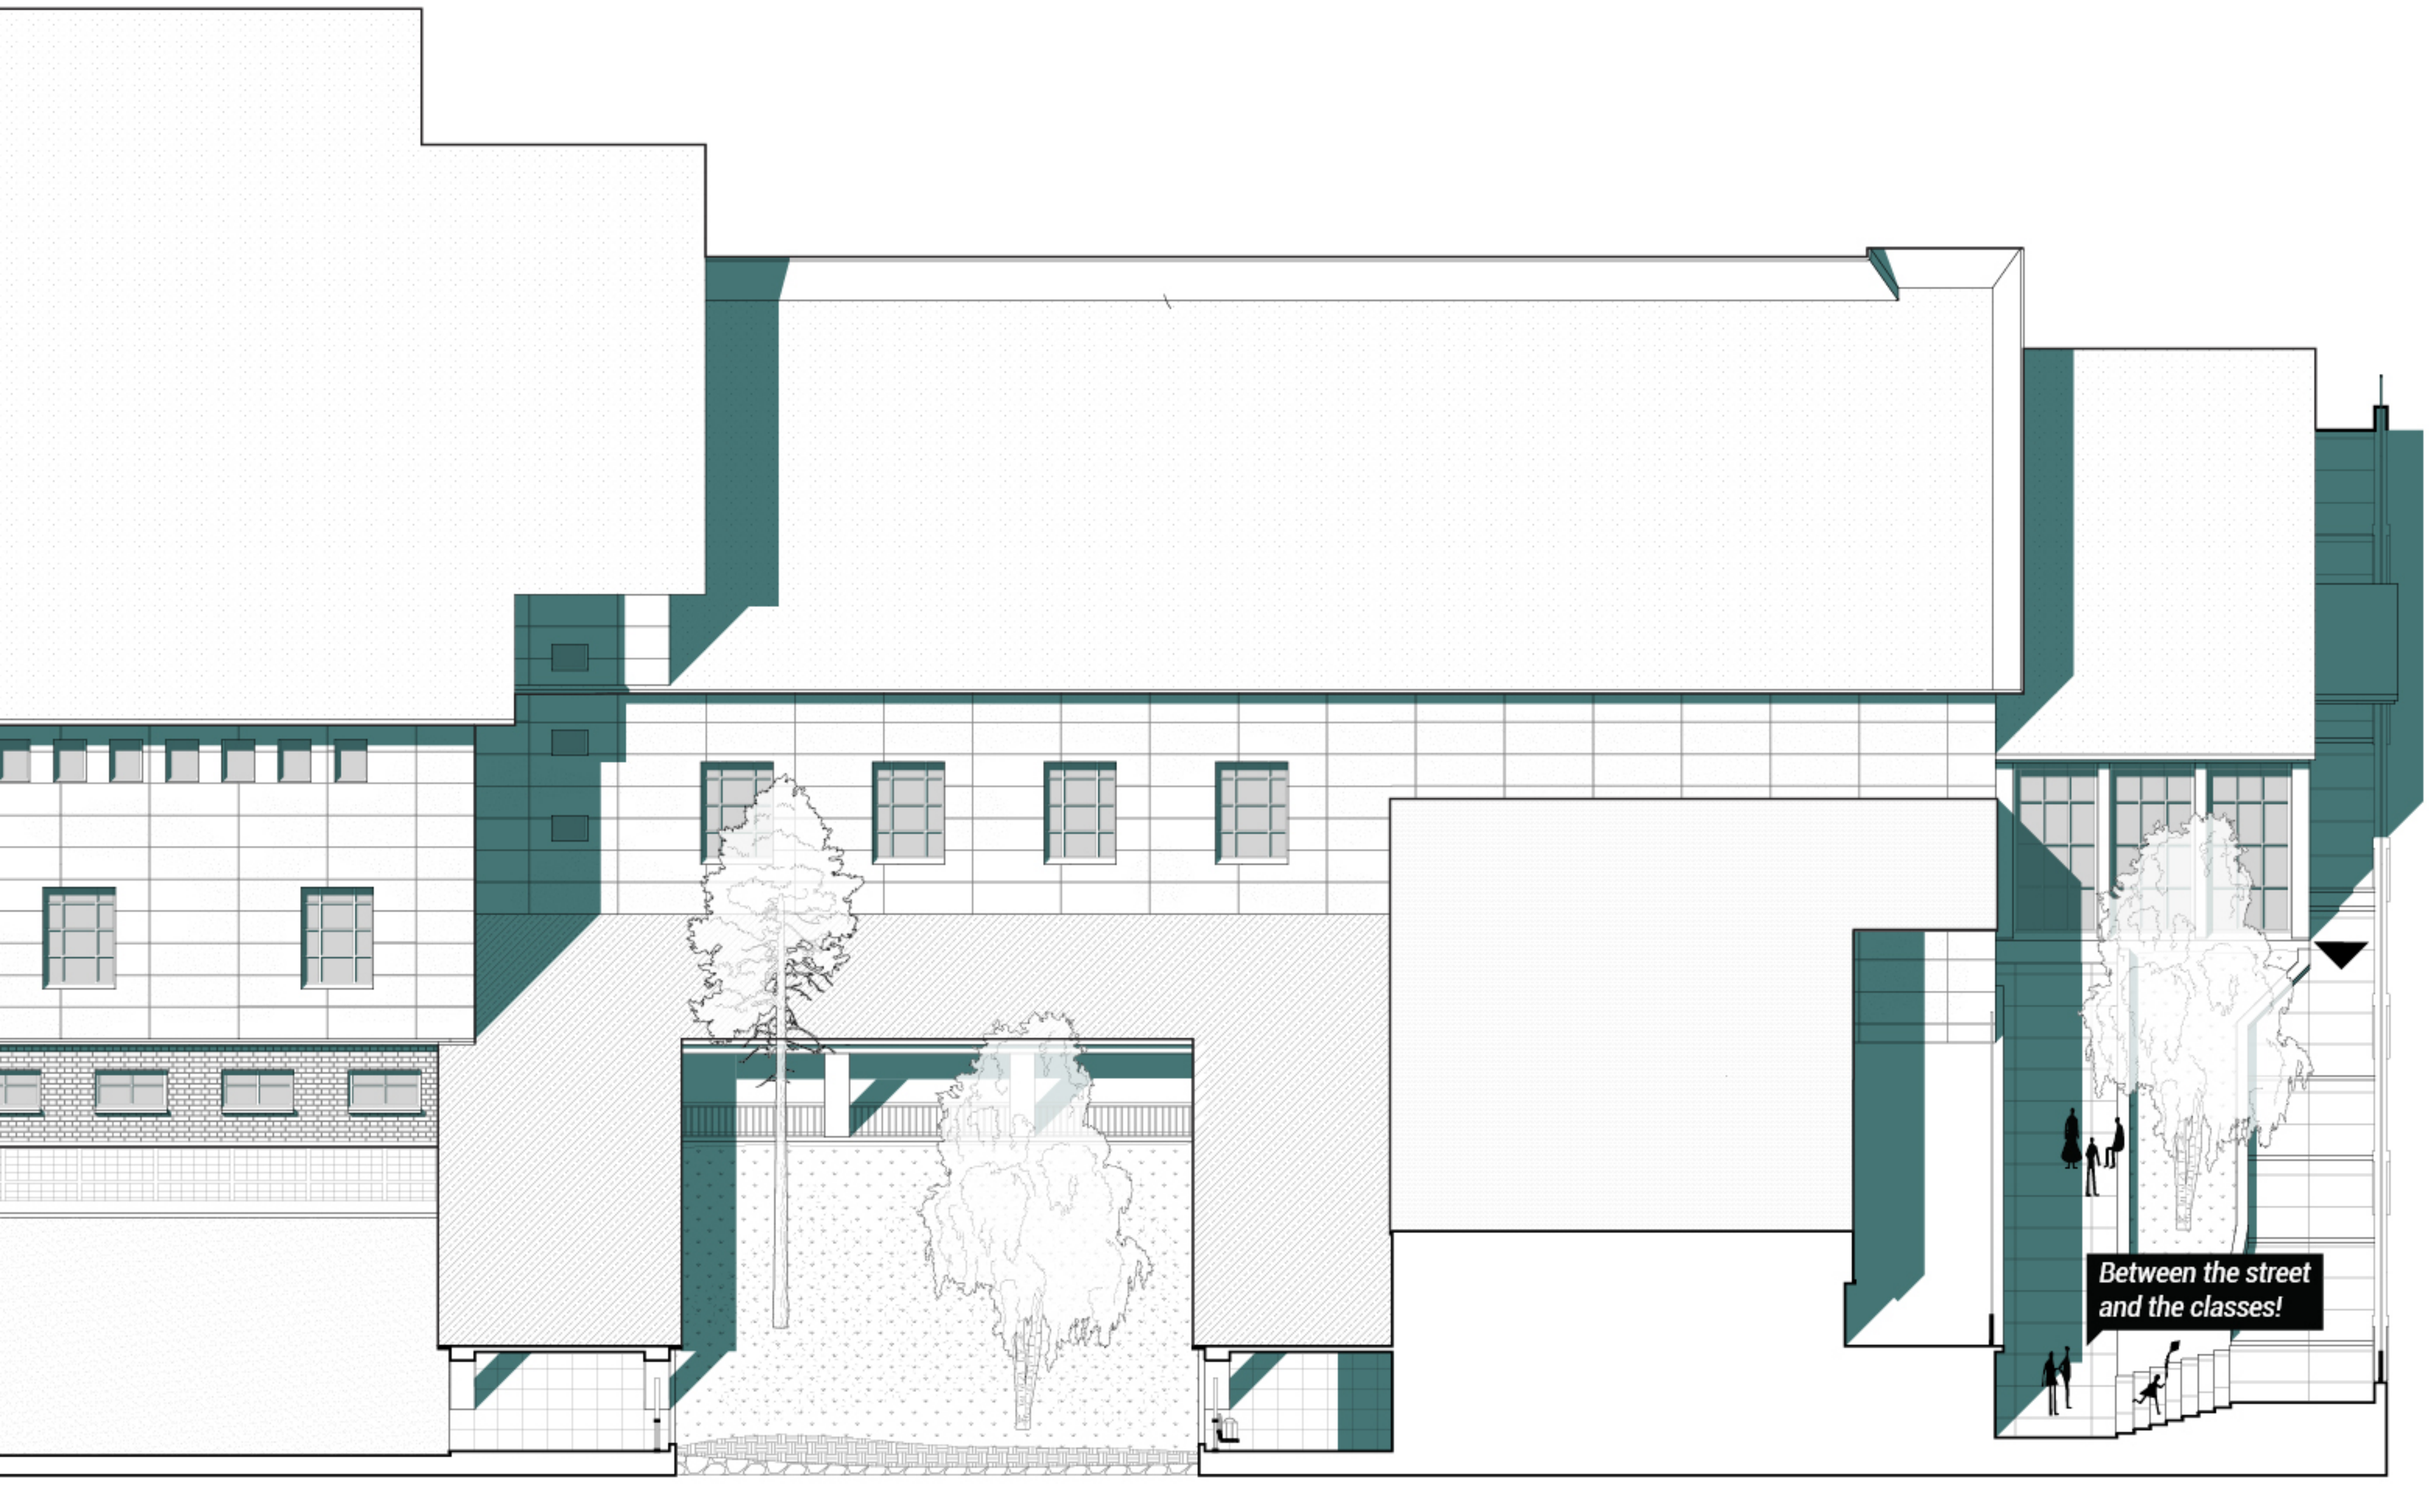

Watching everything and chillin'!

Lunch between trees on the grass!

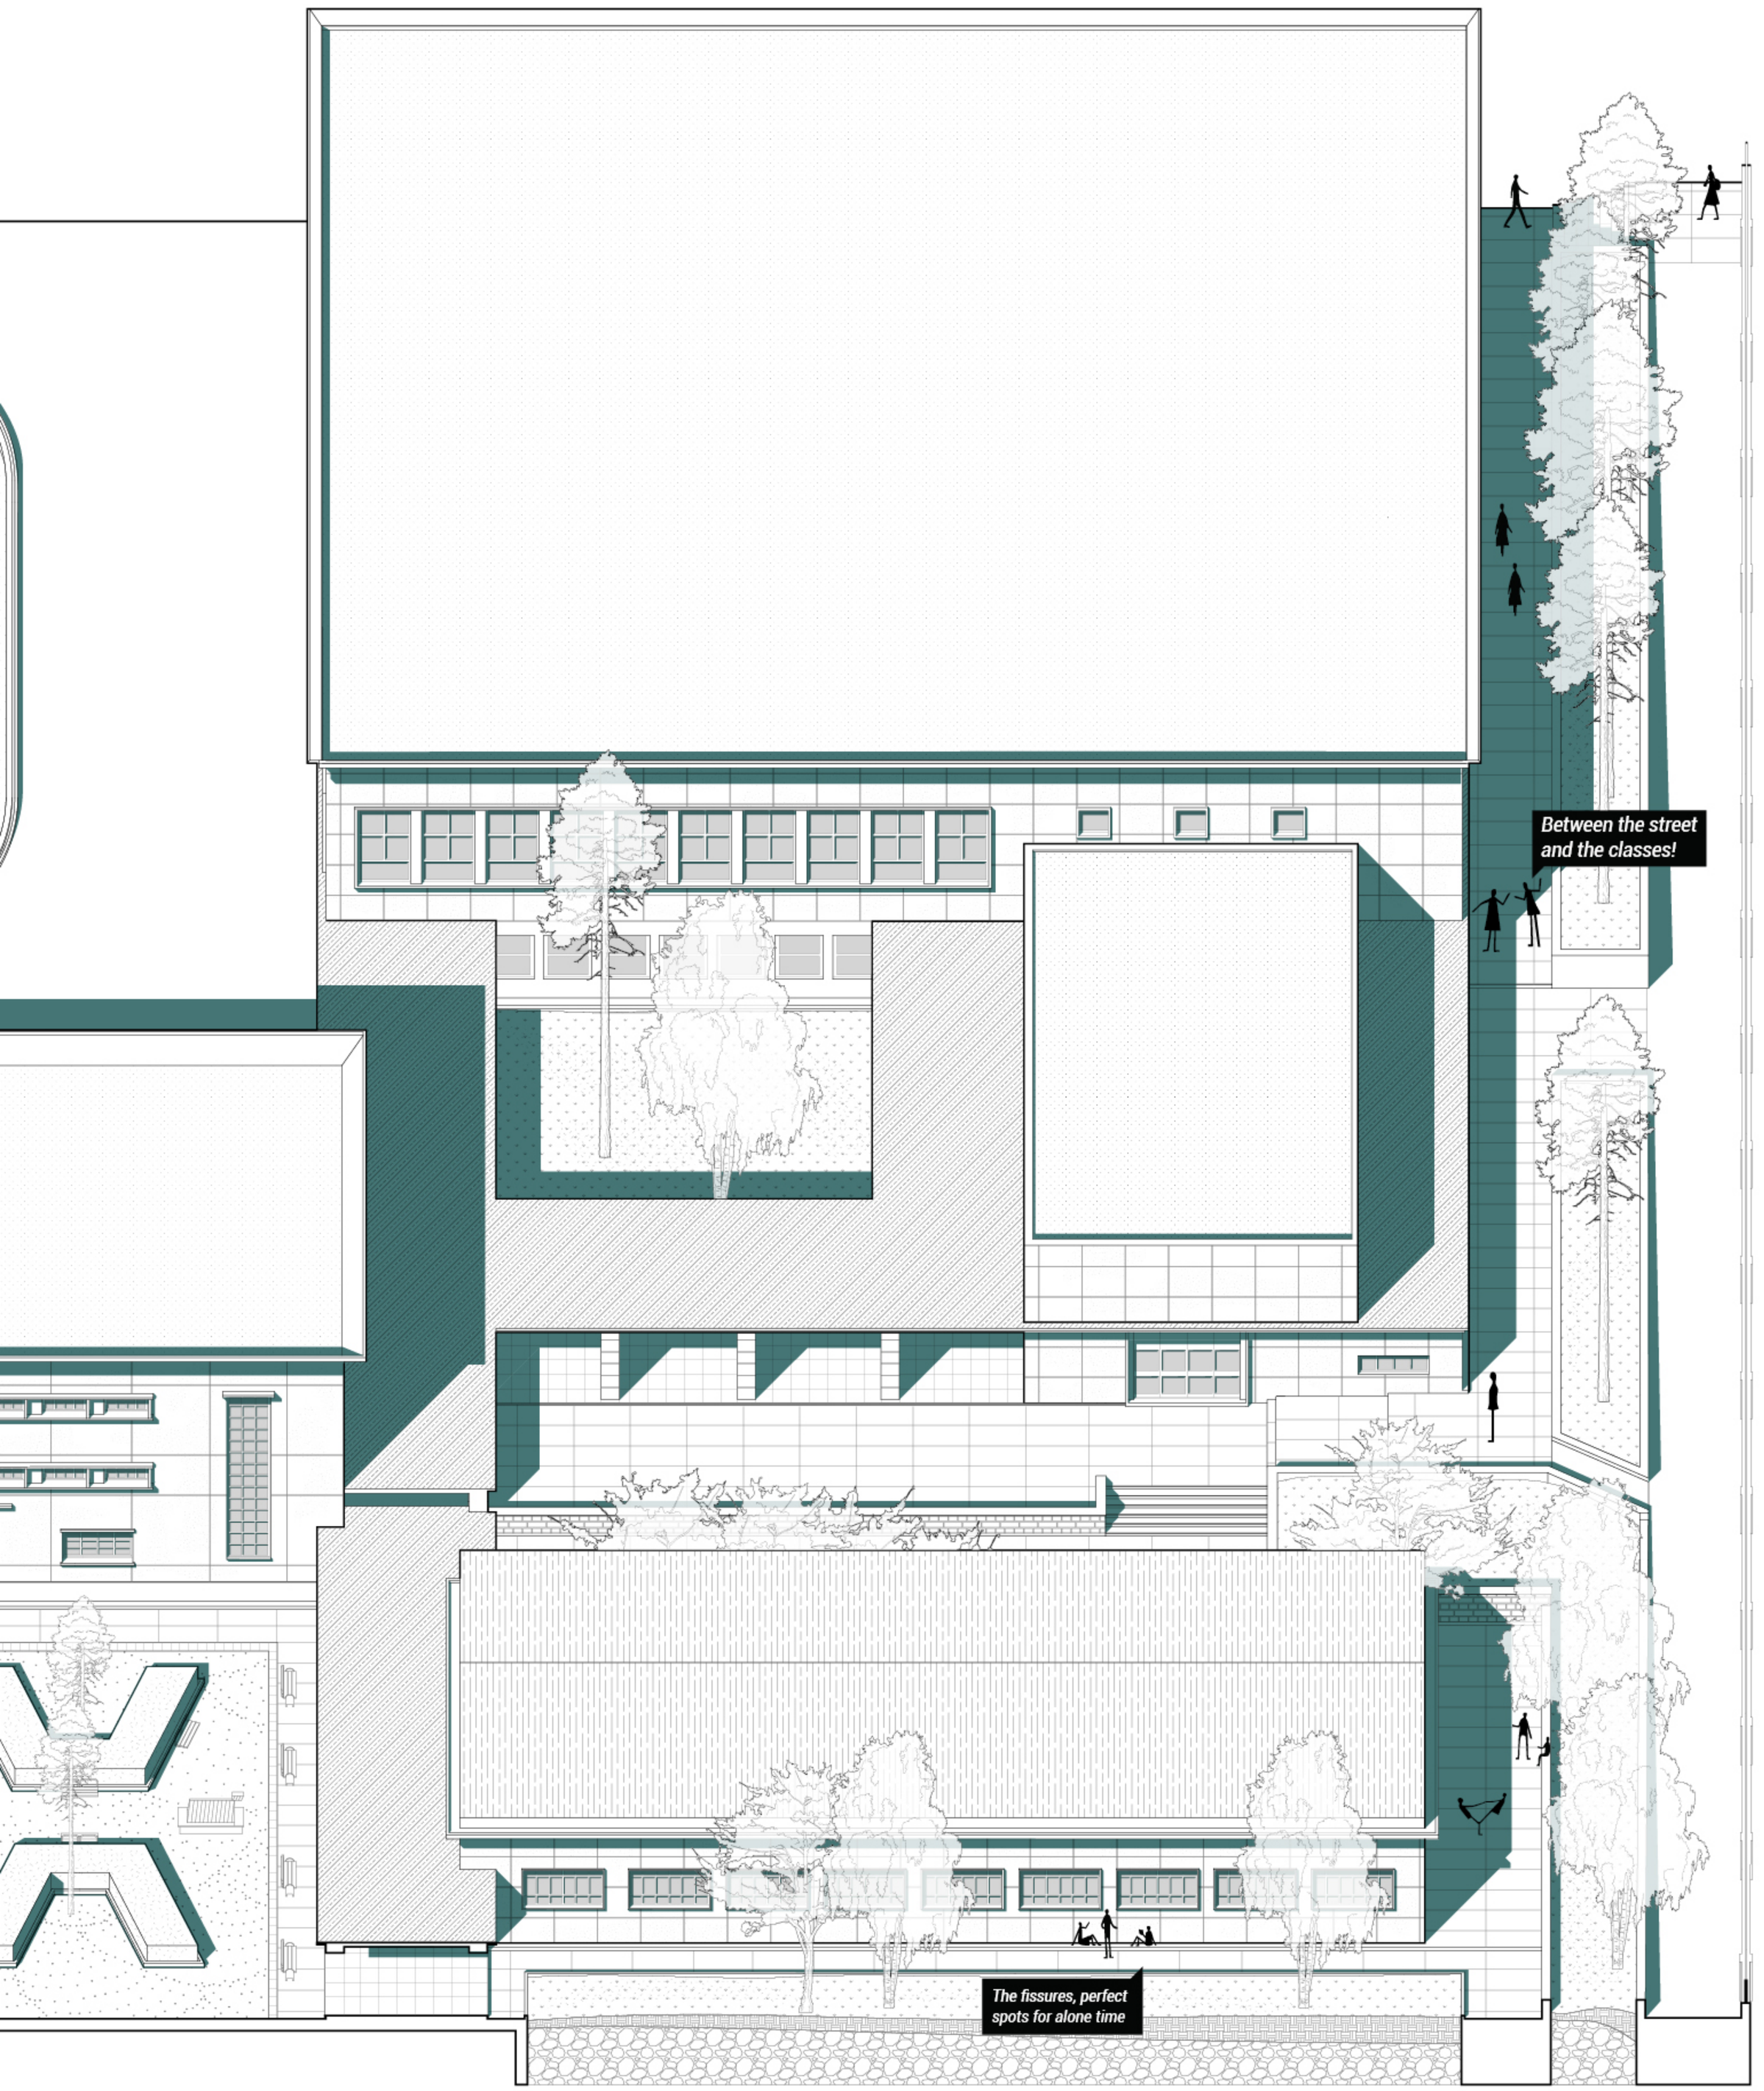
Between the street and the classes!
The fissures, perfect spots for alone time

by the educational institution. The student population offers much more demographic homogeneity in comparison to the general public, and the lack of more complex urban settings and factors such as traffic or different surveillance types enables a more reliable study opportunity for reducing foci of uncertainty to top-down systematic surveillance methods. In this regard, the Faculty of Fine Arts at the University of Tehran serves as a contained laboratory for studying these dynamics, a microcosm where surveillance systems, spatial constraints, and adaptive navigation intersect in observable, modellable ways.

The Faculty of Fine Arts

The Faculty of Fine Arts originated as the School of Fine Arts, established in October 1940 under commission of the Ministry of Culture and Arts, initially operating at Marvi High School. Shortly thereafter, its workshops were integrated into the University of Tehran (UT), [Figure 2.3, 2.4] one of Iran's oldest educational institutions, first relocating to the basement of the Faculty of Engineering in 1941. The school's transformation into a formal faculty began with its annexation as a higher school, culminating in its official designation as the Faculty of Fine Arts by law in April 1949, when it moved to its dedicated site within the university campus. This early phase saw rapid physical expansion, with the core complex designed by Roland Marcel Dubrulle between 1941

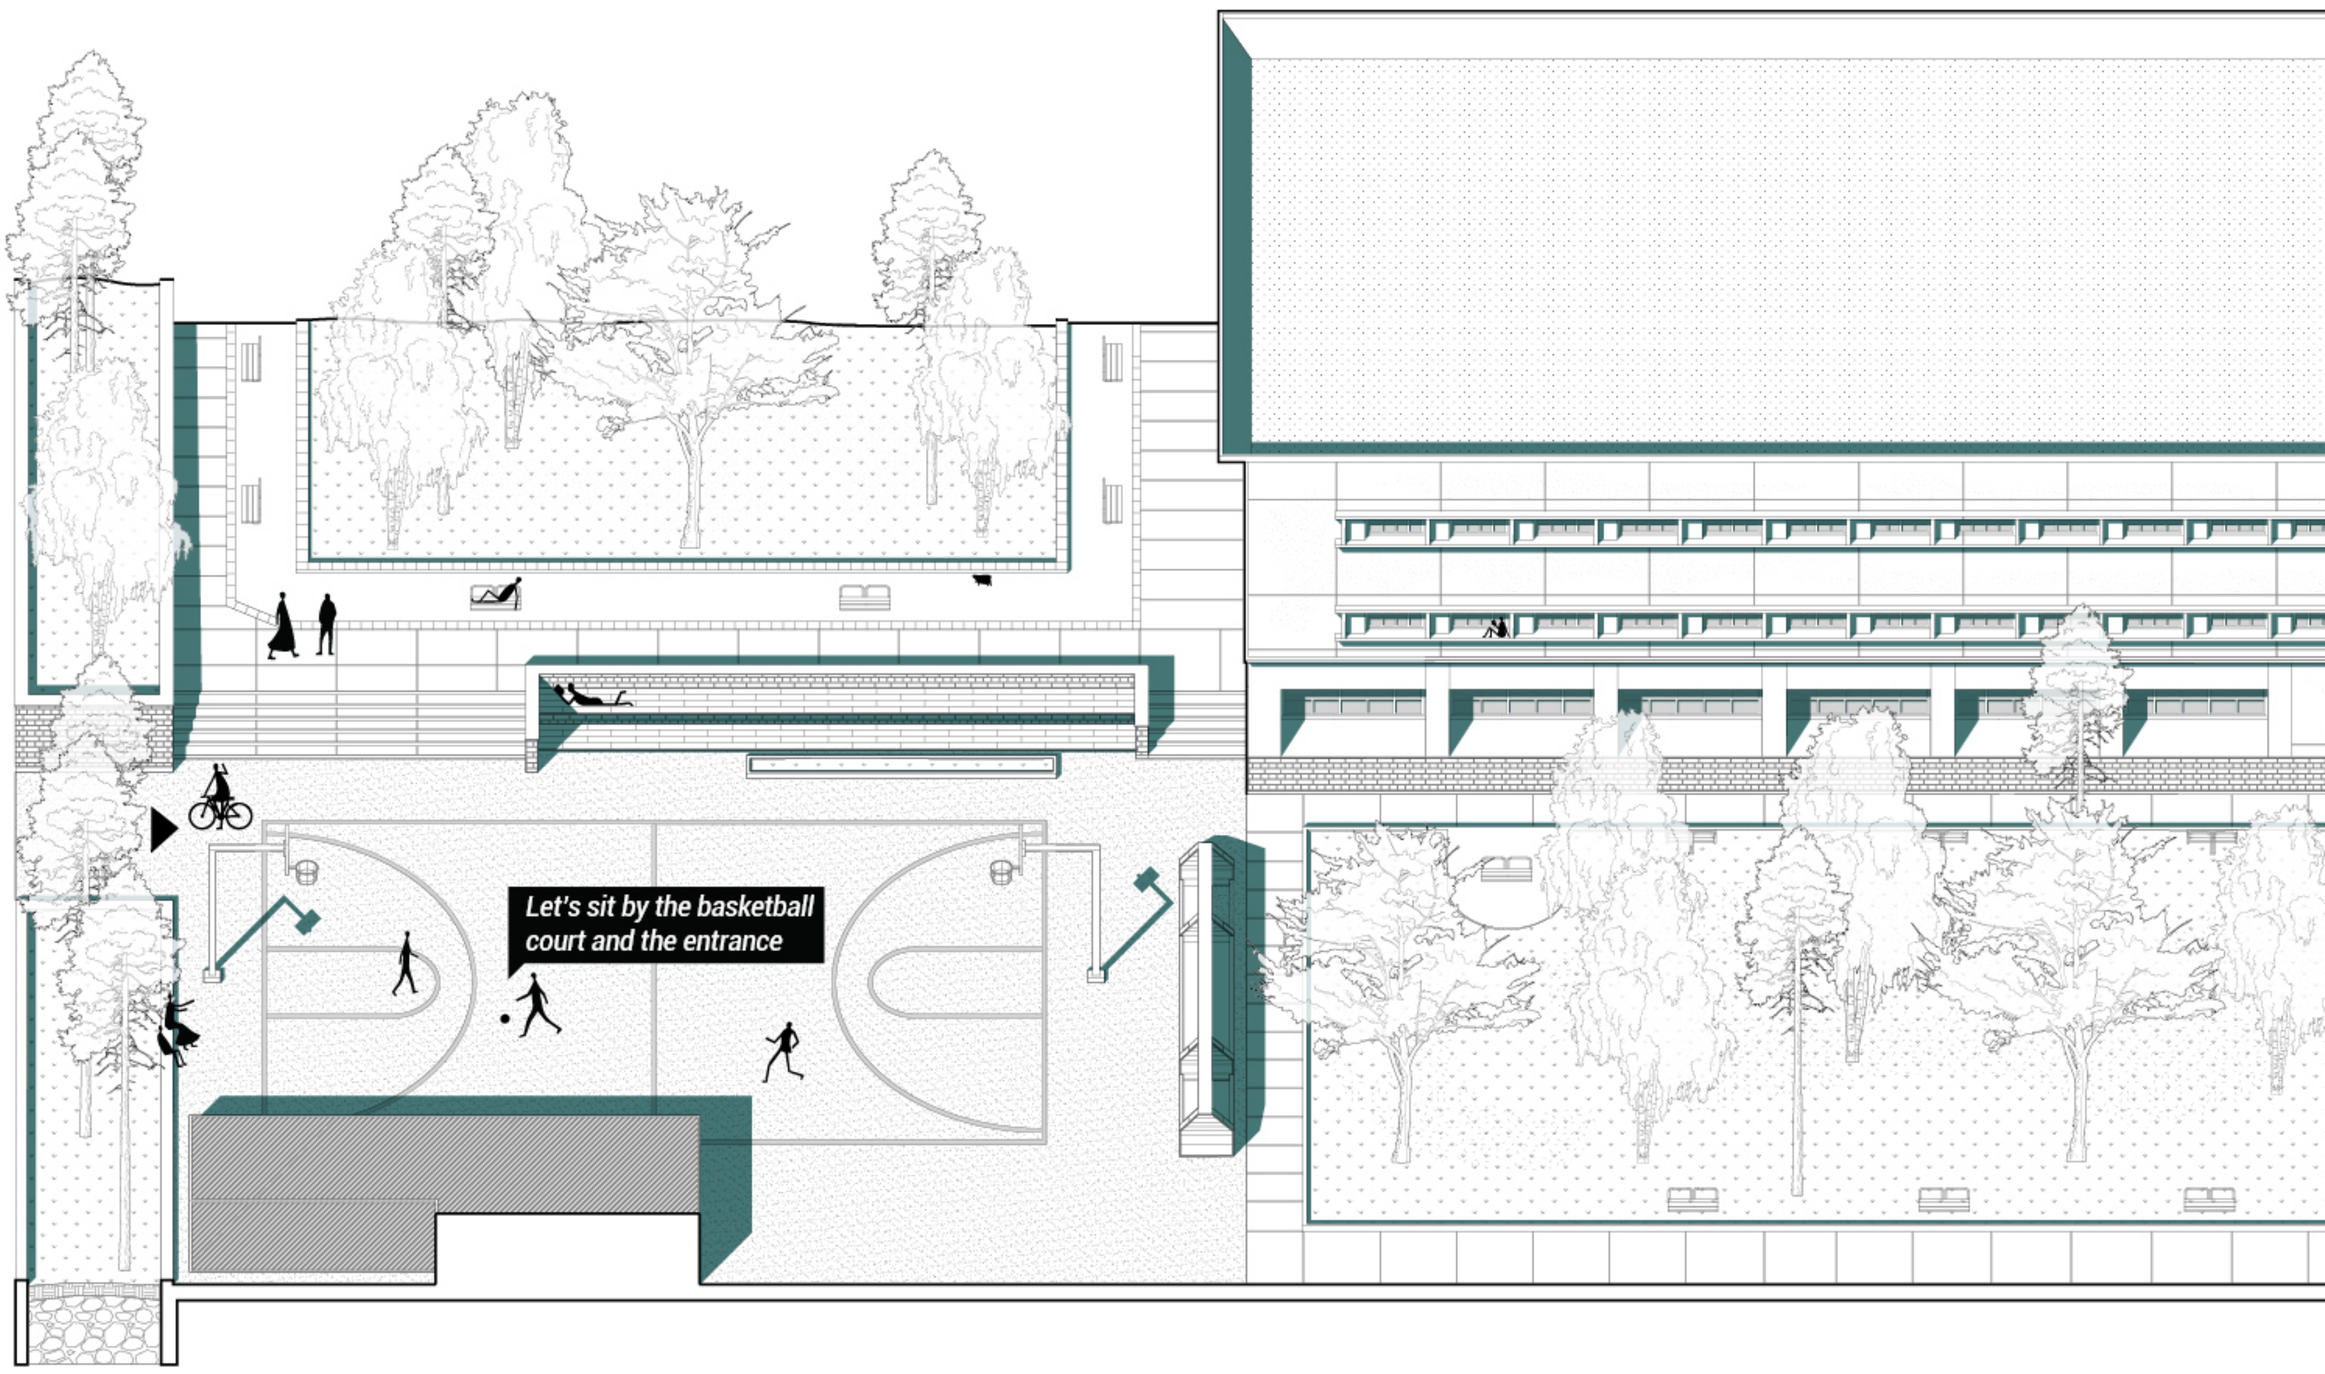

and 1948, evolving to include architecture studios, painting and sculpture workshops, exhibition spaces, and a theater. The main library and additional music and theater classrooms, completed in December 1966, were the final major additions.

The Faculty occupies a distinct rectangular site—approximately 100 by 200 meters—on the southeast side of the University of Tehran campus. Architecturally, it comprises eight interconnected blocks strategically organized around courtyards and green areas. This design utilizes open spaces to ensure natural light, ventilation, and open views throughout the complex, shown in parallel sections of figure 2.7. Key facilities include the amphitheater, main library, specialized studios and workshops, and administrative offices. The faculty is bordered by the Faculty of Literature to the north, Qods Street to the east, Enqelab Street to the south, and green spaces to the west, defining its position within the larger urban and university environment. [CAOI]

2.5 Personal narrative

The experience of surveillance is demonstrably widespread, necessitating collective and technological forms of adaptation. Citizens actively develop and utilize digital tools, mobile applications, social media, messaging platforms, to instantly report locations and details of surveillance equipment and patrols.

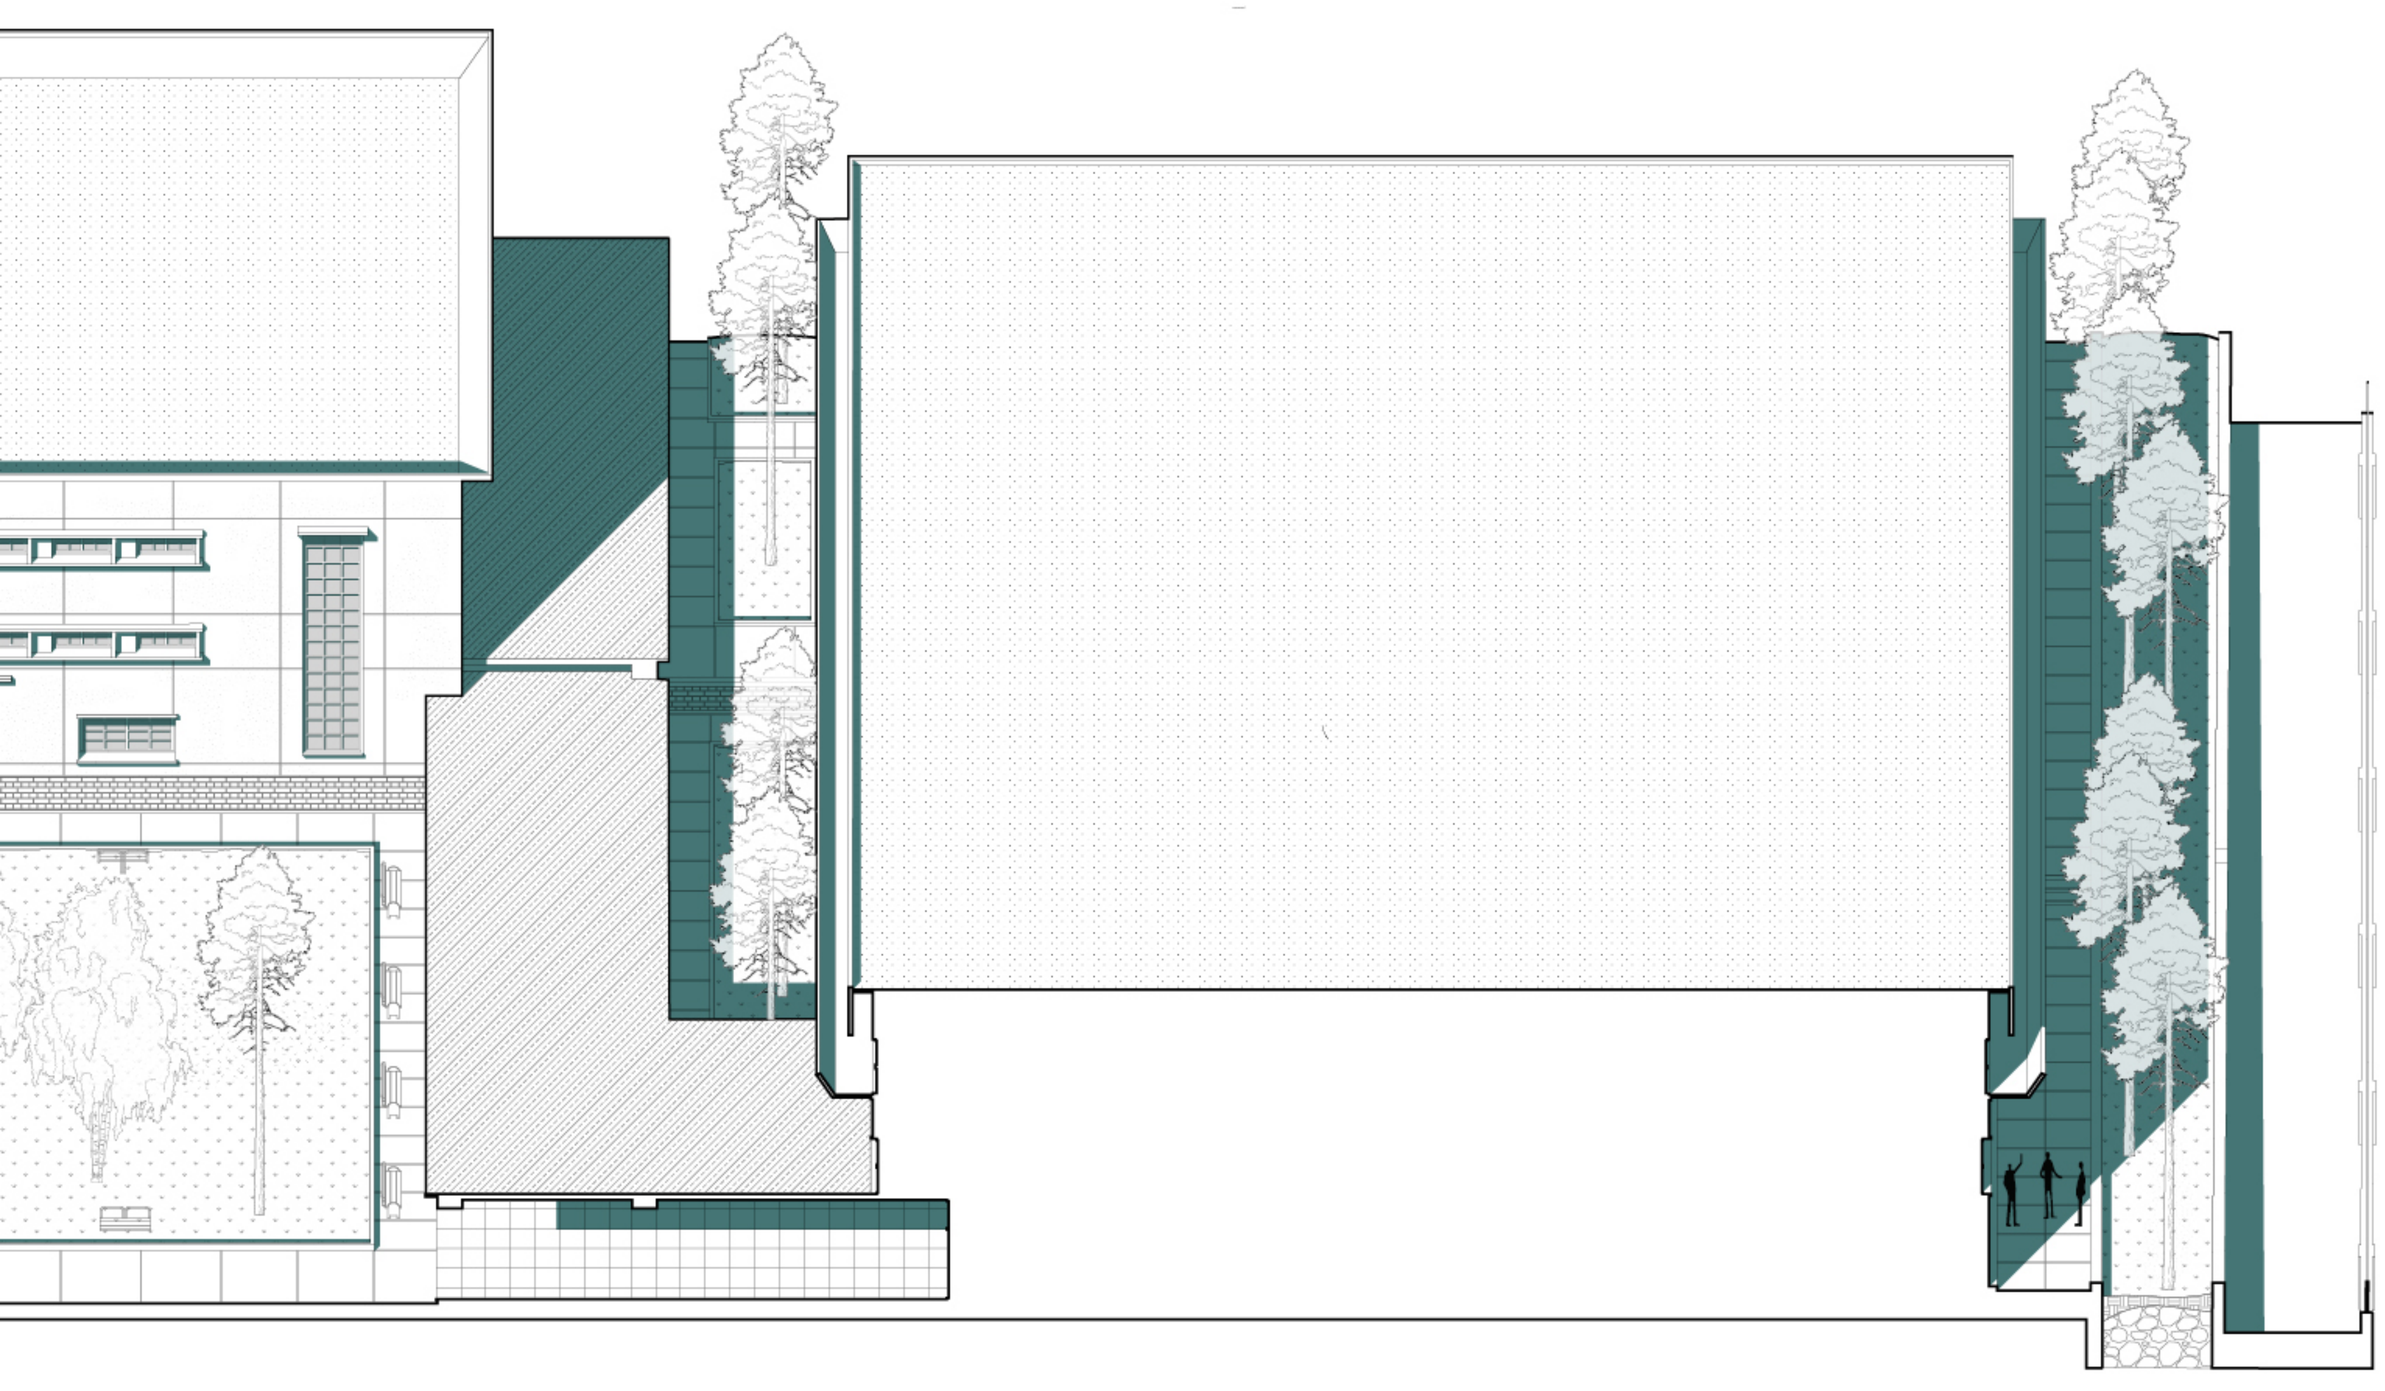

This digital coordination extends to immediate, personal warnings, where people advise each other in passing to avoid specific roads or junctions where patrols have been recently sighted. This indicates the necessity of a constant, real-time information network for navigating public space.

I studied in the Faculty of Fine Arts, Department of Architecture, from 2016 to 2021. Unfamiliar not only with the educational space but also with Tehran's urban setting, as it is not my hometown, I responded to this high-constraint environment by adopting systematic spatial avoidance. After my first year and several confrontations with morality police, I began avoiding main streets and junctions where the possibility of being seen was high. I stopped using subway stations, which, like educational spaces, are easily controlled with limited exits and entrances. I stayed home increasingly, avoided walking when possible, and deferred basic activities. I found comfort in a smaller world where stress was minimized. Eventually, I moved to Italy, cutting my country off the map.

The impact of this system generates internalized behavioral shifts that persist even after leaving the environment. Though I have lived in Italy for two years, I still sometimes wonder unconsciously if I am safe to be in the subway station, or if I am in the men's section of the bus, though no such division exists here. These moments, when I feel the impulse to change direction upon seeing men in suits or police officers, still remind me of feeling threatened or unwanted, making me acutely aware of my body as superseding my individuality. This unconscious self-awareness is a residue of the surveilled urban experience.

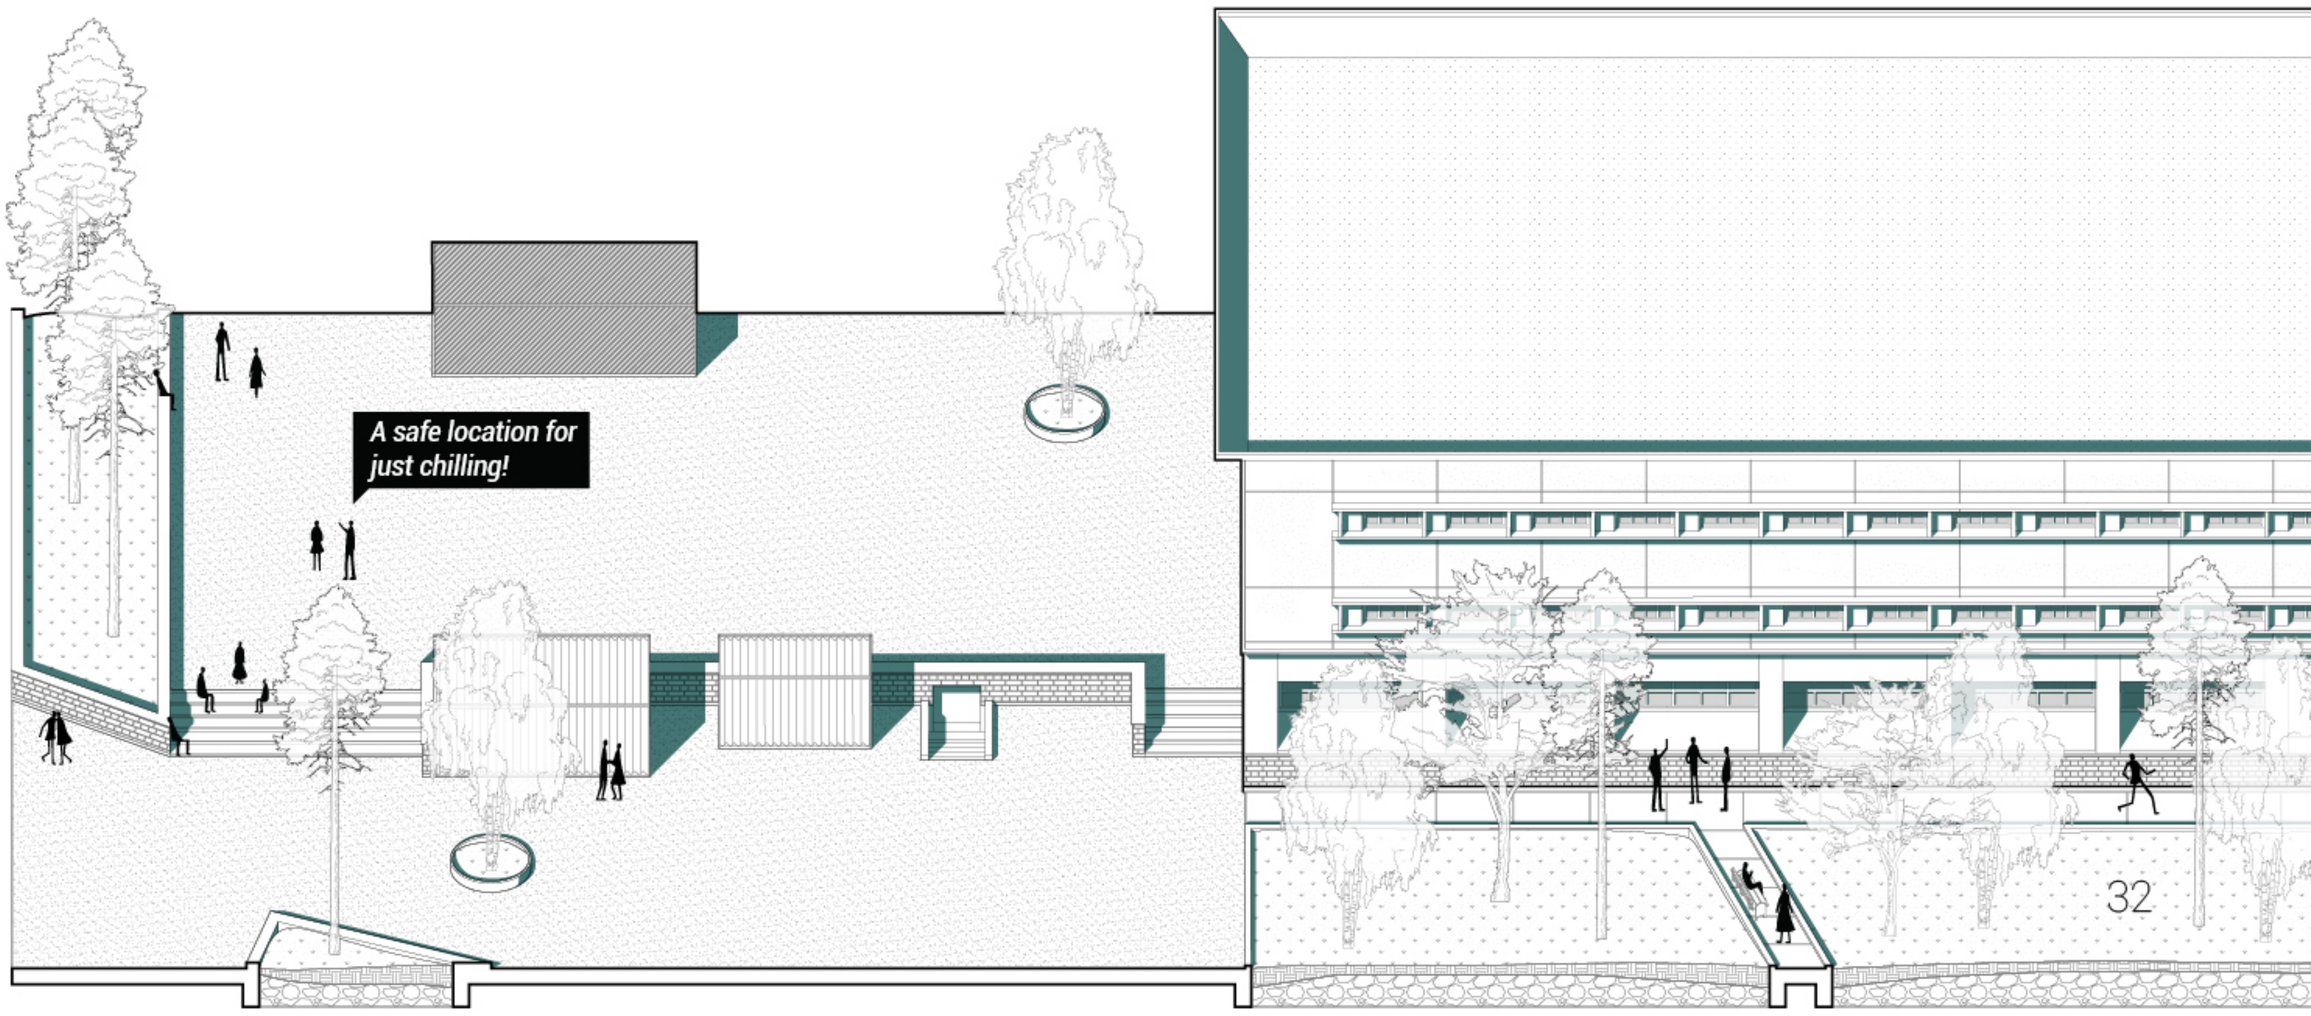

Most of the time in Iran, I experienced stress in public spaces: being asked to leave banks or stores, being denied subway tickets or services, or feeling my heart pound while passing morality police officers at junctions, main squares, or subway station entrances and exits. Retrospective discussions with Iranian friends and former colleagues revealed many shared narratives—the fear, humiliation, and instability felt during simple walks in the city, the withdrawal from public spaces, and the craving for free, unconscious wandering. These conversations crystallized the central question: What is our city?

This inquiry emerged from understanding that state-imposed constraints had fundamentally fragmented and distorted the urban landscape, replacing the right to free movement with a reality defined by calculated risk and spatial limitation.

Mapping Surveillance and Adaptation

On campus, over time, my friends and I, consciously or unconsciously, identified certain safe spots. Interestingly, these were not uniquely our choices; groups of students gravitated toward the same locations to eat, smoke, and talk. More notably, the most popular spots often lacked urban furniture, and we typically sat on steps designed around buildings. However, it must be noted that within the University of Tehran, oppression was more relaxed compared to urban space. Nevertheless, even the possibility of a confiscated student card, being deprived of study rights or compromise of them, and persistent fear was enough to induce spatial consciousness.

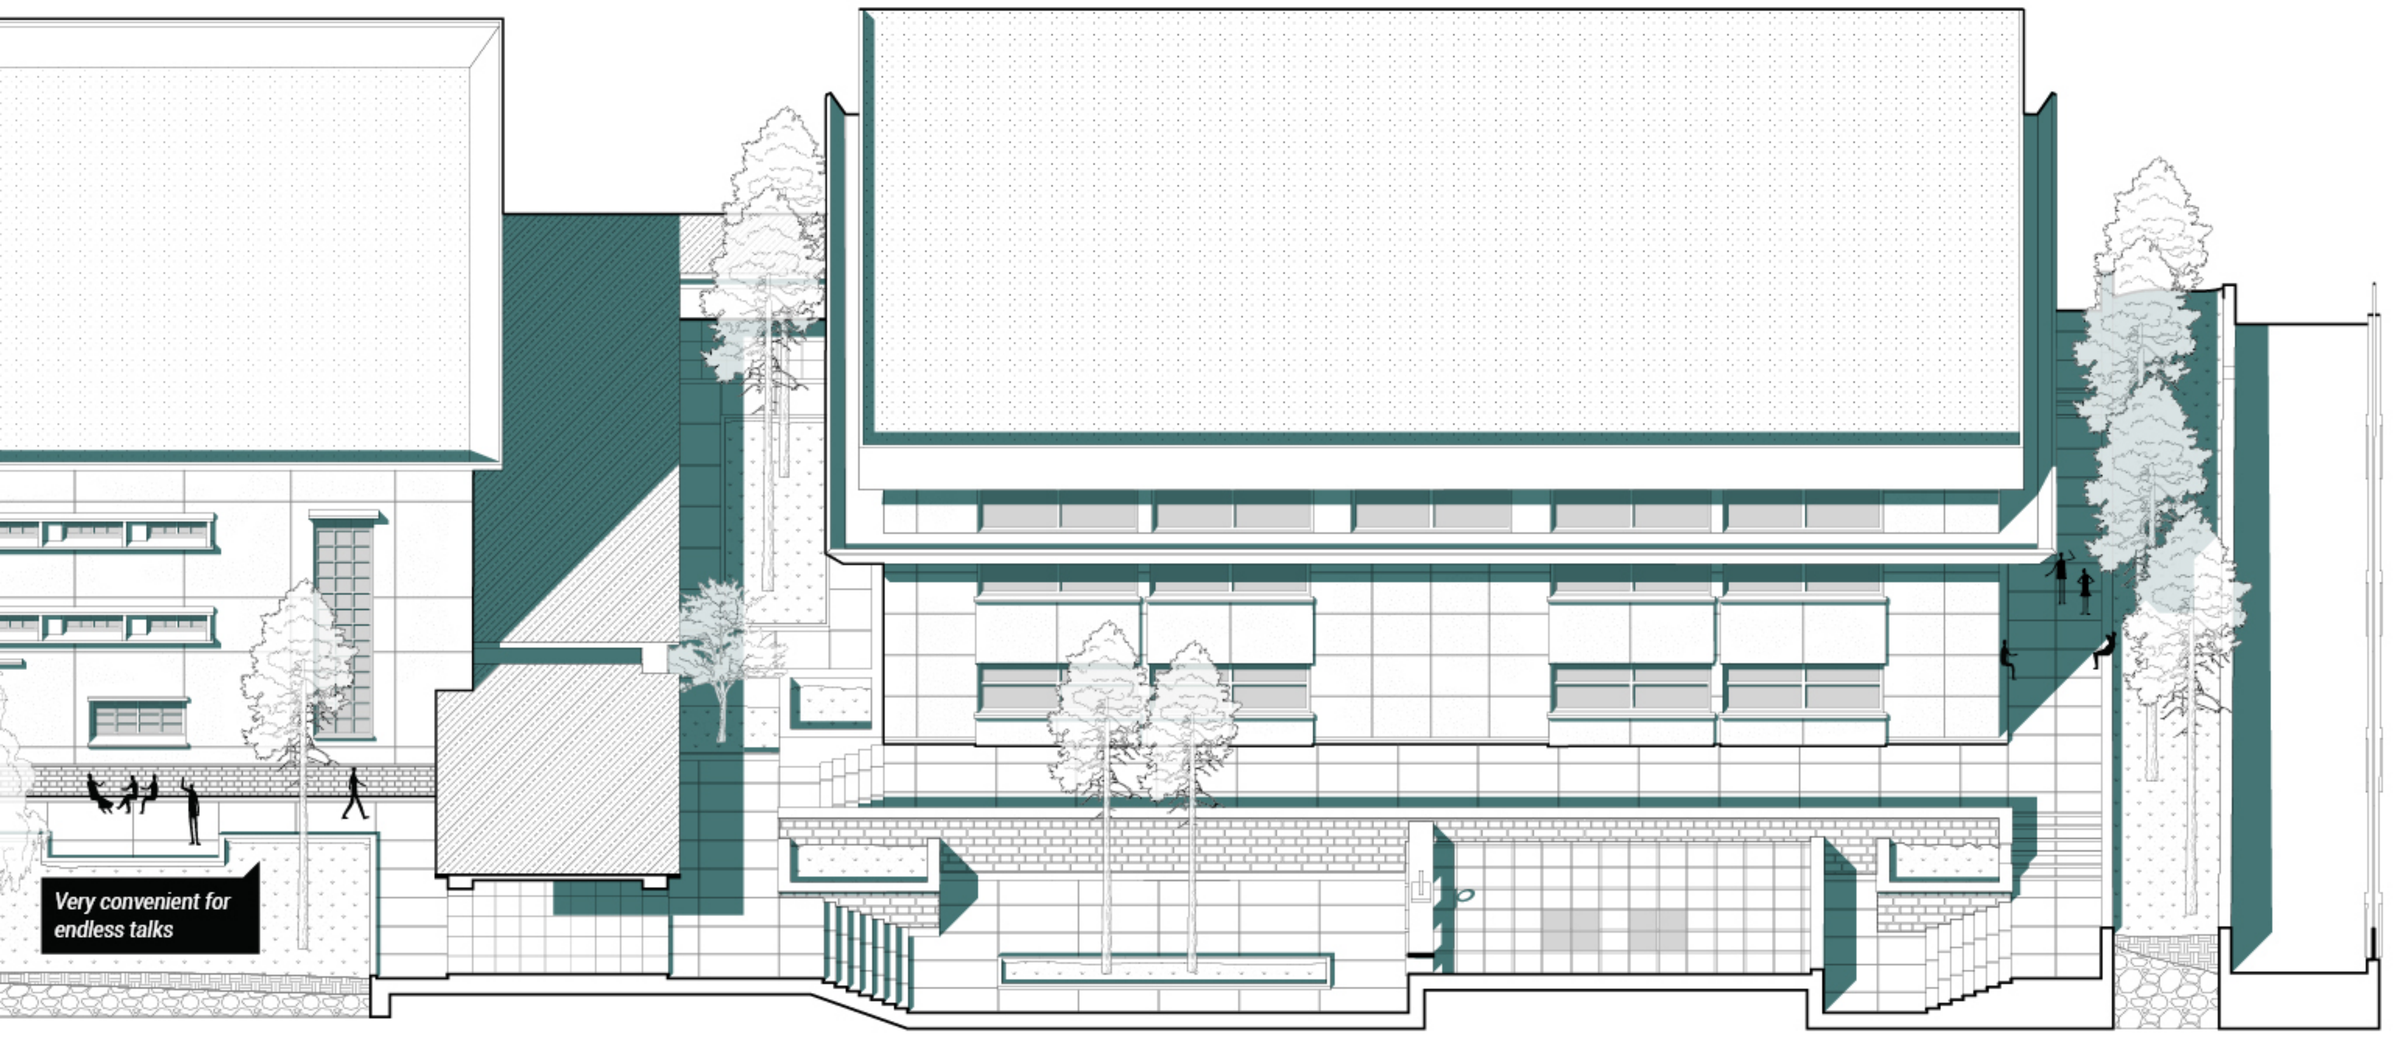

Security Cameras

Patrol Pattern

Gathering Spots

Green Area

Fences

Figure 2.6 *Map of individual safe spots for discretionary time [Author]*

The faculty typically employs two security officers who patrol along certain patterns periodically. Main circulation paths and primary public spaces are most frequently patrolled, and these same locations are also monitored by security cameras. Using a base map, I located and defined the surveillance system in the faculty, peer-mapping with three of my previous colleagues. I asked them to map the surveillance system independently to verify data accuracy, followed by a final check with photographs and visual materials. The surveillance equipment locations and patrol patterns are documented in the following maps. [Figure 2.5]

Additionally, I asked the mentioned colleagues individually to highlight the areas they considered safe for discretionary and leisure time when they were alone. [Figure 2.6] results are shown in the subsequent figures. I used the outcome of this peer-mapping as a validation tool for the simulation results in the next steps, to determine whether the patterns identified by the simulation correspond to lived spatial adaptation, thereby testing whether surveillance infrastructure alone can explain the emergence of these 'safe zones.

My account of the surveillance system dates to 2016-2021. As of 2025, conditions have intensified significantly. As one source reported: "The number of cameras has increased drastically. Before, they were only in the corners and main spots; now they are almost everywhere. The place is not like when we were students; it has lost its old vitality."

Figure 2.5 *Plan of Faculty of Fine Arts and peer-mapped surveillance system [Author]*

2.6 Spatial Fragmentation as Lived Reality

This context establishes the conditions this thesis investigates: how systematic threat reshapes spatial perception and memory. The surveillance system described here does not merely restrict movement; it fundamentally alters how space is experienced, remembered, and navigated. Women develop a perceived map of the city, a version fragmented not by physical barriers but by accumulated stress, learned avoidance, and psychological erosion of spatial access.

The theoretical frameworks of Foucault's panopticon and docile bodies find concrete manifestation in these lived experiences. The perfect operation of disciplinary power lies not in constant enforcement but in internalized control, the self-surveillance that makes women architects of their own restriction. The Faculty of Fine Arts, as a bounded case study site, offers the opportunity to trace these abstract principles through specific spatial patterns, to model how surveillance infrastructure correlates with adaptive navigation, and to quantify the cognitive fragmentation that results from systematic threat.

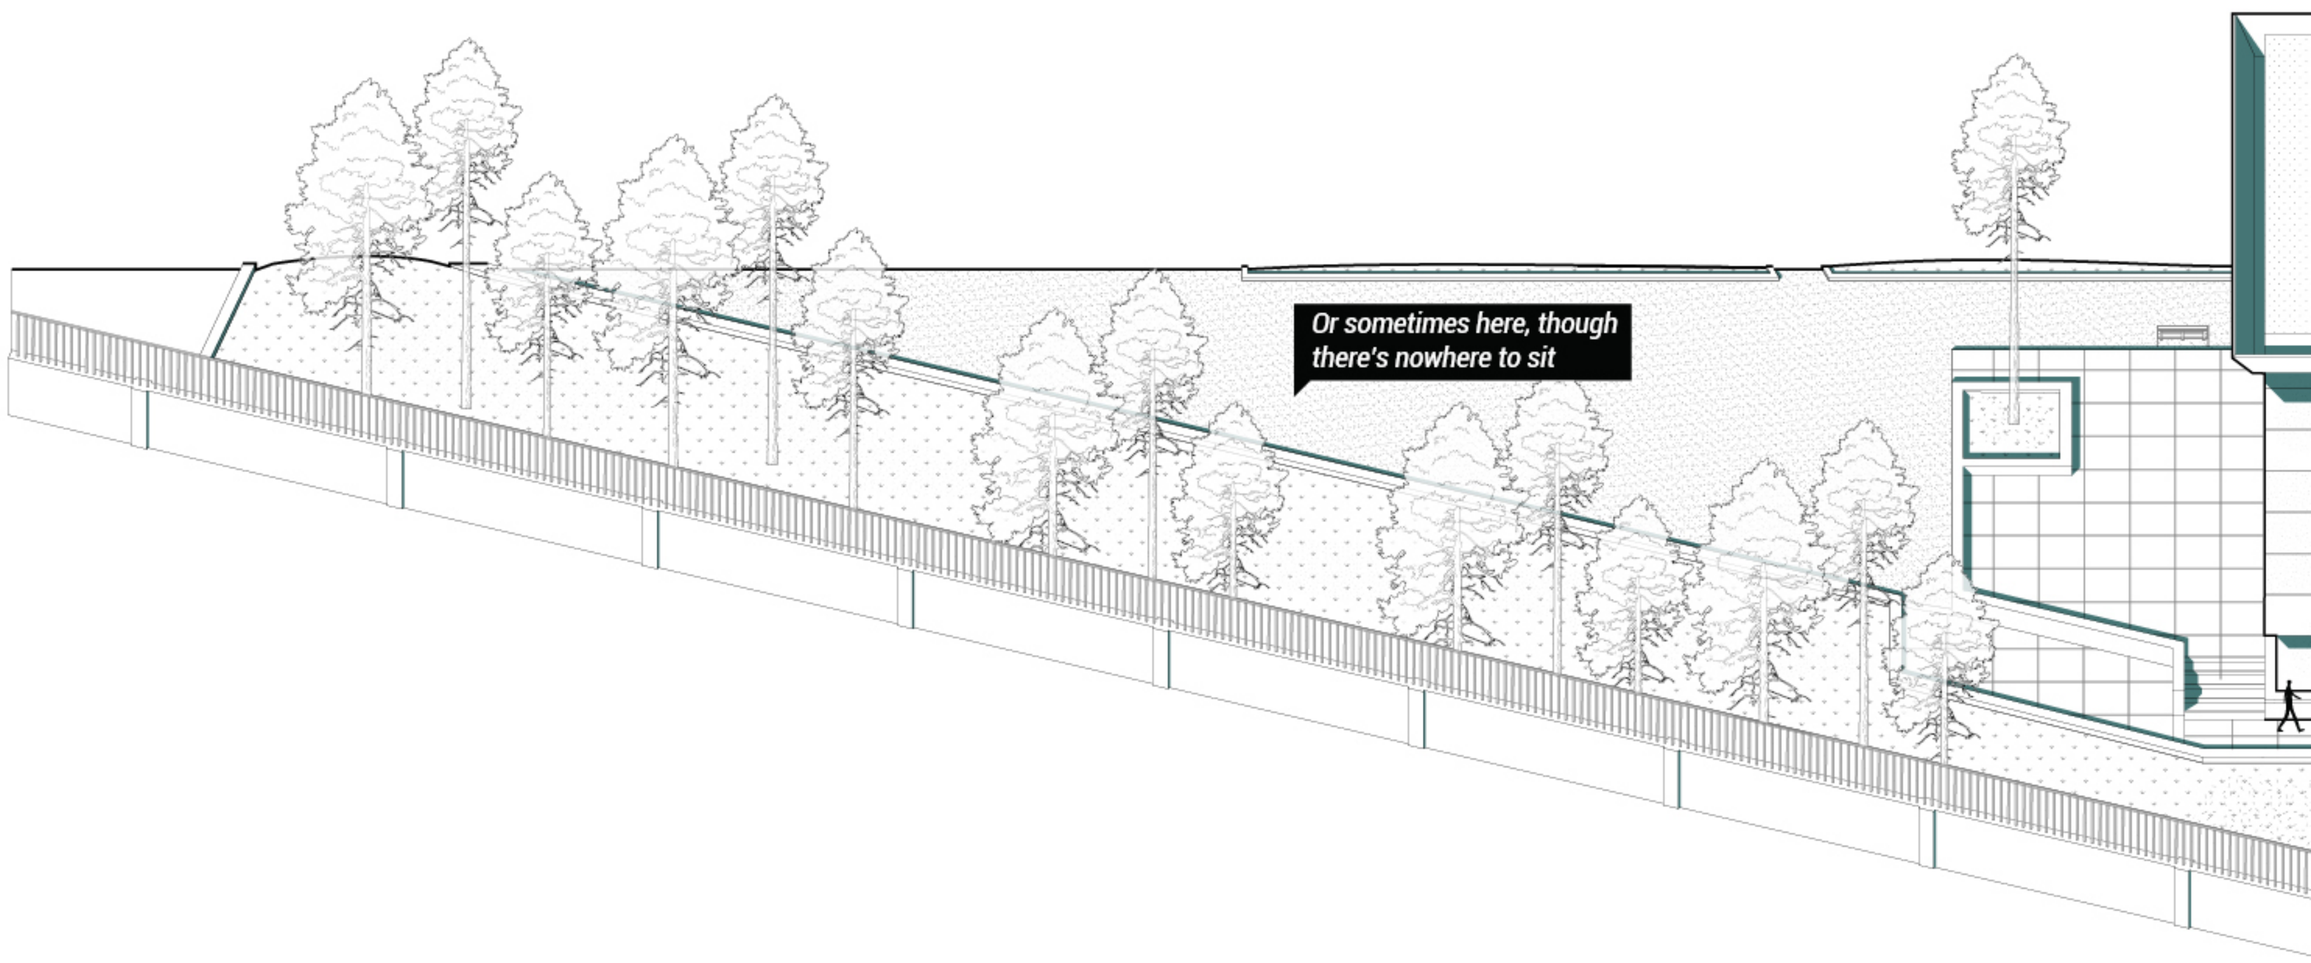

Figure 2.7 *a-e. Parallel Sections of the Faculty [Author]*

The following chapter establishes the theoretical foundation for examining how systematic threat reshapes spatial cognition—the mechanisms of perception, memory, and decision-making that will be operationalized in the computational model. By grounding these theoretical concepts in the lived reality documented here, this research aims to make visible the invisible architecture of control: the ways systematic oppression reshapes not only where women can go, but how they understand and remember the spaces they inhabit.

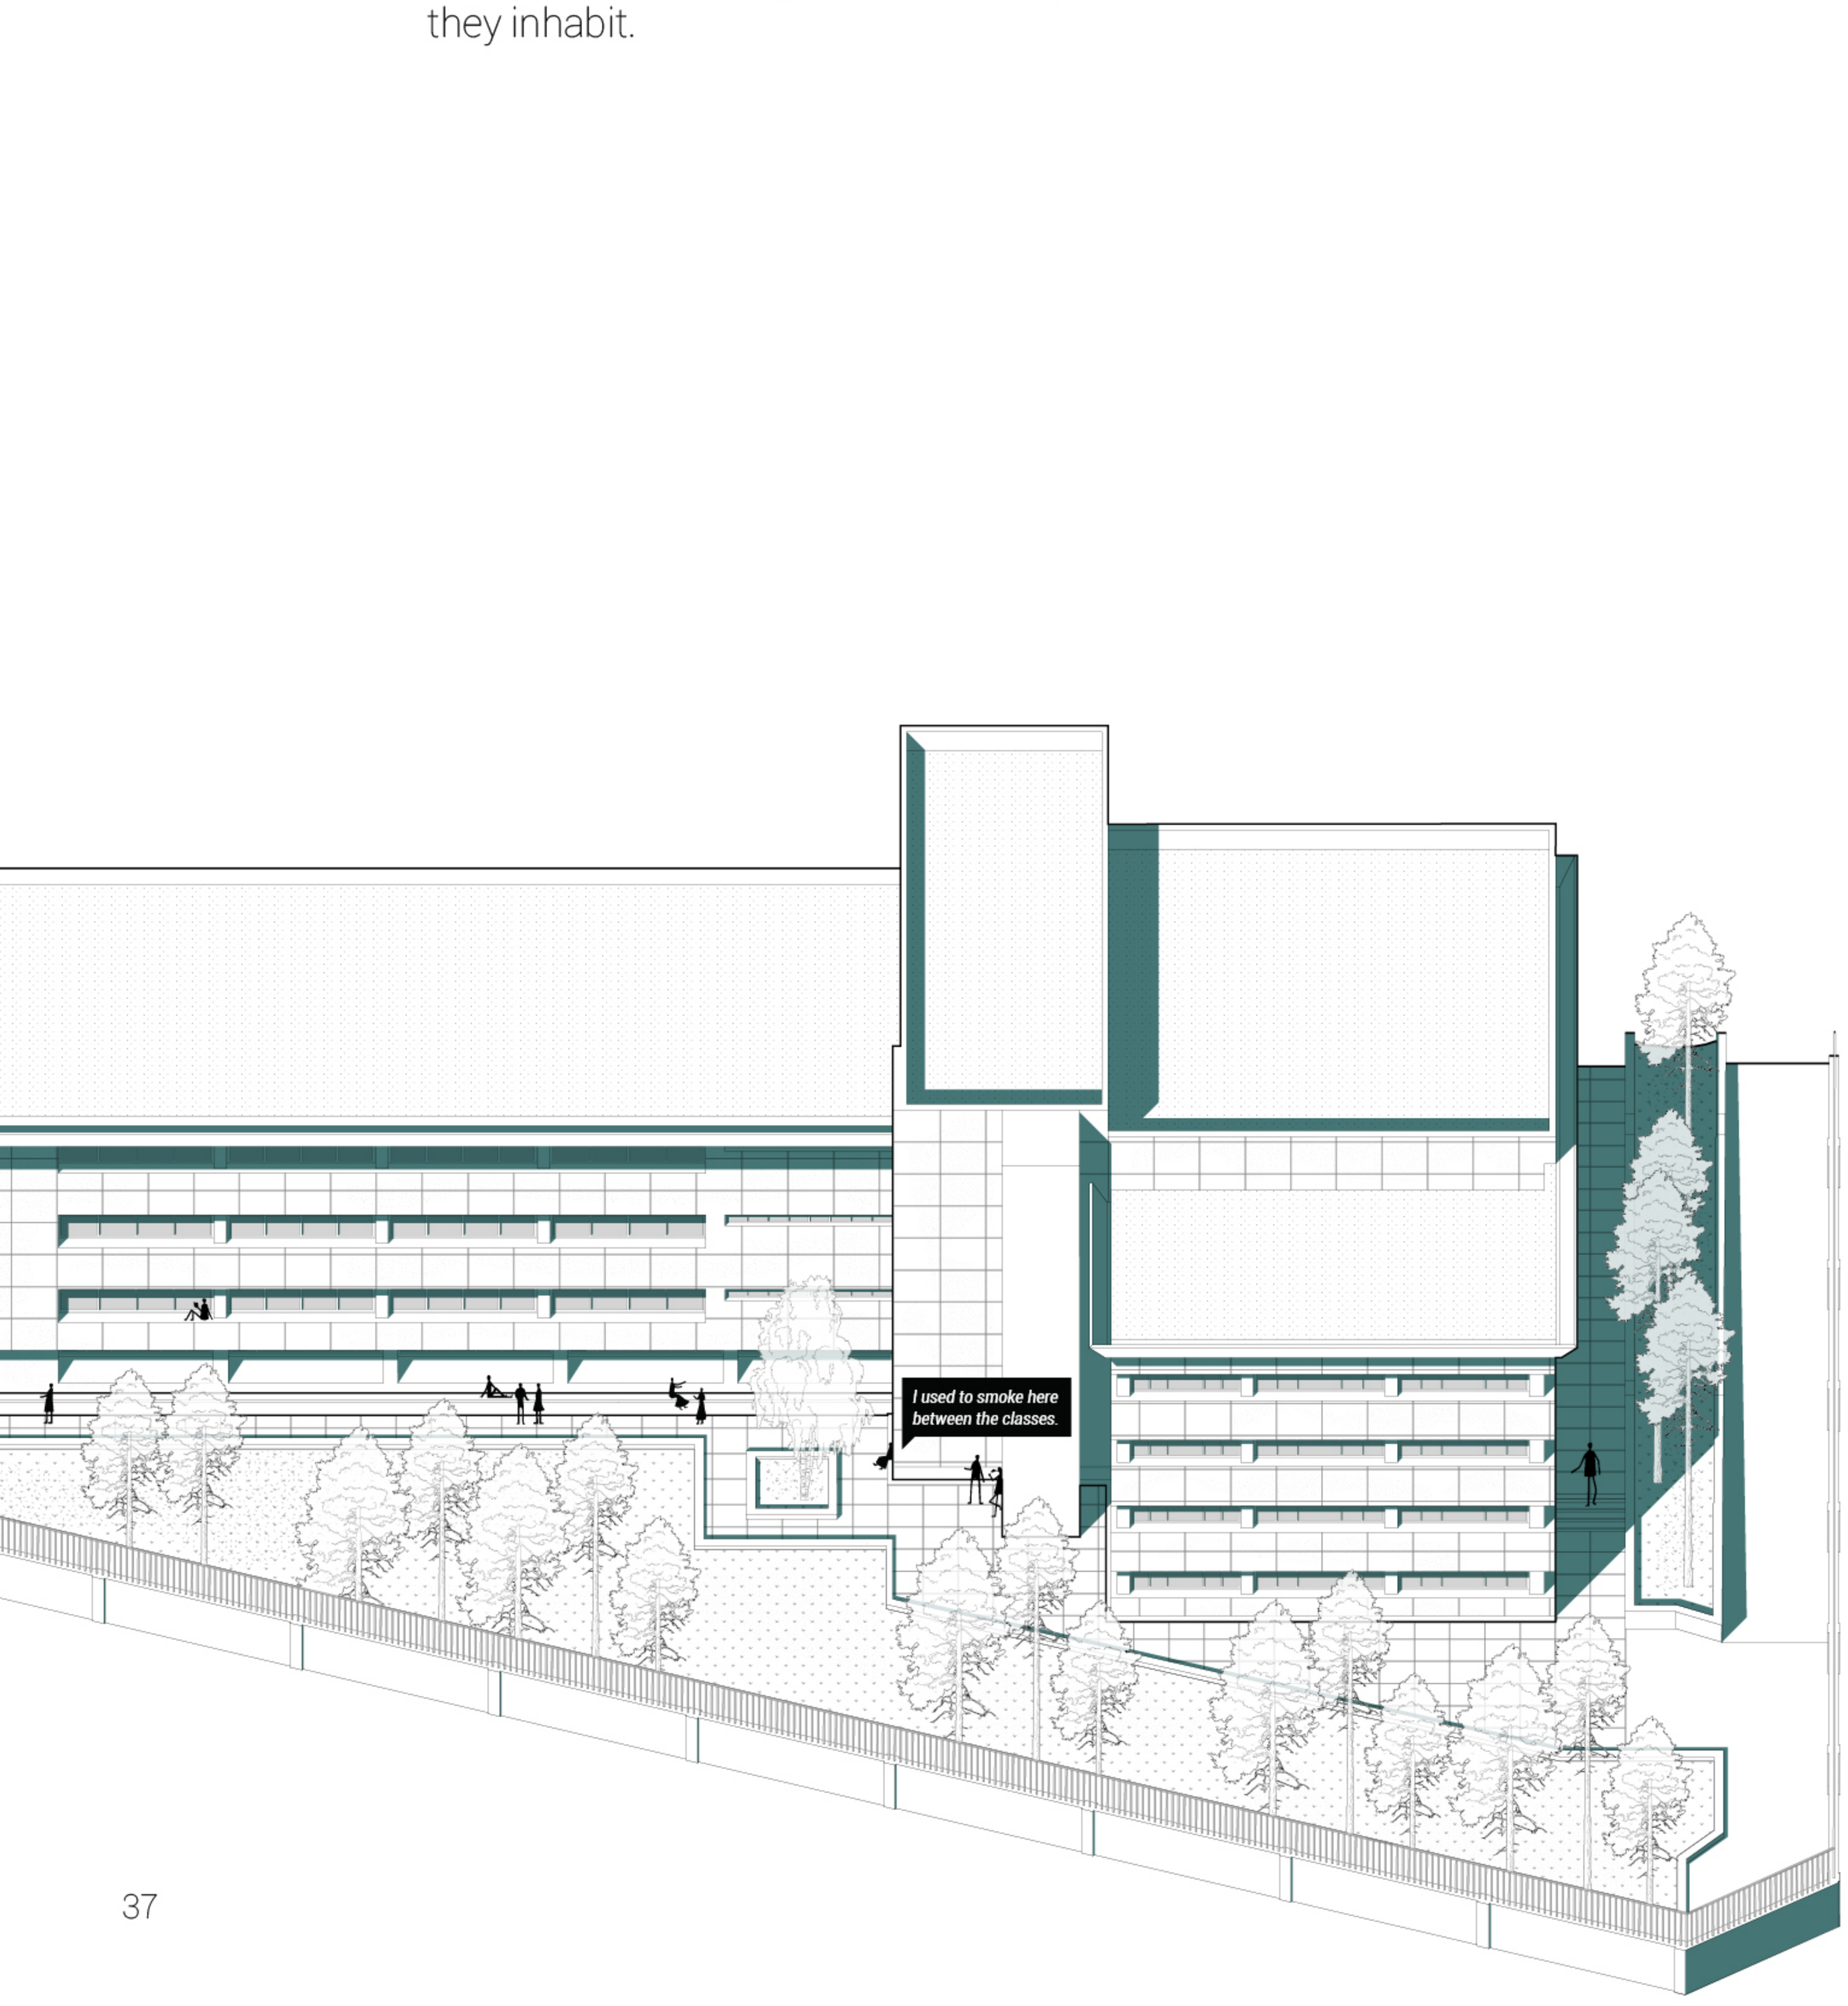

03 | Theoretical Framework

Theoretical Framework
Spatial Cognition Under Threat

This chapter explains how surveillance produces psychological erasure—the phenomenon where physically accessible space becomes cognitively unavailable through accumulated threat experience. We trace this mechanism from the conceptual definition of space and its agency to the specific cognitive mechanisms—perception, memory, and adaptation—that are modeled in the subsequent methodology.

3.1 Defining Space: Dichotomies and the Borderland

Space is where we end and "the world" begins, but the whereabouts of the border is a source of constant dichotomy between the perceived, the represented, and the true space, igniting many thoughts across various fields. As in this thesis, space may be divided into two primary categories:

Physical Space: Attributed to the external world, independent of minds, and often absolute (independent of objects, like a container) and Euclidean. This space acts as the starting map in the model.

Psychological Space: Attributed to the mind, existing only through perception and cognition, and often relative (a set of relationships between objects).

The comprehension of psychological space is defined as learned space,

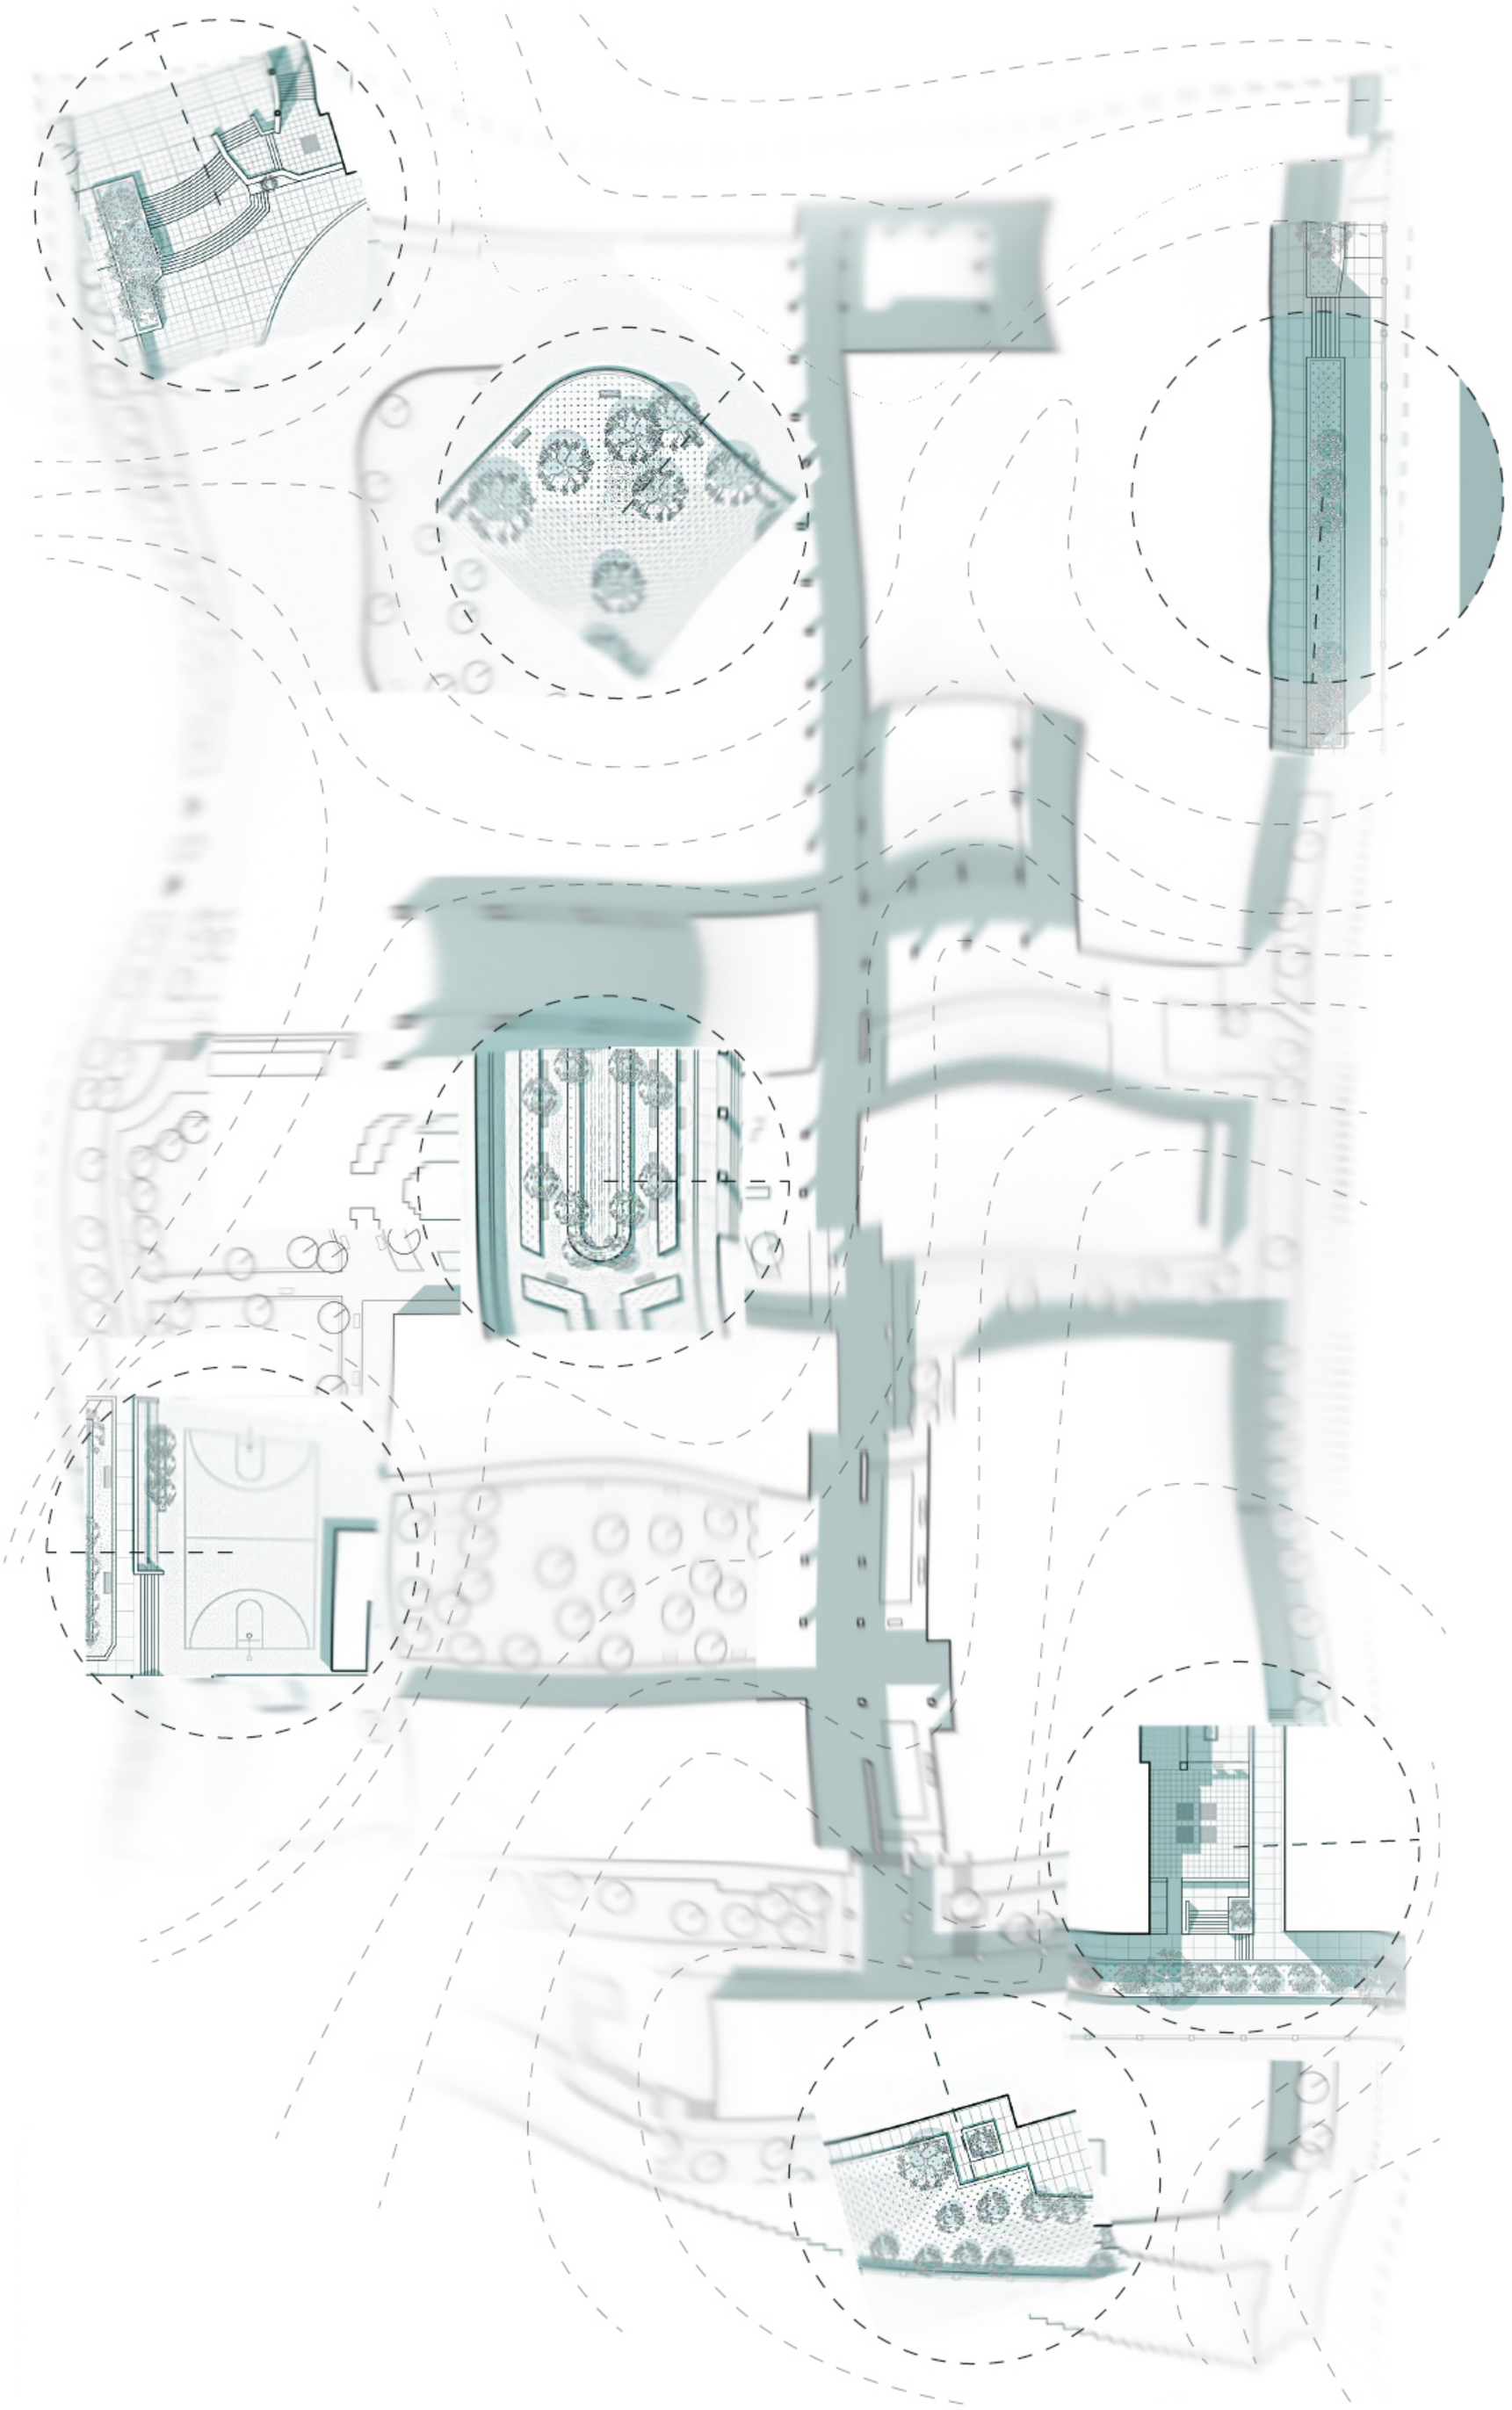

Figure 3.1 *Patchy cognitive structure: "While a person may have vivid memory of their house, the university, and the cafe, the environment in-between is nothing but a blur—a fast continuity of images." [Author]*

the abilities gained through interaction and experience, excluding innate neurological functions. For example, our rabbit's navigation between its den and the carrot lot is learned behavior, while its startle response to sudden sounds is instinctive. The environment we know is a result of this individual learning, a result of interactions and memories of space. [A. Mallot, 2024]

Critically, the representation of learned psychological space is non-Euclidean, as it is not a geometrically constant environment. Furthermore, it is non-continuous, formed as a patchy structure lacking a single, reliable coordinate system. While a person may have vivid memory of their house, the university, and the cafe, the environment in-between is nothing but a blur—a fast continuity of images—and the orientation of each remembered place differs. [Figure 3.1]

Physical space translates to psychological space through a set of events. The university is a learned space as the person has interacted with its physical form. In other words, the physical landscape internalizes through these interactions between the user and the choices and stories it offers them. The individual in the real setting confronts not mere solid geometry, but memories, decisions of other agents, feelings, and perceptions. In the eye of the individual, the composition of all these elements can be represented as one single system.

3.2 Space as an Agent: The Terrain of Uncertainty

The concept leads to the idea of space itself as an agent. While space in its raw, physical sense does not decide, when considered as interacting with an individual with all its beings, systems, and temporality, it functions as an active participant in shaping events for that individual. This structure is named Terrain of Uncertainty in this thesis, abbreviated as ToU.

The ToU is not synonymous with either physical or psychological space alone, but represents their integration in lived experience—the complete navigational field an agent moves through, containing environmental geometry, accumulated memory, and active forces (other agents, temporal changes, emerging affordances).

The ToU comprises the complete experiential environment: the agent's spatial memory, current perceptions, and the active agents operating within that space. The ToU is not merely a representation of memory but a dynamic field containing both remembered associations and present constraints or goals. In this sense, the ToU possesses agency: through the agents it contains, it actively shapes the navigator's choices and constrains movement.

Agency and the Feedback Loop: Through countless interactions, the ToU acquires agency—not as an independent consciousness, but through the active elements it contains (threat agents, other navigators, temporal changes). It enables or constrains the agent's actions in a complex, adaptive feedback loop. This redefines the landscape from a container to an active composition

Figure 3.2 *The fragmentation of ToU and the options it offers: top. Primary situation, a uniform landscape of navigational options, middle. An individualized ToU, with scattered fragmentation due to individual's experiences, bottom. The fragmented ToU showing a clear axis of fragmentation, suggesting existence of a clear focus of uncertainty shaping ToU. [Author]*

of events, organisms, objects, and temporality. Decisions are made as the ToU offers a combination of memorized and learned actions and real-time perception of affordances.

Memory and learning: As Tim Ingold notes, "To perceive the landscape is therefore to carry out an act of remembrance," indicating that the environment is "pregnant with the past." The ToU integrates this spatial memory with real-time perception. While memory itself is passive (stored associations and learned patterns), the ToU as a whole possesses agency through the active elements operating within it, and changing affordances that constrain choice in the present moment.

Affordances: Following Gibson's concept of affordances, the landscape acts upon us through the possibilities and constraints it offers. The same narrow street affords shade and refuge during a warm day but becomes threatening and constrictive at night—the physical geometry remains constant, but its affordances transform with temporal context and the agent's learned associations. In a surveillance environment, affordances are continuously modified by the presence and movement of threat agents. These restricting forces are termed Foci of Uncertainty in this thesis.

Operational Definition: In a surveillance environment, this agency manifests by modifying the perception of risk. The ToU for our agent is driven by clear Foci of Uncertainty—the threat agents propagating stress and restricting navigation. The ToU carries the accumulated memory traces of constraints and hostile elements, acting upon the navigating individual by modifying the affordances of the environment. [Figure 3.2]

To adapt and survive within this environment, the agent must constantly map the ToU by perceiving and processing the risk. This essential task of perception and planning leads us directly to the study of Spatial Cognition. The following section details the internal cognitive systems—perception zones, dual memory structures, and reinforcement learning—which will be operationalized in the agent-based model to demonstrate how surveillance fragments an agent's right to inhabit and remember their world.

3.3 Spatial Cognition: Perception, Memory, and Planning

Spatial cognition focuses on cognitive processes that occur during spatial behaviors: way-finding, spatial planning, spatial reasoning, building and object manipulation, or communication about space. It is the most widespread type of cognition in animals and is known to be the first part of cognition that takes shape. The idea of space is so deeply rooted in the human mind that it is used not only for external objects but also for abstract ideas. From talking about time (e.g., "In the coming days") to remembering words and sequences, we tend to map ideas and locate them, even virtually. [A. Mallot, 2024]

The simplest form of how we interact with the environment, learn, and adapt is stimulus-response behavior. The nervous system receives a stimulus (e.g., Hearing footsteps at night), and it prompts a response (changing route). As Tolman says, it's like turning a switch on and seeing the light. This is the simplest form of behavior, and it cannot be used broadly, since a certain stimulus would not always lead to a certain response. Usually, the receipt of a stimulus will lead to a chain of "inner states," after which the agent will show a response (even if not visible). The inner states include goals, beliefs, whims, percepts, or even sudden insights of the agent.

To understand how the agent adapts to the ToU, we must examine the internal processes of spatial cognition—the loop of perception, planning, and action, where stimuli and inner states are combined in perceiving our movement, the relative location of self and objects, and the layout of the space. [A. Mallot, 2024; Tolman, E. C., 1948]

Spatial Perception and Peripersonal Space

The process begins with the receipt of spatial perception cues from the surroundings. This receipt is limited to the agent's location and the stimuli it receives at the moment, its memories of past experiences (what has been learned), and its capacity to receive and process them. Peripersonal space is a main concept in this dynamic.

Peripersonal Space: The receipt of stimuli is local and occurs within the peripersonal space—the immediate surrounding of the person. This space is an active, egocentric representation that is continuously updated with movement, containing the representation of "self."

ABM Operationalization: Peripersonal space directly correlates with the Perception Zone in the model, defining the limited extent, duration, and heading-direction dependency of information intake. Two main representations of peripersonal space used in that regard are isovists (a map of the visible portion of the peripersonal space) and local charts. [Figure 3.3.]

3.4 Memory Systems and Cognitive Maps

Memory is where the temporality of experienced space finds meaning and where planning takes place. We categorize memory as Iconic (very short-lived memories of received stimuli, like the silhouette of the sun after looking at it), Working, and Long-Term. [A. Mallot, 2024]

1. Working Memory (WM): WM is the intermediary between stimuli, long-term memory, and planning an action—it is the information we retain while performing an act, such as when passing through a crowd, and we record the location of nearby people to avoid colliding with them. Usually, shortly after finishing the task, this memory is forgotten.

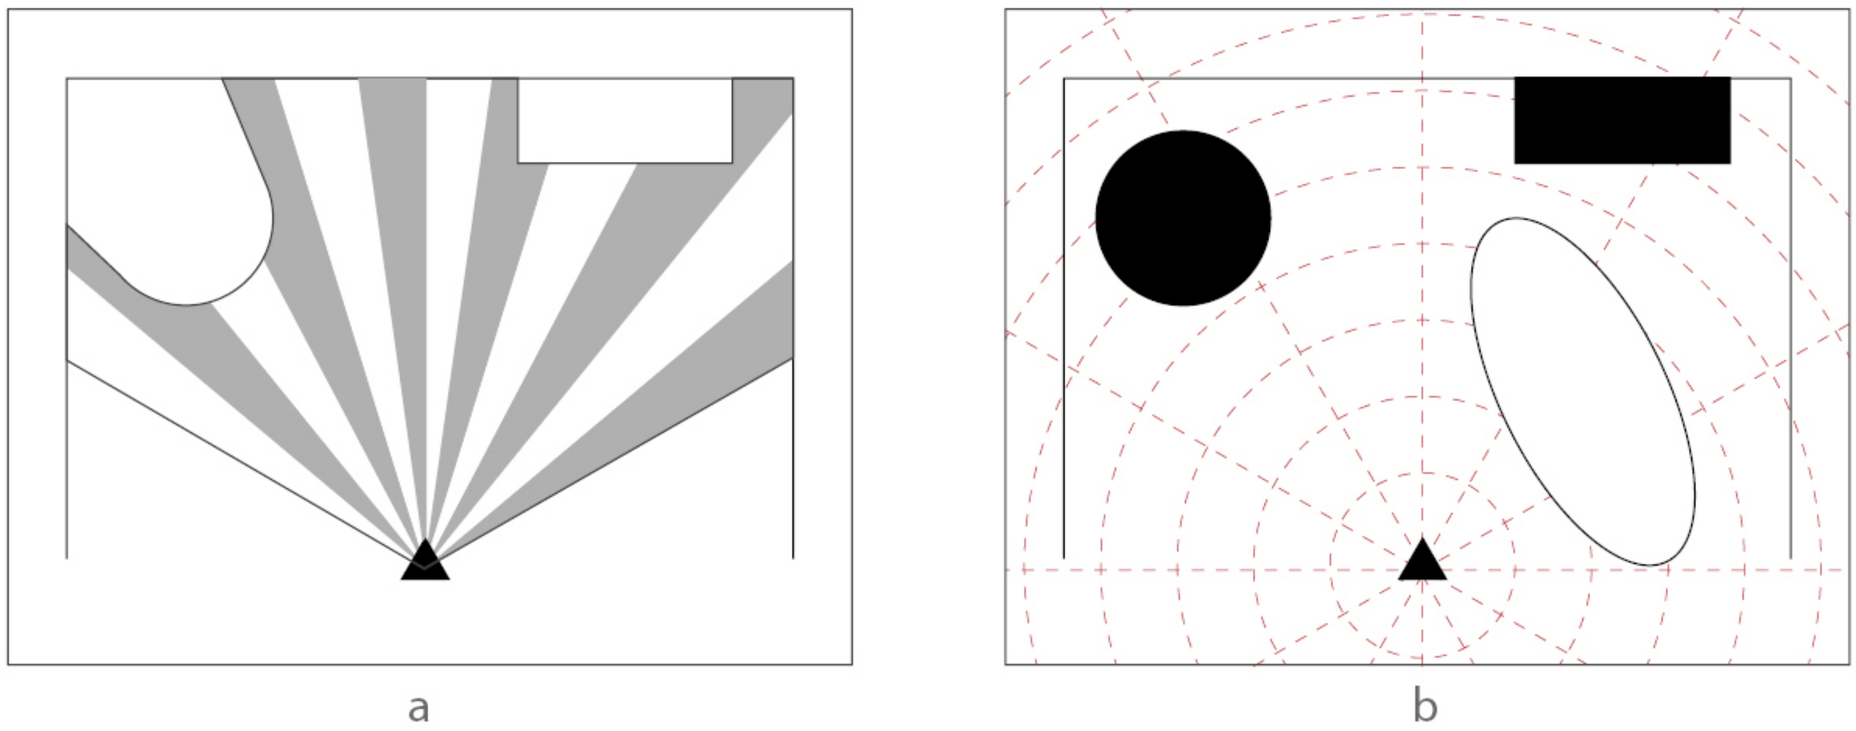

Figure3.3 a. Isovist, a map of visible portion of the floor, b. Local chart with egocentric coordinates of peripersonal space and the configuration of objects [A. Mallot, 2024]

Working memory is limited in capacity, other than being transient. To minimize the problem, the data we perceive is accepted as "chunks" of data. For example, while a limited number of words can be remembered, by changing them into more complex units of meaning, like several sentences, they will be easier to remember.

As for spaces and objects, it's easier to recall the locations relatively, by creating a visual narrative of the events, the components of the scene, and recording them, then using this general memorized imagined space as a source of comparison. The imagined space is not a mere image and consists of remembered landmarks, objects, but also memories, smells, etc., associated with the location. Furthermore, not only will it differ upon different recalls, but it is not a replica of the scene.

Furthermore, working memory is where learning takes place through rehearsal. Learning and modification or storage of behavioral mechanisms lead to plastic changes in the brain, which are called memory traces and shape the structure of long-term memory.

1.a. Local Charts and Orientation: WM is responsible for creating a local chart of the environment—a representation of peripersonal space that is locally allocentric but selects its content egocentrically, often having an intrinsic reference axis aligned with the salient properties of the environment [A. Mallot, 2024]. This chart is the basis for immediate planning and updating the perceived environment. [Figure 3.4]

1.b. ABM Operationalization: The WM is operationalized in the planning and navigation mechanism of the agent: by combining the memory and the perception and inner states, the agent analyzes the ToU for navigational decision-making. The safest orientation is chosen based on this analysis, and movement takes place.

2. Long-Term Memory (LTM): Through learning, the local charts move to long-term memory, retrieved upon necessity, and loaded into WM for planning.

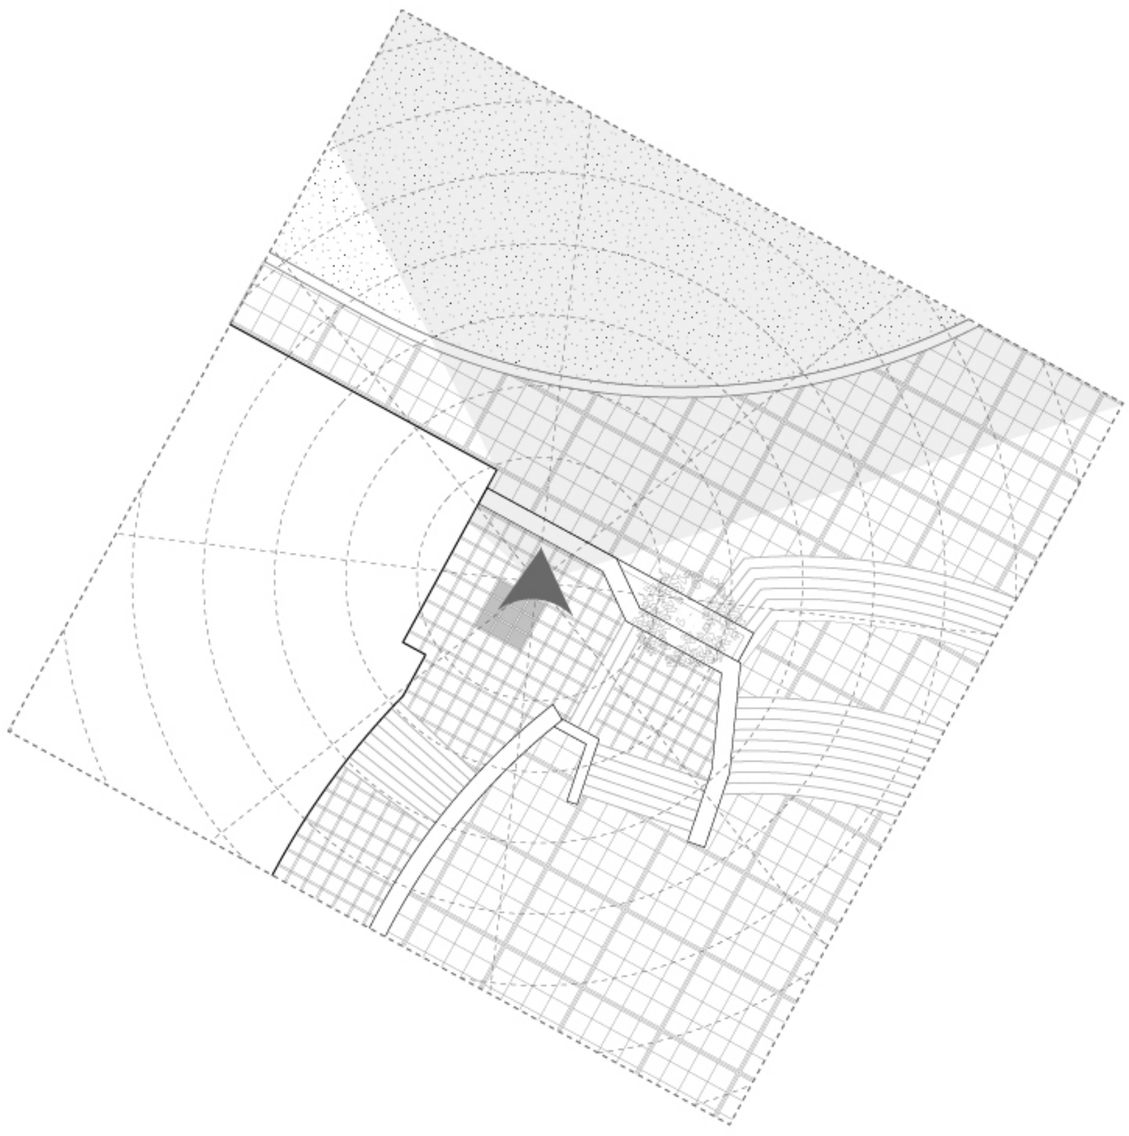

a

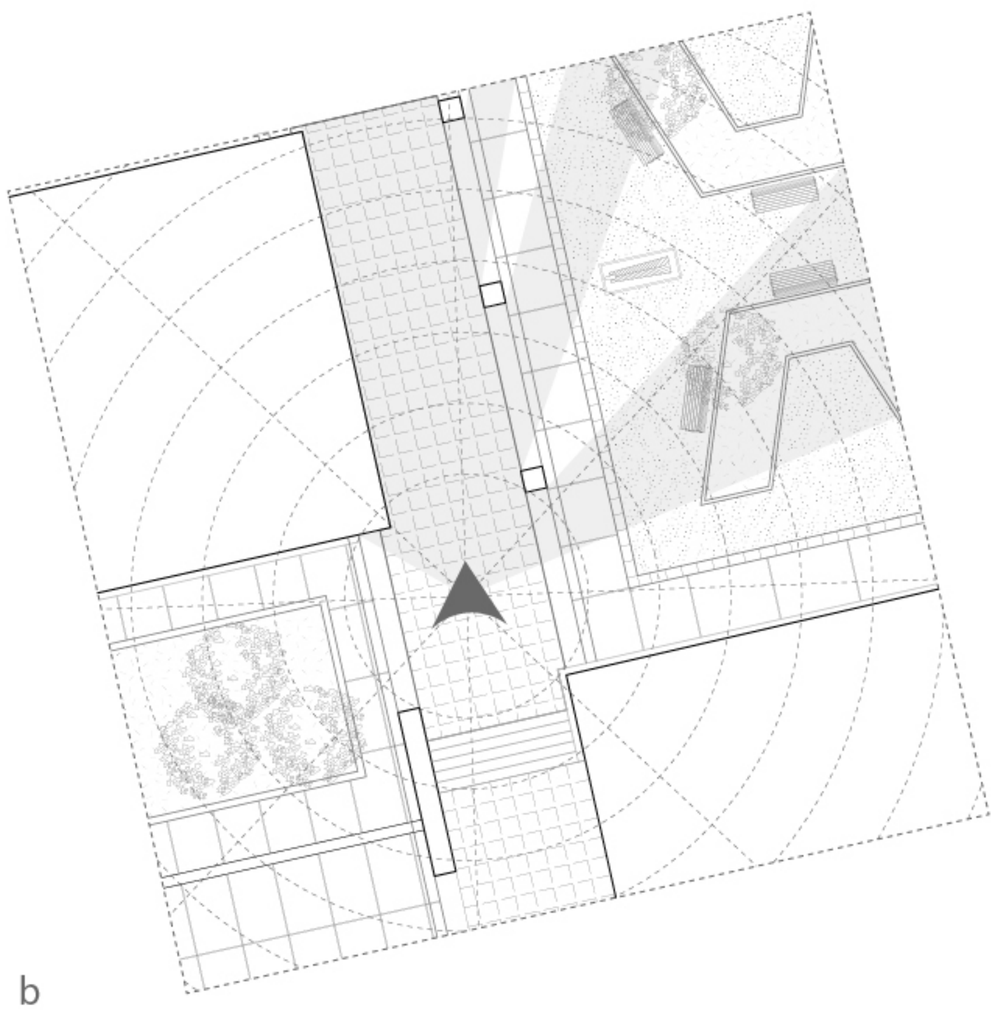

b

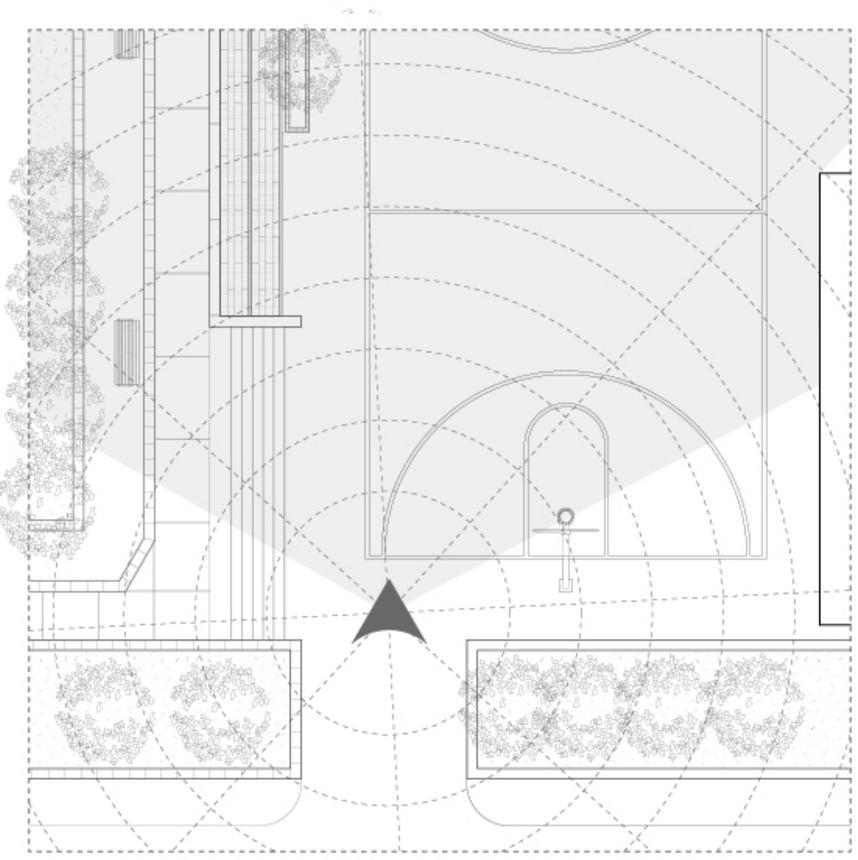

c

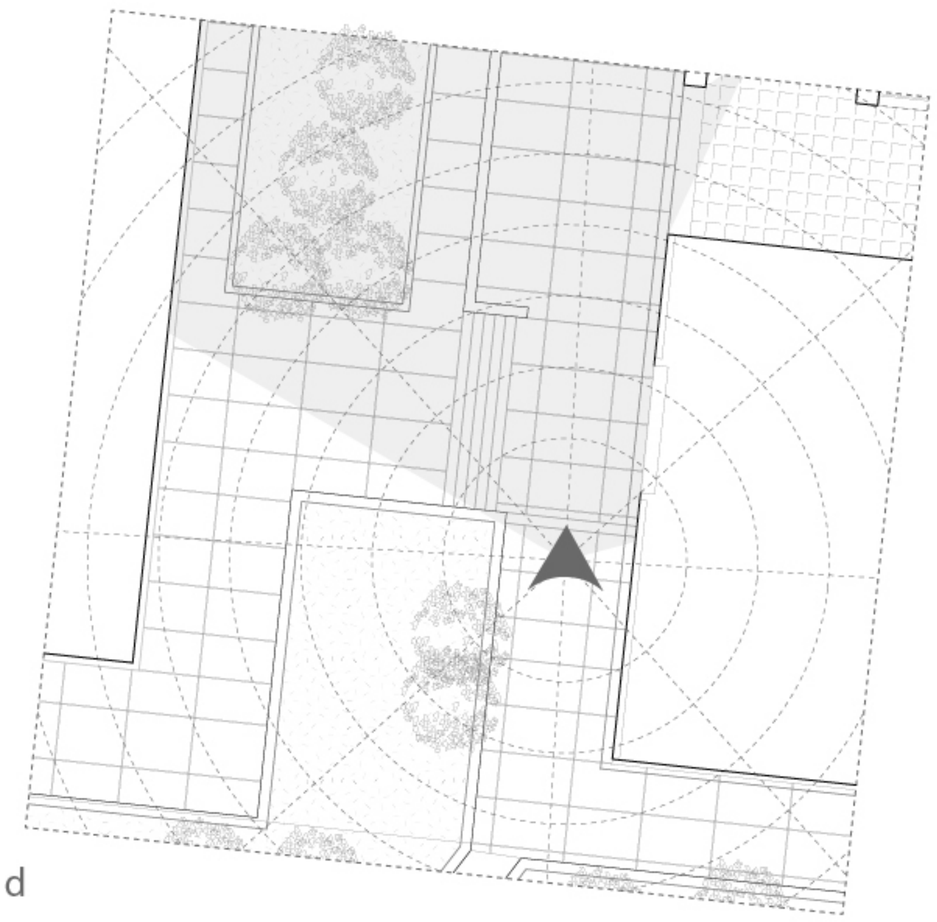

d

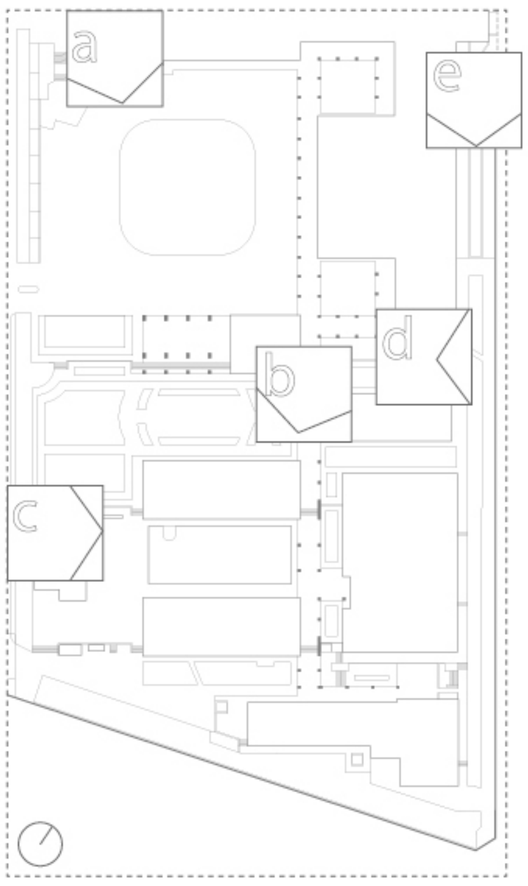

Legend

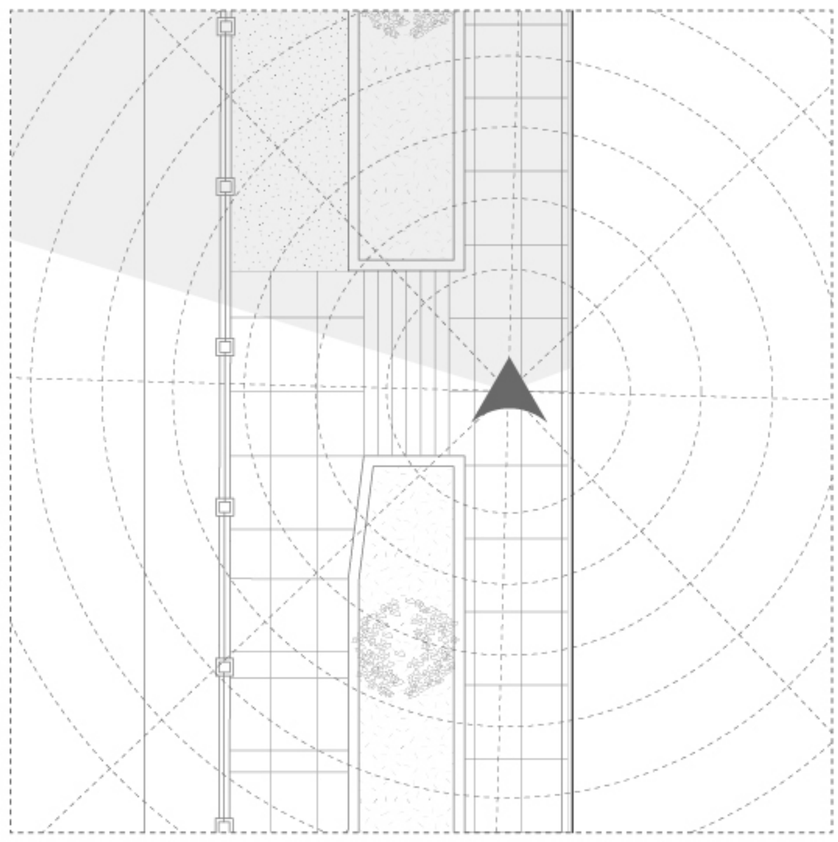

e

Figure 3.4 *The local charts are memorized in long-term memory, each remembered along a memorized orientation. a-e: samples of local charts in Fine Arts Faculty, rotated to align with the remembered orientation. Legend shows true north.*

Remembering routes and places relies on several factors, such as the distinctiveness of landmarks, the informational value of cues, their relevance to the current task, and how often they are encountered. Associations between states, actions, and outcomes are built through these cues. Over time, these associations enrich memory, allowing it to go beyond simple route-following toward a more flexible representation of space—a "widening of the map." In other words, LTM progresses from rigid Route Memory (sequential, fixed-goal paths likened to "strip maps") to flexible Cognitive Maps. [A. Mallot, 2024]

2.a. Cognitive Map (Tolman)

The structure of memory built from state, action, and new state associations goes beyond simple stimulus–response chains. Non-declarative memory governs habits and routines: automatic reactions to immediate needs. Declarative memory, by contrast, is rule-like and can be expressed in if–then form—for example, if going to the piazza, then turn right. This form allows for flexibility, the invention of new routes, and the transfer of knowledge to novel situations. [Figure 3.5]

The cognitive map is the flexible, declarative structure built from accumulating associations between states, actions, and outcomes. It allows navigation toward multiple goals and is dynamic, permitting multiple paths and decisions. Crucially, its richness is behavioral, not spatial, dependent on the number of alternatives and expected outcomes available to the agent.

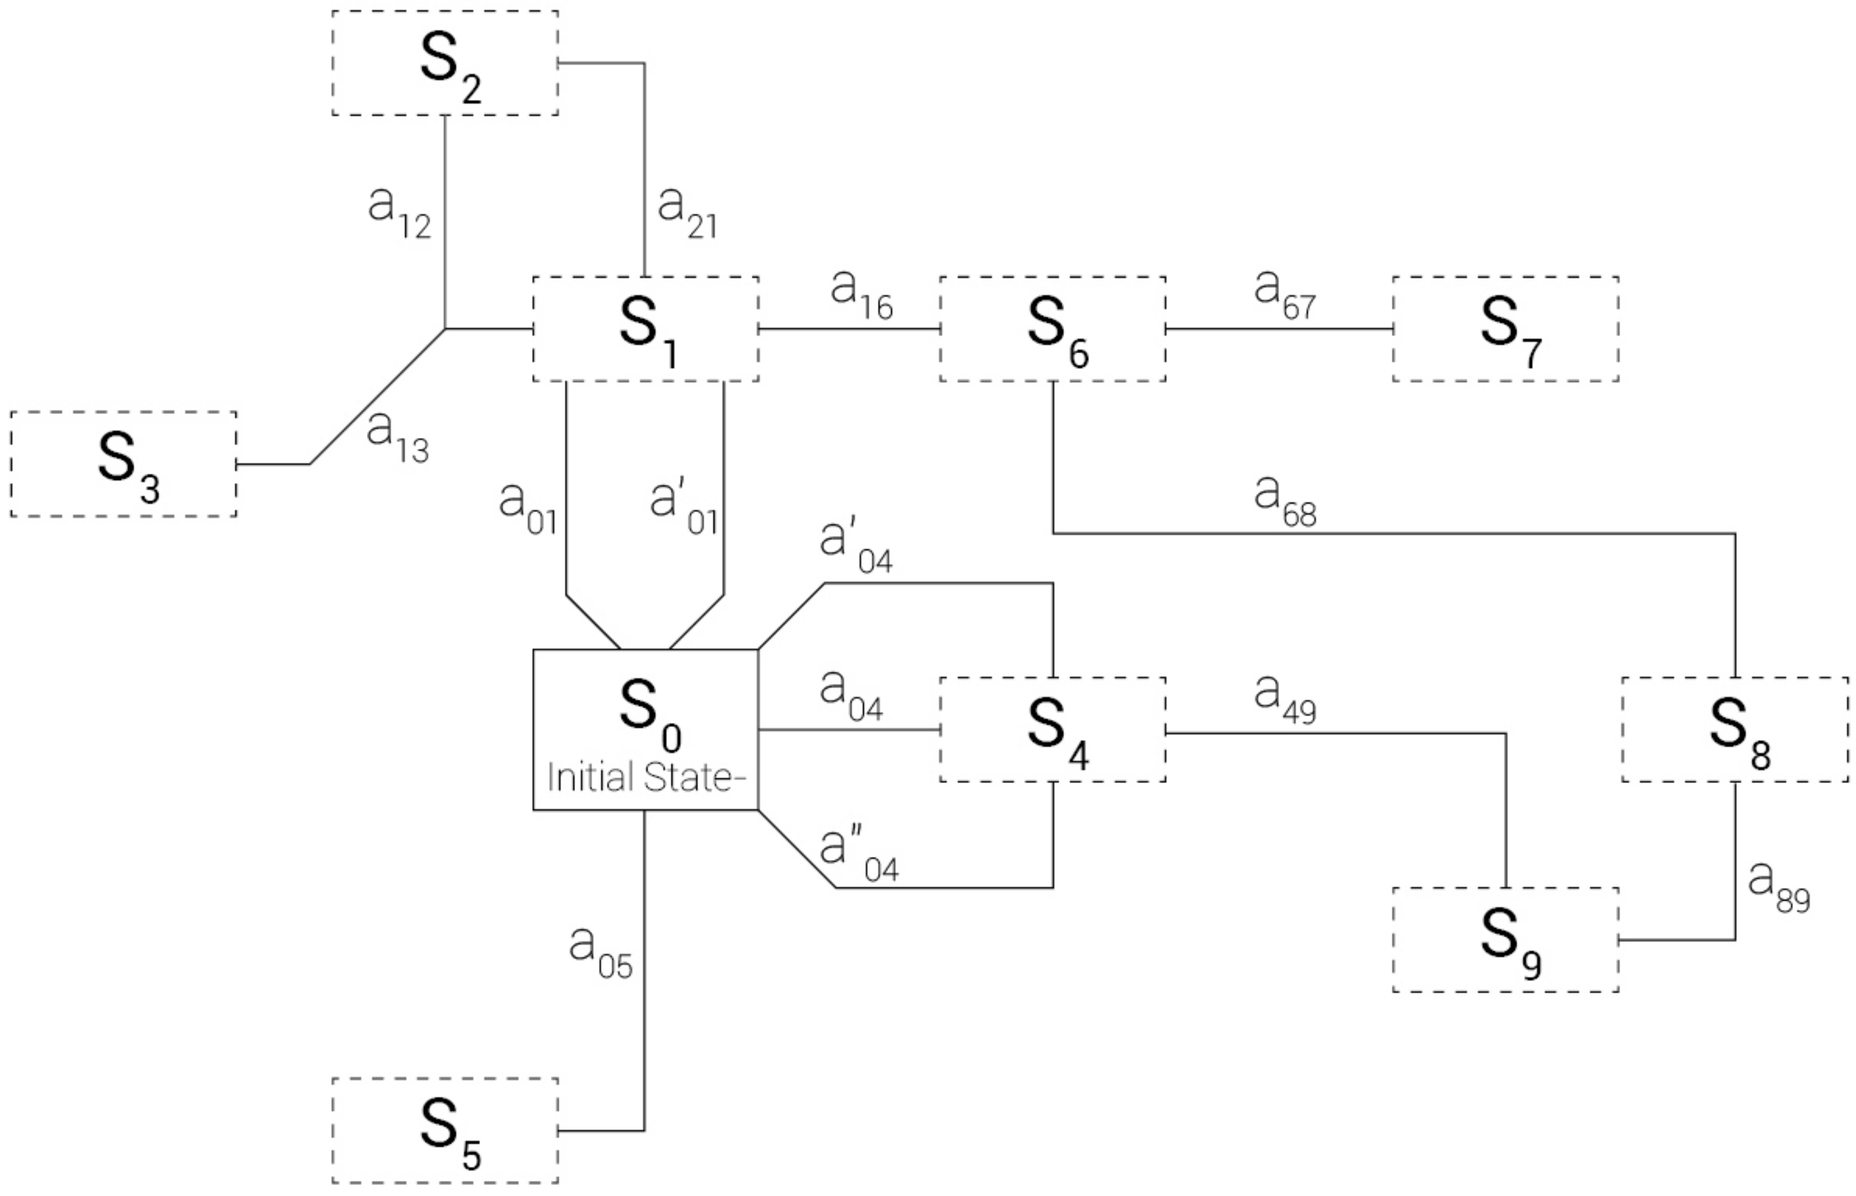

Figure 3.5 *Cognitive map: state-action associations leading to the next states, a declarative structure of long-term memory showing learned spatial behavior options [A. Mallot, 2024].*

In this sense, someone may have a far richer cognitive map of their home, built through accumulated associations and possibilities, than of foreign places they have visited superficially. Expansion of the map is therefore behavioral, not spatial; it depends on the number of alternatives and expected outcomes, not on distance traveled.

To conclude, cognitive maps enable higher-level planning and adaptive behavior. They represent not the physical layout of the world but the agent's knowledge of possibilities—an inner map of states, actions, and potential trajectories. The cognitive map's flexibility depends critically on how new information is acquired and integrated: the process of learning. [A. Mallot, 2024; Tolman, E. C., 1948]

3.4.1 Learning

Learning is the adaptive change of behavior based on experience that reshapes LTM. Learning plays a key role in the so-called widening of the map; as more action alternatives are acquired and linked to predictable outcomes, the repertoire of behaviors grows.

Latent Learning: enriches the map by storing potential, non-obligatory options.

Reinforcement Learning (RL): RL assigns value to state–action pairs through incentives (rewards/penalties). This process is central to the thesis:

In threatening environments, whether predation in nature or systematic surveillance in controlled spaces, persistent exposure to threat acts as a powerful negative incentive. This negative RL signal assigns an extreme negative utility or cost to exposed states, causing the agent to not merely deprioritize a route, but to actively psychologically erase that path from the set of viable options in the cognitive map. This psychological erasure—the removal of physically accessible paths from the agent's set of considered options—is the mechanism of fragmentation. The space remains physically navigable but becomes cognitively unavailable. In other words, the familiar, unconscious learned space will be lost, and turns to an uninhabitable land in which the main factor of navigation is the probable threat.

These learning mechanisms operate within a dynamic memory system that determines what knowledge persists and what fades. Under surveillance, this system becomes the site of fragmentation. [A. Mallot, 2024; Tolman, E. C., 1948]

ABM Operationalization of LM and Learning

Long-term memory stores the learned space. In a loop of rehearsal of

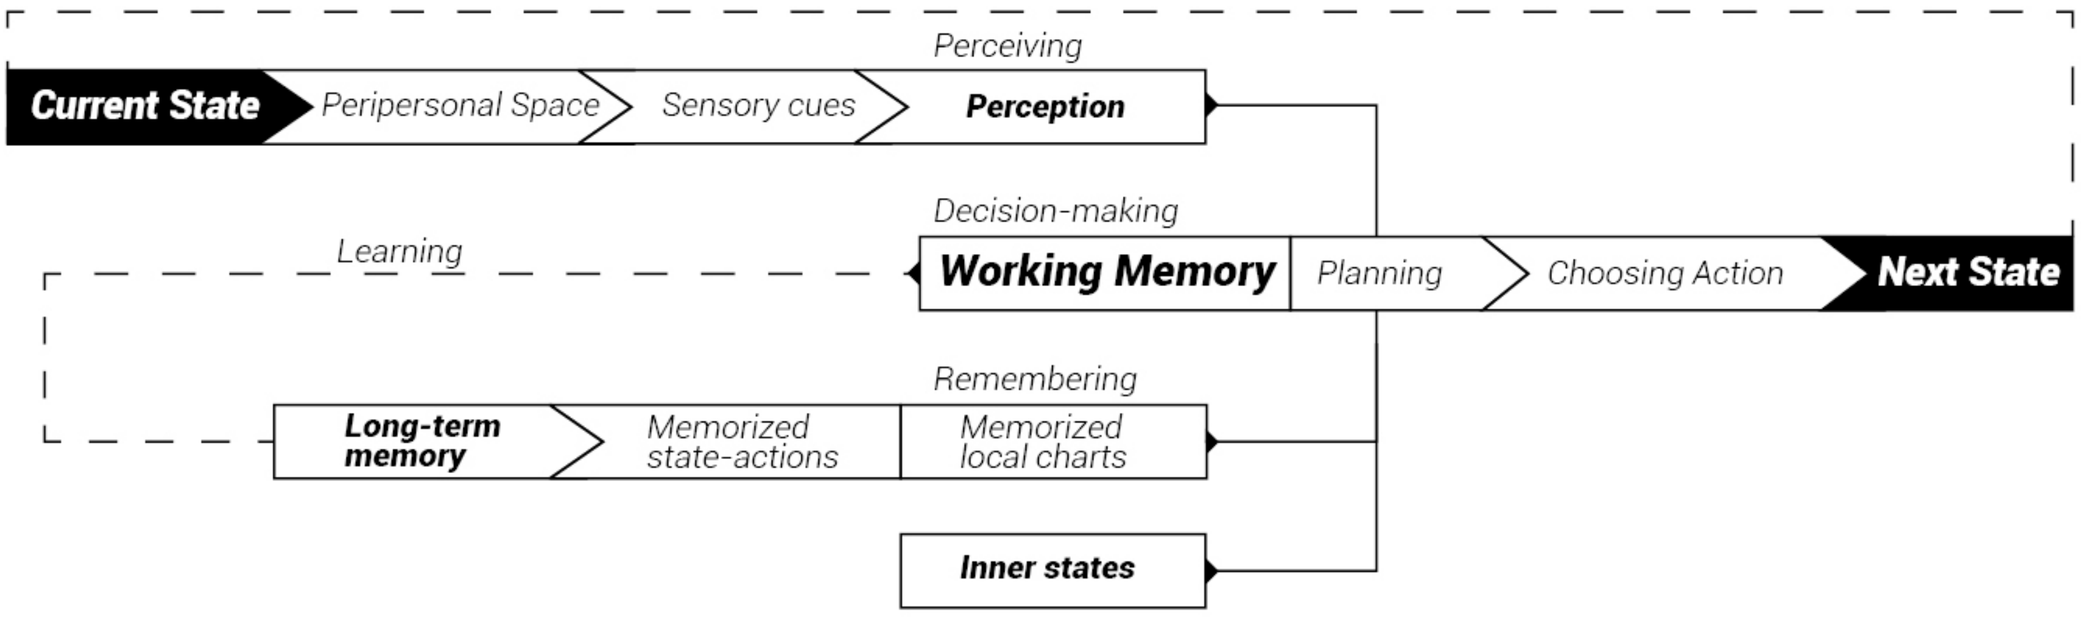

Figure 3.6 *The loop of decision-making in spatial behavior [A. Mallot, 2024].*

navigational actions, the agent learns the avoided and safe areas with a certain rate from the memories. A navigational cognitive map takes shape as edges are added and learned or removed and forgotten. Not only the navigated nodes, but all the remembered nodes with low avoidance will be parts of this map as possible states. The latent learning leads to storage of detail memory, or familiarity with the setting, and the reinforcement learning leads to avoidance memory.

The *breaking point* is defined as the learning curve flattens, when the agent has reached the maximum adaptation with the ToU. At this point, the inhabited, safe areas are mostly concrete, and it is possible to deduct quantifiable results.

3.5 Adaptation, Memory Dynamics, and Fragmentation

The agent's ability to survive in the ToU is governed by the dynamics of its internal state, particularly the acquisition and decay of knowledge and the assignment of value. The navigation of the agent in this system is discretionary spatial behavior—wandering and lingering without prescribed destination. In designing the simulation of the ToU in interaction with the agent, the following mechanisms are involved.

A. Memory Dynamics: The Dual System and Decay

To capture how the ToU is adapted, the agent maintains a dual memory system, distinguishing between neutral knowledge and critical threat assessment:

Memory of Details (Detail Memory): stores the neutral knowledge, such as the visual arrangement and geometric configuration of a node.

Memory of Stress (Avoidance Memory): stores the emotional or punitive value associated with a node due to past exposure to threat. Crucially,

the agent might not remember the visual details but still retain the high negativity associated with the experience.

This knowledge is not static. Memory Decay ensures the map is not omniscient and introduces temporality into the agent's knowledge. As time passes, the resolution of rarely-visited states is lost, and their details are forgotten. The adapted rate of this decay will be defined mathematically in the methodology, but theoretically, it confirms that knowledge is dynamic and requires continuous reinforcement to be maintained.

B. Information Intake and Knowledge Fidelity

The agent's knowledge acquisition is locally constrained:

Peripersonal Space Radius: This defines the maximum radius within which the agent receives cues and updates their local chart. It sets the literal bounds of immediate awareness, outside of which knowledge relies entirely on memory.

Information Intake Zone (FOV): Perception is further constrained by the Field of View, operationalized as a limited, heading-dependent cone, reflecting limits imposed by human attention and occlusion.

Knowledge Degree & Error Rate: Information received within the intake zone is subject to distance-based distortion and an Error Rate, acknowledging the subjective nature of perception.

C. Value Assignment and Stress Propagation

The core of the ToU lies in value assignment—the breaking of spatial uniformity.

- The Stress Intensity: This is a factor used to evaluate the suitability of a node for exploratory behavior, dependent on the Foci of Uncertainty. Nodes closer to a threat have higher stress.
- Actual vs. Perceived Stress: Actual Stress is fixed by the objective distance to the surveillance source. Perceived Stress is the agent's internal assessment, which is influenced by the node's visibility, the decay rate of Memory of Stress, and the Error Rate in the current perception.
- The Avoidance Memory: The learned factor resulting from high remembered stress, forming the basis of decision-making and the Reinforcement Learning process.

Tyranny of Value

In ordinary navigation, cognitive maps balance multiple values: efficiency, social connection, aesthetic pleasure, and safety. Under persistent surveillance, the stress dimension progressively dominates all others. The agent's cognitive map doesn't shrink spatially but becomes tyrannized by a single factor that

overrides competing values. This is the critical distortion: not loss of knowledge, but loss of optionality as threat-avoidance subordinates all other navigation criteria.

D. The inner states

To introduce heterogeneity and modulate the agent's response to learned stress, a single critical factor governs the decision-making process: the Avoidance Threshold (TA).

> Avoidance Threshold: This is a defining, individualized property that determines the maximum level of Perceived Stress an agent is willing to tolerate before rejecting a state.
>
> Individualizing Factors: The agents exploration can also be modified to prioritize factors such as familiarity, avoidance, or perception. However, in this thesis the changing individualizer is avoidance threshold.

Agents with a high TA are more risk-taking (analogous to courage), proceeding along high-stress paths. Agents with a low TA are more hyper-vigilant and conservative, rapidly eliminating areas from their viable navigation space.

The TA acts as the decisive variable for fragmentation, as it determines the breaking point at which the agent's learned avoidance behavior overrides their exploratory action, leading to the collapse of navigational options.

3.6 Operational Bridge: From Theory to Simulation Variables

The theoretical framework is translated into a structured computational model by mapping each core concept onto a specific ABM variable or mechanism:

Theoretical Concept	ABM Operational Variable/ Mechanism	Role in Fragmentation
Physical Space	Graph Structure (Nodes & Edges)	Defines the physically fixed environment.
Peripersonal Space	Perception Zone	Governs limited, egocentric information intake and occlusion.
Cognitive Map (Details)	Spatial Knowledge Memory Layer	Stores declarative route and location detail.
Memory Decay	Exponential Decay Function	Determines the rate at which knowledge is forgotten over time.
ToU Agency / Stress	Stress Propagation Function	External input defining the negative incentive (risk) assigned to nodes.
Reinforcement Learning	Avoidance Memory Layer	Stores the learned threat value that overrides spatial detail.
Latent Learning	Detail Memory	Familiarity with space
Individuality	Avoidance threshold	Modulates the risk-acceptance thresh-old for decision-making.
Fragmentation Result	Breaking Point Analysis	Measures the quantitative collapse of the agent's viable navigation options.
Decision-Making	Weighted Utility Calculation	Combines all factors to determine the next action

Table 3.1 Operational bridge: translating theoretical framework into operational ABM concepts. [Author]

04 | Methodology
ABM

Methodology
Agent-based Modeling

This section details the computational architecture and formal rules used to translate the theoretical framework into an operational Agent-Based Model (ABM). The model simulates the dynamic fragmentation of the protagonist's cognitive map as they navigate the Terrain of Uncertainty, ToU. This exploratory simulation models discretionary spatial behavior, wandering and lingering without prescribed destination, rather than obligatory movement to functional spaces (classrooms, studios, cafe). This methodological choice isolates stress-driven navigation from necessity-driven movement.

ABM is chosen for its ability to model emergent behavior (fragmentation) from the local decisions of heterogeneous agents. The simulation is developed in NetLogo, utilizing its capacity for spatial modeling (patches) and agent interaction. Complete model specification following the ODD protocol (Grimm et al. 2020) is provided in Appendix A.

This experiment is an exploratory simulation. The formulas and parameters were calibrated through iterative testing to produce memory decay and stress accumulation rates that correspond to the timescale of semester-long spatial learning, creating a plausible dynamics. While specific values are subjective, the qualitative behaviors (initial exploration, gradual convergence, stress-driven avoidance) are robust across reasonable parameter ranges. The sensitivity analysis to set the parameters is provided in Appendix B.

4.1 The Agents

The simulation involves three distinct agent types:

- The Protagonist (Rabbit): The main actor following profile of female student described in chapter 2, responsible for navigation, decision-making, and memory processes.

- The Foci of Uncertainty: Patrols & Cameras: These are the hostile agents, or "predators," serving as the attractors of the stress representation. They are mobile (Patrols) or static (Cameras).

- The Environment (Patches): The environment is represented as a grid of spatial patches, each associated with values that change based on the interactions with other agents.

4.1.1 The Environment and patch Variables

The initial assumption is a uniform environment, which gradually becomes a personalized ToU as patches are experienced and navigated. The environment is a translation of the case study, information encoded in color-coded patches. Each patch is assigned a set of ToU variables and states:

Variable	Type	Role in Fragmentation
Passable?	Boolean (True/False)	Define the graph structure and occlusion for perception (isovist).
See-through?	Boolean (True/False)	Distinguishes visual occlusion (walls) from physical barriers (fences)
Time Since Last Visit (Tlast)	Numerical	Factor for calculating Memory Decay
Detail Memory(Md)	Numerical (0-1)	Stores neutral spatial knowledge.
Avoidance Memory (Ma)	Numerical (0-1)	Stores learned negative utility, Long-term memory of accumulated stress
Real-Time Stress (Sr)	Numerical (0-1)	Objective stress propagated by nearby threat agents.
Perceived Stress (Sp)	Numerical (0-1)	Protagonist's current perception of risk
Visit Frequency	Numerical	Number of times a patch is seen
Presence Frequency	Numerical	Number of time a patch is navigated
Knowledge Degree (K)	Numerical (0-1)	Fidelity of the information perceived

Table 4.1 The variables of environment patches. [Author]

4.1.2 Stress Propagation: Hostile Agents

Threat agents (Patrols and Cameras) define the ToU's hostility. They possess a Stress Propagation Radius and a stress intensity that can be modified to simulate different surveillance intensities. The result is an update in Actual Stress (Sr), which is the objective stress value propagated by a threat source to the patch, inversely related to the distance from the source.

This value is always present in the environment but is only sensed when the threat enters the protagonist's Perception Zone, updating the Perceived Stress (Sp) of the patch. The mobile threat has a movement pattern defined by the color-coded environment. An overview of threat agents' variables is shown in the following table:

Variable	Type	Role in Fragmentation
Stress Intensity	Numerical (0-1)	Strength of stress propagation
Propagation radius	Numerical	Radius of view-cone
FoV Angle	Angle (0-360°)	Angle of view-cone
Movement Pattern	Set of patches	Patrol navigation pattern (for mobile threats)

Table 4.2 *The variables of threat agents. [Author]*

4.1.3 Protagonist

The Protagonist is the decision-maker. The agent senses the stress and the familiarity of the patches in its peripersonal space, and by calculation of stress and memory decides the best patch to move to. The mechanism replicated a simplified version of planning: combining stimuli and LM in WM to reach a decision. The variables assigned to the protagonist are as follows:

Variable	Type	Role in Fragmentation
Heading	Angle (0-360°)	Current facing direction
Avoidance Threshold(Ta)	Numerical (0-1)	Individual risk tolerance (key individuality parameter)
Error Rate (Pe)	Numerical (0-0.1)	Perception error/misinterpretation factor
Goal-Patch	Coordinates	The destination patch selected in the decision-making algorithm.
FoV Angle	Angle (0-360°)	Angle of view-cone
FoV radius	Numerical	Radius of view-cone

Table 4.3 *The variables of protagonist. [Author]*

4.2. The Mechanisms

Four integrated algorithmic components govern the protagonist's decision-making.

4.2.1 Perception and Information Intake

The protagonist senses the environment within their Perception Zone, which the Peripersonal Space Radius and the Field of View cone define.

- The occlusion mechanism: First, the isovist of the protagonist is defined, as all the patches in the view cone that are in line of sight, excluding the patches that are not visible at the moment.

- Knowledge Degree: The fidelity of information of a patch is a direct product of distance from the agent, reflecting the decrease in fidelity of information as distance increases.

- Perceived Stress Calculation: Upon entering the Perception Zone, the protagonist calculates the subjective risk for patch i:

  - *SP(i) = Sr(i) * K(i) * (1 + random-float (2 * Pe) - Pe)*

  - Where Sp(i) is the currently sensed stress, K is the knowledge degree, and Pe is the perception error, which adds a stochastic factor to account for misinterpretation.

4.2.2 Memory Dynamics

The two memory layers are updated based on the Time Since Last Visit:

Memory Acquisition: When the protagonist occupies a patch, memory of the occupied patch and the patches in the view cone increases. MA is updated through Reinforcement Learning (RL) based on experienced perceived stress SP.

$$MA(i) = MA(i) + \alpha * SP(i) * K(i),$$

Where α is the learning rate. Avoidance Memory also moderately increases for patches with visible stress within the perception zone.

MD is updated in the same manner, without the effect of SP.

$$MD(i) = MD(i) + \beta * K(i)$$

Where β is the learning rate for memory of detail.

Memory Decay Function: Memory decays exponentially over time, confirming that knowledge is dynamic and requires continuous reinforcement to remain current.

$$M = M0 * (1 - (time\text{-}last \wedge cd))$$

Where cd is the decay constant, allowing for different rates between Cd-avoidance and Cd-detail (e.g., Cd-avoidance < Cd-detail to model persistent fear). The decay is fast at the beginning and slow at the end, and as the frequency of visiting a patch increases, the decay becomes slower and decreases. [J.E.R. Staddon, I.M. Chelaru, J.J. Higa b]

4.2.3 Decision-Making

The protagonist's decision to move operates in a two-phase cycle, moving beyond simple node-to-node comparison to a broader environmental assessment.

Attractiveness (A) Calculation (Weighted Utility): Since the agent is wandering without a specific goal, attractiveness balances baseline familiarity against avoidance cost.

$$A(i) = w_Base \times MD(i) - w_Avoid \times MA(i) - w_Perc \times Sp(i) + 1/(visit\text{-}frequency + 1)$$

Where w_Base and w_Avoid and w_Perc are weighting coefficients defining the agent's general motivational priorities. When MA > TA, Attractiveness collapses to 0. To avoid the trap of staying forever in the familiar patches, the patches that have not been visited have higher attractiveness (1/(visit-frequency + 1)). This formulation operationalizes the '*tyranny of value*' hypothesis: as MA increases, the avoidance term dominates, causing other criteria (familiarity, curiosity) to collapse to zero once MA > TA.

A. Phase 1: Area Assessment (Exploration)

The protagonist first assesses the surrounding area to determine the viability of current direction for a safe navigation.

View Cone Scanning: The agent assesses current heading direction, calculating the average Attractiveness of the visible patches along with number of safe patches. If the direction proves to be dangerous with a low attractiveness or a high number of dangerous patches, the agent rotates and sample a new heading direction.

Best Direction Selection: The agent selects the best direction that maximizes the attractiveness.

B. Phase 2: Local Movement (Selection)

The agent commits to the optimal adjacent patch j within the best direction.

Threshold Check (Psychological Erasure): The agent filters the best patches in the selected direction j, and chooses the best patch as the goal-patch.

Final Movement Selection: The agent moves one step toward the selected goal-patch. If no goal-patch exists or the agent gets stuck, it attempts to move randomly and find a new goal-patch. A degree of randomness is added to the movement to ensure more human-like behavior.

This two-step process formalizes the idea that the agent's macro-level decision (which way to turn) is based on a wide scan of the ToU, while the micro-level decision (the exact step) is based on local node utility within that chosen, safer direction.

4.2.4 The Fragmentation Mechanism (Avoidance Threshold)

The Avoidance Threshold (TA) is the key to fragmentation, acting as a filter in both the Area Assessment and Local Movement phases. Any path that exceeds this subjective tolerance is psychologically erased, directly constraining the agent's available route options.

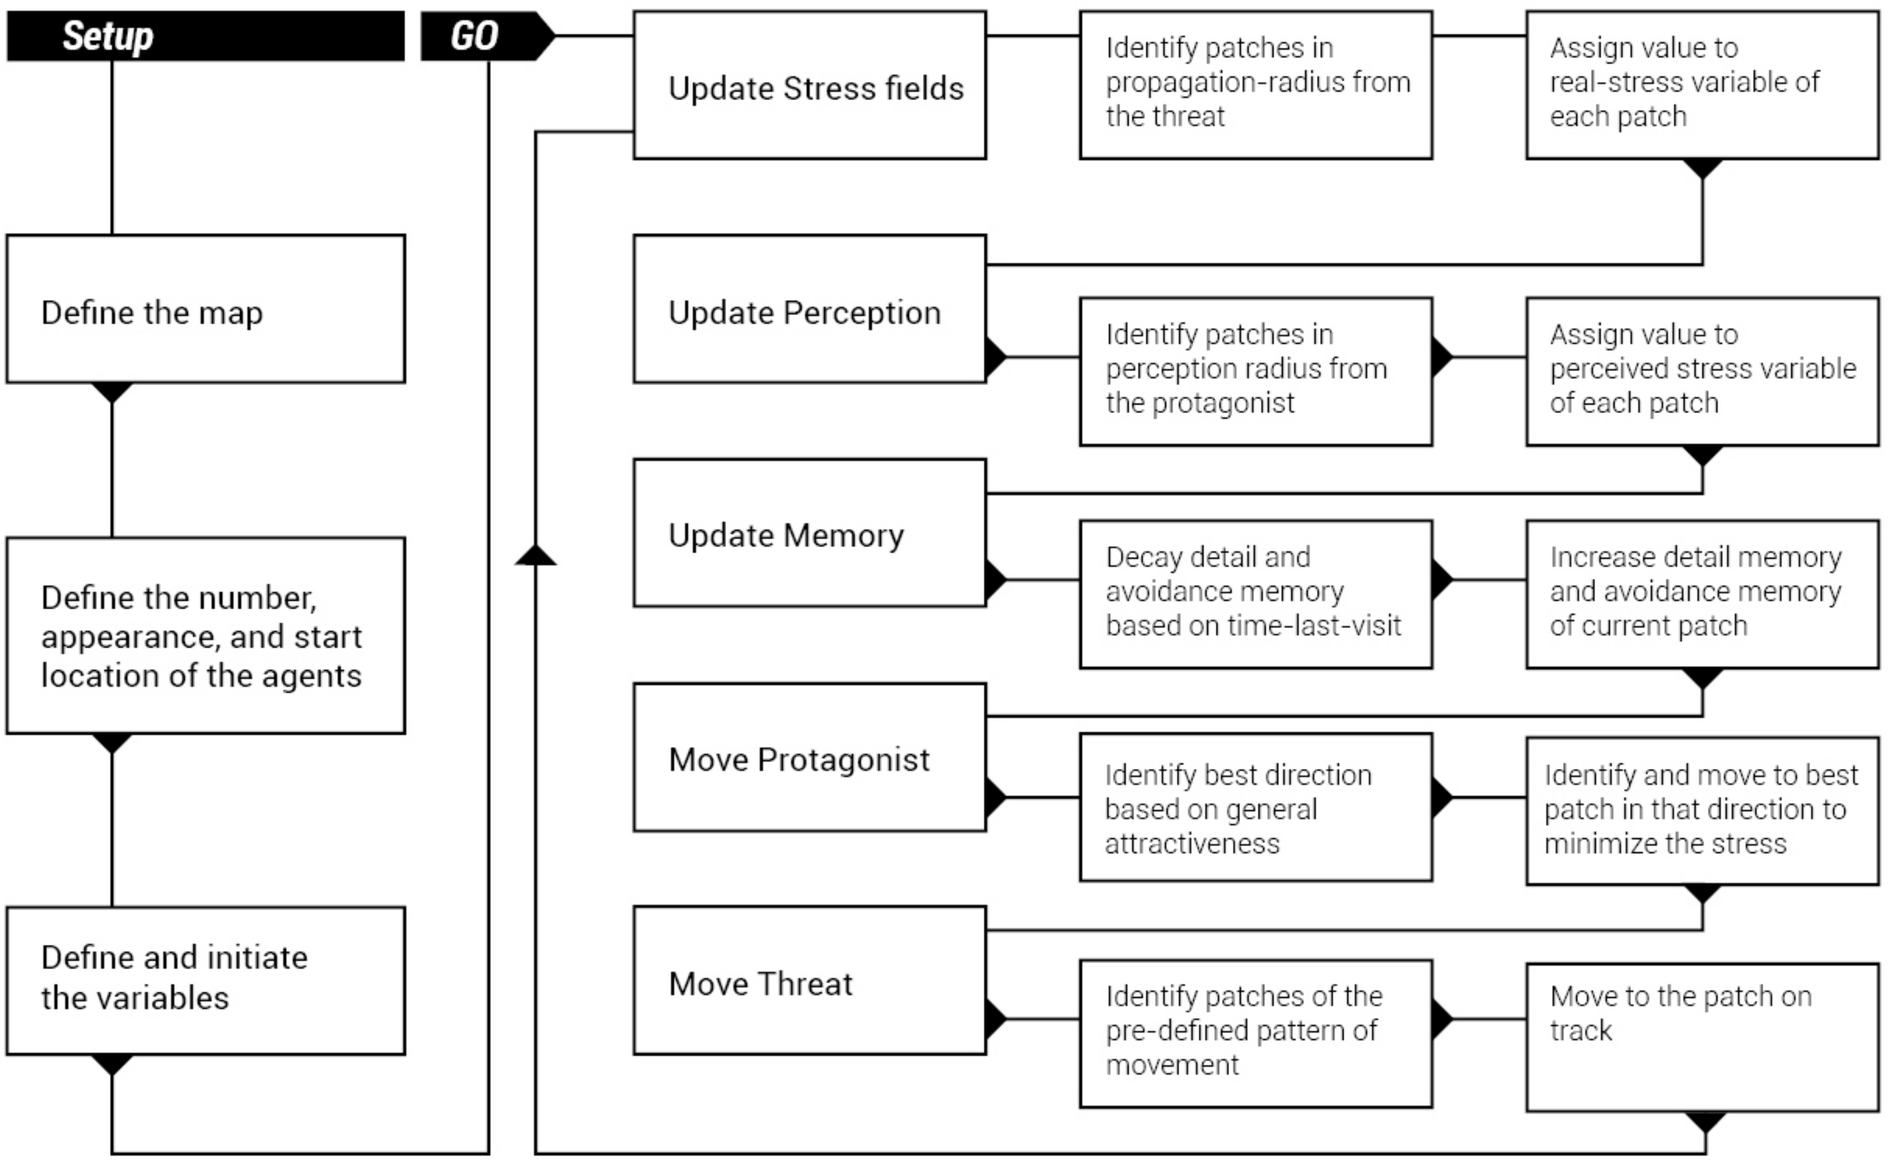

Figure 4.1 *The loop of main commands in the simulation. [Author]*

4.3 The Parameters

The parameters introduced in the model were set to following values:

Parameter	Value in the Model	Description
FoV Radius	30 m	Maximum sensing distance
FoV Angle	120 °	View-cone angle
Detail Decay Constant	6 - 8	Fast decay rate
Avoidance Decay Constant	8 - 10	Slow decay rate
Learning Rate	0.05 , 0.1	Avoidance memory acquisition rate
Base Weight	-0.2	Familiarity preference, set negative to encourage exploratory behavior
Avoidance Weight	1	Avoidance effect in decision making
Perception Weight	1	Perception effect in decision making
Avoidance Threshold	0.2-0.9	Individual risk tolerance (varied)
Error Rate	0.1	±10% perception noise

Table 4.4 Parameters of the simulation. [Author]

4.3.1 Sensitivity analysis

To determine the suitability of the parameters to match the heightened risk perception reported in female students' experience, I performed a series of sensitivity analysis runs, the most important of which are:

1. Decay constants: in the first scenario, using respective numbers 2 and 4, the decay is very fast, disrupting the process of memory accumulation. By increasing the parameters to 8 and 10, decay is much slower and almost not effective on the learning process. Values 6 and 8 prove to be matching the effectiveness of decay in the simulation conditions. For more details go to appendix B.

2. Learning rate: In the 0.01 , 0.02 scenario, learning is too slow and inefficient, however the results would not change, though the simulation time needed to reach the breaking point increases significantly. By increasing the learning rates to 0.1, 0.2, learning happens faster than ordinary, making double visits of a patch almost non existent, although the results show the same pattern of avoiding the threats.

3. Stuck threshold: The values are not much effective on the results, however, if the threshold is too high, it can produce unwanted detail memory solely because the agent is trapped and has no options.

4.4 The Model: Simulating the Faculty of Fine Arts

Following the definition of the ABM, the model is applied to the environment of the case study, the Faculty of Fine Arts of the University of Tehran. The site of the faculty is used to define environmental variables, which play a fundamental role in determining the perception zone, patrol patterns, and threat locations. The site is simplified to its main components: The visibility and navigation occluders (main obstacles that limit visibility), the navigational obstacles that won't affect visibility (fences, elevation changes), and the passable patches. In addition to this, the patrol movement pattern and the location of security cameras are also parts of this translation of the environment. The environment is simplified and introduced to the site as a raster image, and the information is encoded as colored pixels.

Component	Type	Role in Fragmentation
Site Map	Color-coded patches	Patches colored in white(passable), black (buildings) and gray (obstacles, such as fences)
Surveillance Locations	List of patch IDs	Fixed camera positions
Patrol Routes	Color-coded patches	Set of red-colored patches

Table 4.5 *Translation of Faculty of Fine Arts into simulation patches. [Author]*

The simplified faculty map is shown in the figure. The map was reduced to the main occluders and obstacles to reduce the unwanted noise of interaction between the agent and edges/corners. The Faculty of Fine Arts was discretized into a 117 × 210 patch grid (≈1m per patch), yielding 14,806 passable patches after excluding buildings and obstacles. Each passable patch connects to up to 8 adjacent neighbors, forming a navigation graph of potential routes.

The simulations ran for 10000 samples under different scenarios, one control scenario of free navigation and three scenarios involving the protagonist with variegated values of avoidance threshold. Upon completion, the patches with the highest visit frequency and detailed memory were detected, effectively clustered around the gathering places mentioned before.

Figure 4.2 *The prototype of the model. a: the green(the protagonist) and the red(the threats) start to move, leading to the formation of environmental memory, and the intensity of the memory is shown by the whiteness of the patch, b: the memory map after a while, c: the map of frequently visited patches(red) and detail memory of the environment (decreasing from light to dark), d: the map of avoidance memory, decreasing from black to blue. Although the prototype, the results show the fragmentation of space for the protagonist. [Author]*

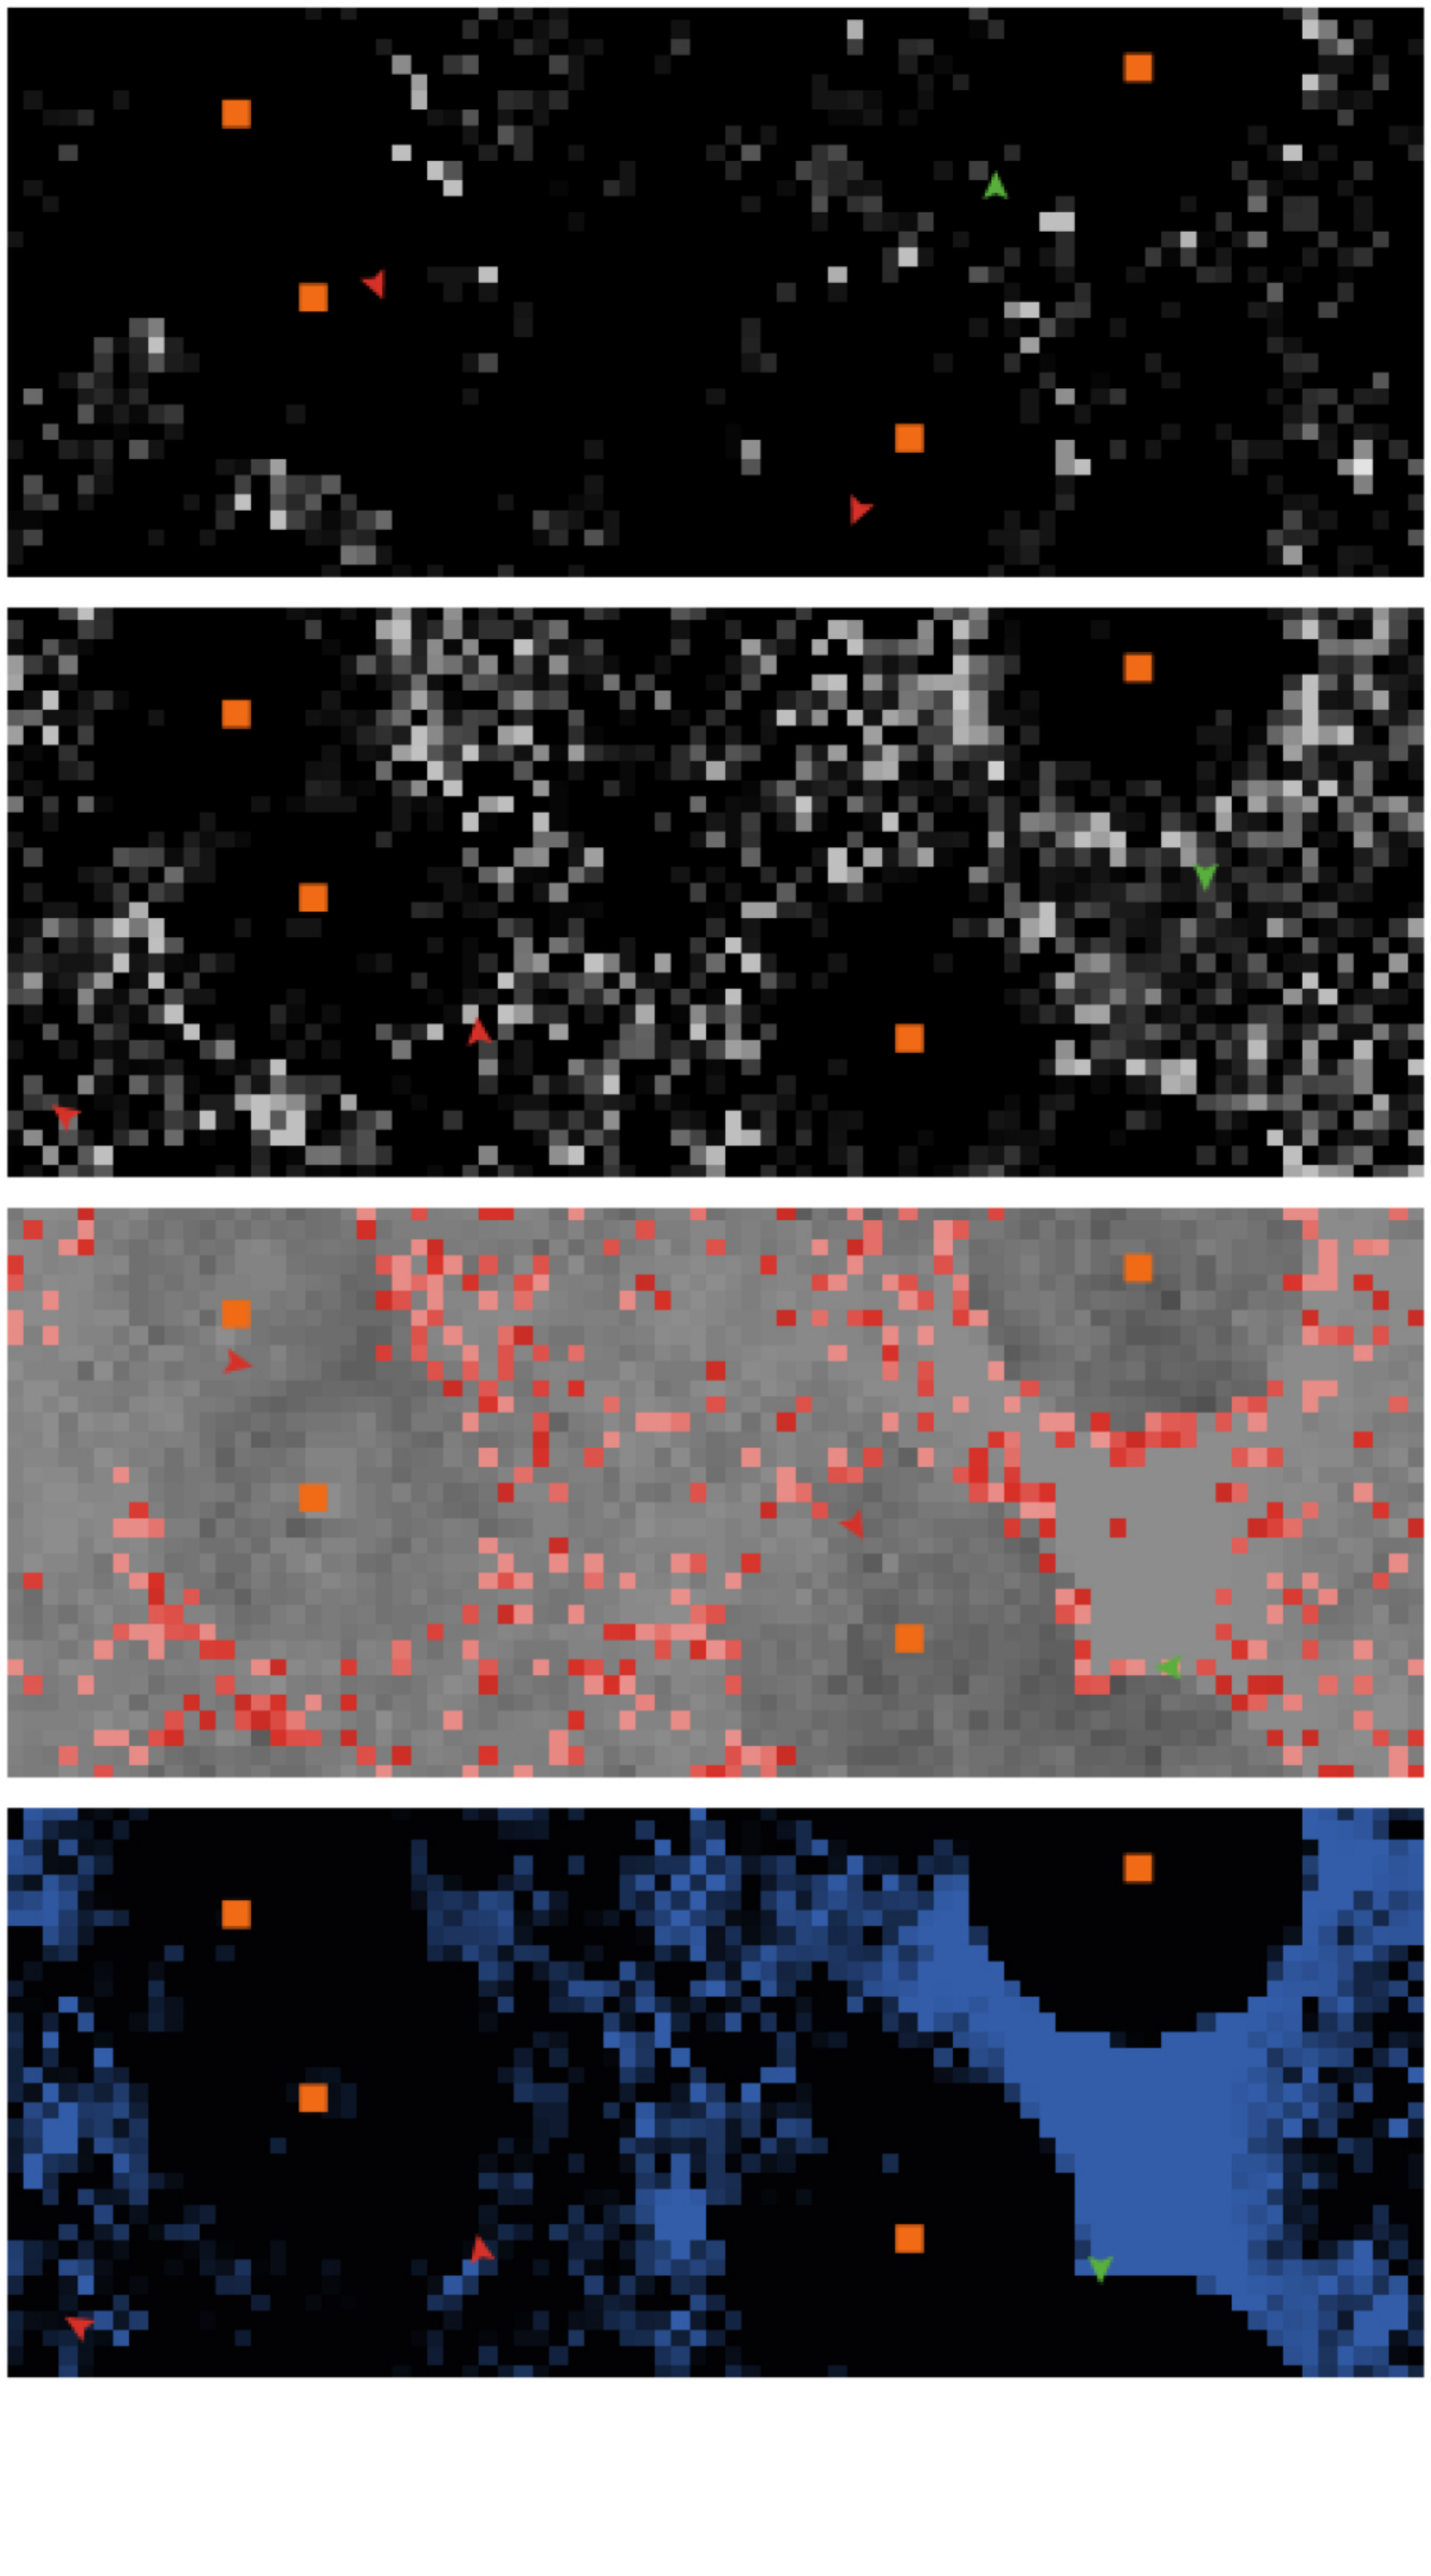

a

b

C

d

The patches were then visualized as nodes of various sizes based on how memorized and visited they are, and through pruning the patches with low detail memory and scarcely visited, the analysis showed a significant reduction in the number of viable edges for the agent, and the accessible world collapsed at least in half. It is noticeable that some places are not at all visited, although open, since the stress levels at the entrances are higher than accepted levels for the agent, even though the agent has a 10 percent chance of wandering randomly.

The decrease of options is also seen in the views, as highly stressed areas are more navigated in routes, sequential and narrow, and safe spaces are rather a collection of accessible, rather equally memorized patches.

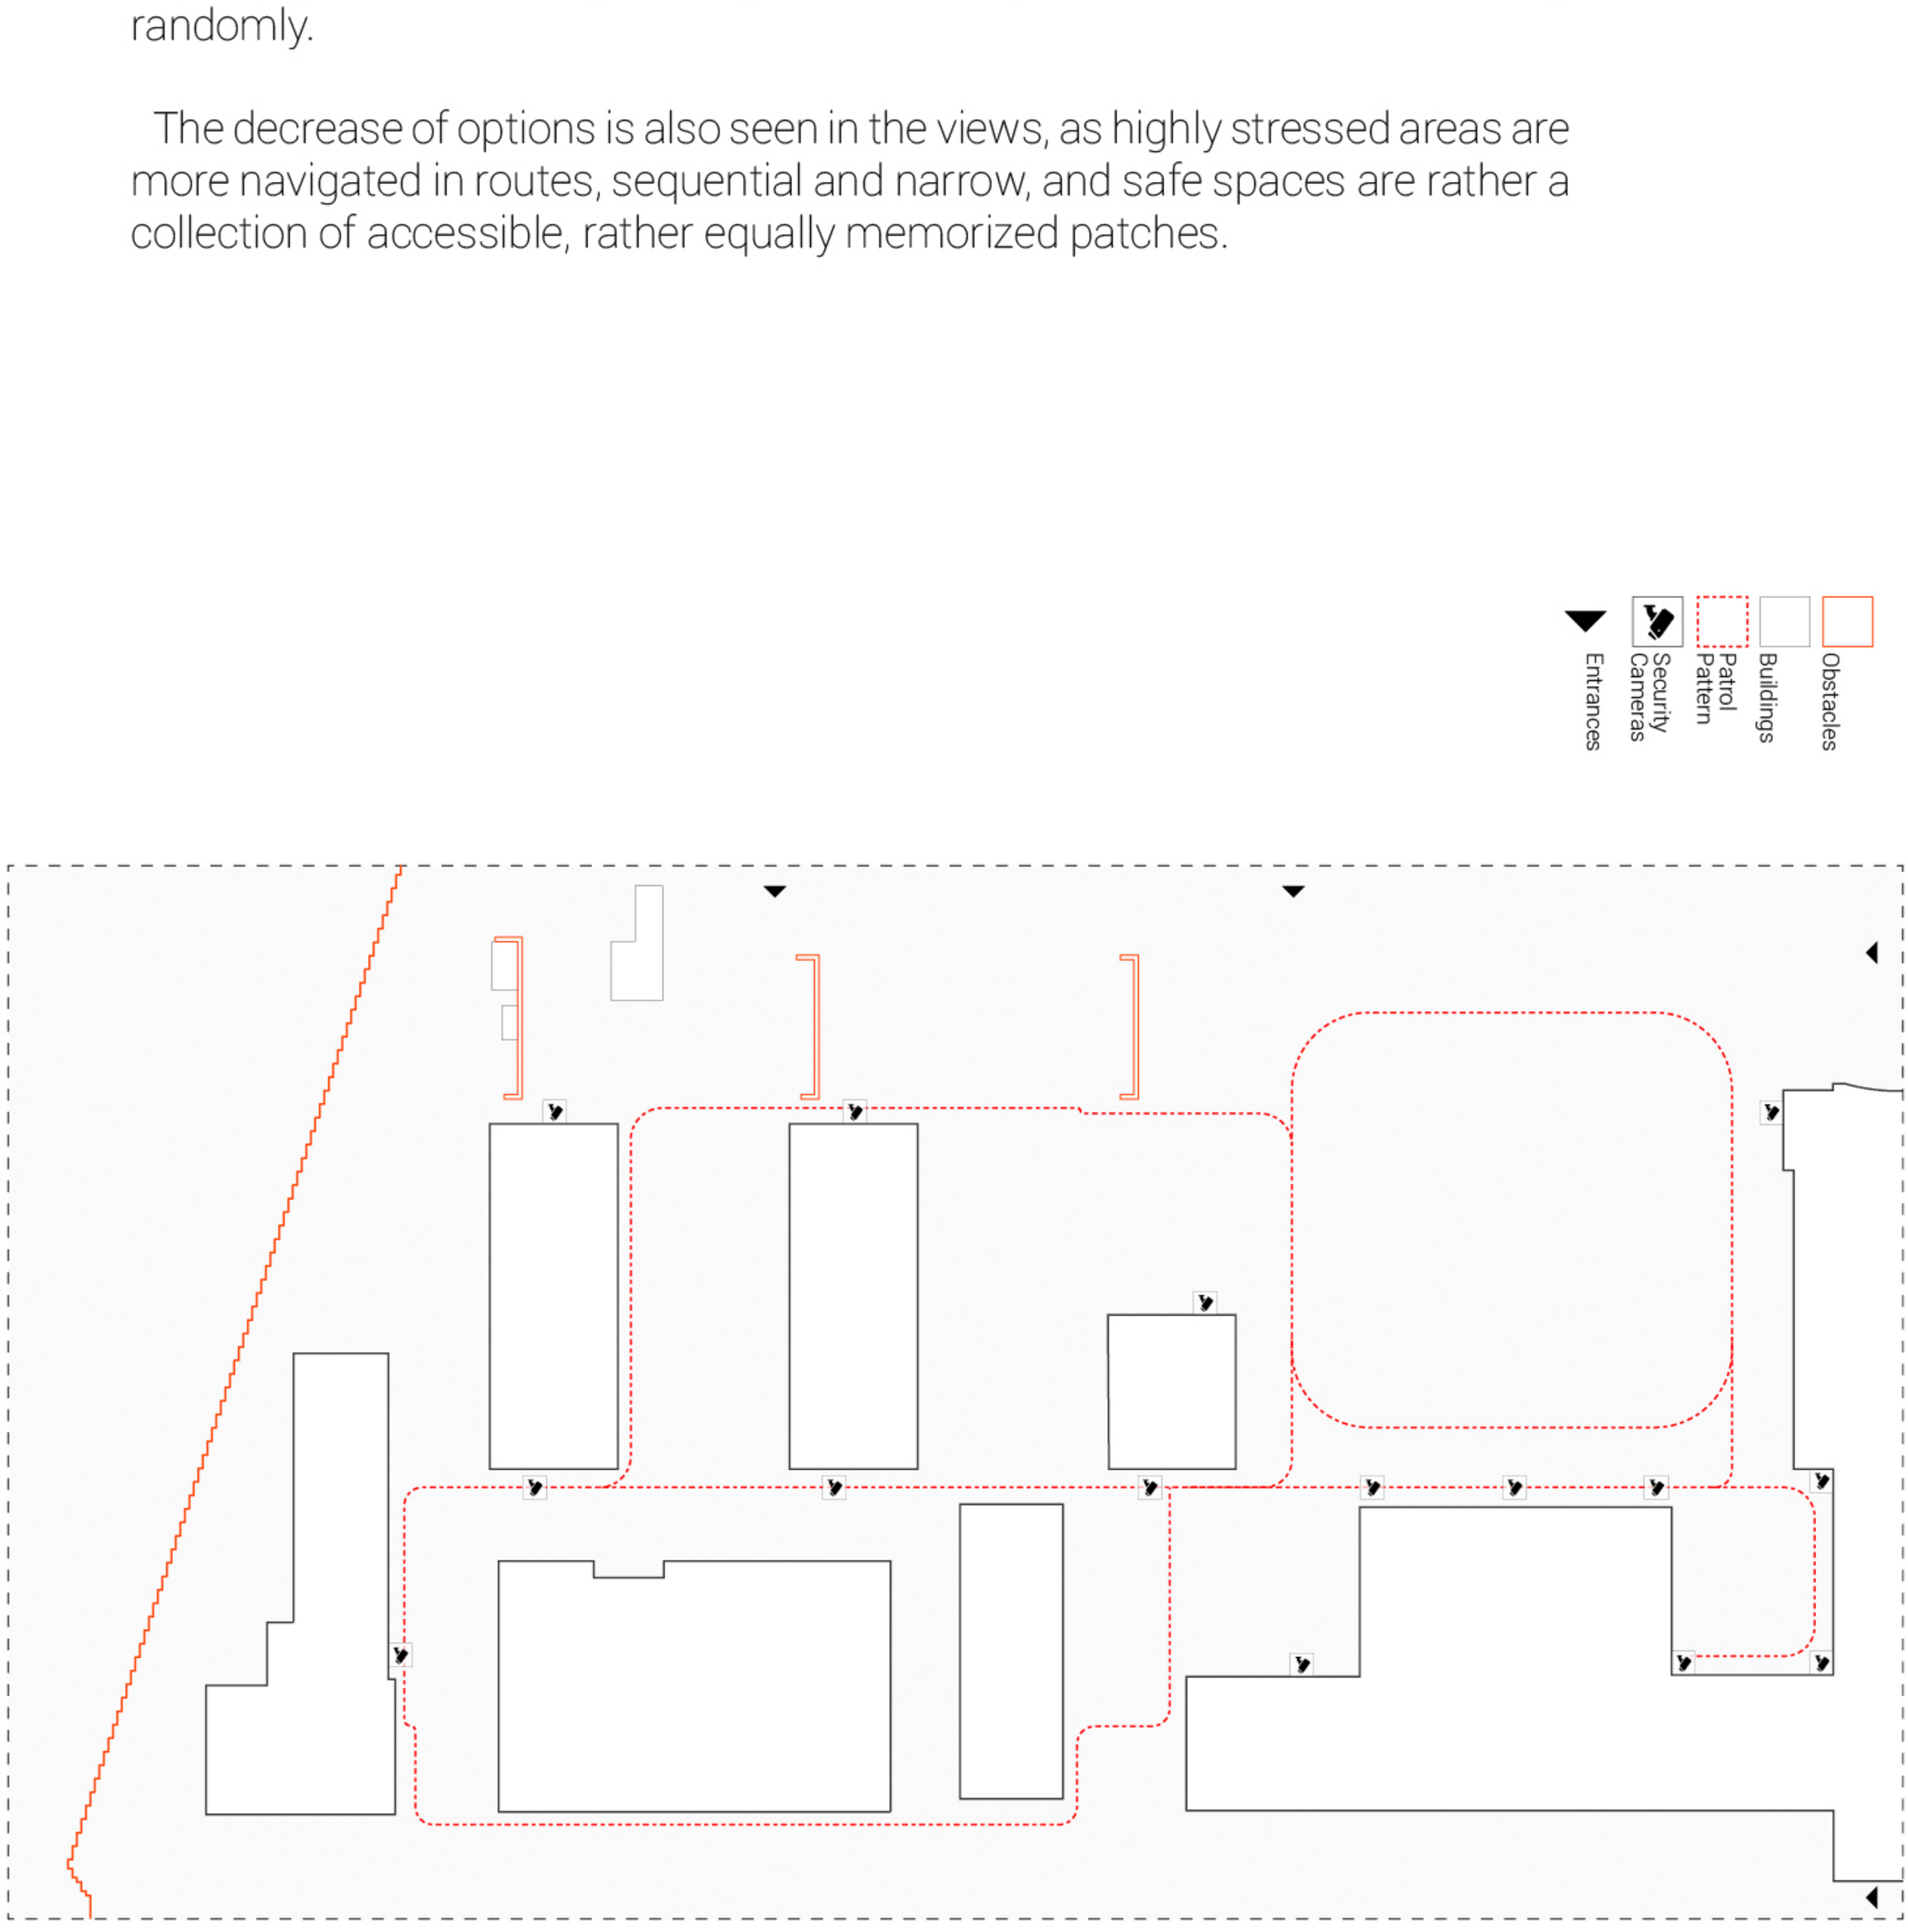

Figure 4.3 *Translation and simplification of Faculty to environment patches in Netlogo. [Author]*
Figure 4.4 *Map of individual safe spots for discretionary time [Author]*

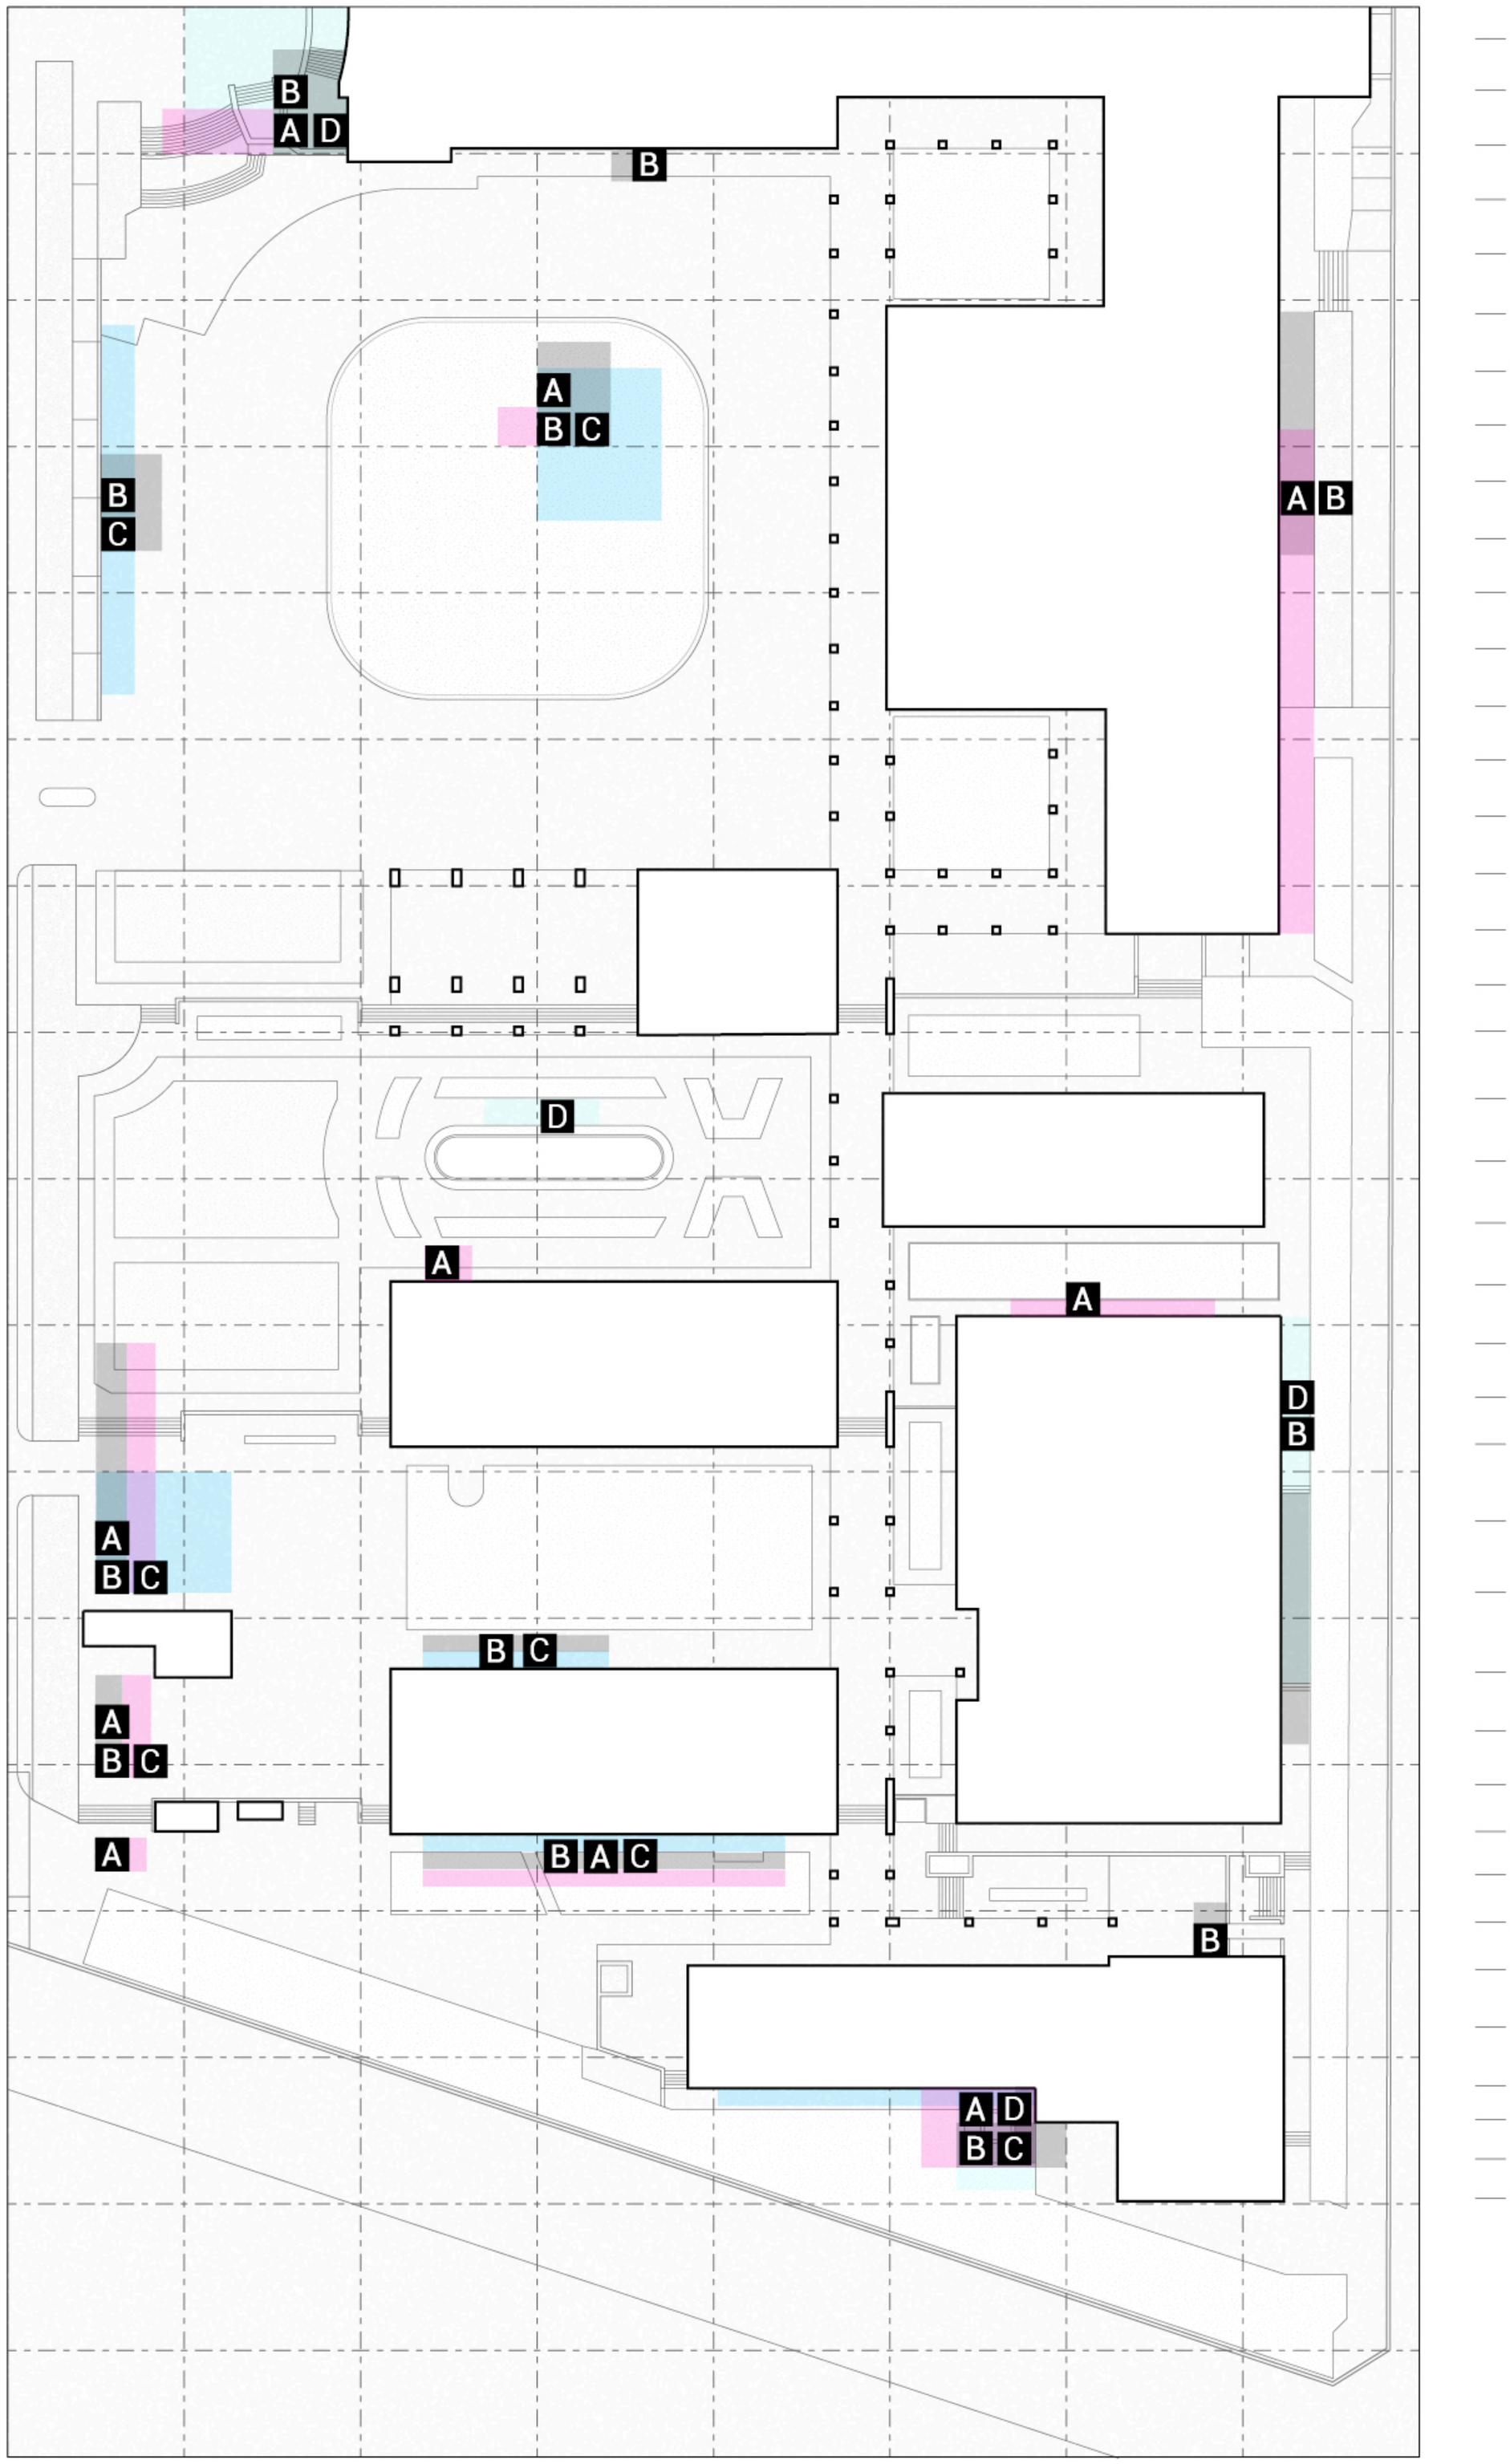

0 5 10 20m

05 | Results

Results

The fragmented landscape

This chapter presents the results of the experiment across four scenarios with varying avoidance thresholds. The first section provides quantitative results under each scenario, followed by comparative visualizations of the experienced landscape. What emerges is not a linear relationship between risk tolerance and spatial constraint, but a more troubling pattern: surveillance fragments the cognitive map in ways that cannot be predicted simply by measuring fear.

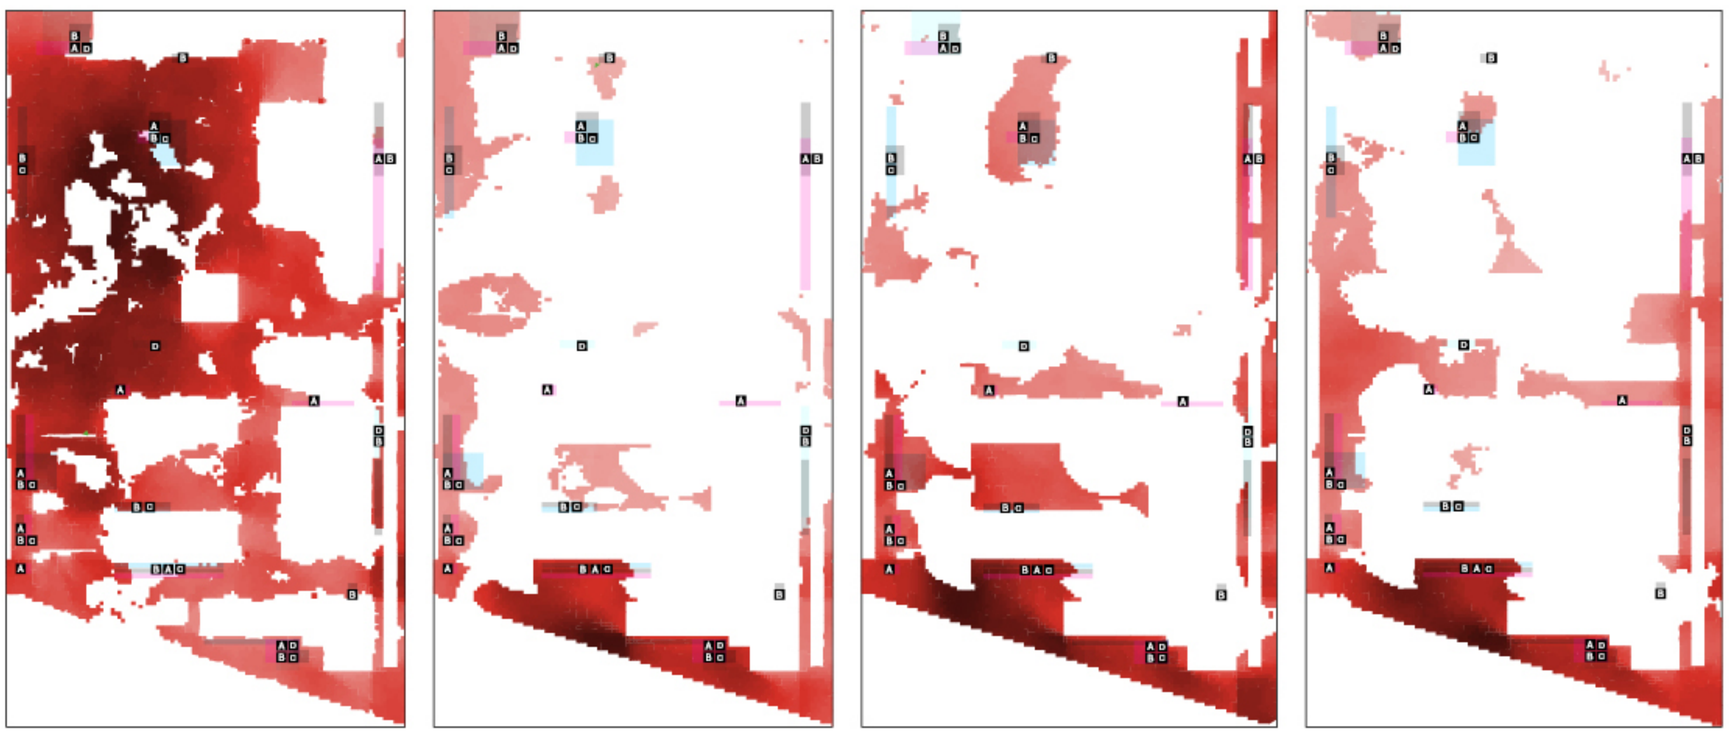

Figure 5.1 *Four scenarios of fragmentation based on 4 individual characters: a. Control scenario, the privileged explorer, b. Mildly brave with a TA of 0.4, c. Risk-avoider with a TA of 0.2, and d. The brave with a TA of 0.9. The intensity of red indicates the intensity of familiarity with an area, and fragmentation is vividly visible in all of the surveilled cases. [Author]*

5.1 First scenario: Control, the privileged explorer

In this scenario, stress intensity propagated by the security system is set to zero, establishing a control condition in which the agent moves freely. The agent wanders and explores without constraint.

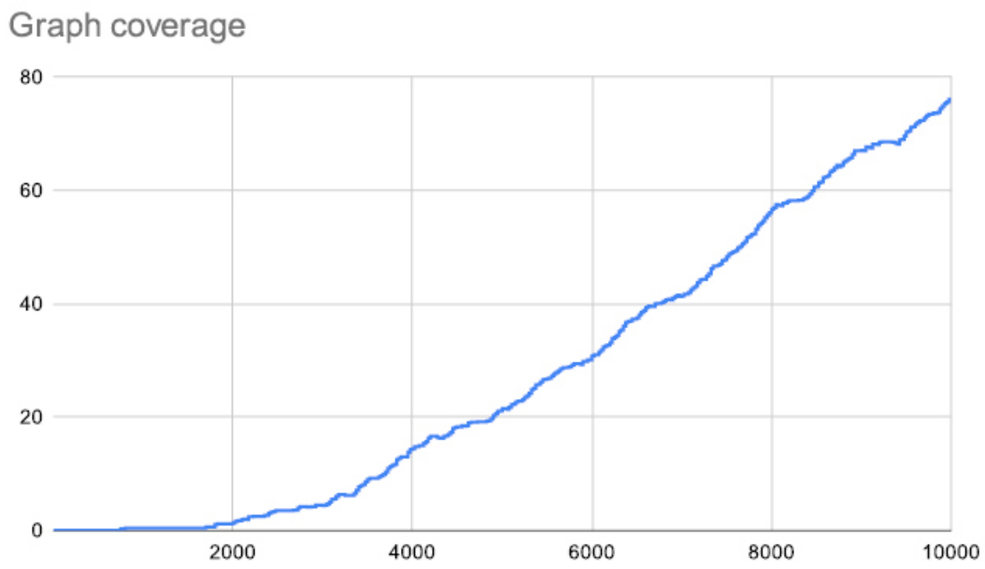

Graph 5.1 Graph coverage over time shows a semi linear increase over time, ceasing solely due to timeout.

The exploration graph shows consistent exploration rates, with graph coverage—the percentage of known and explored area relative to total accessible area—constantly increasing. When the simulation stopped at 10,000 ticks, graph coverage had reached 76.1%, and the memory maps show diffused, nearly equal intensity across the entire environment. Total movements indicate the agent made navigational choices more than 80% of the time, rarely freezing or becoming paralyzed by indecision.

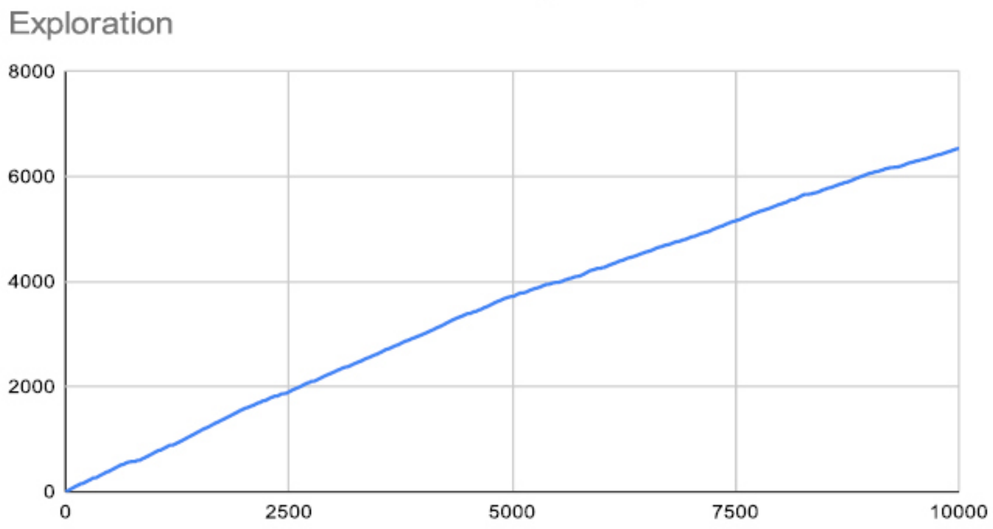

Graph 5.2 Exploration rate is increasing in an almost linear way, ceasing solely due to timeout.

Graph Structure		
Active Nodes	Nodes used and remembered	11262
Pruned nodes	Nodes not used/ not memorized	13754
Graph coverage	Active nodes / total passable nodes	76.1%
Navigation Usage		
Nodes navigated	Nodes used for navigation	6537
Edges used	Navigated edges	8070
Total movements	Total movements	8514
Avg degree	Average edges of a passable node	0.72

***Table 5.1** Control Scenario - Graph structure and navigation statistics over the simulation. [Author]*

The baseline is clear: without surveillance, the agent maintains a distributed cognitive map, explores continuously, and moves decisively.

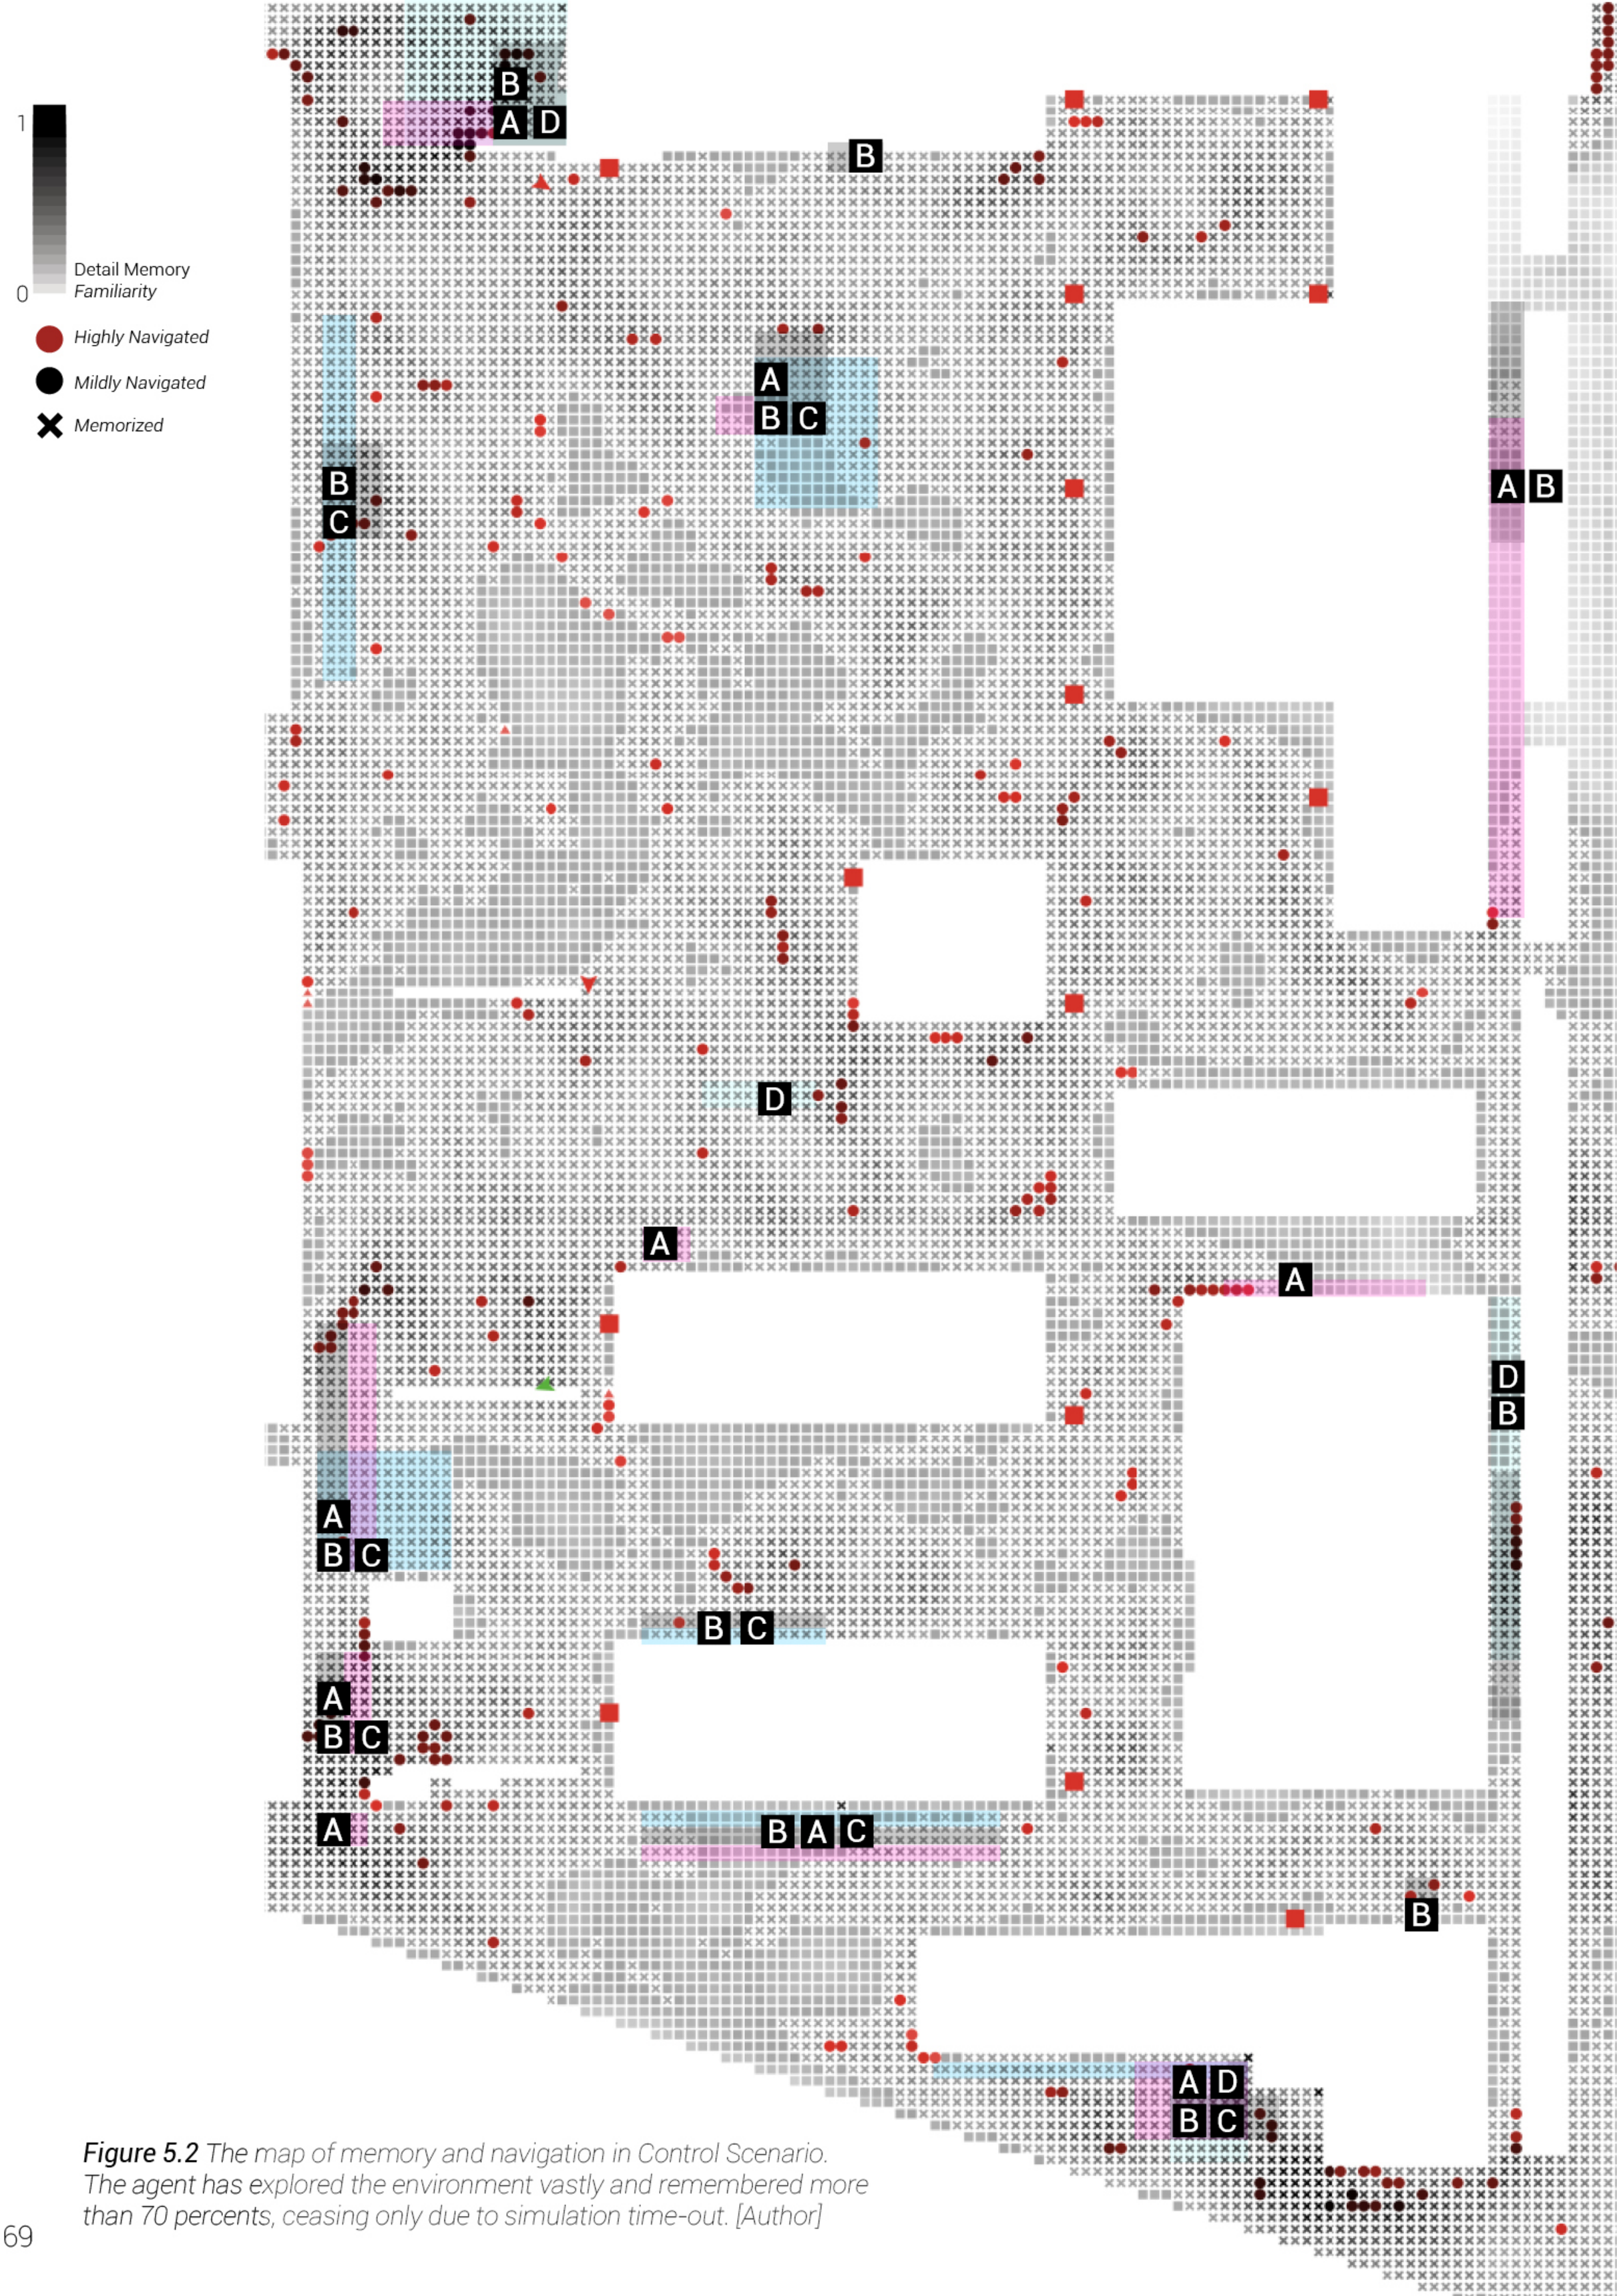

Figure 5.2 *The map of memory and navigation in Control Scenario. The agent has explored the environment vastly and remembered more than 70 percents, ceasing only due to simulation time-out. [Author]*

1

0

Detail Memory
Familiarity

Not Navigated

Navigated

Figure 5.3 *The map of detail memory in Control Scenario. There is no meaningful correlation between the mapped safe spaces and the memorized patches [Author]*

5.2 Second scenario: mildly brave explorer

The second agent has an avoidance threshold of 0.4, meaning it can reluctantly move through areas with mild avoidance memory but will avoid patches that exceed this threshold. The results show a drastic collapse: graph coverage drops by 50.2 percentage points to just 25.9%. The exploration graph reaches its climax well before the final sample, indicating the agent stopped encountering novel areas and began circulating within the safe zones it had already identified. More telling is the shift in navigation topology. The quantity of navigated nodes and edges drops significantly (-42% and -22% respectively), yet total movements decrease only slightly (-9%). This produces a 127% increase in average degree, from 0.72 to 1.64, confirming that the agent is no longer exploring efficiently but instead clustering and repeating paths within a constrained set of familiar, memorized patches.

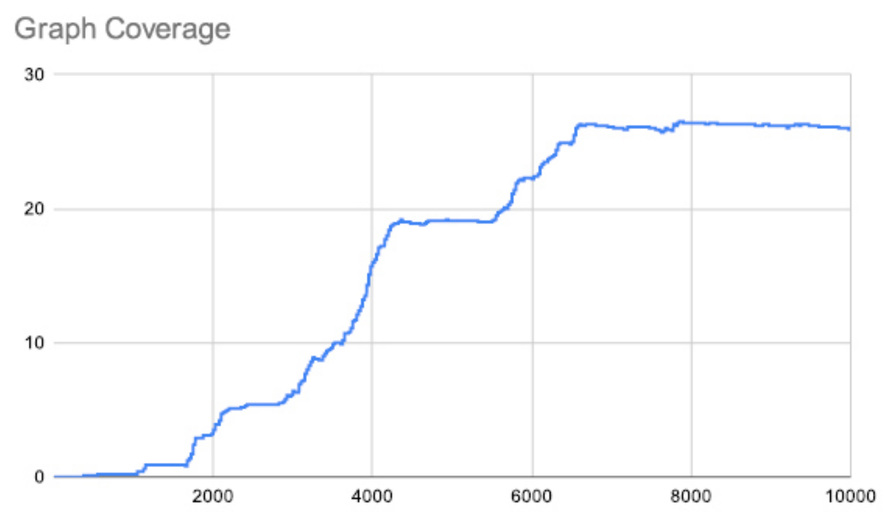

Graph 5.3 *Graph coverage over time shows the major expansions of the navigational graph and its flattening after around 8000 runs at ~25.9%.*

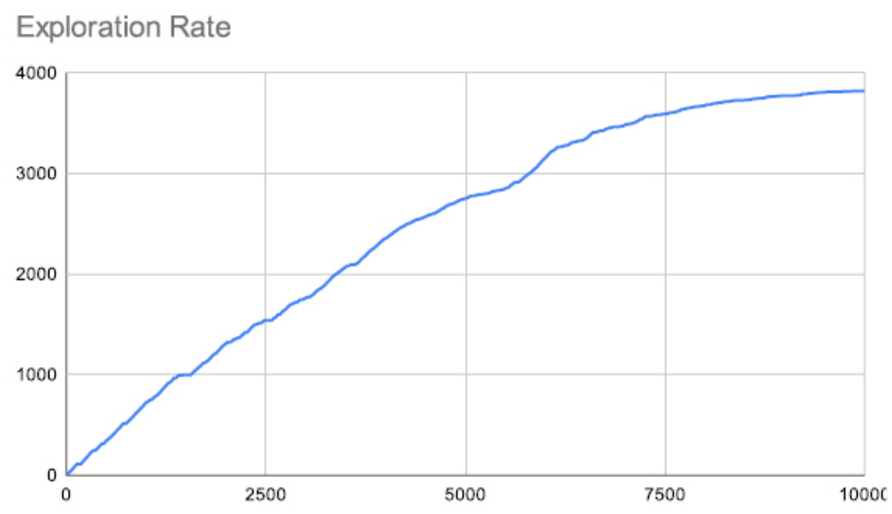

Graph 5.4 *Exploration over time shows the rate of navigation in new areas in the simulation. The flattening of the curve indicates breaking point and end of inhabitation.*

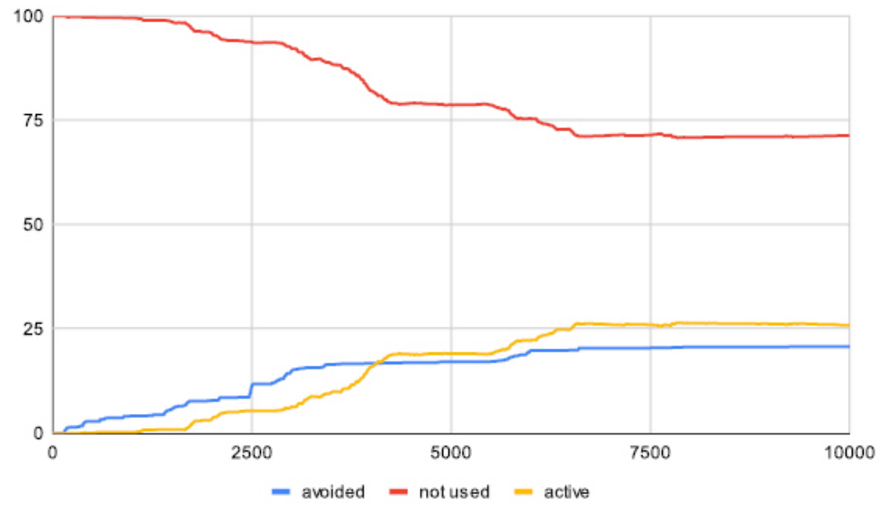

Graph 5.5 *Percentage of avoided, active and never-visited areas over time. For this agent, active patches are slightly more than the avoided, and the existence of Inaccessible void is evident.*

Graph Structure		
Active Nodes	Nodes used and remembered	3841 (-66 %)
Pruned nodes	Nodes not used/ not memorized	21175 (+53%)
Graph coverage	Active nodes / total passable nodes	25.9% (-50.2%)
Navigation Usage		
Nodes navigated	Nodes used for navigation	3822 (-42%)
Edges used	Navigated edges	6315 (-22%)
Total movements	Total movements	7743 (-9%)
Avg degree	Average edges of a passable node	1.64 (+127%)

Table 5.2 Mildly Brave Scenario - Graph structure and navigation statistics over the simulation. [Author]

The Emergence of the Inaccessible Void

A critical finding appears in these numbers. The graph shows that patches fall into several categories: avoided (pruned due to high avoidance memory), actively used, and a third category—patches that are neither avoided nor used. These are the Inaccessible Void: patches that remain physically passable and psychologically acceptable (avoidance memory below threshold) but have become structurally cut off from the agent's reduced navigation network. They are not banned; they are simply unreachable through the remaining viable paths. The agent cannot get there anymore, even though nothing directly prohibits it. In this scenario, 10,965 patches are pruned from the graph, yet fewer than 10,965 patches exceed the avoidance threshold. The difference is the Inaccessible Void—spaces lost not through direct threat, but through network fragmentation.

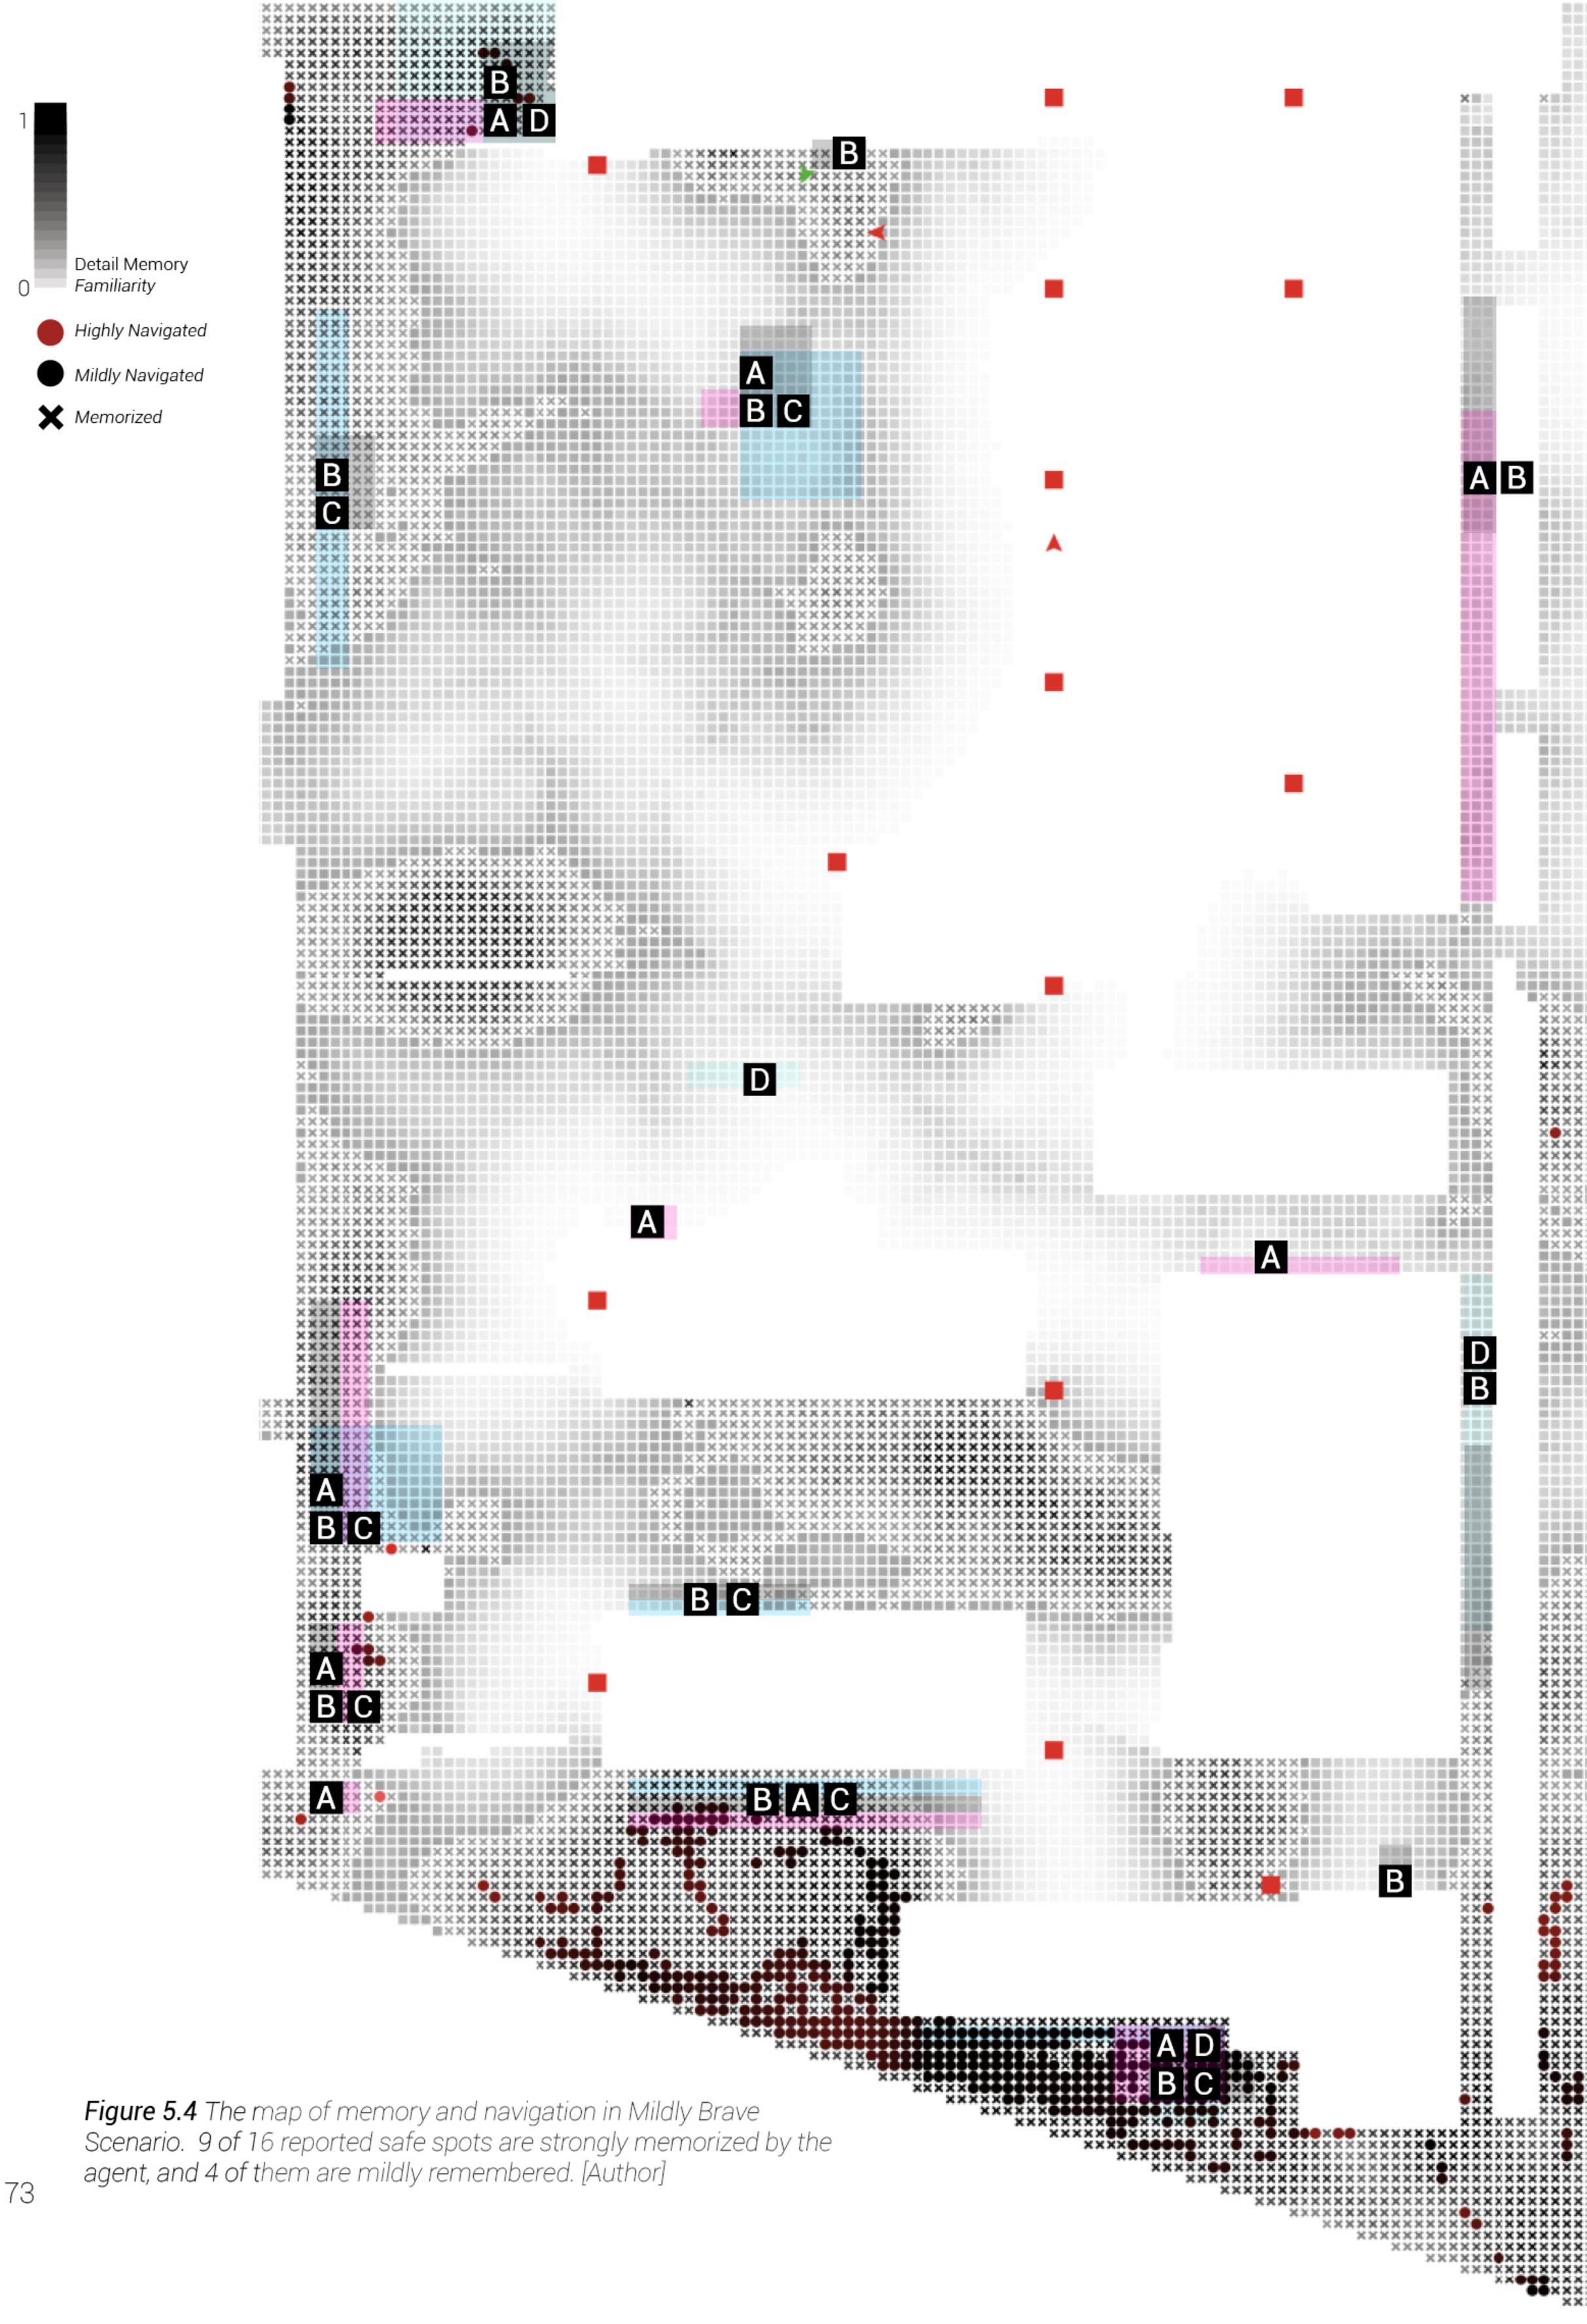

Figure 5.4 The map of memory and navigation in Mildly Brave Scenario. 9 of 16 reported safe spots are strongly memorized by the agent, and 4 of them are mildly remembered. [Author]

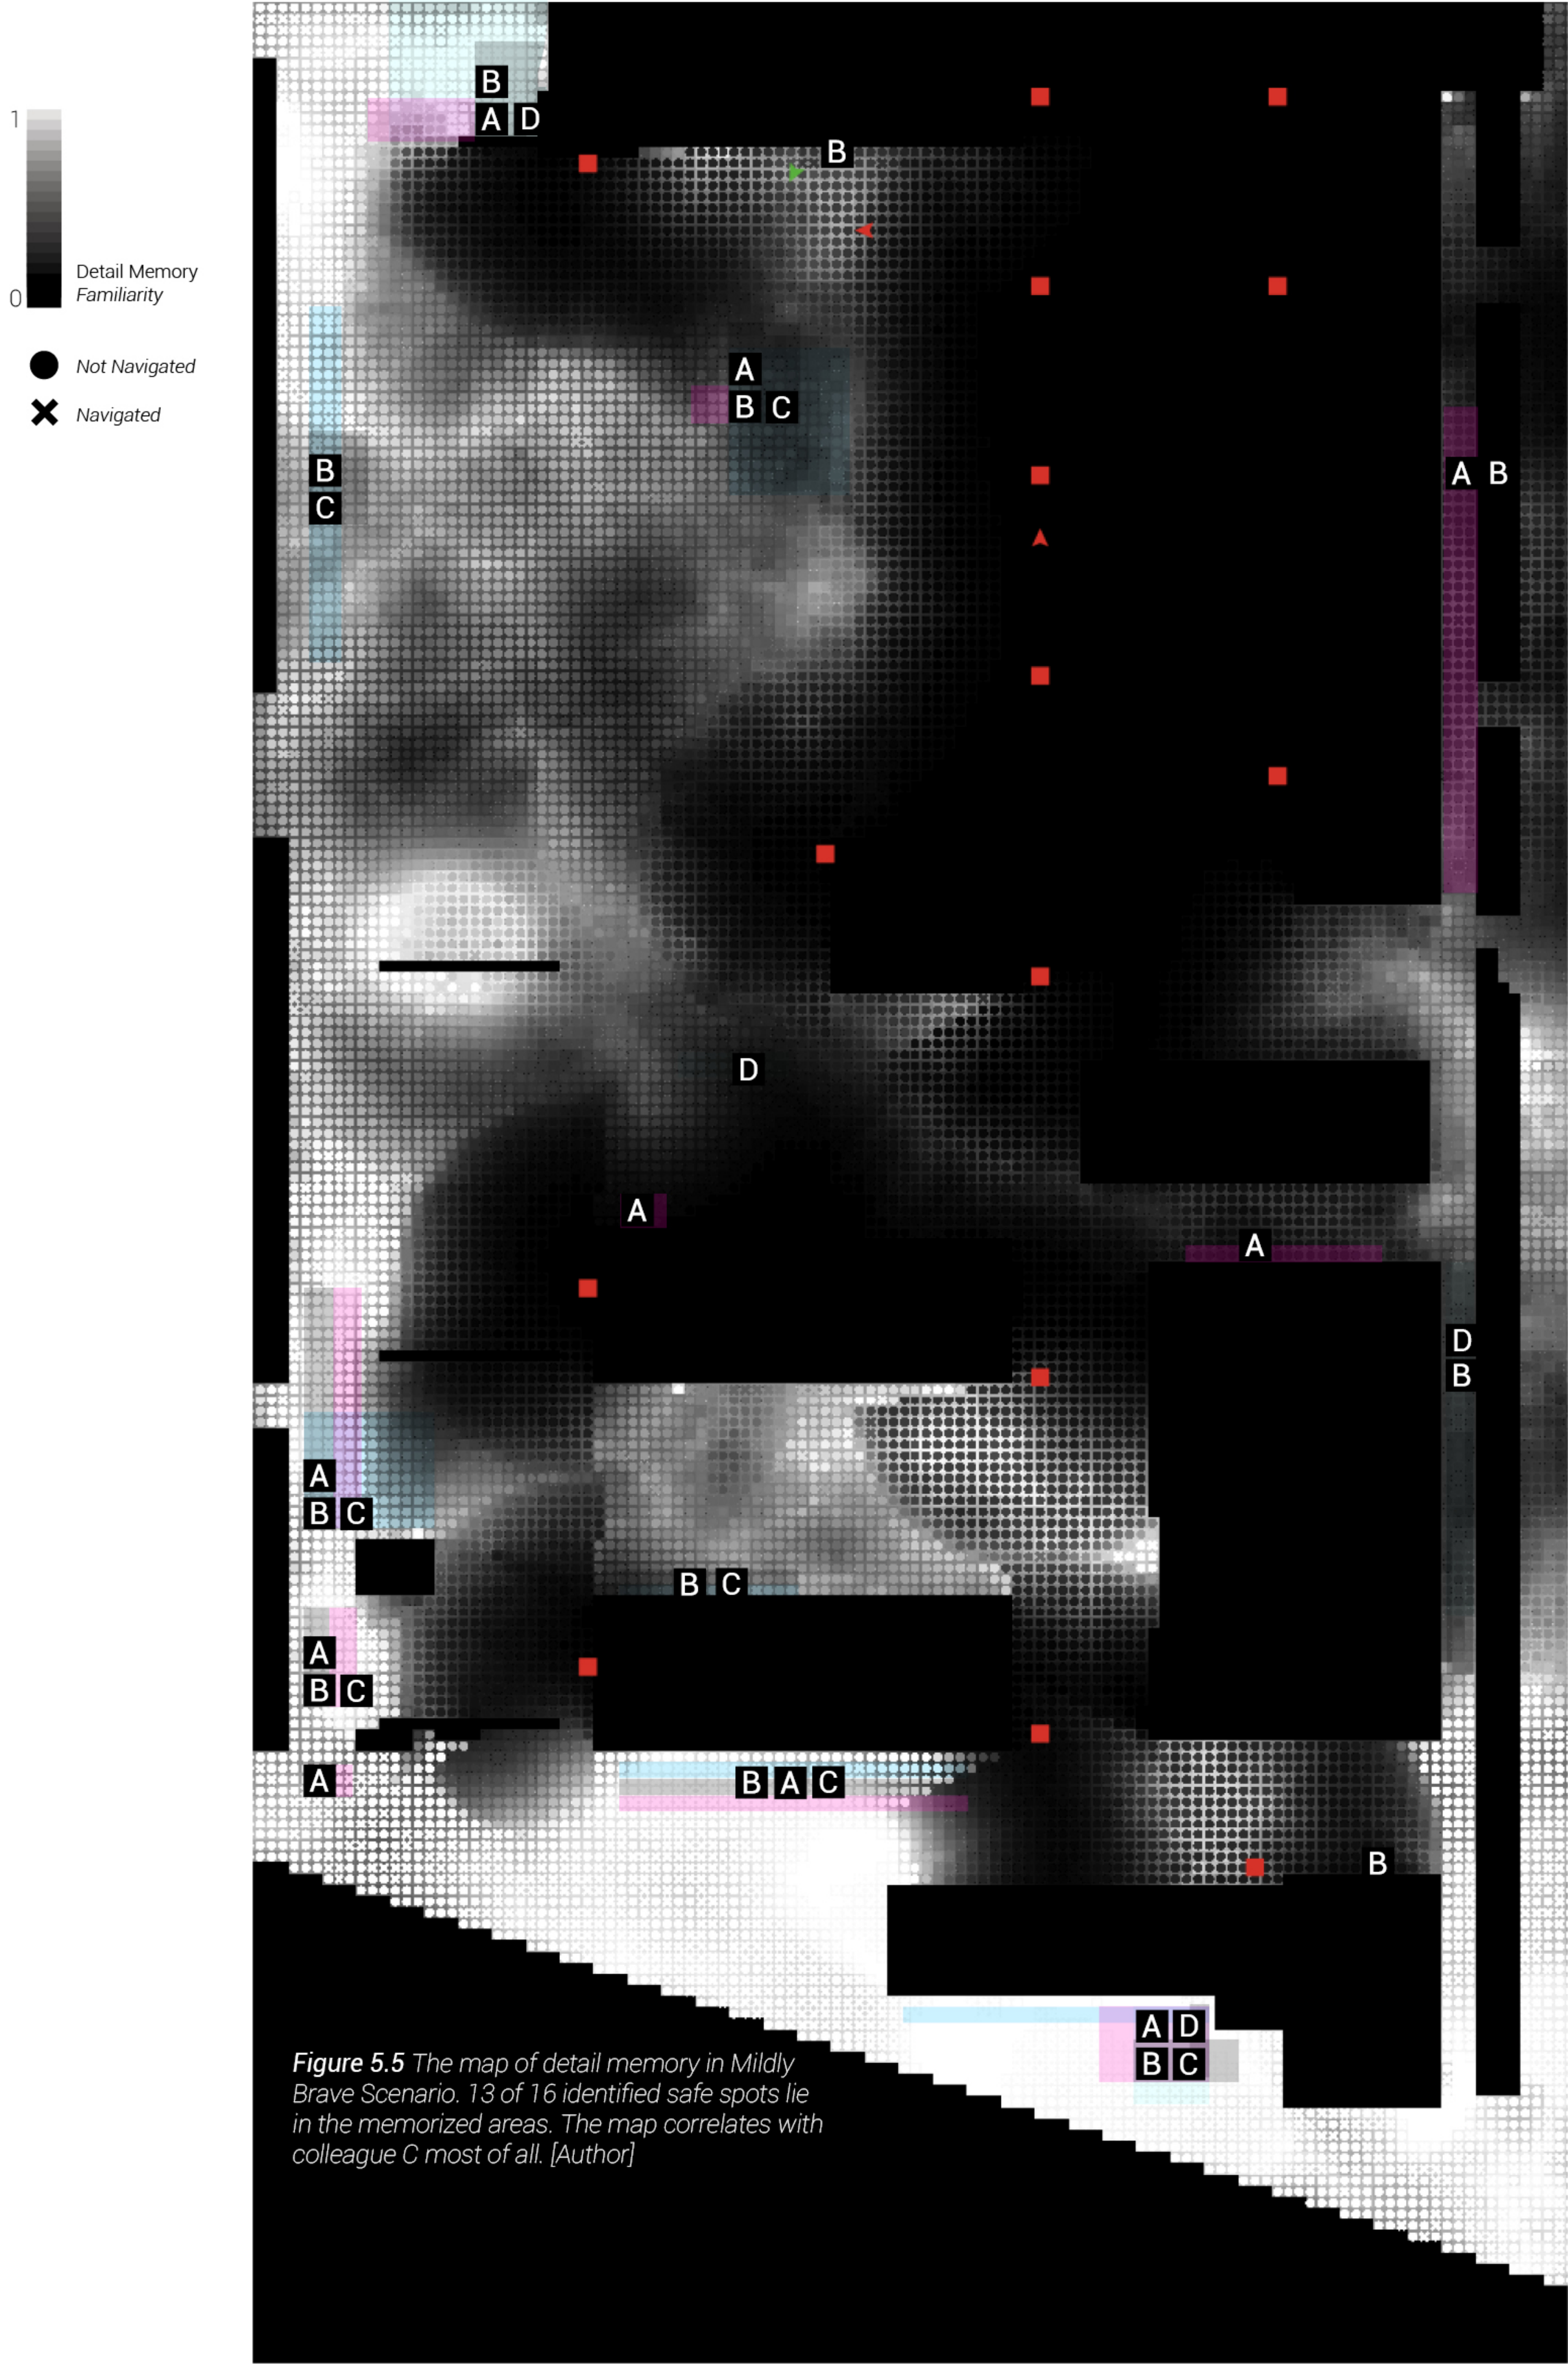

Figure 5.5 *The map of detail memory in Mildly Brave Scenario. 13 of 16 identified safe spots lie in the memorized areas. The map correlates with colleague C most of all. [Author]*

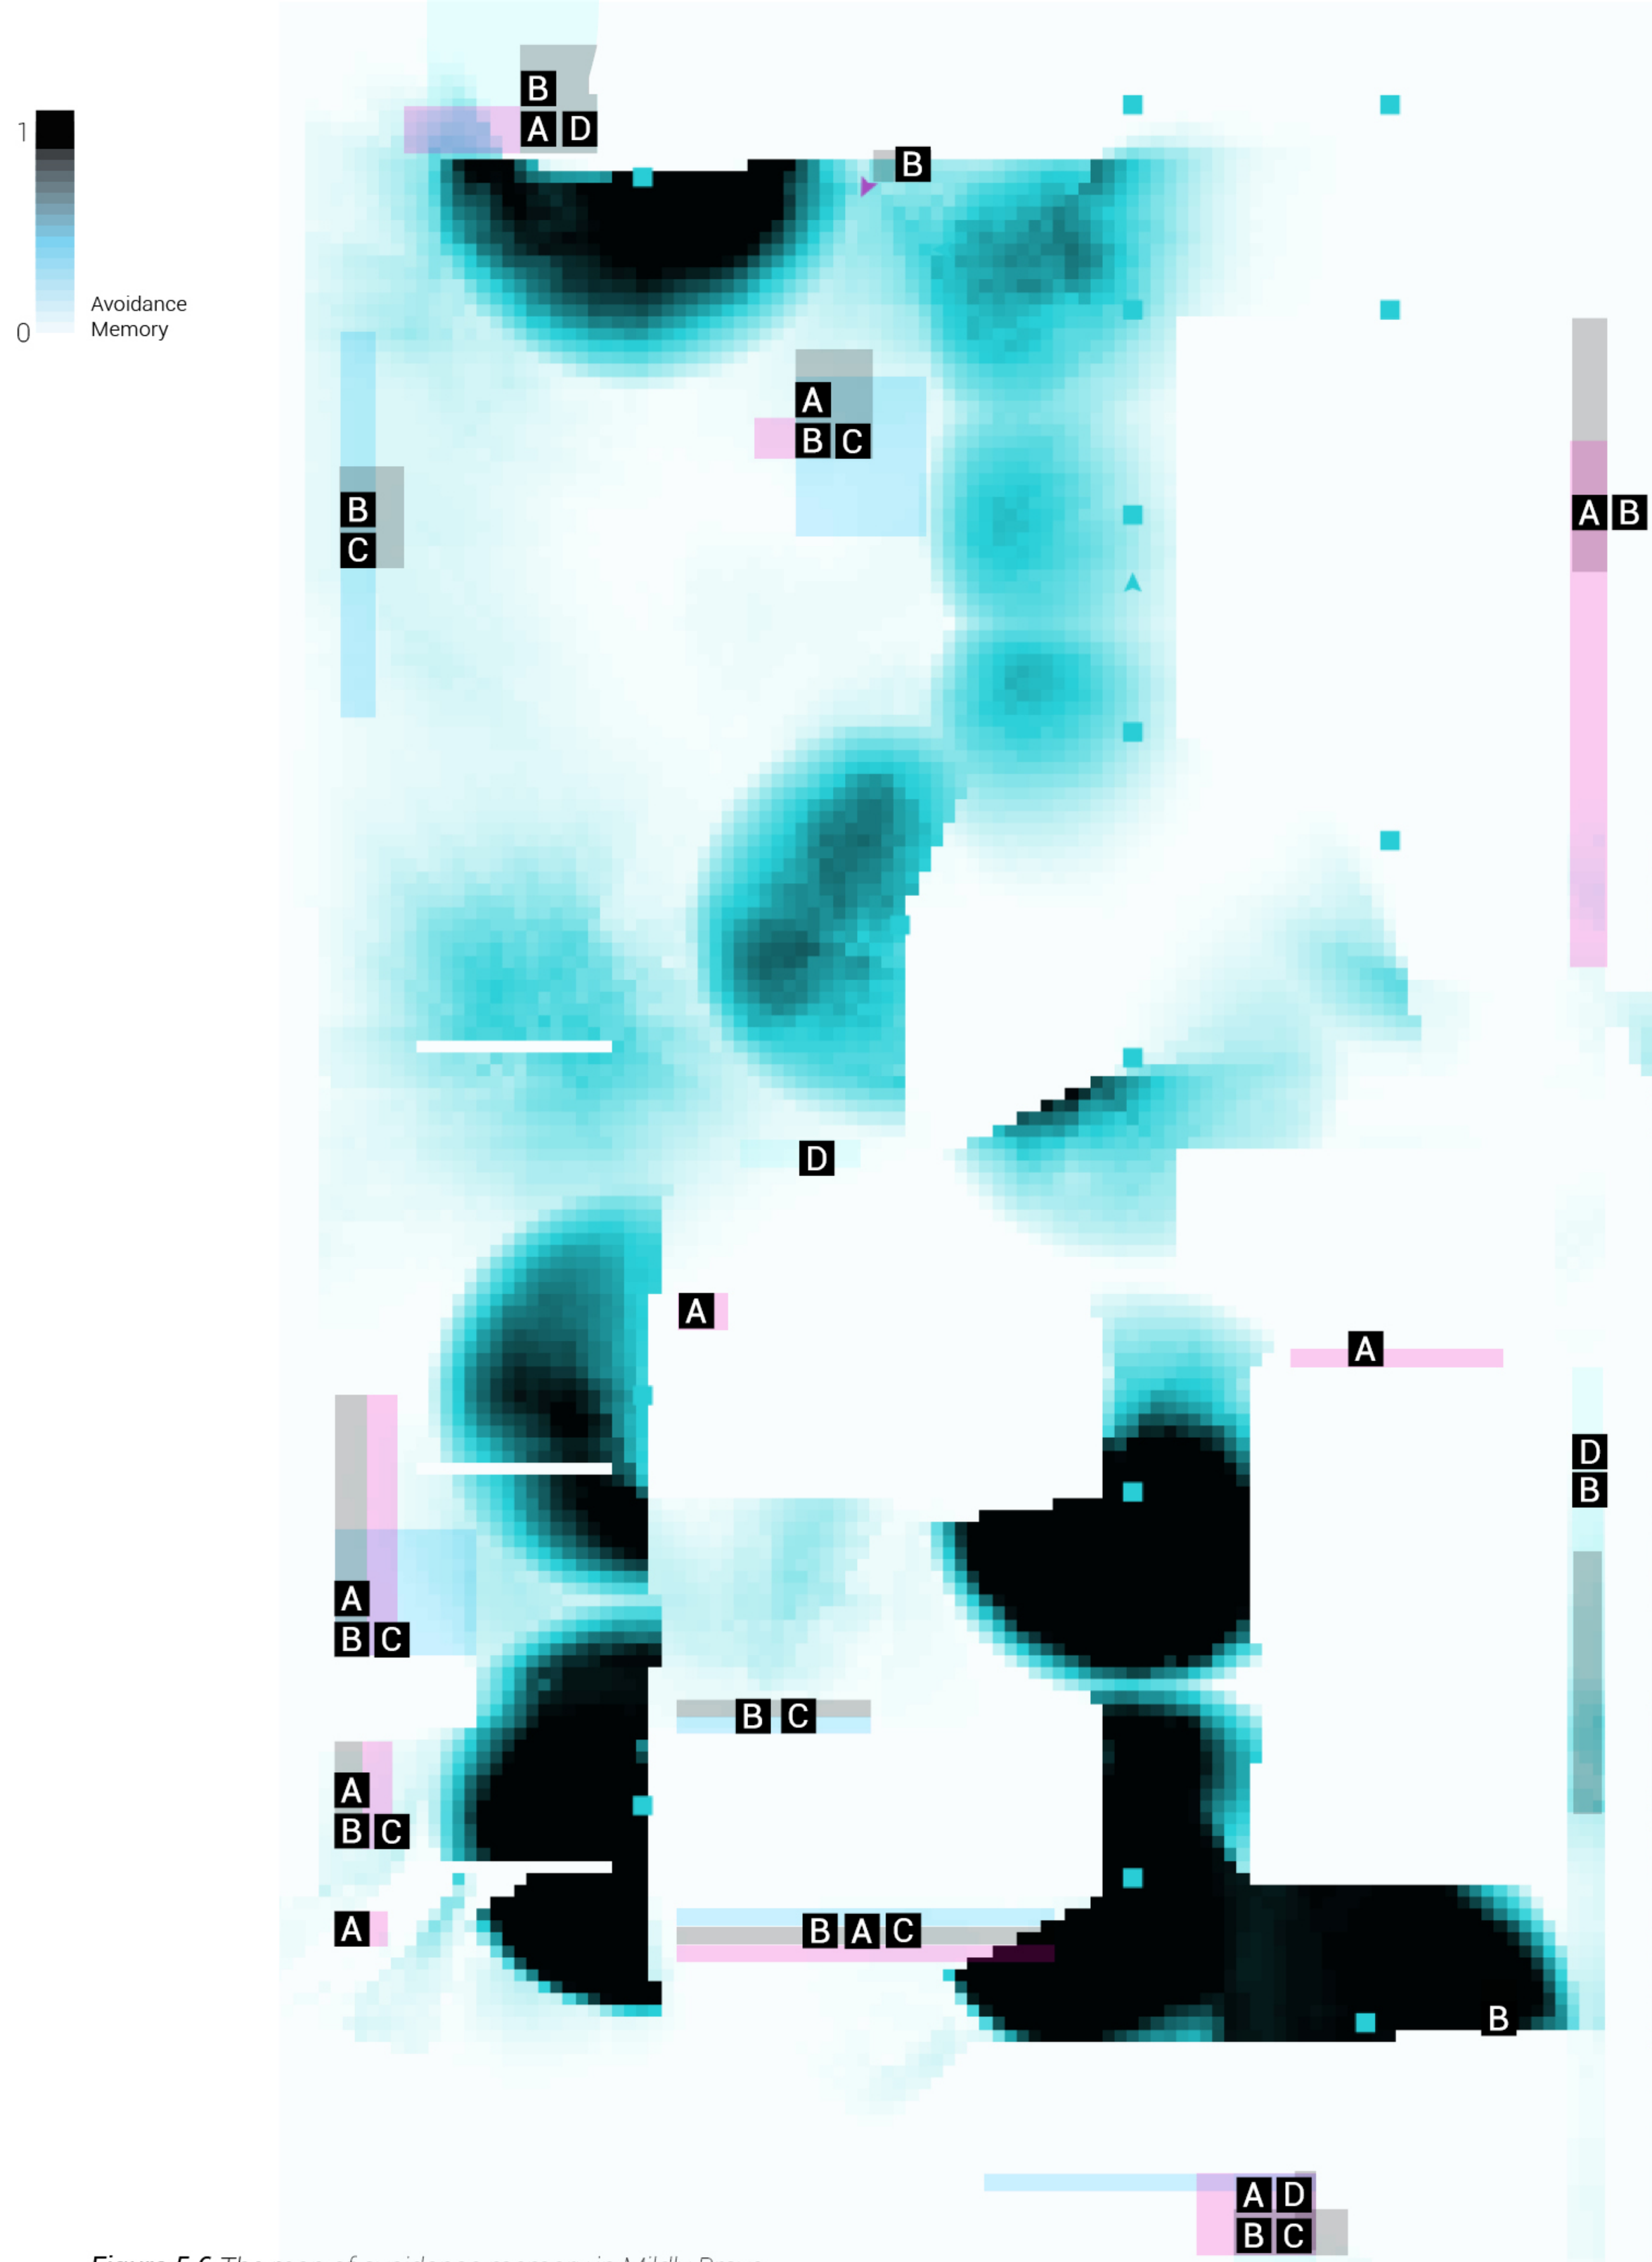

Figure 5.6 *The map of avoidance memory in Mildly Brave Scenario. All but one of the reported safe spots are strongly located in low-stress zones. [Author]*

5.3 Third scenario: risk-avoider

The third agent has an avoidance threshold of 0.2, making it highly sensitive to perceived risk. One might expect this agent to experience the most severe fragmentation. The reality is more complex.

Graph coverage shows nearly the same collapse as the mildly brave scenario (28.8%, a drop of 47.3 percentage points). Active graph nodes are actually slightly higher than in Scenario 2, suggesting the agent maintains access to a comparable spatial network.

The critical difference lies in navigational behavior. While the number of nodes navigated remains similar, total movements drop dramatically—to just 67% of control trials, a −22% decrease compared to −9% in Scenario 2. This agent moves less, hesitates more, and experiences what can only be described as spatial paralysis.

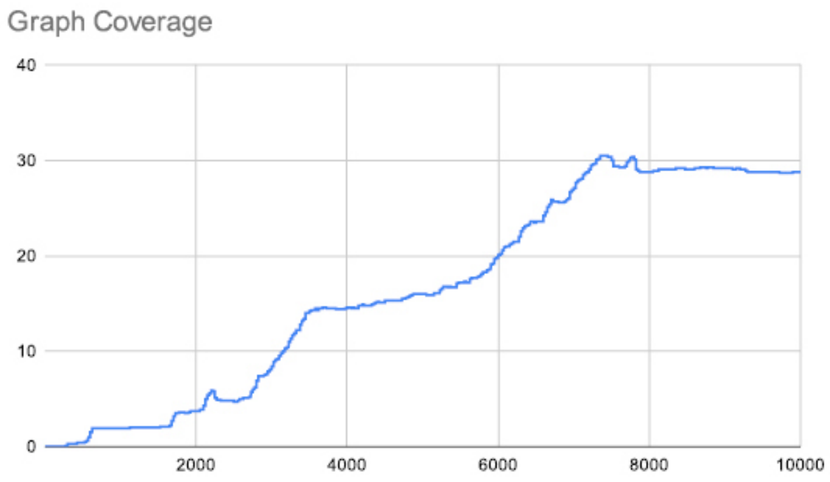

Graph 5.6 Graph coverage over time shows the major expansions of the navigational graph and its flattening after around 7000 runs at ~28.8%.

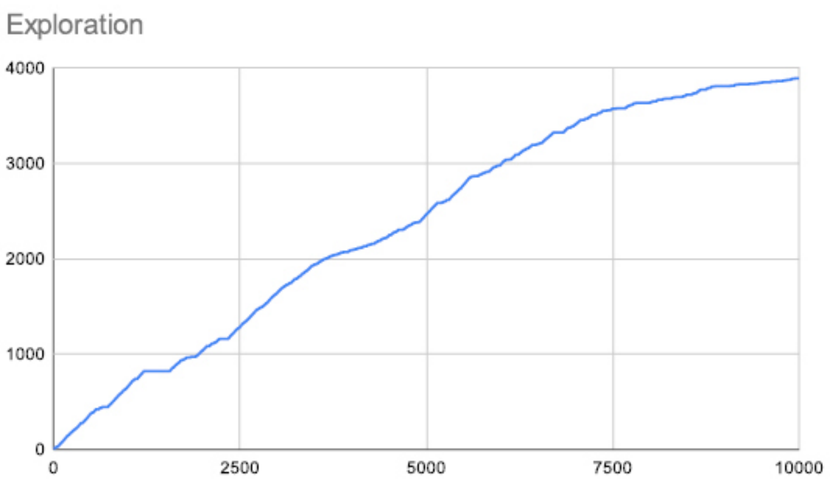

Graph 5.7 Exploration over time shows the rate of navigation in new areas in the simulation. The flattening of the curve indicates breaking point and end of inhabitation.

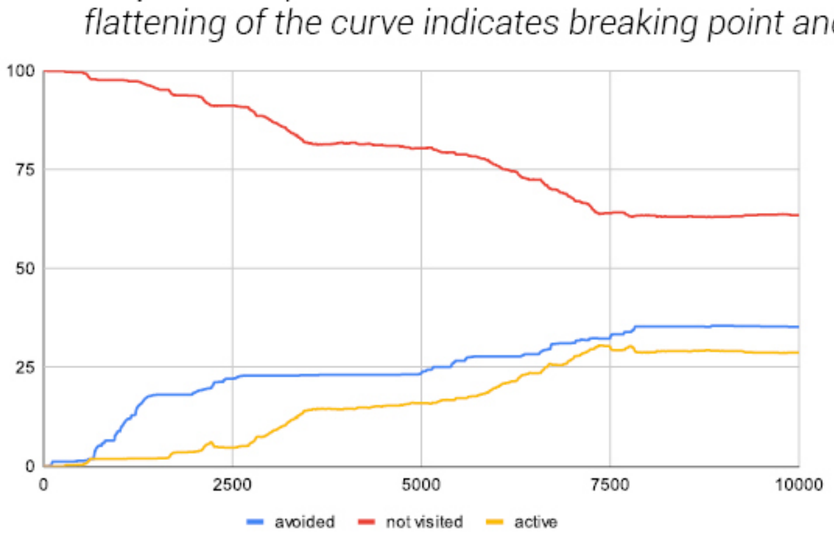

Graph 5.8 Percentage of avoided, active and never-visited areas over time. For this agent, avoided patches exceed the active patches.

This suggests that extreme surveillance doesn't merely redirect behavior—it produces cognitive freezing. Architecturally, this means spaces under heavy surveillance may become effectively unusable even when technically accessible. The agent knows where it could go but cannot bring itself to move.

Furthermore, the number of avoided nodes in this scenario exceeds the number of active nodes—a reversal not seen in other scenarios.

Graph Structure		
Active Nodes	Nodes used and remembered	4263 (-72 %)
Pruned nodes	Nodes not used/ not memorized	20753 (+51%)
Graph coverage	Active nodes / total passable nodes	28.8% (-47.3%)
Navigation Usage		
Nodes navigated	Nodes used for navigation	3892 (-40%)
Edges used	Navigated edges	5656 (-30%)
Total movements	Total movements	6749 (-22%)
Avg degree	Average edges of a passable node	1.33 (+85%)

Table 5.3 *Risk-avoider Scenario - Graph structure and navigation statistics over the simulation. [Author]*

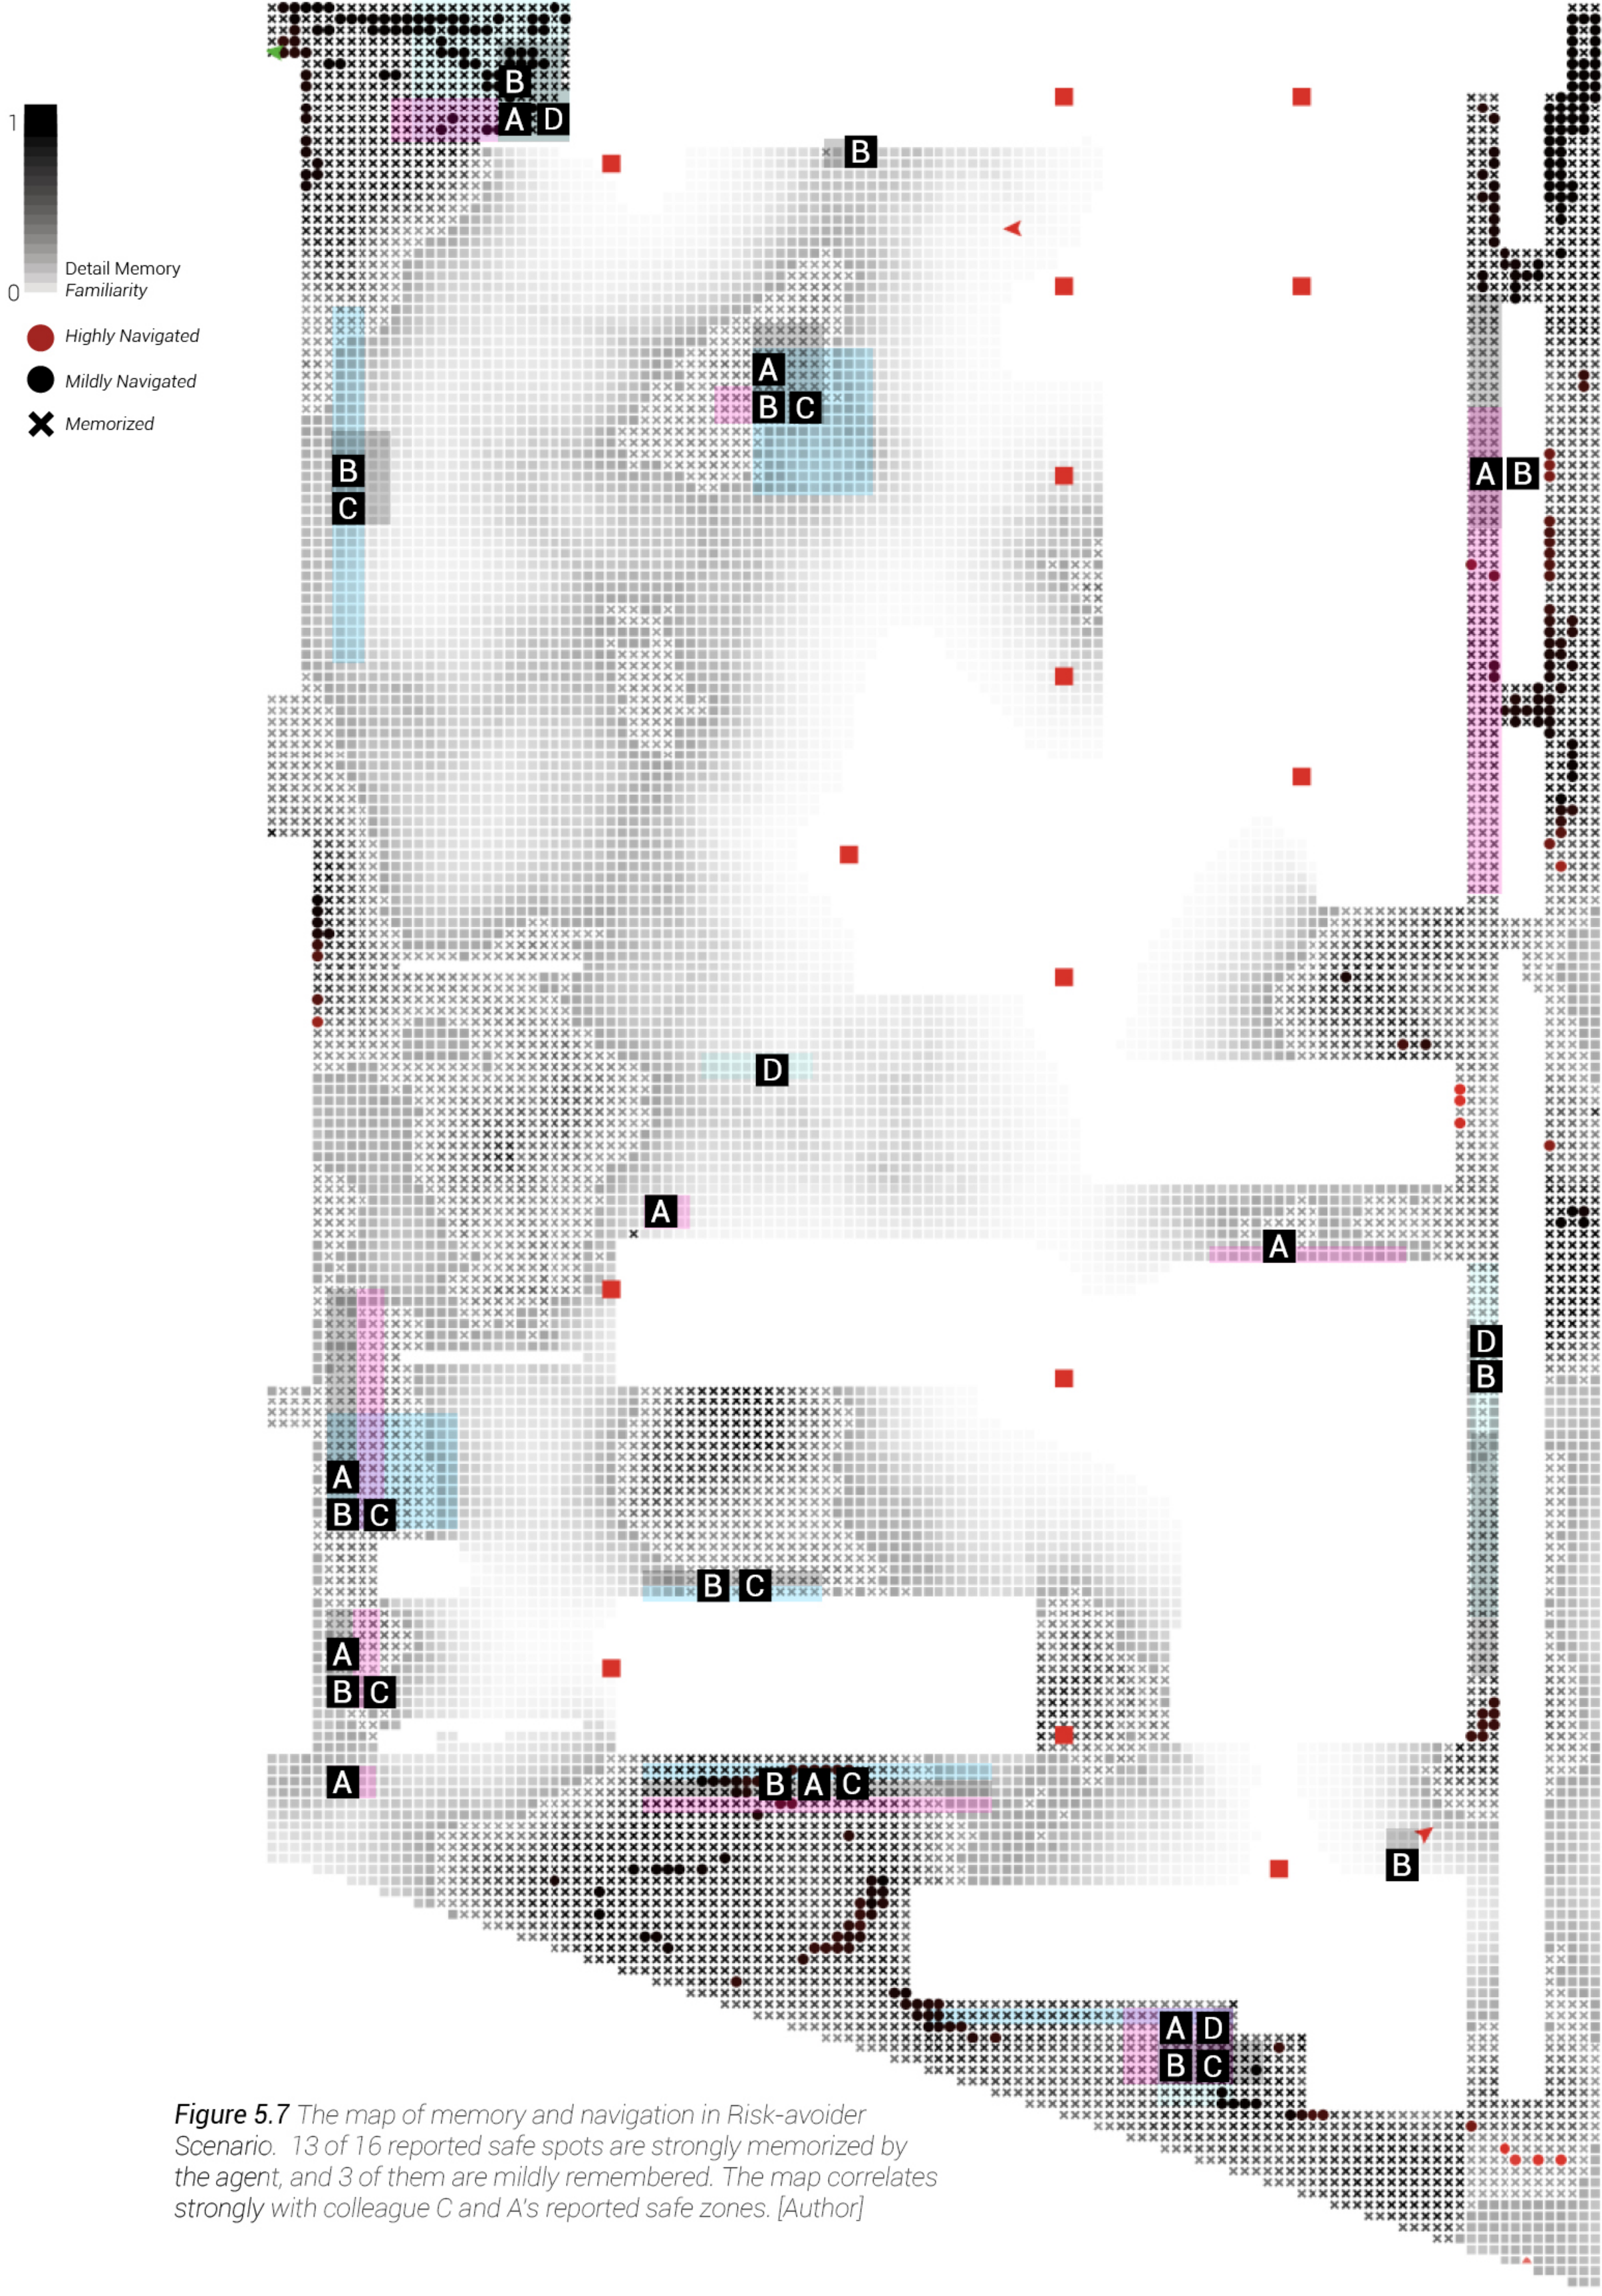

Figure 5.7 *The map of memory and navigation in Risk-avoider Scenario. 13 of 16 reported safe spots are strongly memorized by the agent, and 3 of them are mildly remembered. The map correlates strongly with colleague C and A's reported safe zones. [Author]*

1

0

Detail Memory
Familiarity

Not Navigated

Navigated

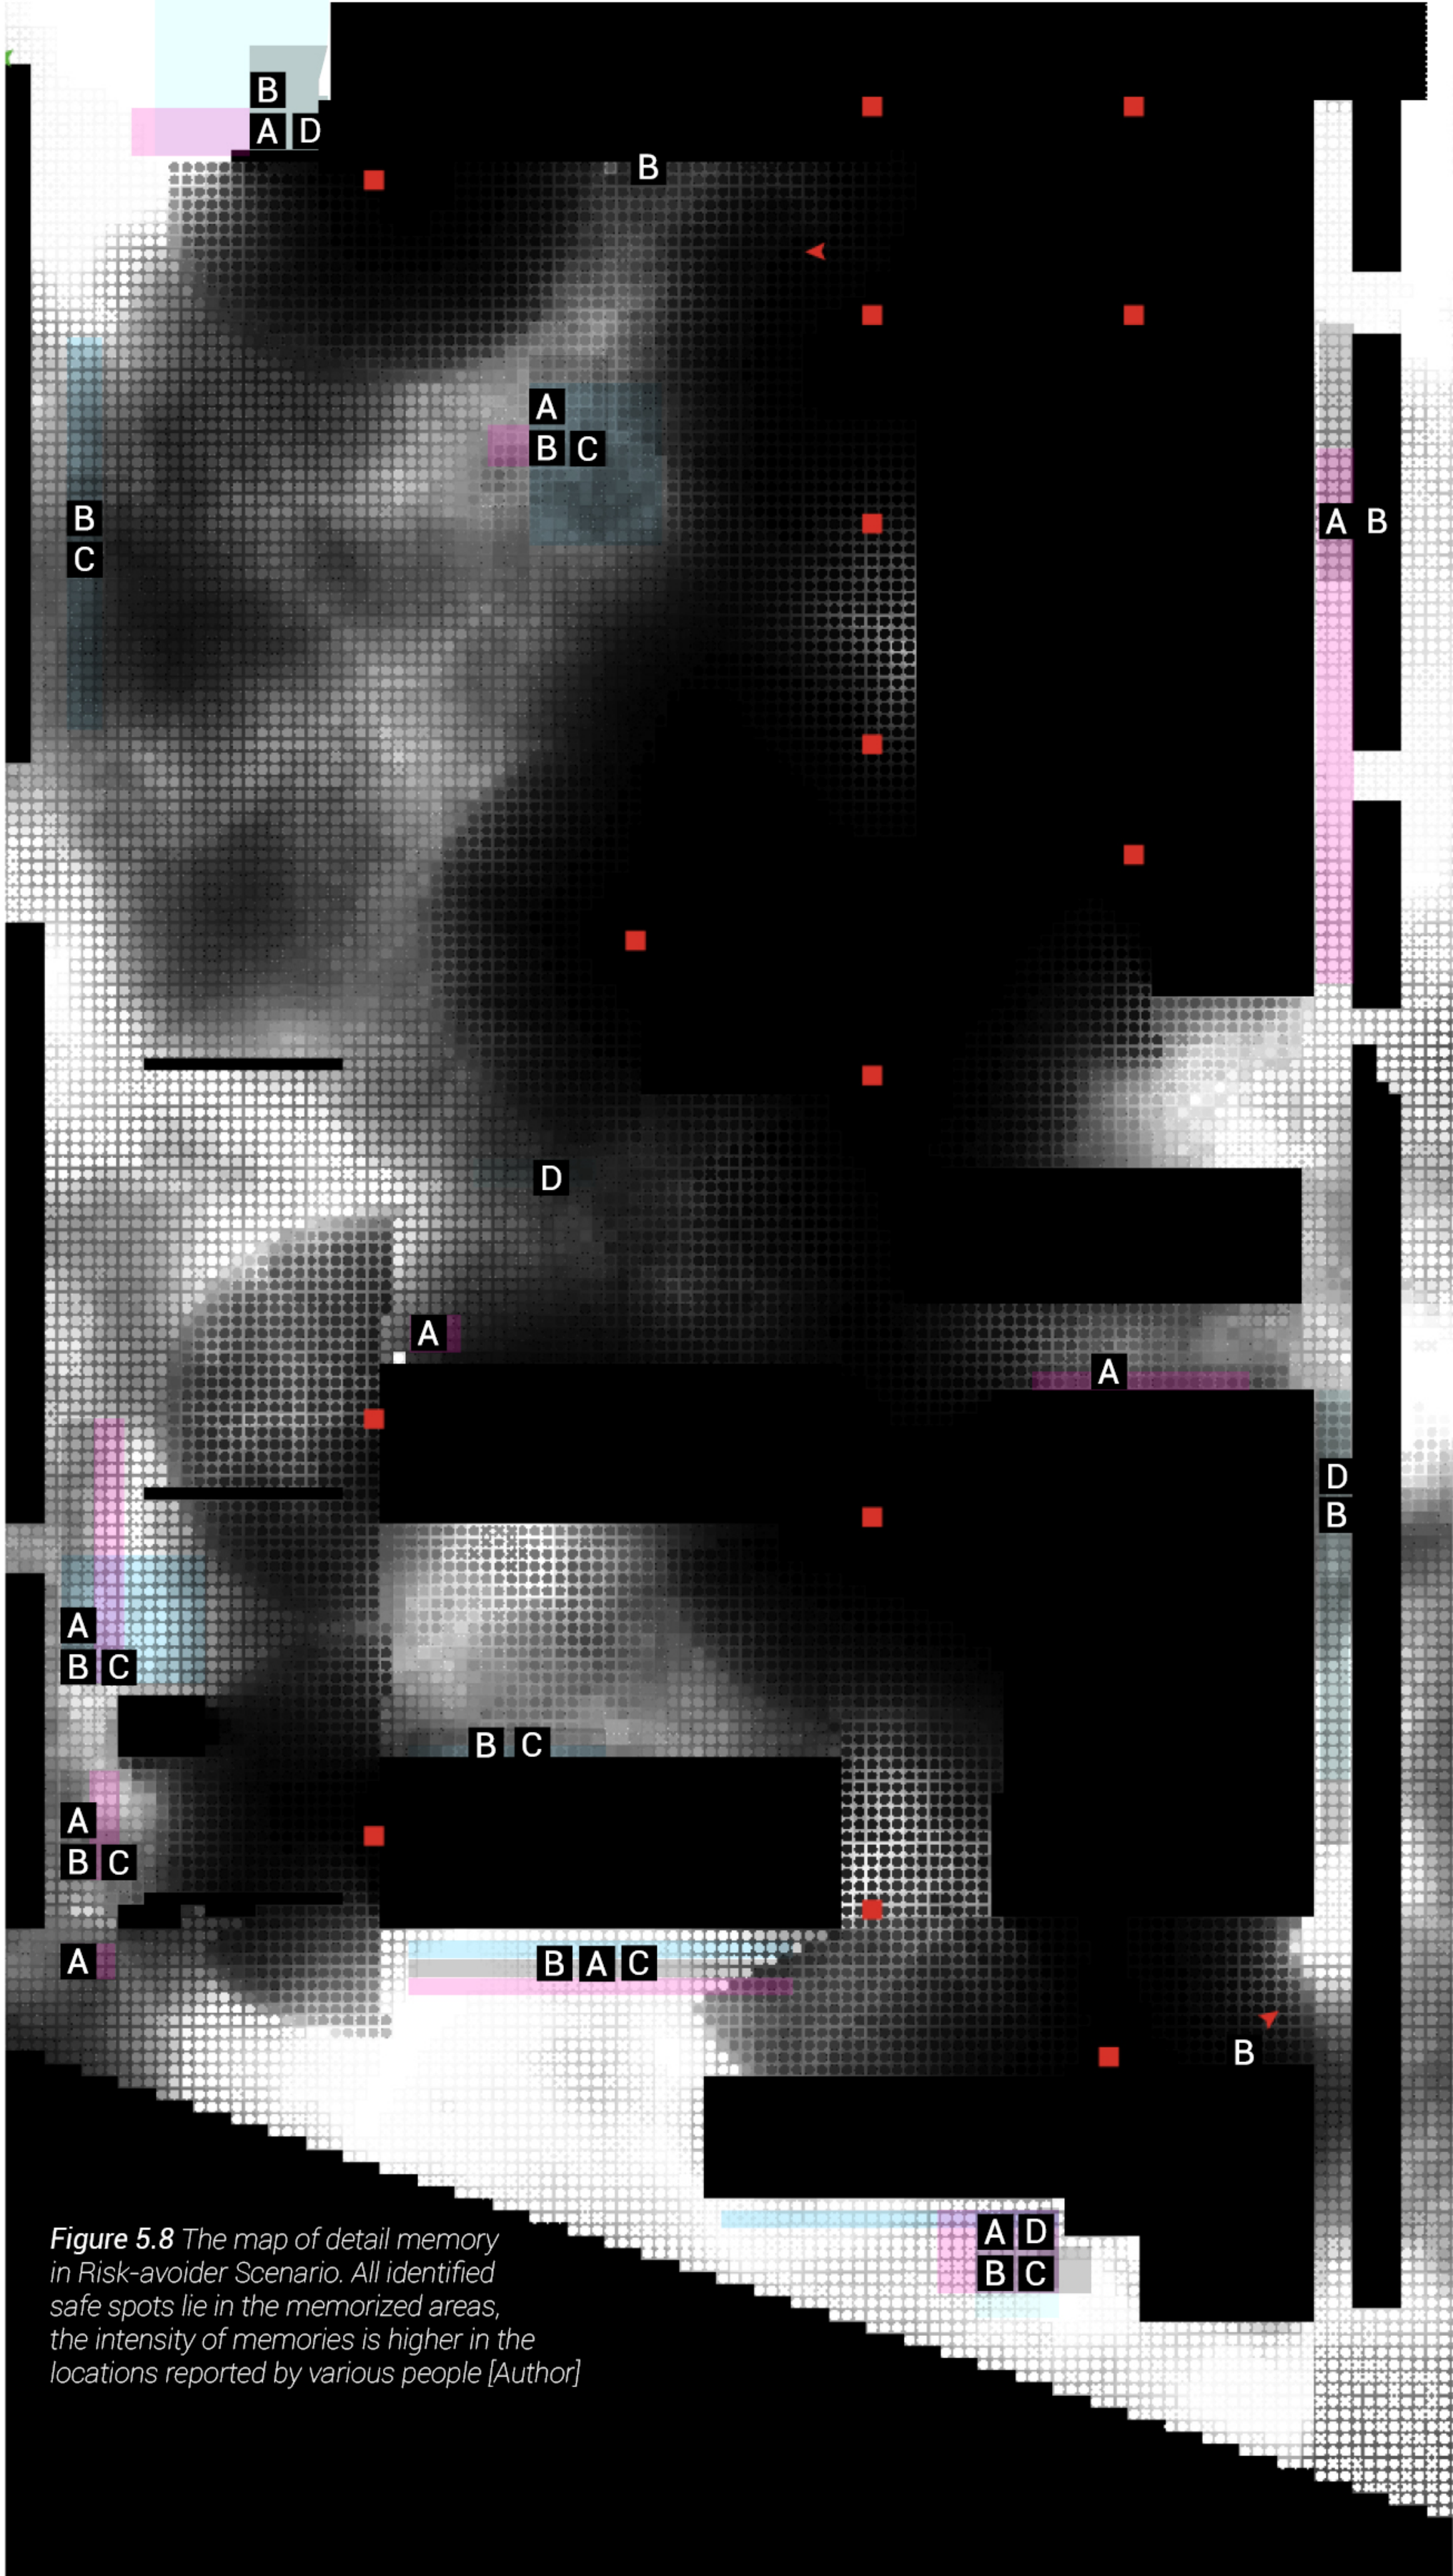

Figure 5.8 *The map of detail memory in Risk-avoider Scenario. All identified safe spots lie in the memorized areas, the intensity of memories is higher in the locations reported by various people [Author]*

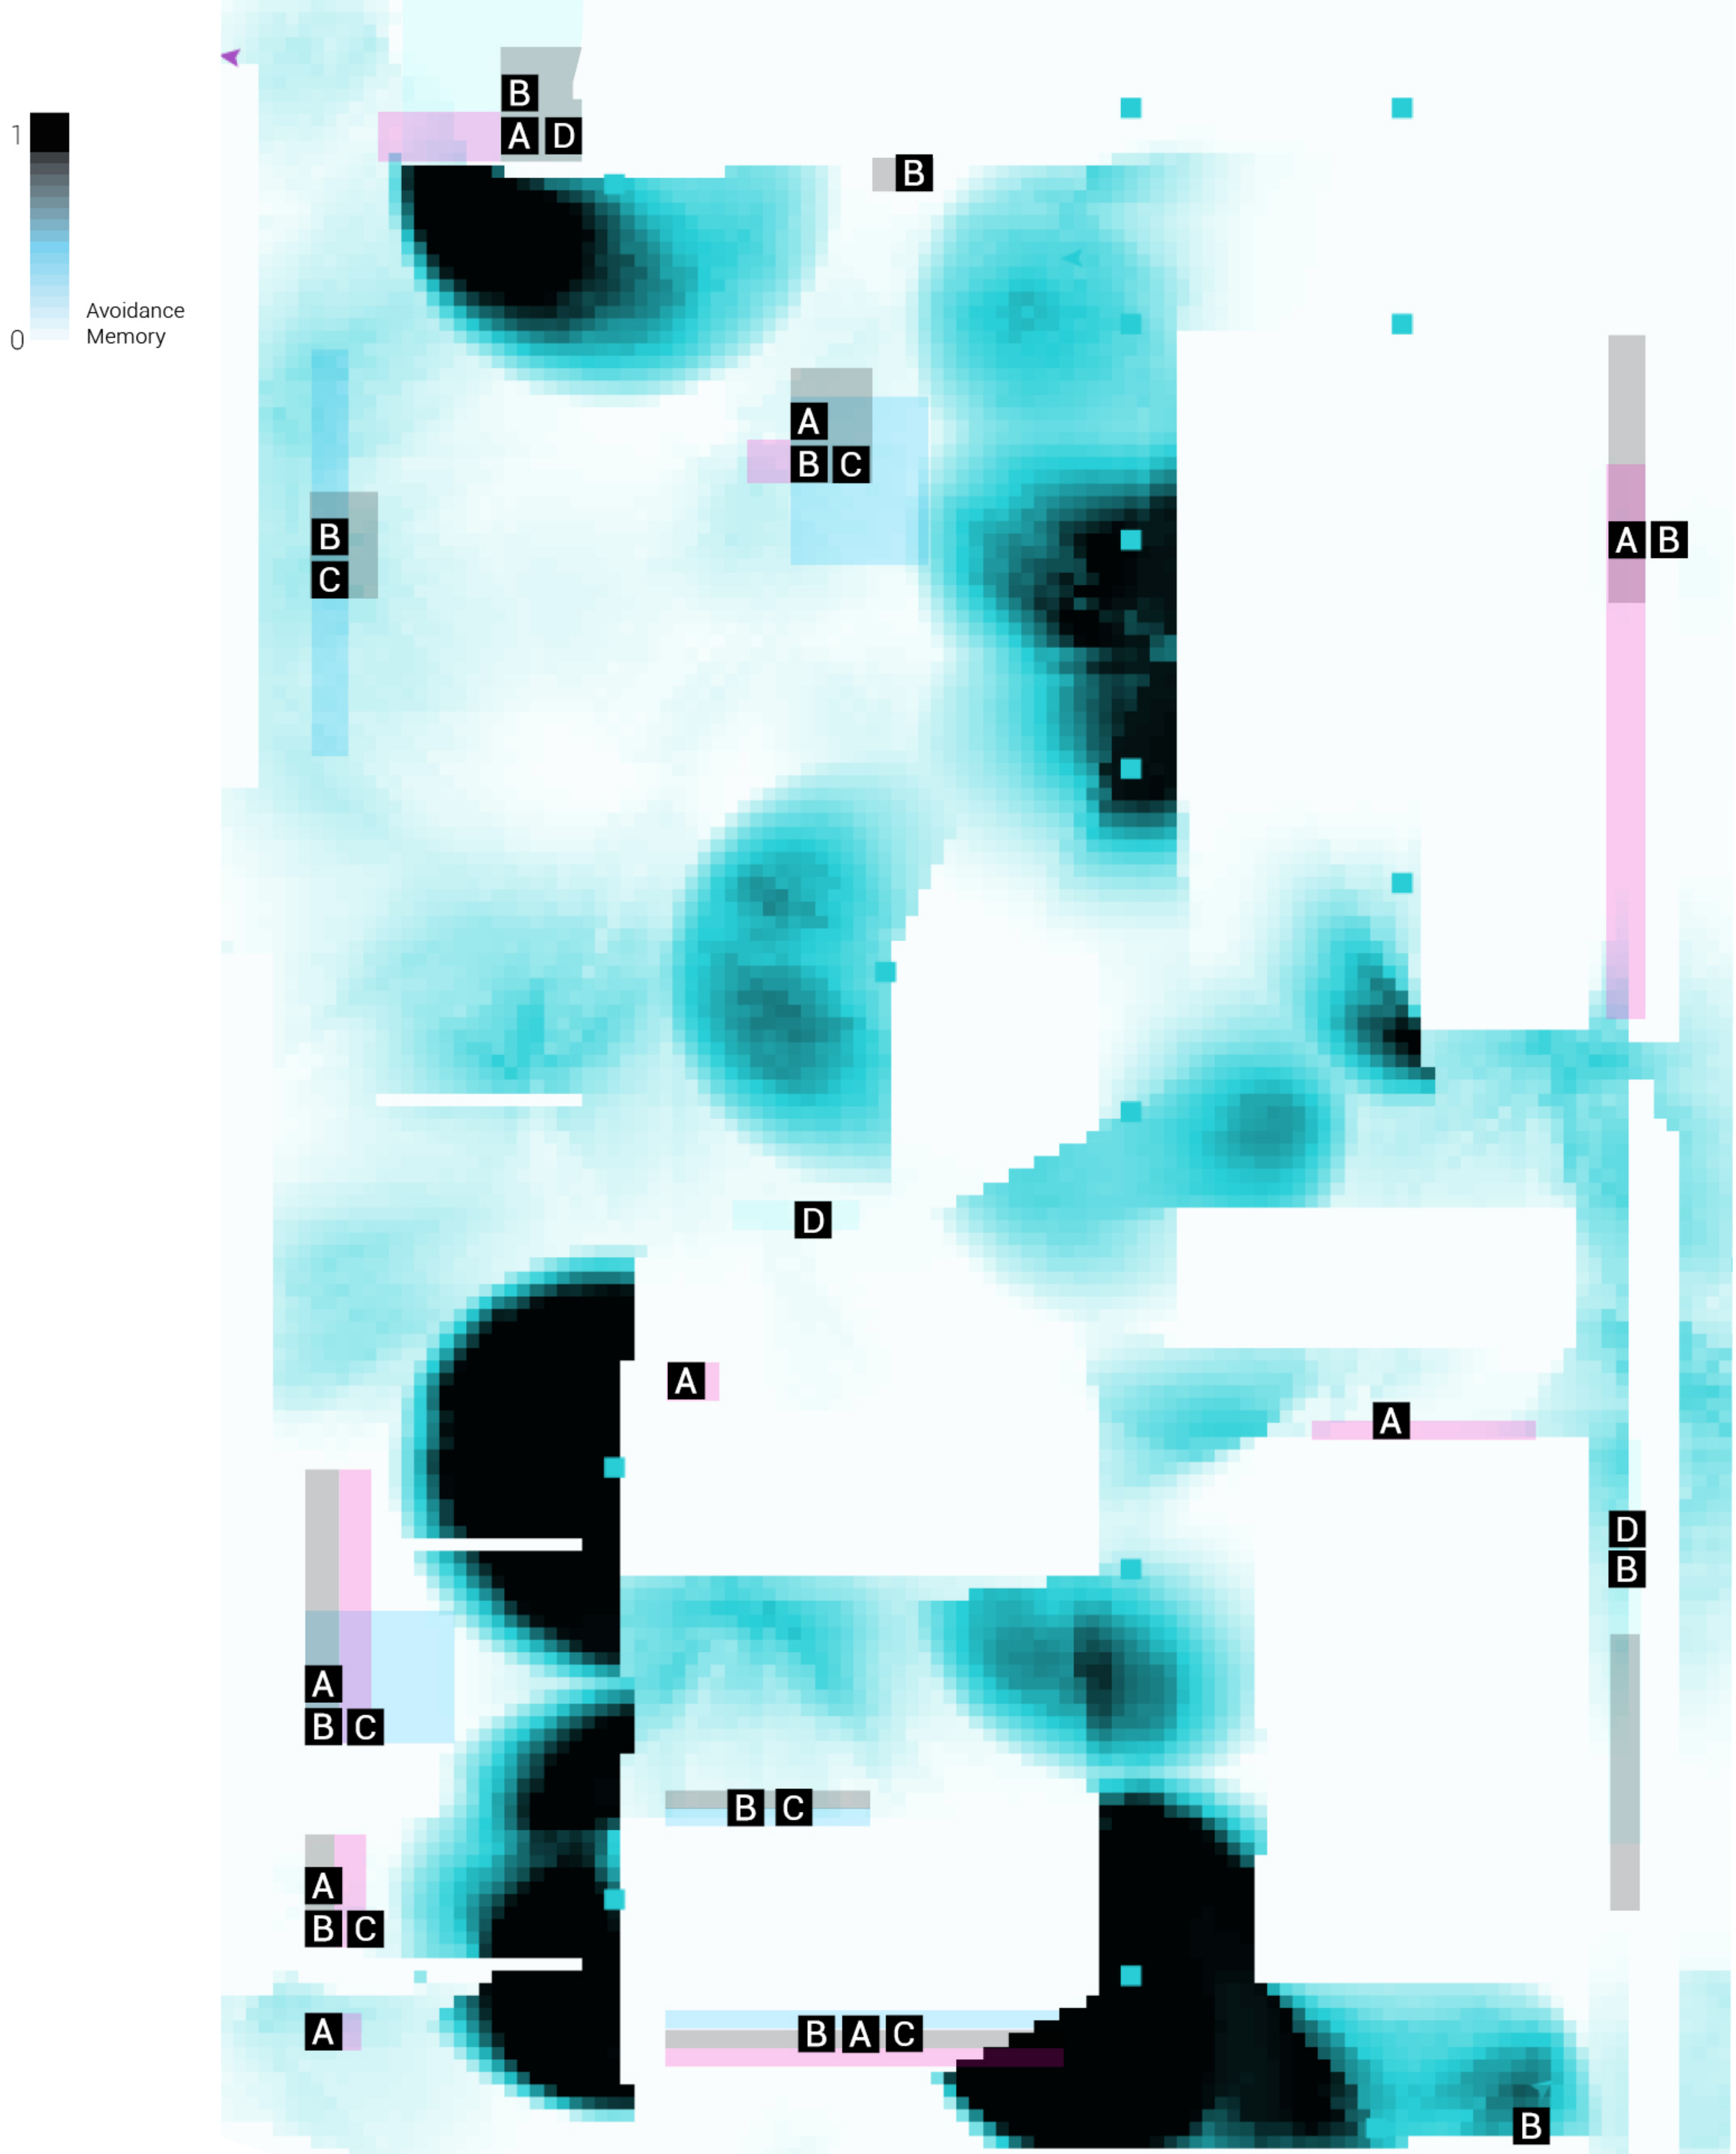

Figure 5.9 *The map of avoidance memory in Risk-Avoider Scenario. All but one of the reported safe spots are located in low-stress zones. [Author]*

5.4 Fourth scenario: The brave

The final agent has an avoidance threshold of 0.9, meaning it will tolerate high levels of perceived risk before avoiding an area.

Graph coverage remains significantly lower than the control scenario but is noticeably higher than in Scenarios 2 and 3—38.5%, a reduction of 37.6 percentage points. Total movements are only slightly less than the control group (−5%), yet the same clustering behavior observed in other scenarios is still visible. Average degree increases by 67% to 1.2, confirming that even a highly risk-tolerant agent engages in repetitive path-following rather than free exploration.

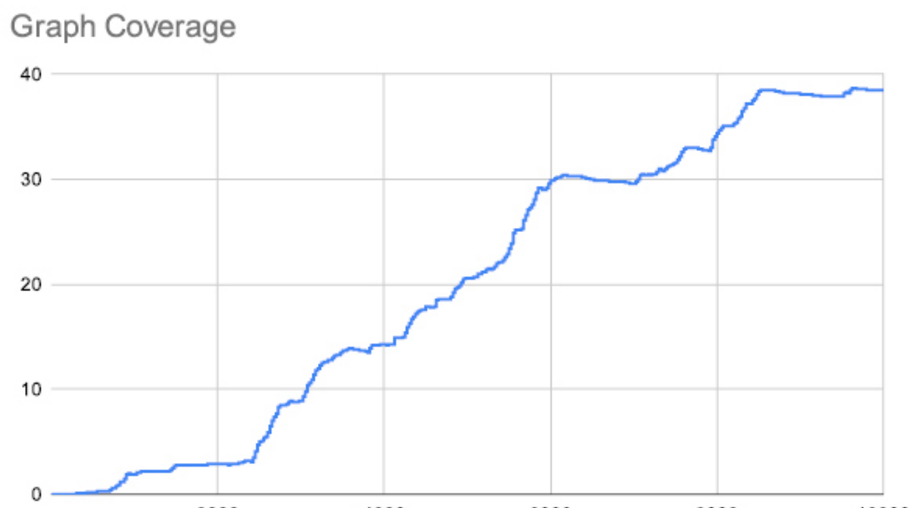

Graph 5.9 *Graph coverage over time shows the major expansions of the navigational graph and its flattening after around 9000 runs at ~38.5%.*

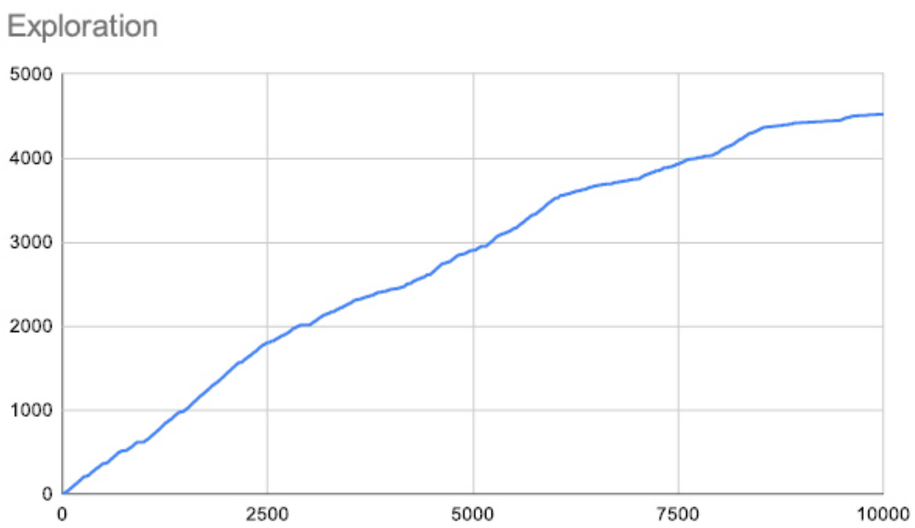

Graph 5.10 *Exploration over time shows the rate of navigation in new areas in the simulation. The flattening of the curve indicates breaking point and end of inhabitation.*

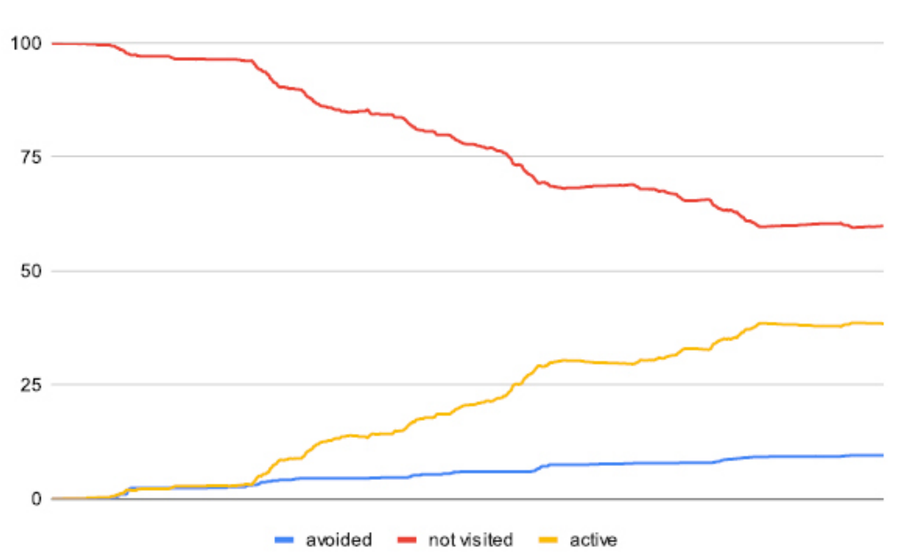

Graph 5.11 *Percentage of avoided, active and never-visited areas over time. For this agent, active patches are significantly higher than the avoided patches.*

GRAPH STRUCTURE		
Active Nodes	Nodes used and remembered	5695 (-49 %)
Pruned nodes	Nodes not used/ not memorized	19321 (+40%)
Graph coverage	Active nodes / total passable nodes	38.5% (-37.6%)
Navigation Usage		
Nodes navigated	Nodes used for navigation	4523 (-31%)
Edges used	Navigated edges	6818 (-16%)
Total movements	Total movements	8117 (-5%)
Avg degree	Average edges of a passable node	1.2 (+67%)

Table 5.4 Brave Scenario - Graph structure and navigation statistics over the simulation. [Author]

The coverage and exploration graphs show stabilization during the run, though the agent slows down more gradually than in Scenarios 2 and 3. The number of avoided patches is significantly smaller than in previous runs, yet the Inaccessible Void persists. Even bravery cannot fully restore the cognitive map.

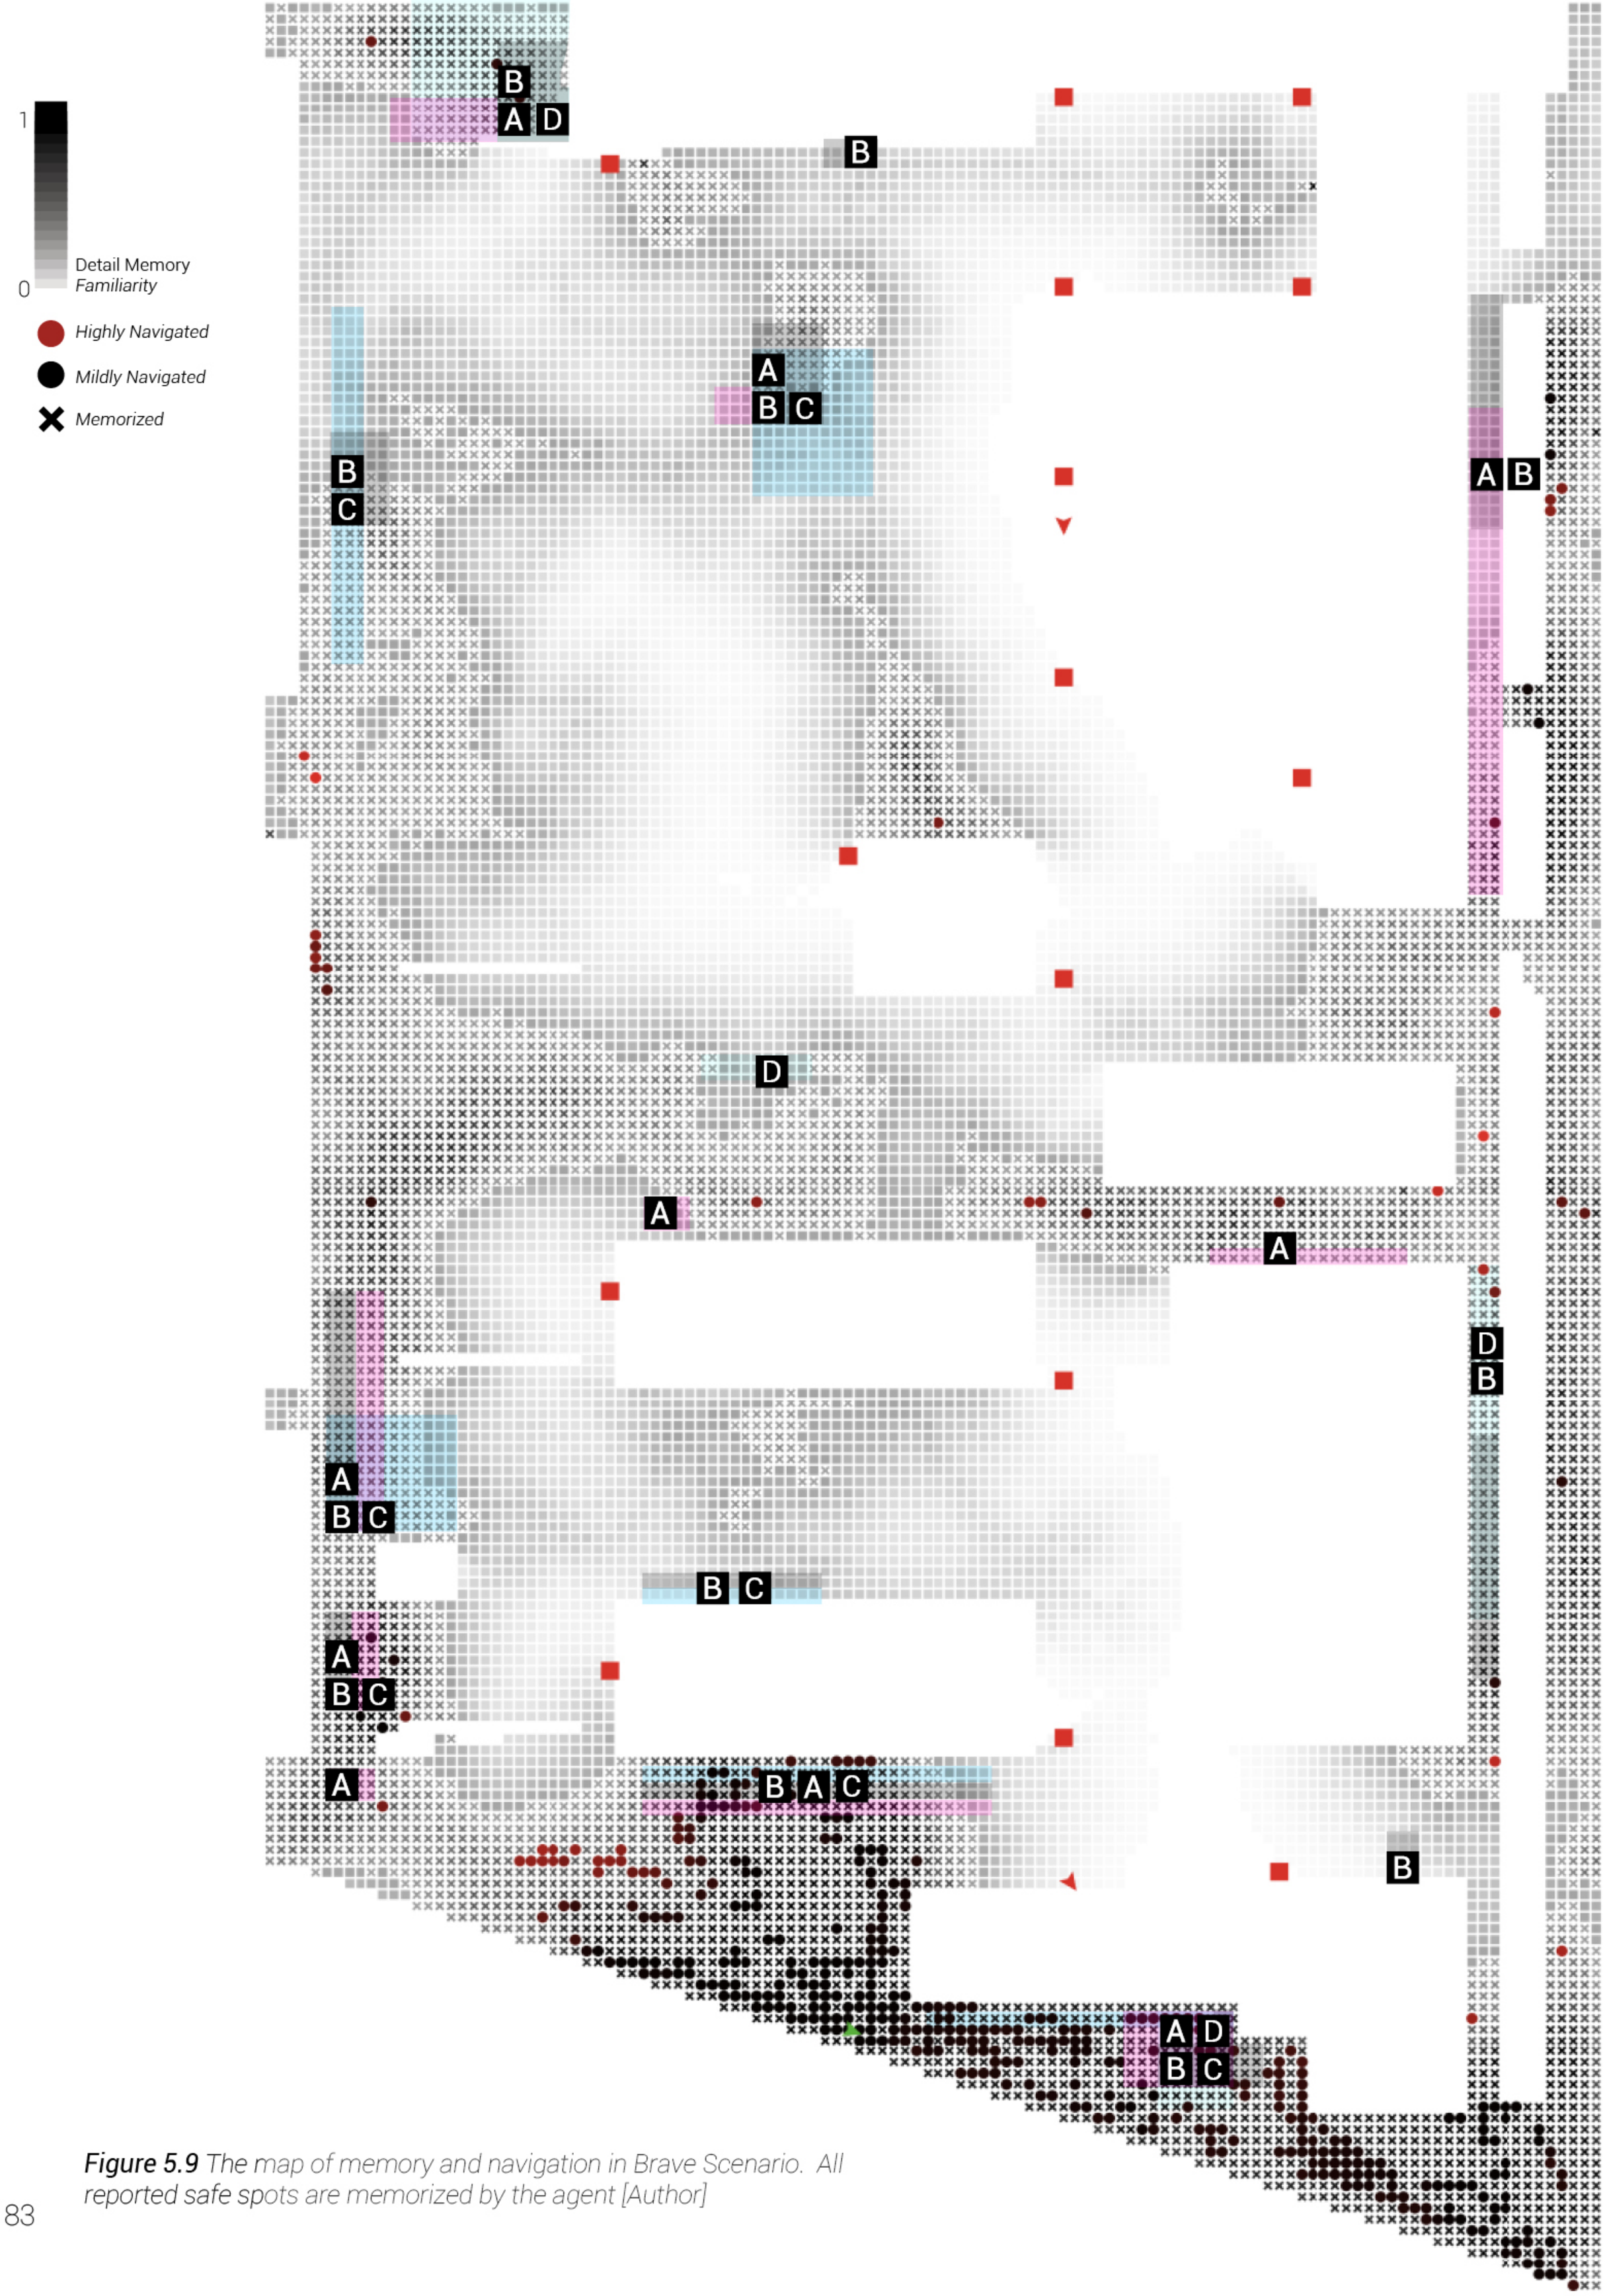

Figure 5.9 *The map of memory and navigation in Brave Scenario. All reported safe spots are memorized by the agent [Author]*

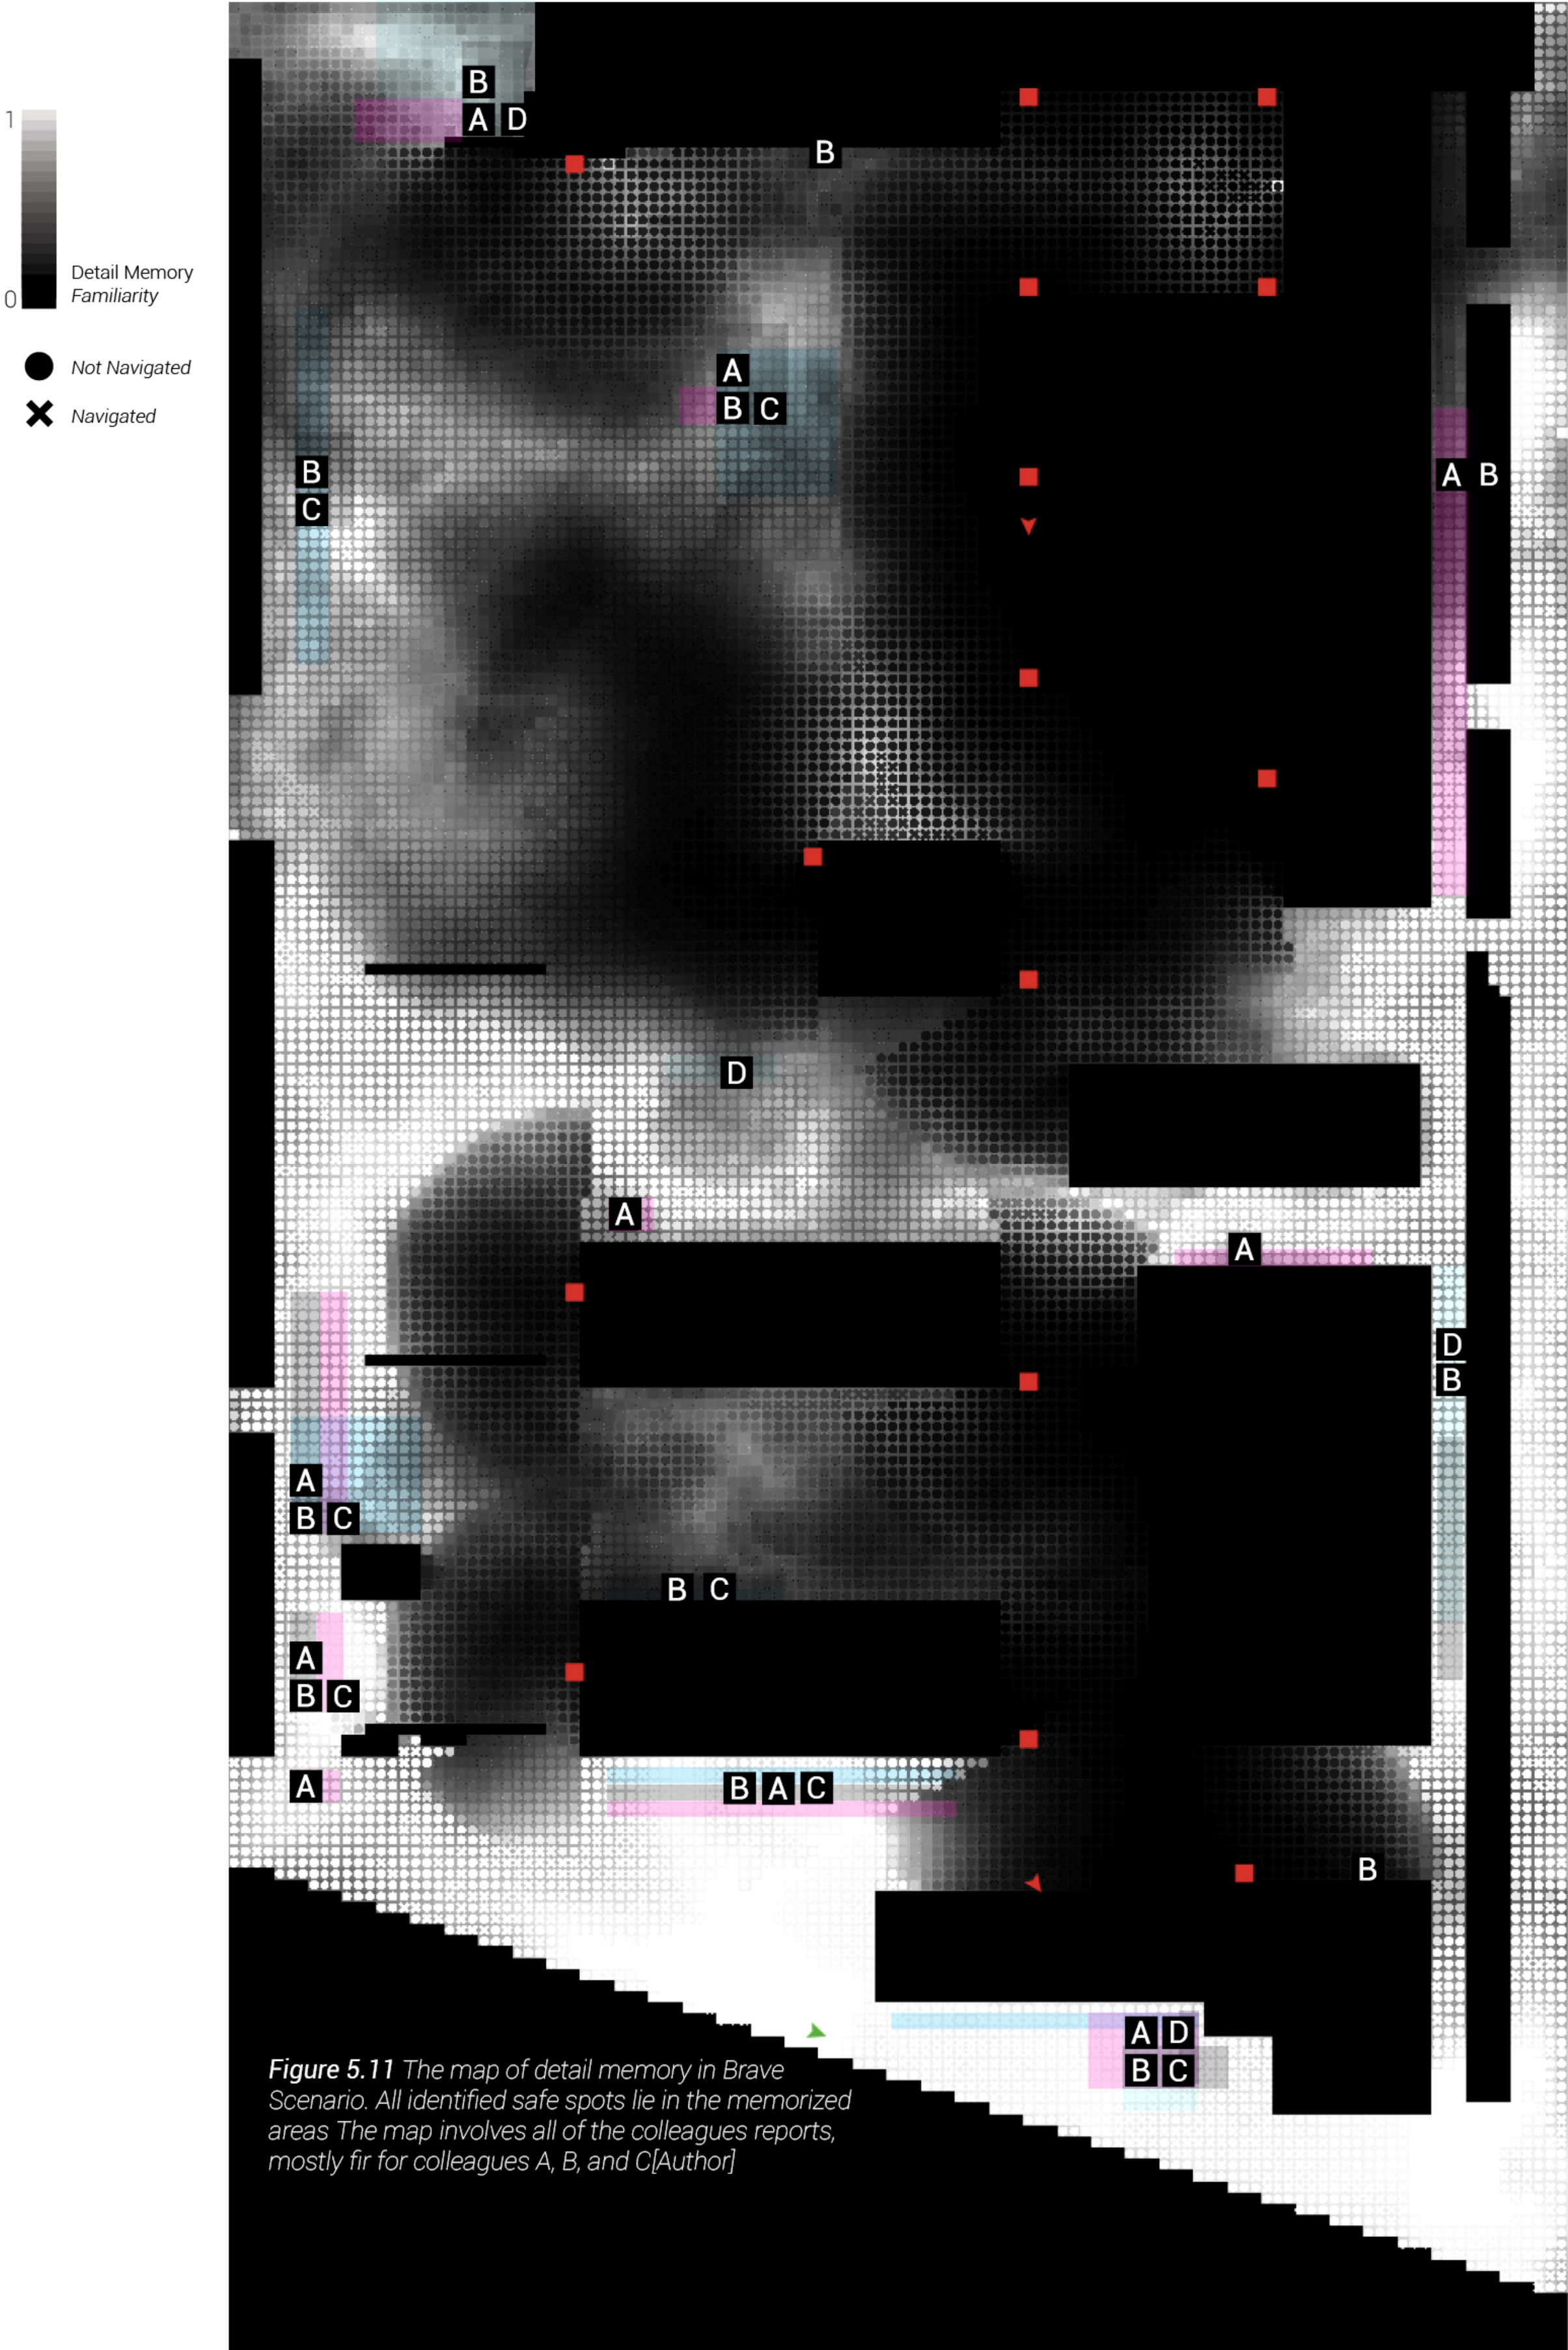

Figure 5.11 *The map of detail memory in Brave Scenario. All identified safe spots lie in the memorized areas The map involves all of the colleagues reports, mostly fir for colleagues A, B, and C[Author]*

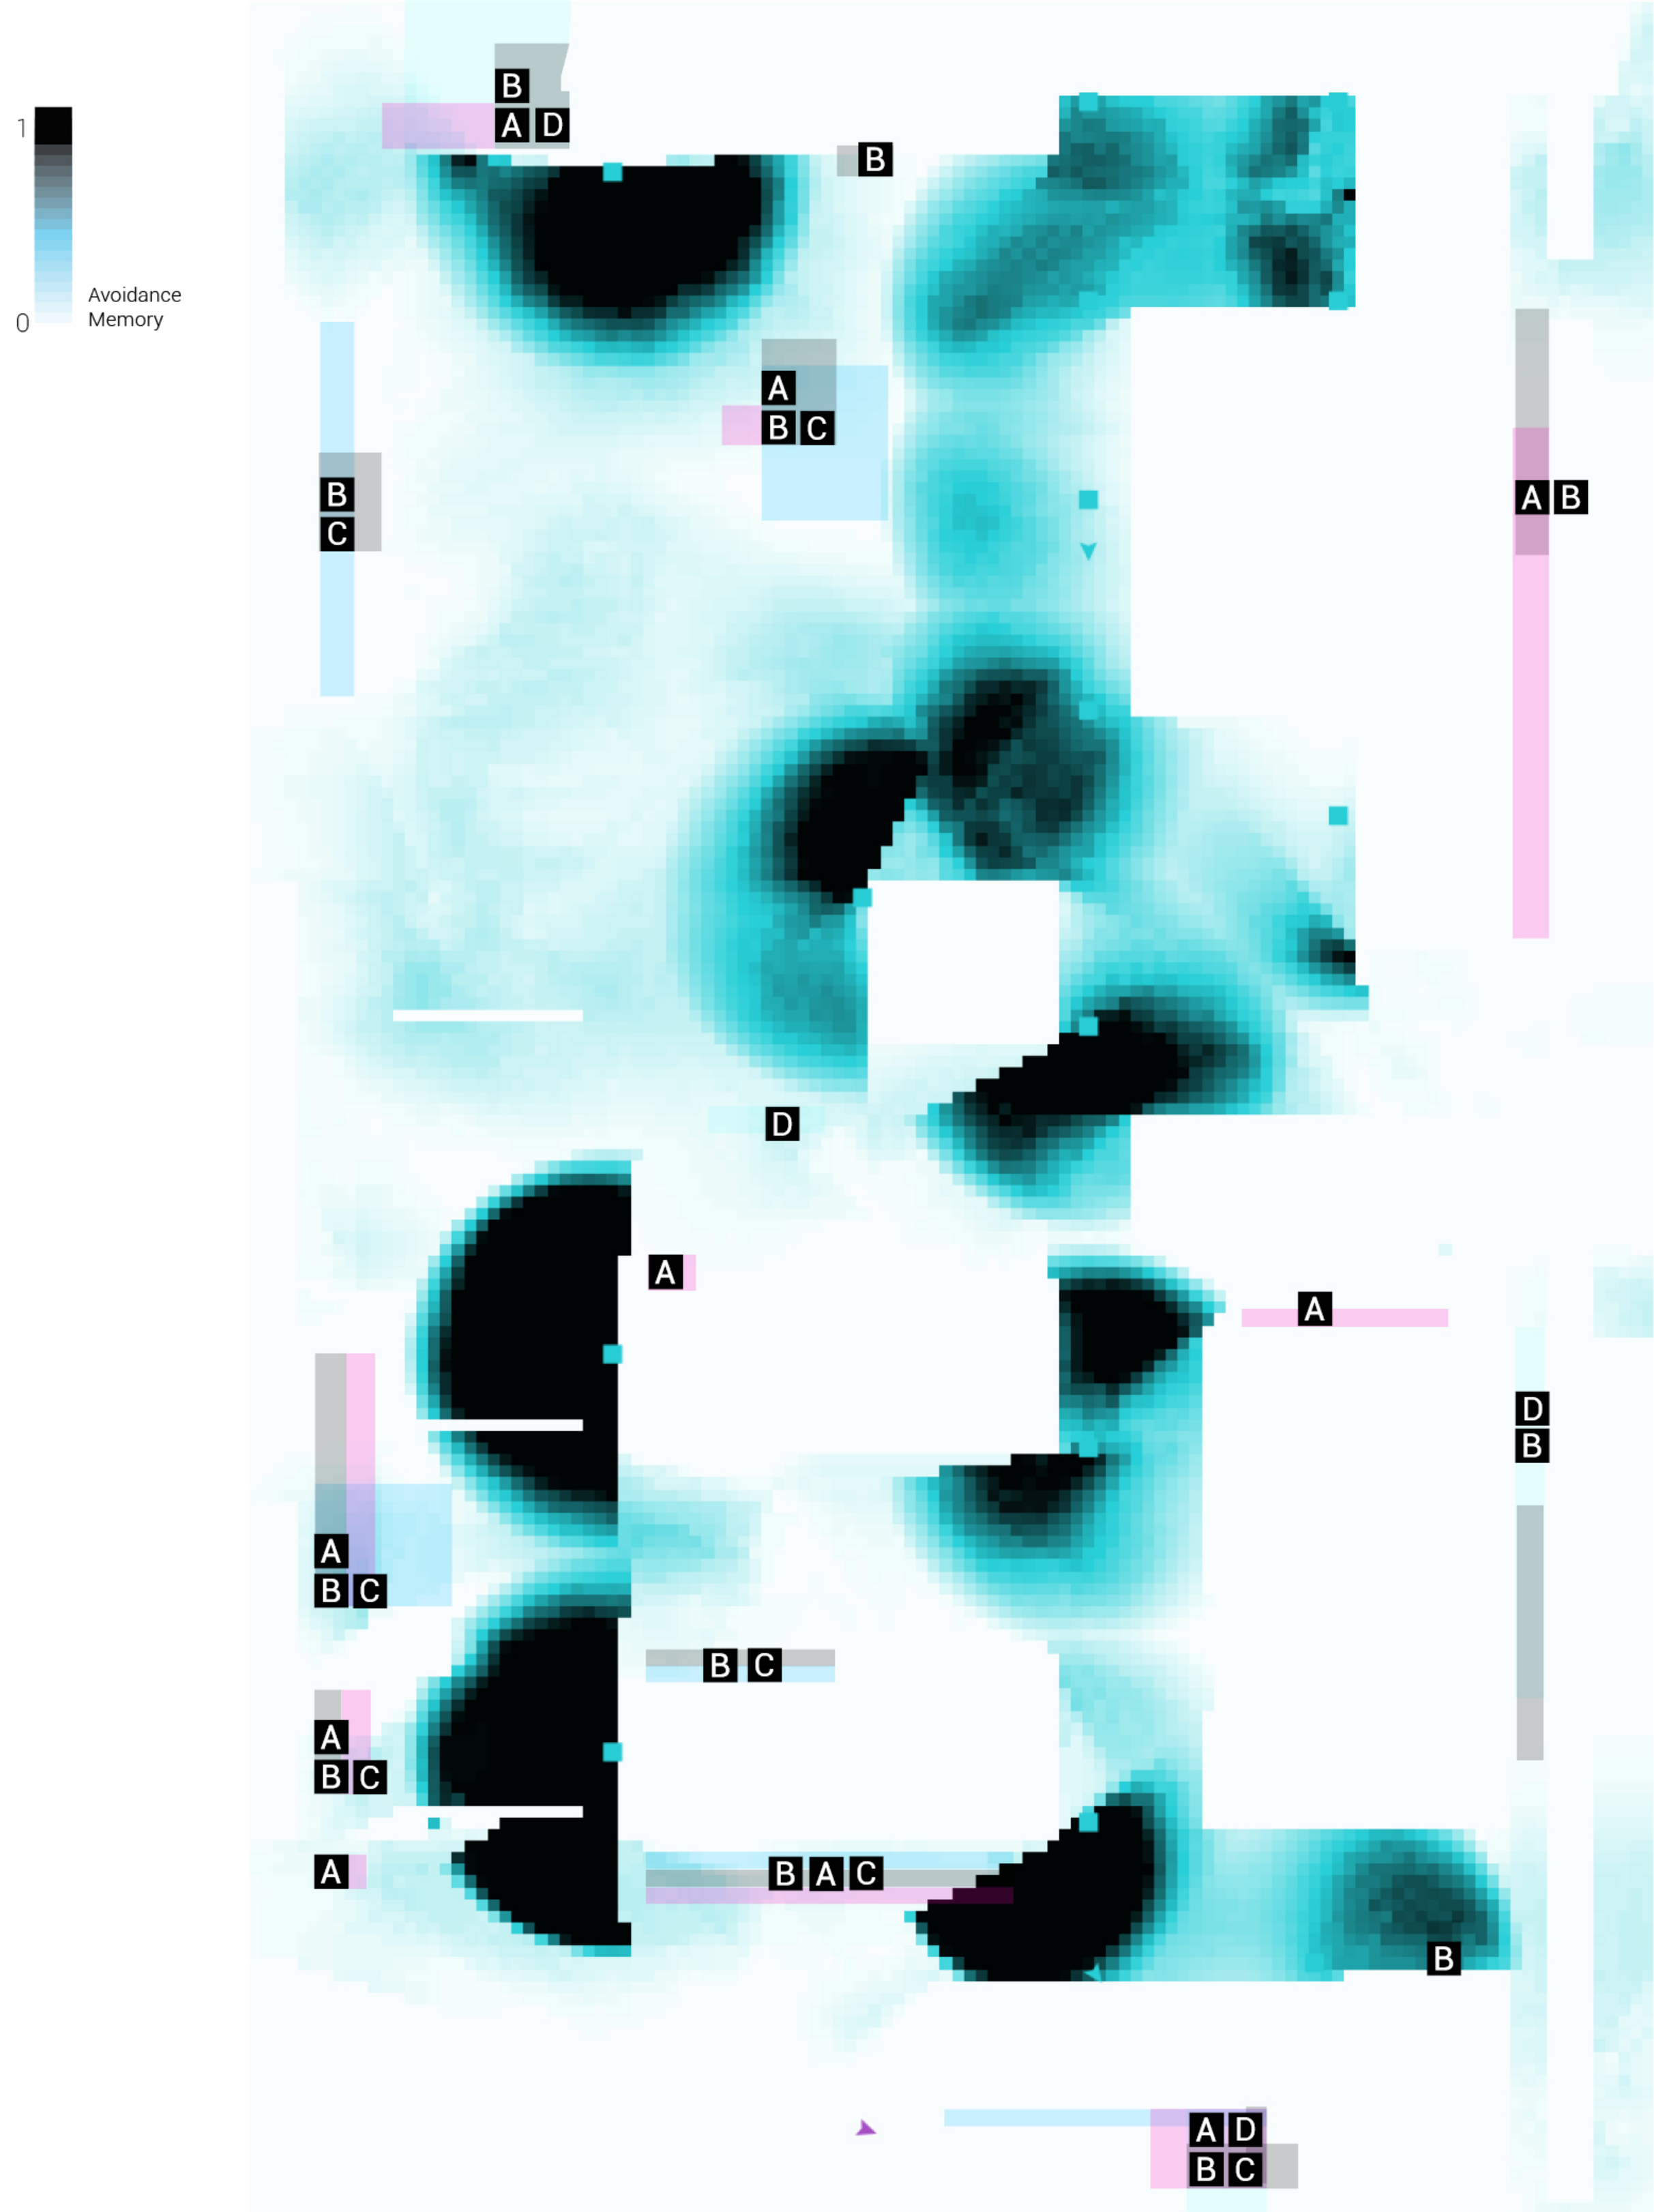

Figure 5.12 *The map of avoidance memory in Brave Scenario. All but one of the reported safe spots are located in low-stress zones. [Author]*

5.5 Conclusion and Discussion

To conclude, the behavior of agents in different scenarios can be summarized as follows:

Metric	Control	Mildly Brave	Risk-avoider	Brave
Graph Coverage	76.1%	25.9% (-50.2%)	28.8% (-47.3%)	38.5% (-37.6%)
Movements	8514	7743 (-9%)	6749 (-21%)	8117 (-5%)
Avg degree	0.72	1.64	1.33	1.2
Behavior	Diffuse	Clustered	Frozen	Active but constrained

Table 5.5 Comparative Table of all the scenarios. [Author]

5.5.1 The Collapse of the Accessible World

The simulation demonstrates that the accessible world collapses by at least half across all stress scenarios. Critically, this fragmentation is not linear but peaks around the mildly brave agent, revealing a complex relationship between risk tolerance and spatial constraint.

Non-linear Collapse: The fragmentation of the cognitive map does not necessarily progress linearly with increasing risk sensitivity. Instead, the most severe behavioral constraint appears in the mildly brave agent (Scenario 2), not the most risk-averse.

Graph coverage drops most dramatically for the mildly brave and Average degree peaks in the same scenario (+127%), indicating maximum clustering and repetitive movement. However, Total movements collapses in risk-avoider scenario.

This suggests that the relationship between fear and spatial behavior is not monotonic. Moderate risk sensitivity produces the most fragmented navigation patterns—agents brave enough to keep moving but cautious enough to avoid large portions of space. The most risk-averse agents freeze, reducing their movement but paradoxically maintaining slightly more active nodes in memory. The brave agents retain more territory but still exhibit clustering.

Quantitative Collapse: The average reduction in graph coverage reaches approximately 50%, validating the fragmentation hypothesis proposed at the outset.

Behavioral Fragmentation: Under stress scenarios, the average degree of

active nodes increases significantly—from 0.72 to 1.64. This confirms the agent is not moving efficiently through space. Instead, it clusters and repeats paths within a limited set of safe, highly memorized zones. The total number of viable edges has dropped, forcing the agent into repetitive circuits.

The Breaking Point and Withdrawal: All constrained agents experience flattening of the exploration curve. After this breaking point, the agent's spatial memory is effectively fragmented, quantified with the collapse of cognitive map in comparison to the control scenario. This fragmentation is accompanied by a drop in total movements. The Risk-Avoider agent shows the most severe behavioral constraint, with total movements dropping to 67% of baseline trials. This constitutes evidence of cognitive paralysis.

Fragmented Route Typology: The visualizations confirm that highly stressed areas are navigated through sequential, narrow routes, while safe areas become collections of equally memorized, highly interconnected patches. This is the fragmented cognitive map made visible.

The Inaccessible Void: While the avoidance threshold directly governs the psychological erasure of high avoidance-memory patches, a third, subtler category of non-use emerges: the Inaccessible Void. These are patches that remain physically open and cognitively acceptable but become structurally cut off from the agent's reduced navigation network.

The numbers are telling. In the Avoidance Threshold = 0.4 scenario, the total number of unusable patches (14,806 - 3,841 = 10,965) significantly exceeds the number of patches where avoidance memory directly surpasses the threshold. This difference is accounted for by the Inaccessible Void—patches rendered unusable not through direct stress, but because their surrounding network edges were erased.

This demonstrates that surveillance acts as a structural contaminant, causing loss of usability far beyond the immediate point of threat. The poison spreads through the network topology itself.

5.5.2 Individual Experience and Long-Term Convergence

The experience remains closely tied to the individual's trajectory through space. Learning occurs based on events encountered during exploration, meaning one agent may never interact with a mobile threat in the same location where another agent has developed maximum avoidance memory. However, over the long run, these differences diminish. The mildly brave and risk-avoider scenarios show only slight variations in their final states, suggesting that surveillance eventually homogenizes experience regardless of initial disposition.

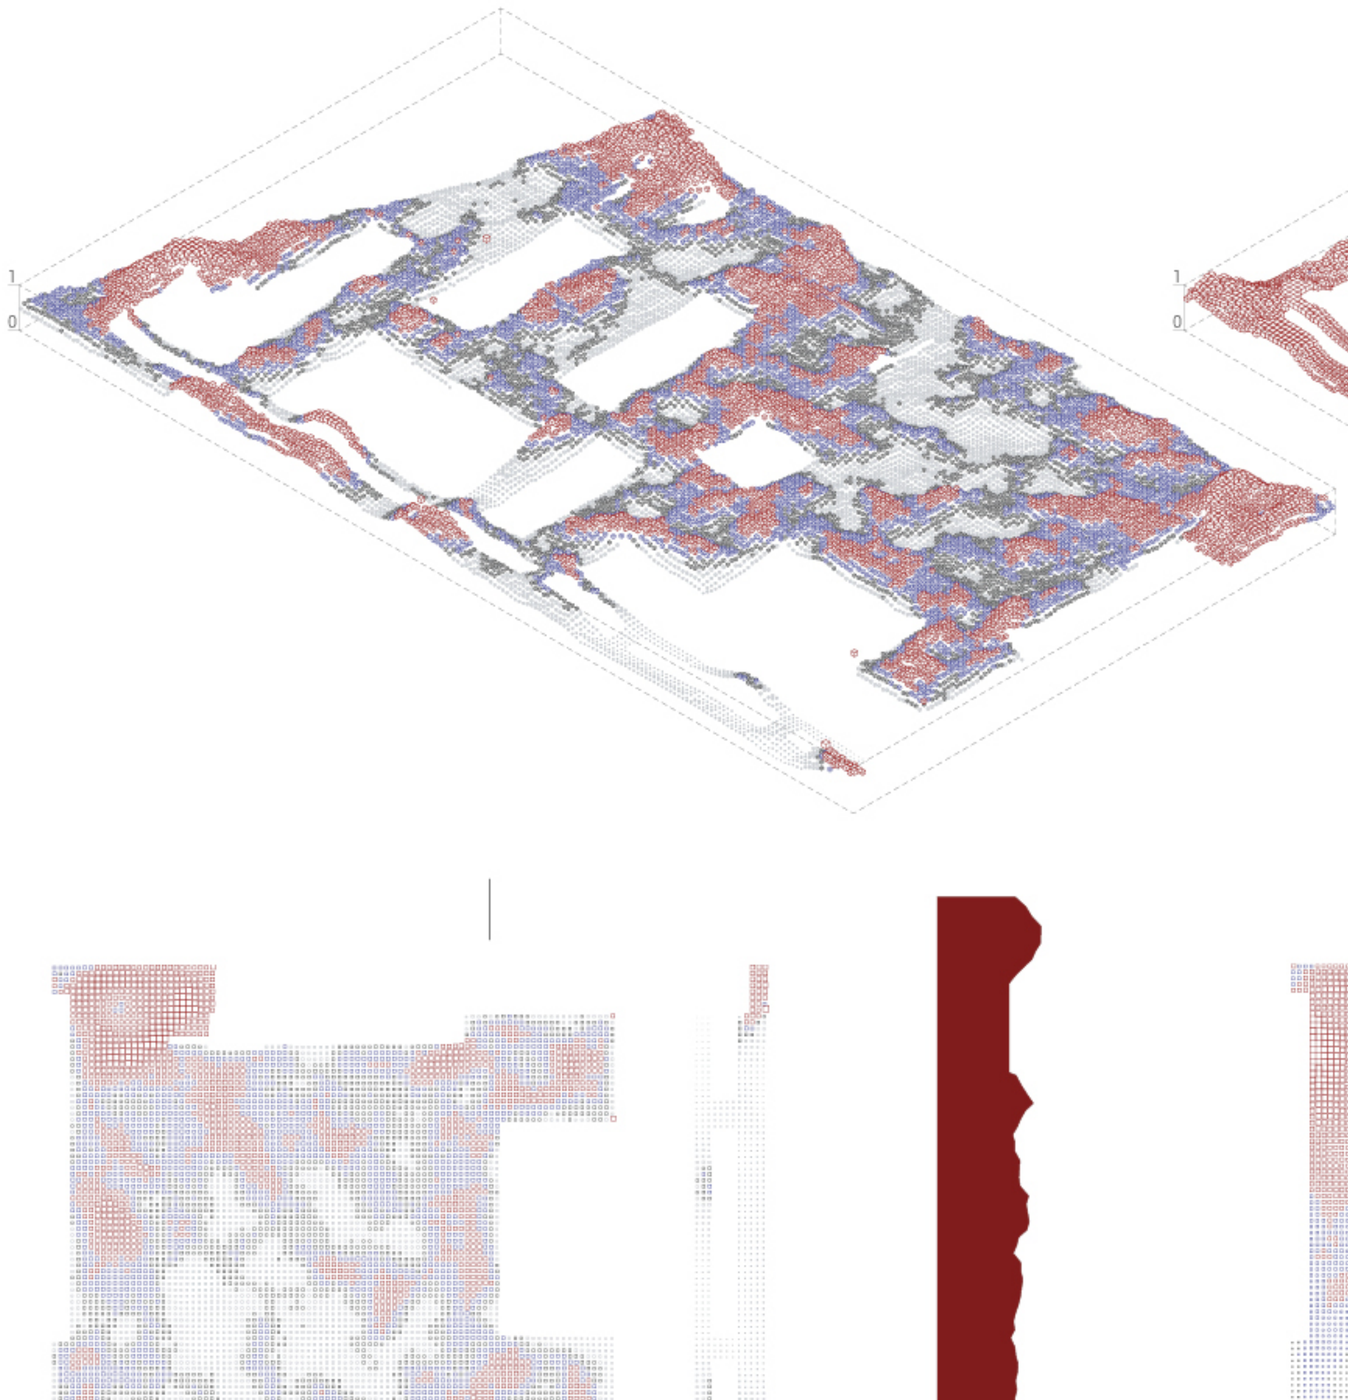

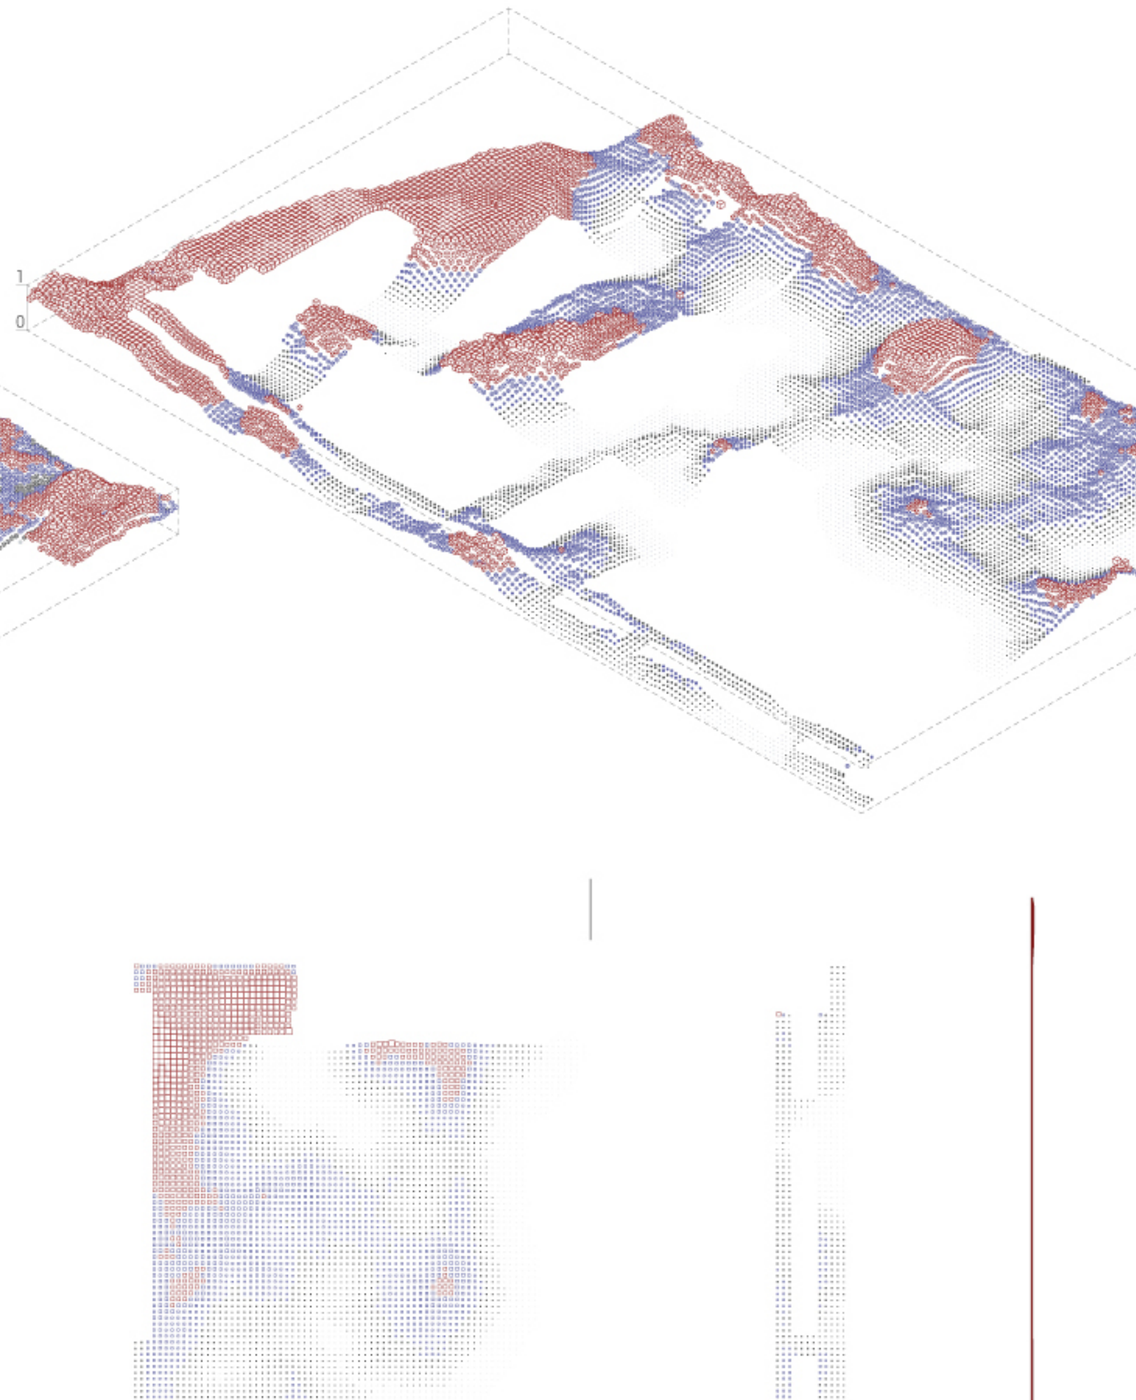

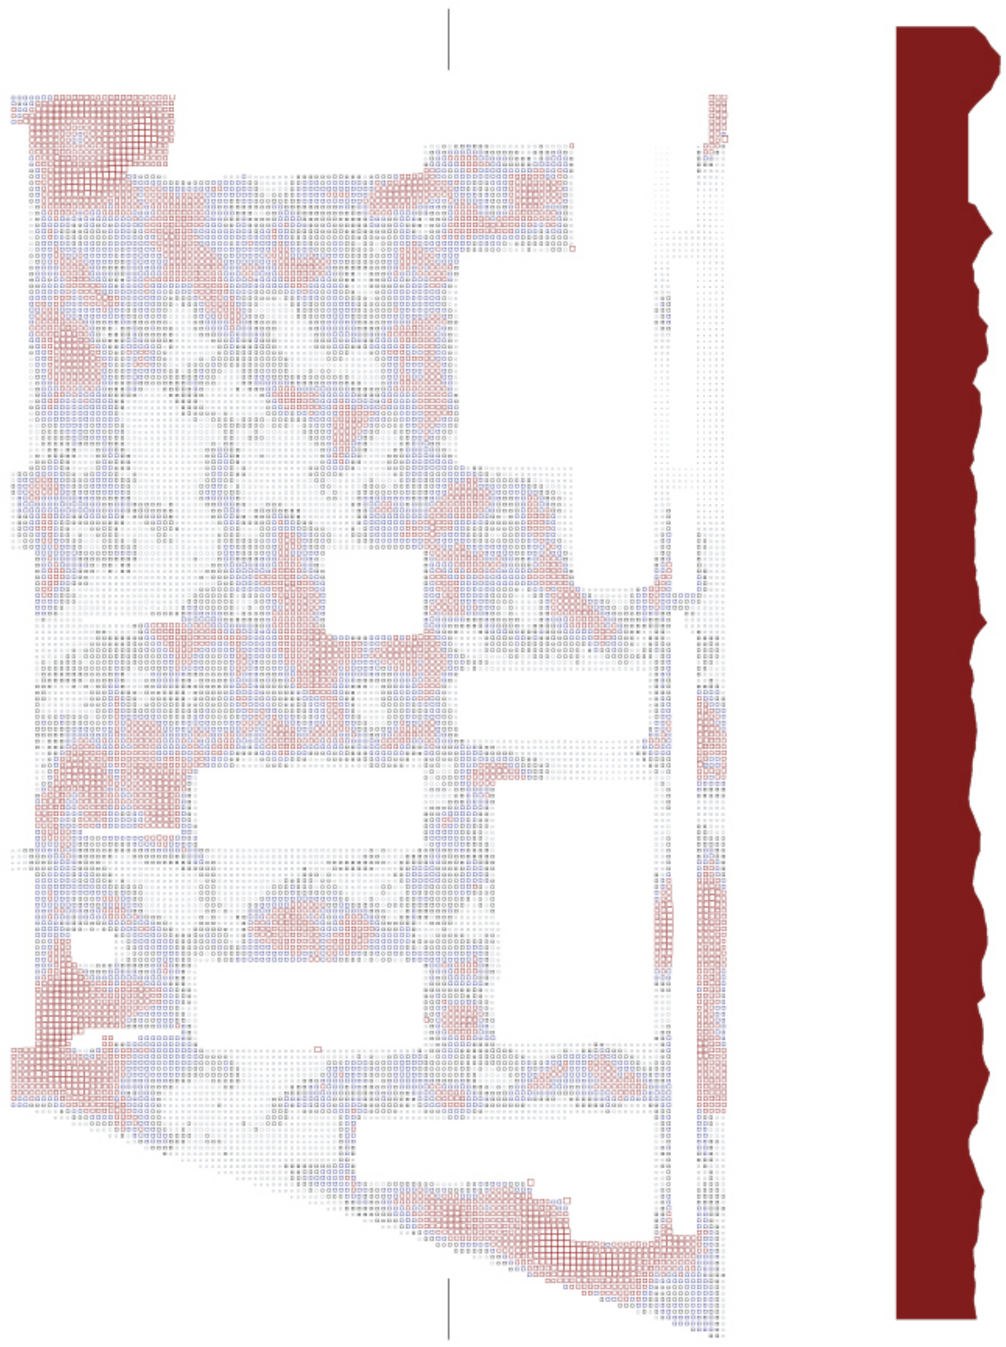

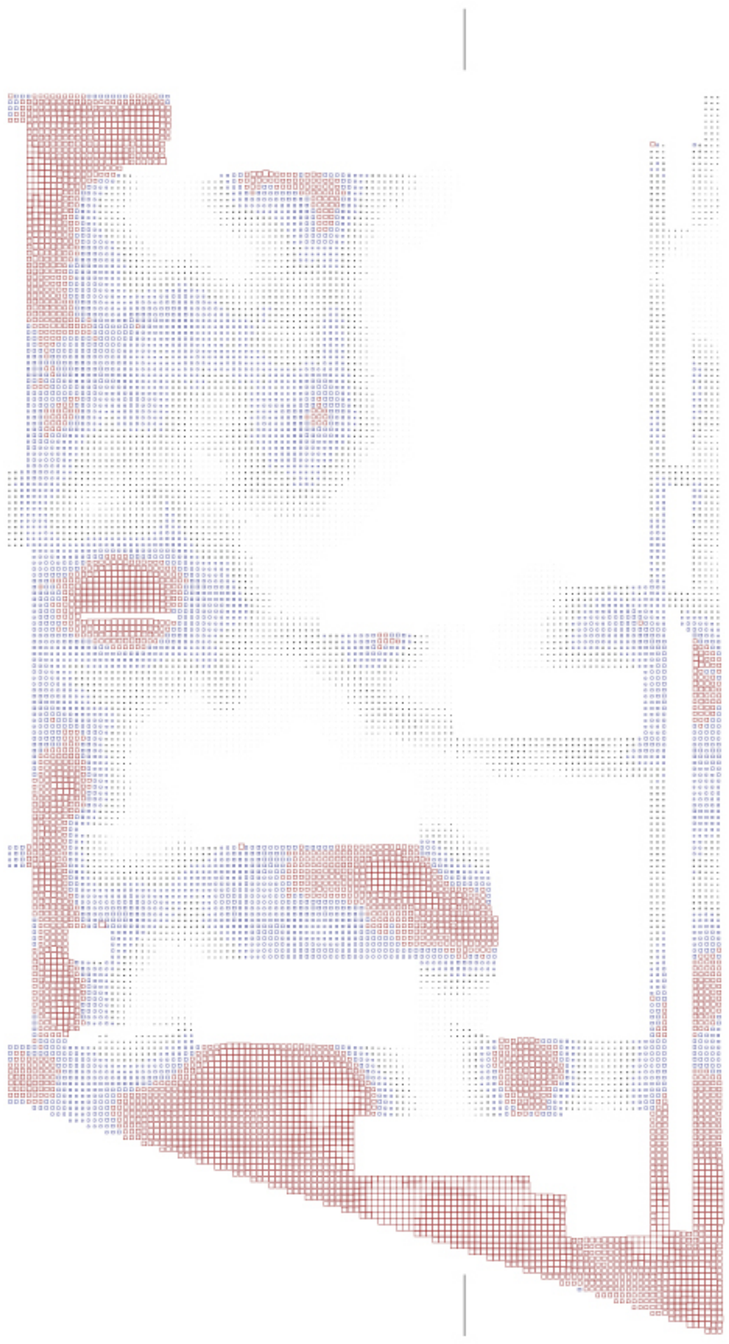

Control Scenario

The privileged explorer

The privileged explorer exhibits a diffuse, high-intensity memory landscape with remarkably uniform distribution. Both isometric and sectional views reveal consistent memory formation across the terrain, characterized by low deviation and gently modulated peaks. This baseline condition demonstrates unrestricted spatial cognition, where knowledge accumulates evenly without systematic avoidance patterns.

1st Stressed Scenario

Mildly brave explorer

Peripheral clustering emerges as the defining characteristic, with memory intensity concentrating at the edges of accessible zones. The sectional analysis reveals dramatic variation—sharp peaks in familiar territories juxtaposed with steep valleys in avoided areas. This high deviation signals the beginning of fragmentation, where navigation compresses into limited corridors while vast portions of space fade from cognitive maps.

- *High*
- *Medium*
- *Low*
- *Very low*

Intensity of Detail memory (Familiarity)

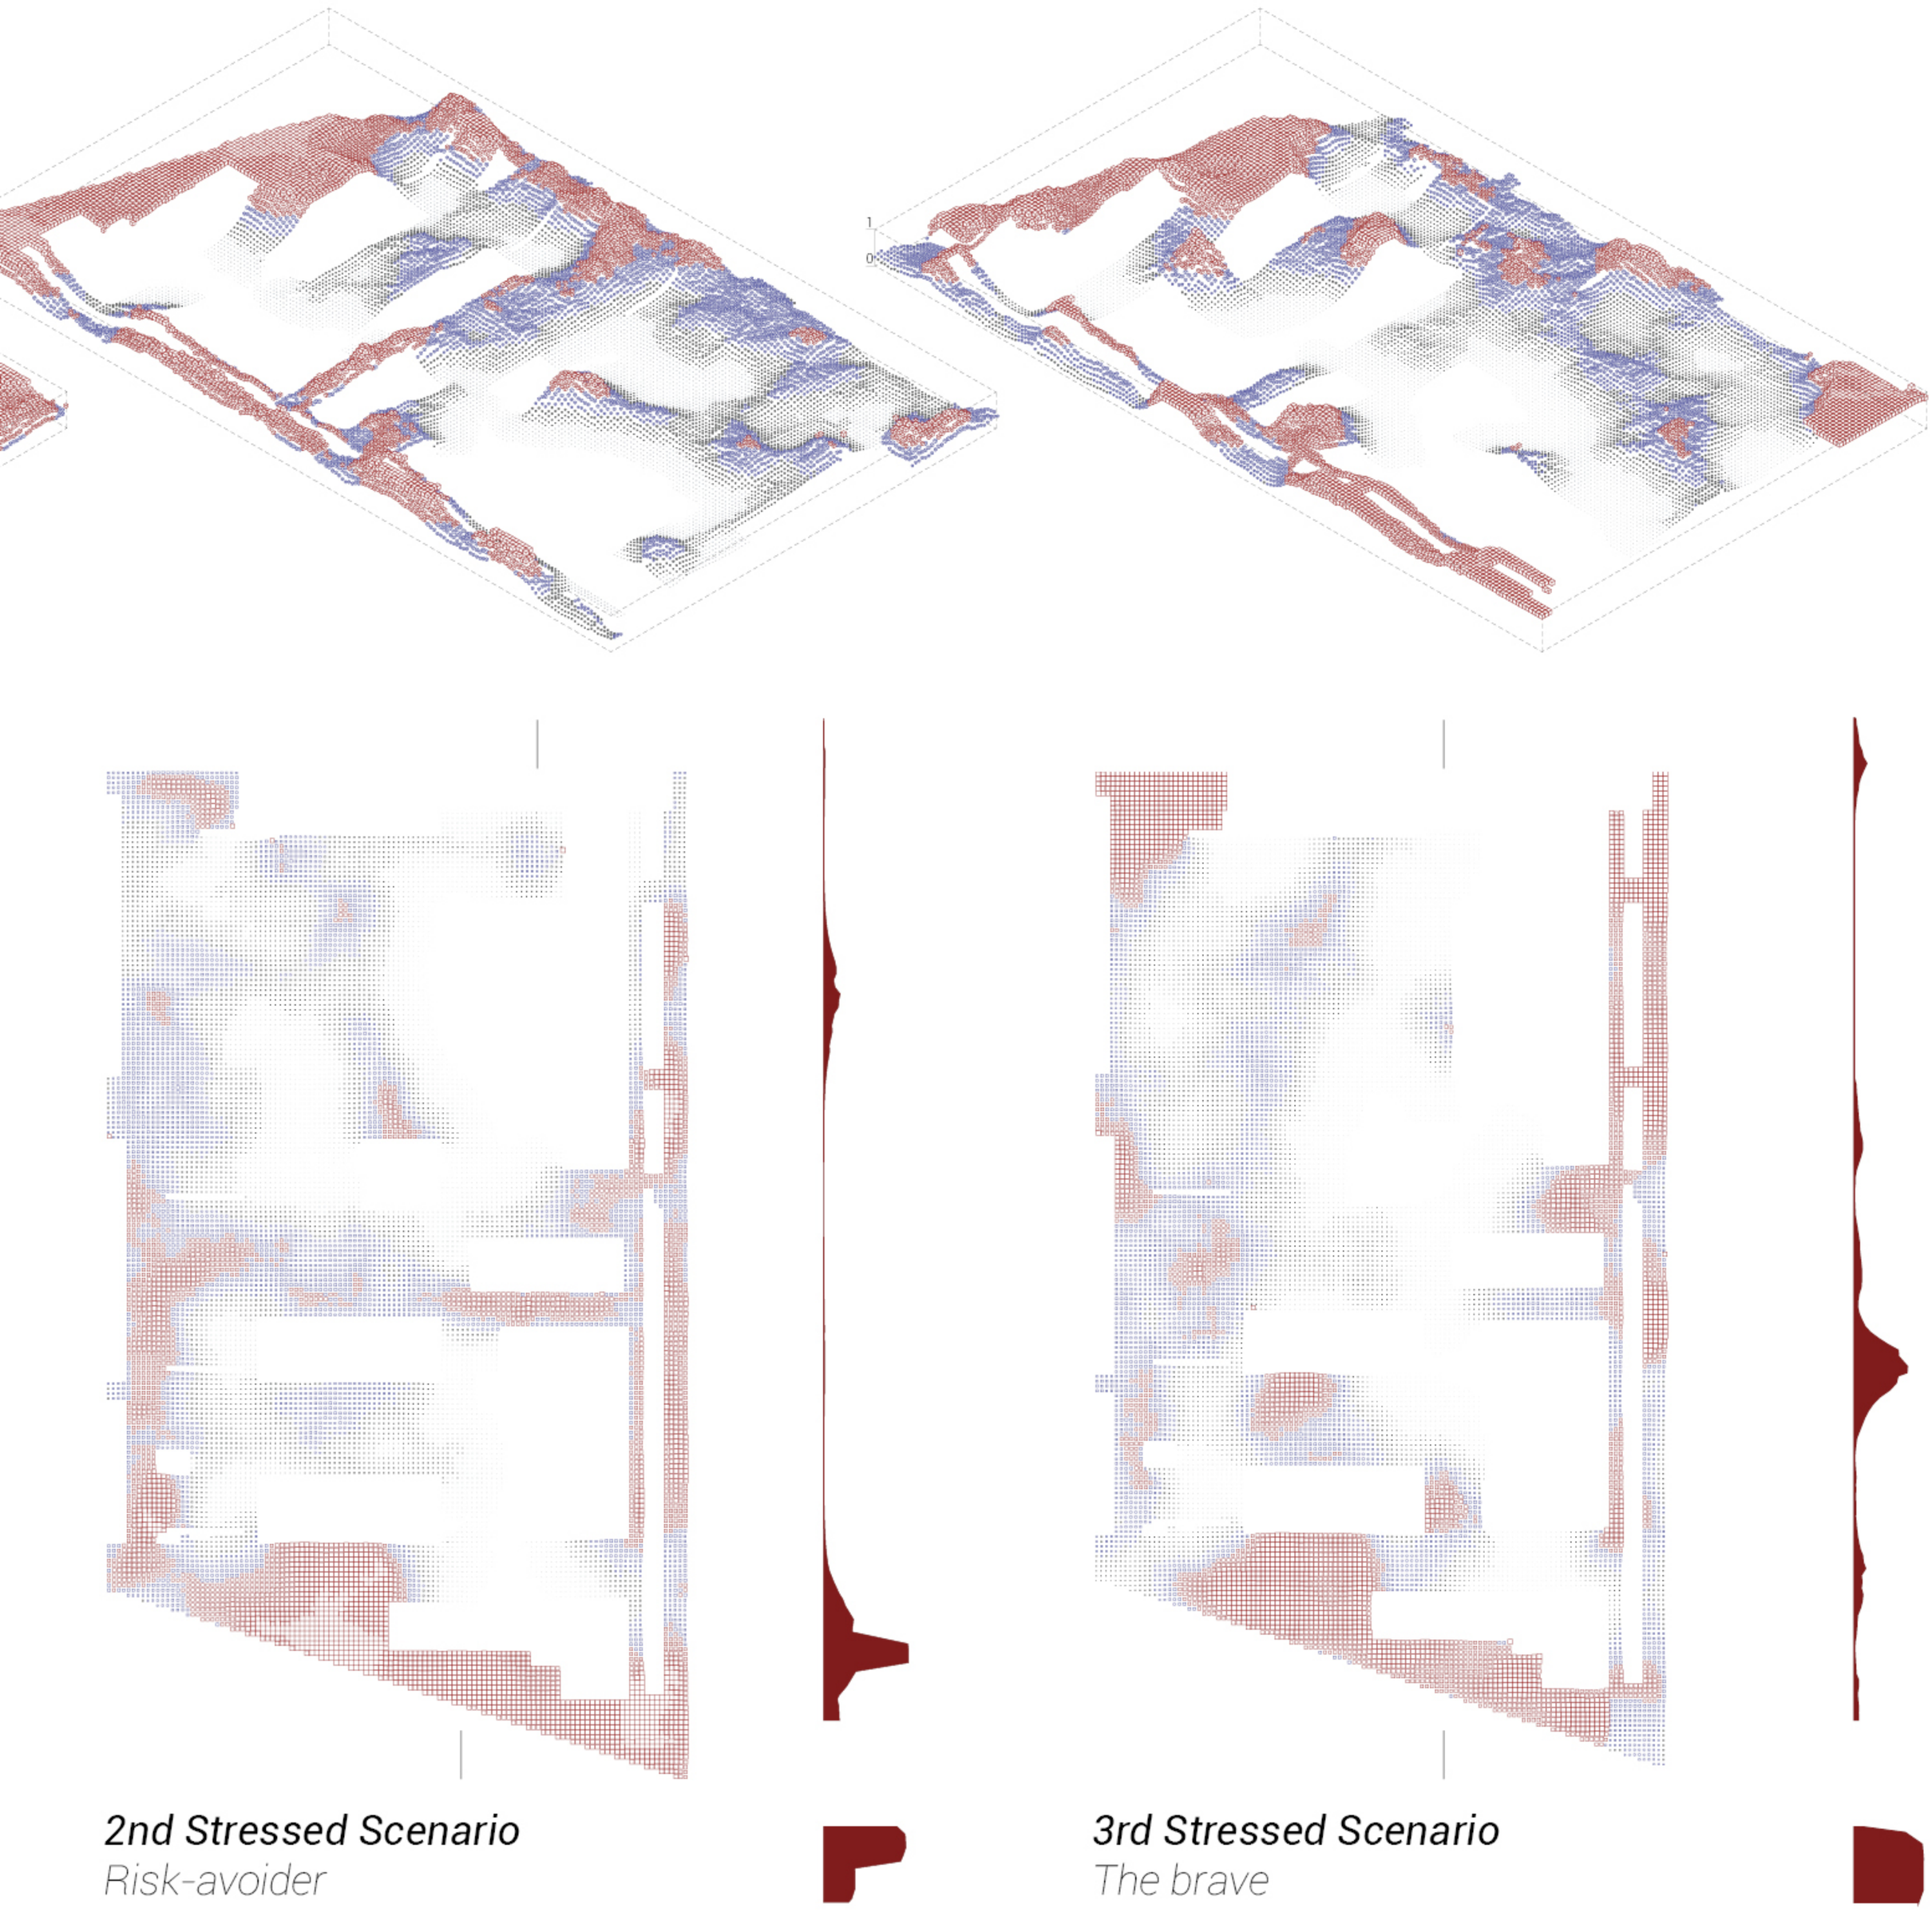

2nd Stressed Scenario
Risk-avoider

The risk-avoider exhibits clustering patterns similar to the mildly brave explorer, yet with unexpectedly intensified peaks and more abrupt transitions. The memory landscape becomes increasingly polarized—intensely familiar safe zones contrasted against near-void territories. This scenario reveals how heightened threat sensitivity produces not just avoidance but cognitive erasure, leaving remnant islands of habitable space.

3rd Stressed Scenario
The brave

The brave explorer maintains clustering patterns analogous to moderate stress scenarios, but with a critical distinction: the memory floor rises perceptibly. While peak intensities remain concentrated in safe zones, the baseline memory retention across the terrain shows modest improvement, suggesting persistent exploration despite elevated threat levels. This represents constrained yet resilient spatial cognition.

Figure 5.13 *The fragmentation of memory at the end of the simulation for all of the scenarios: the changes of detail memory is shown in different visualizations, showing peripheral clusters of familiar spaces in the stressed scenarios. [Author]*

5.5.3 Comparison with the Reported Map

All simulations subjected to stress propagation show clustering patterns similar to those generated by myself and my colleagues during the original mapping exercise. There are inconsistencies due to model limitations—the simulation does not account for small-scale spatial settings or temporal changes in surveillance. For example, the gardens provide well-hidden safe spots. Nevertheless, the agent's safe zones match the safe spots we actually used to a great degree.

This alignment suggests that, consciously or unconsciously, we were following a similar navigational mechanism. It further demonstrates that surveillance reduces the cognitive map to a single dimension: the level of threat. Individuality and inner states lose their effect on navigational decisions. As chunks of the urban environment are removed from individual memory, the individual disappears from these parts of the city.

5.5.4 The Emergence of Concepts

The emergence of the Inaccessible Void and the loss of options to the point of paralysis are among the most interesting findings of this experiment. These are concepts I did not notice prior to this work, however visible they were in lived experience. As navigational options collapse under threat, spaces that are not themselves surveilled are cut from the agent's memory and possibilities. The protagonist not only loses optionality but lives in a state of constant doubt and fear, leading to a slow shrinkage of the inner landscape and degradation of urban experience.

However, I do not conclude that the behavioral options of the protagonist are necessarily compromised in all dimensions. While unjustly deprived of navigational options and urban space, the constrained inhabitation leads to clusters of patches with high detail-memory. As familiarity with remaining spaces increases, small-scale decisions appear: painting graffiti, sitting on a window sill, appropriating neglected corners. Spatial oppression may paradoxically intensify certain forms of place-making within the residual territory.

5.5.5 Limitations of the Model

While the simulation reveals the fragmentation of individual spatial memory under surveillance, it has several limitations:

Single Agent Simulation: Collective dynamics and social learning significantly affect individual decision-making. These features are not included here, to simplify the experiment and isolate the studied parameters.

Static Threat Patterns: The simulation does not consider the adaptability of threat agents. Patrols and cameras are pre-defined by movement patterns or fixed locations, whereas real surveillance adapts and learns.

Simplified Environment: Spatial details, small-scale obstacles, and visual occluders are not fully represented in the simulation.

Homogeneous Risk Baseline: In reality, the spatial variation of threat is heterogeneous and layered. The real environment is not a uniform space with identical memory values, but an already heterogeneous terrain of information, history, and social meaning.

Limitations in Data: The data presented here dates from 2016-2021 and is auto-ethnographic. While reinforced by peer mapping exercises, it remains limited in temporal scope and perspective.

Single Case Study: The argument could be strengthened by studying additional case studies across different contexts and populations.

5.5.6 Future Research

This study can expand in several directions. By adding multiple agents and including collective decision-making mechanisms, modifying stress propagators to exhibit adaptive learning, and conducting comparative studies across different environments, the experimental framework can continue to grow.

Design implications should be integrated into future iterations. By introducing bidirectional stress fields and inverting the gaze—allowing the agent to become a source of discomfort to surveillance—the effectiveness of counter-strategies and emergent resistance patterns can be studied. The simulation could also allow agents to create temporary changes in space through furniture, installations, vegetation, and other territorial markers, observing how territories grow and shrink in response.

Beyond surveillance, the framework developed here—understanding environment as a terrain of uncertainty that stores ever-changing information relative to specific foci of attention, with information revealed through sensing and movement—is expandable to non-human navigation. Robot sensing, autonomous surveying, and other forms of machine perception could benefit from these insights.

5.5.7 Implications for Urban and Architectural Design

The experience of an individual is not often treated as a game-changer in urban design or even architecture. However, the results of this study point

toward a fundamental observation: the architect or urban designer is an enabler of choice, and each wall effectively creates and removes various options from the possibility space. Every design decision is also a decision about what will be remembered, what will be avoided, and what will become structurally inaccessible through no direct prohibition. The next section discusses the lessons learned through this experiment and their implications for spatial practice.

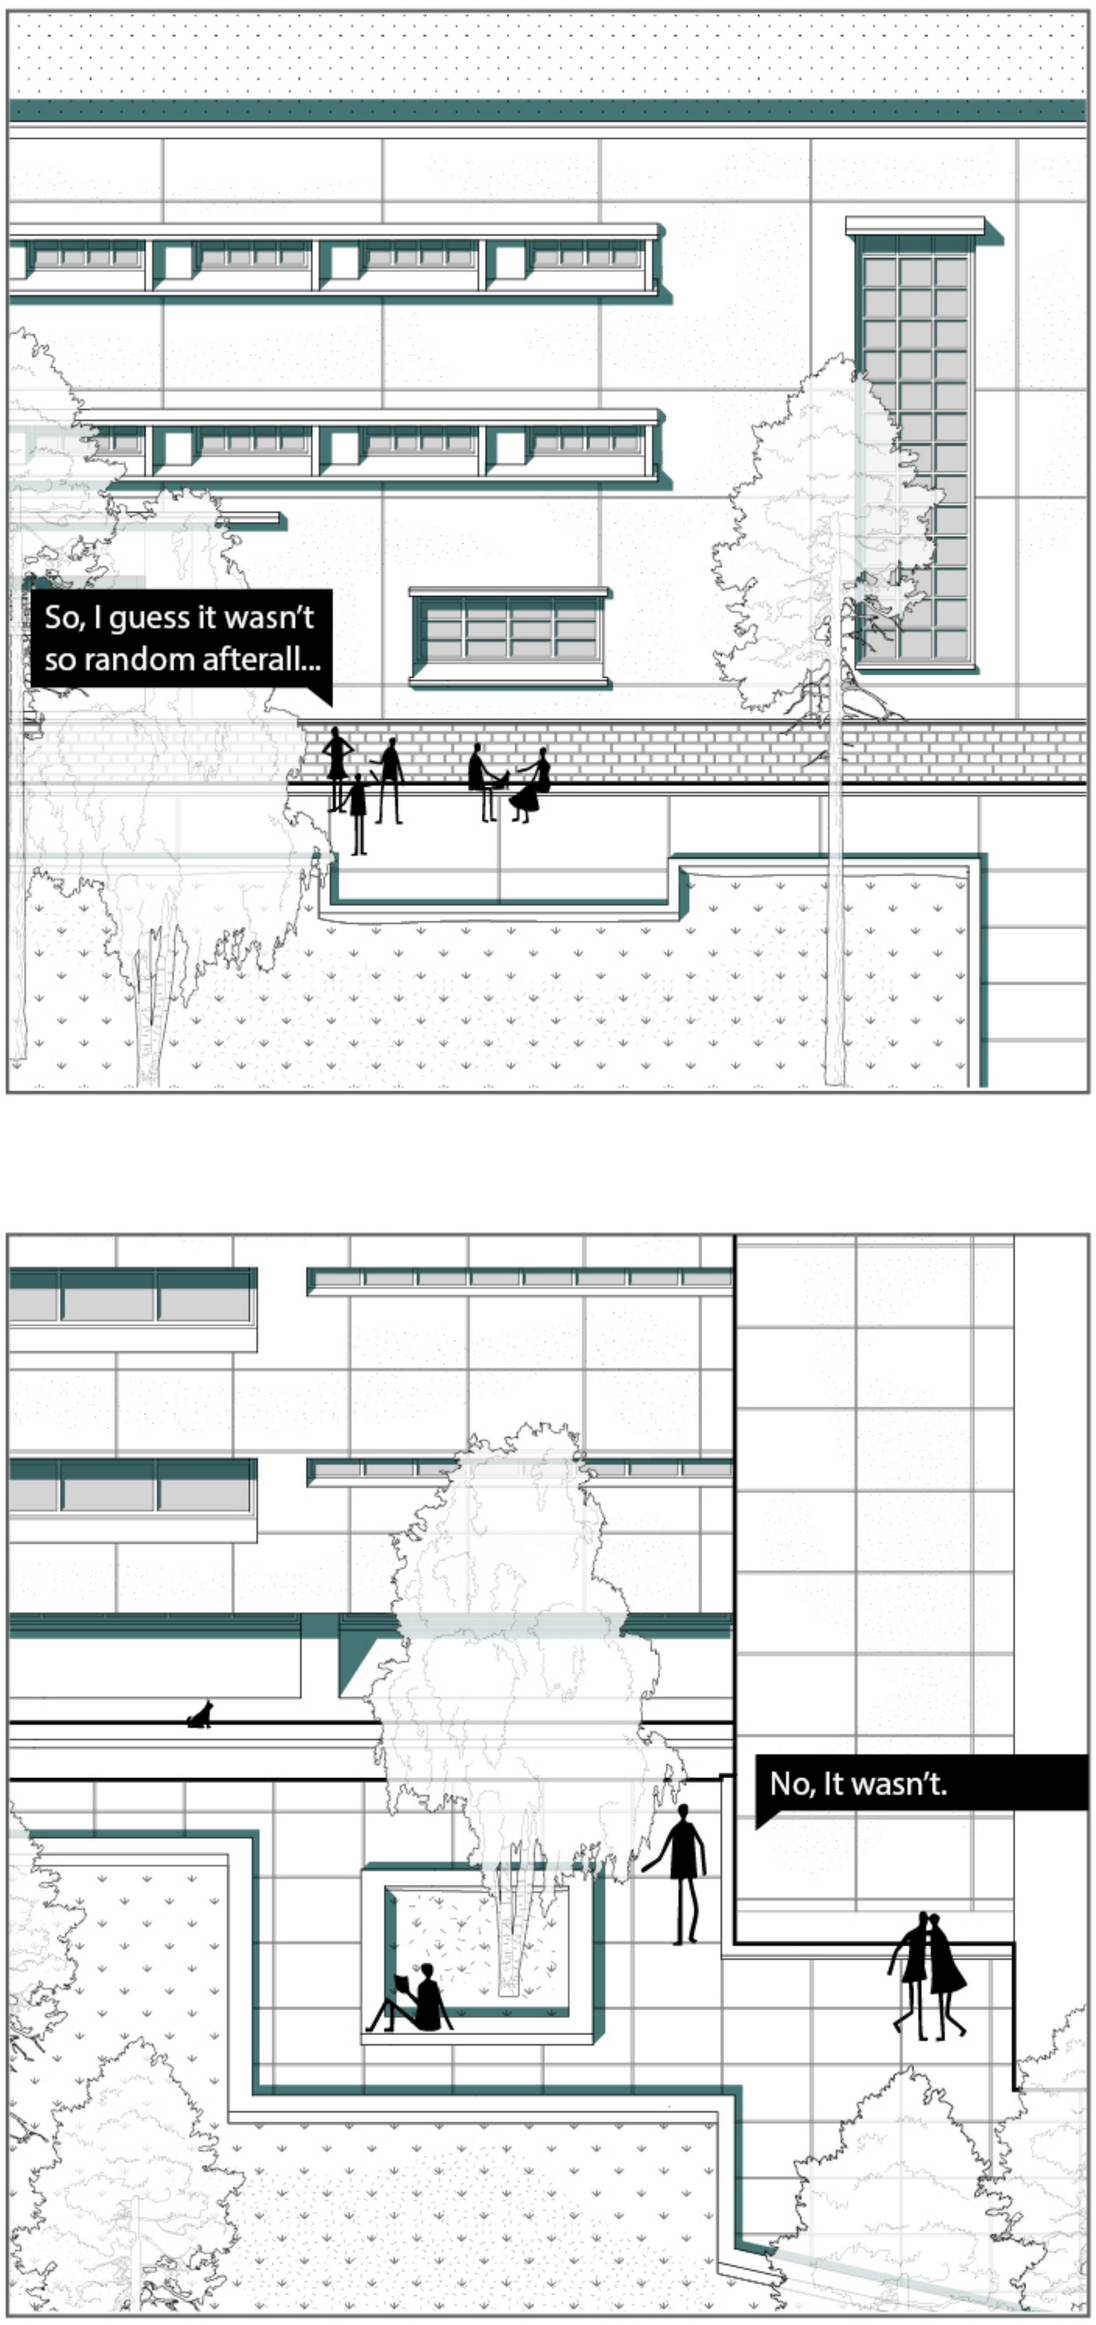

06 | Discussion

Design Implications
Architecture as Friction

The simulation demonstrates that surveillance fundamentally restructures spatial behavior, reducing graph coverage by 50-75% and, at extreme levels, inducing paralysis. Furthermore, the violence induced by surveillance is not merely social, but spatial, reducing the possible world to a one-dimensional ToU and shrinking the experiential environment.

The role of the architect and urban designer in resisting this cannot prevent a determined state power from establishing control. However, architecture can increase the friction, the cost in labor, money, legitimacy, and visibility required to achieve total surveillance.

For example, controlling all space in a centralized plan is far easier than managing the scattered, porous structure of the Faculty of Fine Arts. The faculty's design maximizes user options while complicating the establishment of oppressive control, even as it accommodates necessary security functions. This is not naïve resistance but strategic design: accepting that security exists while refusing to make control easy.

The following sections outline design strategies derived from this study, interventions to maximize friction and preserve cognitive optionality for users under surveillance pressure.

6.1 The Designer's Ethical Crisis: Defending Optionality

The simulation's most significant finding is not just that surveillance reduces spatial access, but that it collapses the cognitive optionality of navigation. Agents don't lose space uniformly; they lose the ability to choose. At an avoidance threshold of 0.2, movement drops to 67% of trials. The remaining 33% is not detours or alternative routes—it is paralysis and withdrawal.

This is the designer's ethical crisis: architecture that enables total surveillance doesn't merely constrain movement. It subordinates all other values, curiosity, sociability, aesthetic experience, spontaneity, to a single imperative: threat-avoidance. The environment becomes one-dimensional.

The designer's obligation is not to eliminate security, which may be necessary and legitimate, but to prevent this tyranny of single-value optimization. The goal is to maintain cognitive optionality even under surveillance pressure, ensuring that individuals retain the capacity to make choices based on multiple, competing values rather than fear alone.

The strategies below outline how architecture can defend optionality without denying the reality of security concerns.

6.1.1 Mutual Visibility (Invert the Gaze)

Surveillance assumes unidirectional observation: the authority sees, the subject is seen. The first counter-strategy is to introduce bidirectional visibility, using the same panopticon logic in reverse. In the simulation, this strategy is equal to decreasing intensity of stress propagated by the threat agents, which can result in tolerable stress levels at some locations even in the presence of the threat agent. [Figure 6.1]

The Mechanism: Watch the watchers

By designing spaces where security actions are visible to a public audience, architecture creates a form of passive deterrence. The presence of potential witnesses—students in balconies, pedestrians on elevated walkways, passersby near transparent fences—decreases the probability of harsh, unrecorded consequences. Confrontations that occur in public view are subject to social and institutional accountability in ways that isolated encounters are not.

This is watching of watchers. It introduces a bottom-up stress propagation that affects threat agents themselves, reversing the asymmetry of the panopticon.

Design Elements

| Maximize the audience: Balconies, terraces, and elevated pedestrian routes that overlook zones of potential confrontation (entrance gates, patrol congregation points, administrative checkpoints).

- **Transparent boundaries:** Fences and barriers that allow visual access rather than solid walls that isolate and conceal.

- **Blend public and private:** Educational spaces, cafes, and informal gathering areas positioned adjacent to security interfaces, ensuring that enforcement actions occur within sight of everyday activity.

Effectiveness and Limits

This strategy leverages the public nature of institutional spaces and depends heavily on societal sensitivity and willingness to intervene. In contexts where public witnesses are indifferent or complicit, mutual visibility loses its deterrent power. However, even in authoritarian contexts, the visibility cost. The labor required to manage optics and justify actions, remains non-zero.

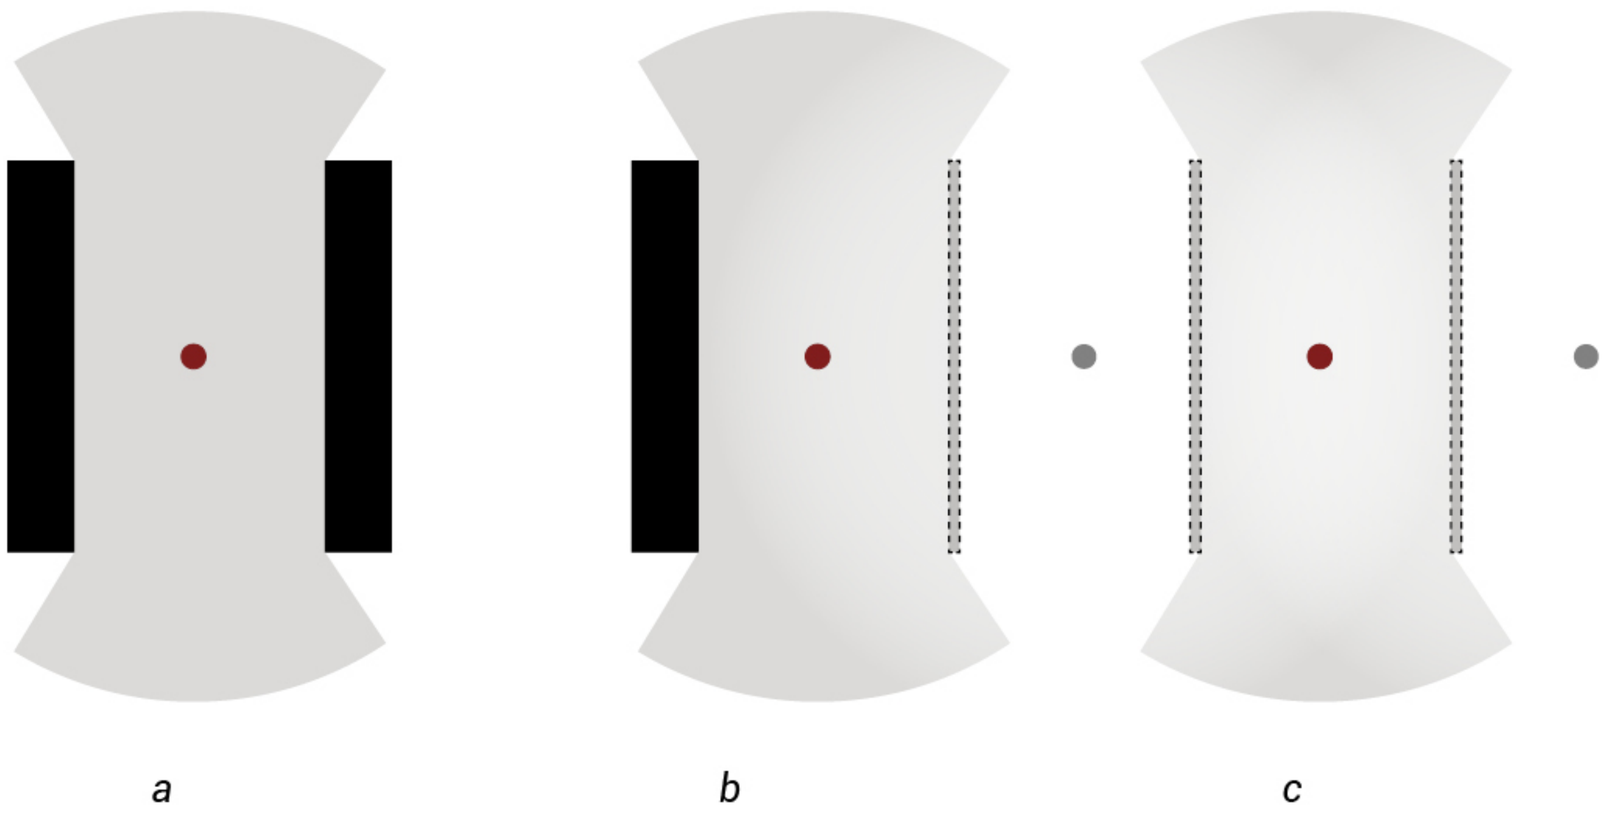

Figure 6.1 *The decreased intensity of propagated stress by introducing public gaze to the design, or inducing stress for the threat agents: a. Passage in blocked on both sides by walls, the propagated stress intensity is the highest possible, b. Passage is visible to public on one side by implementing fences, the counter-stress reduces the stress intensity, c. Passage is open to both sides, maximum use of public gaze as counter-stress [Author]*

6.1.2 Efficient Zoning (Security Without Control)

Not all security is oppressive. Perimeter access control, asset protection, and emergency response are often necessary functions. The challenge is to prevent necessary security from metastasizing into comprehensive spatial control.

Zone Isolation

The Faculty of Fine Arts' scattered pavilion plan is effective because it allows security functions to be localized and compartmentalized. By concentrating security oversight in specific zones (such as the open corridor that connects buildings), the design ensures that most space remains undefined and open.

This is efficient zoning: accepting that certain interfaces (entrances, exits, storage areas) require monitoring while resisting the diffusion of surveillance across all public space. The mechanism is equal to minimizing view-cone of the threat agents, reducing number of patches affected by the propagated stress. [Figure 6.2]

Open vs. Controlled Space

The architectural goal is to maximize the ratio of open, multi-purpose space to explicitly controlled zones. This resists the assignment of high perceived stress to most areas. Spaces designed for wandering, informal gathering, and spontaneous appropriation remain structurally resistant to surveillance because they lack the functional legibility that control requires.

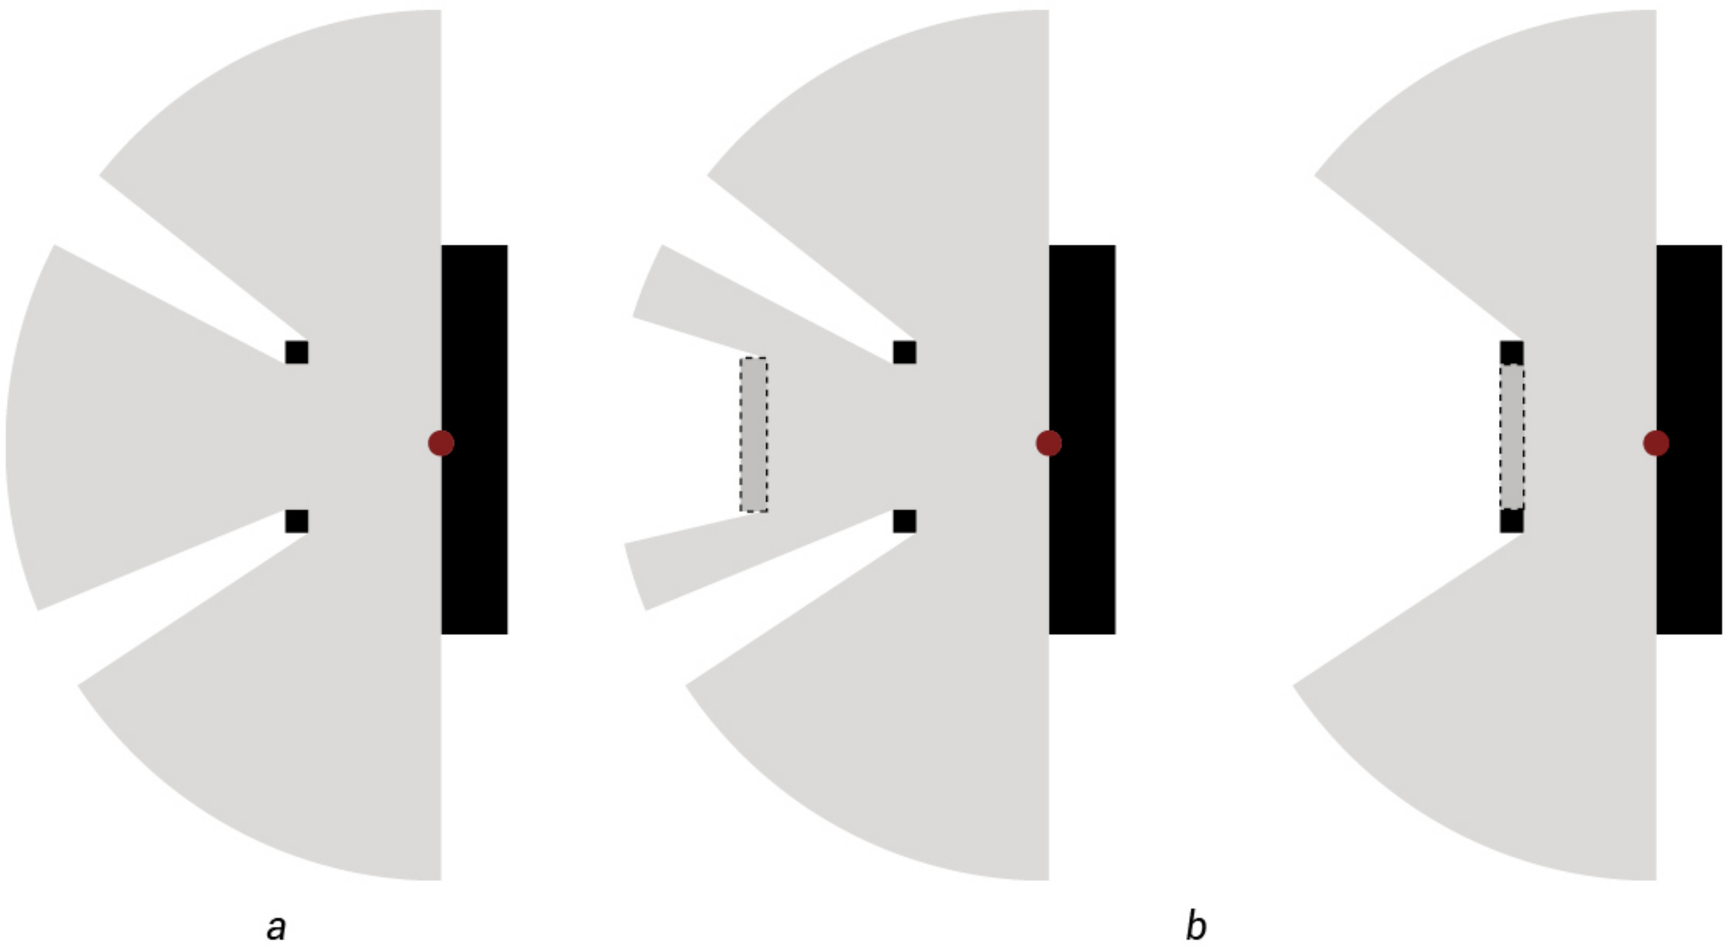

Figure 6.2 *Maintaining security of assets while minimizing control: a. The camera's view cone is almost completely unobstructed, b. Using a semi occluder (e.g. A hedge, an installation, etc) the view cone is reduced to different extents, while preserving the necessary surveillance [Author]*

Efficient vs. Oppressive Structure

An efficient security structure concentrates oversight where it is functionally necessary. An oppressive structure treats all space as potentially threatening and thus subject to observation. The design distinction is clear: compartmentalize security; distribute freedom.

6.1.3 Functional Ambiguity (Resist Legibility)

Control requires predictability. Surveillance systems—whether human or algorithmic—depend on spaces being legible, classifiable, and mappable to expected behaviors. Architecture can resist this through functional ambiguity.

Design for Wandering

Spaces that resist singular classification evade the bureaucratic legibility required for surveillance. A space that can be a study area, a gathering spot, a protest site, or a place to rest does not map cleanly onto control schemas.

Design tactics include:

- **Open, undefined spaces:** Areas without prescribed function that users can appropriate as needed
- **Scalar variety:** Spaces that accommodate intimate gatherings (2-3 people) and collective assemblies (50+), preventing functional fixity
- **Paths without obvious destination:** Circulation that encourages wandering rather than point-to-point transit
- **Temporal layers:** Movable furniture, temporary partitions, and user modifications (graffiti, posters, installations) that allow space to be redefined dynamically

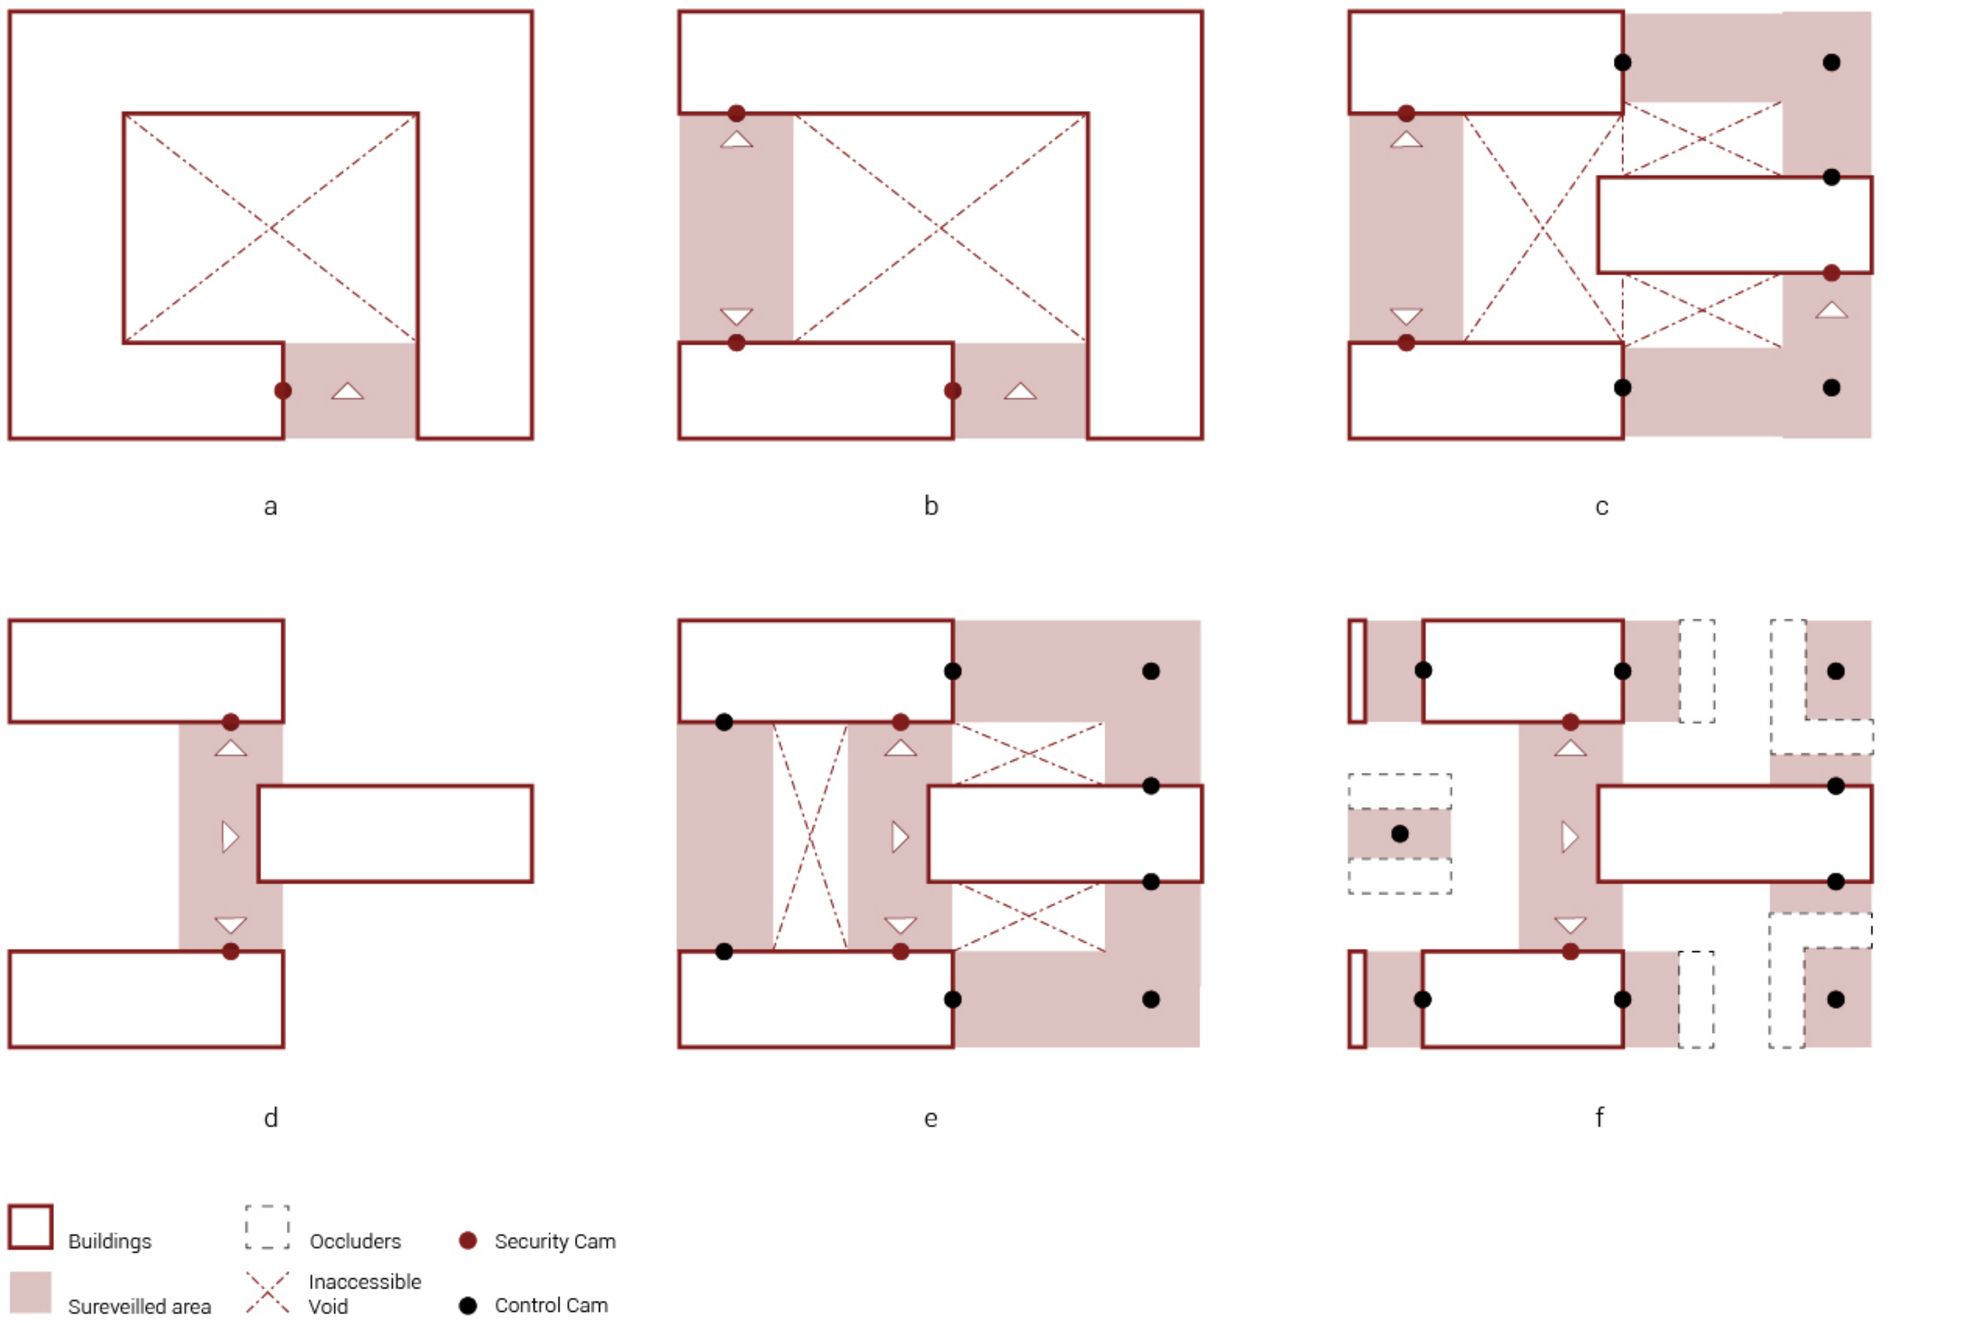

Figure 6.3 *Sample of increasing friction through architectural design: a. The castle, with only one entrance the security and control act as one, minimum friction, b. Opening up the structure, the inaccessible void remains the same, but the friction increases, c. A structure similar to Faculty of fine arts with the entrances placed to minimize friction, d. By changing the access point the minimum surveillance is needed to for security, e. Imposing surveillance of a here requires excessive effort as the friction is very high, and f. By using temporal layers, ambiguous spaces, and other occluders the friction is even higher. [Author]*

In-Between Spaces

The architectural strategy is to introduce mediatory zones—thresholds, courtyards, covered walkways, alcoves—that exist between clearly defined functions. These spaces complicate surveillance by multiplying the number of distinct zones that require monitoring.

A gradient of spaces, from fully public to semi-private to intimate, creates surveillance fatigue: the cost of comprehensive coverage becomes prohibitive. This strategy contributes also to maximizing affordances, and it is connected to the next strategy, both increasing the optionality and decreasing predictability.

Strategic Porosity

Porosity must be designed carefully. Naïve porosity—maximum visual openness—can be co-opted by surveillance, creating perfect sightlines for panoptic observation. The solution is not maximum occlusion but multiple, equally viable sightlines that prevent any single vantage point from becoming critical.

This is why the Faculty of Fine Arts' scattered structure works: there is no single "panopticon point." Achieving total surveillance requires comprehensive coverage across multiple, non-overlapping zones, making monitoring economically and logistically prohibitive.

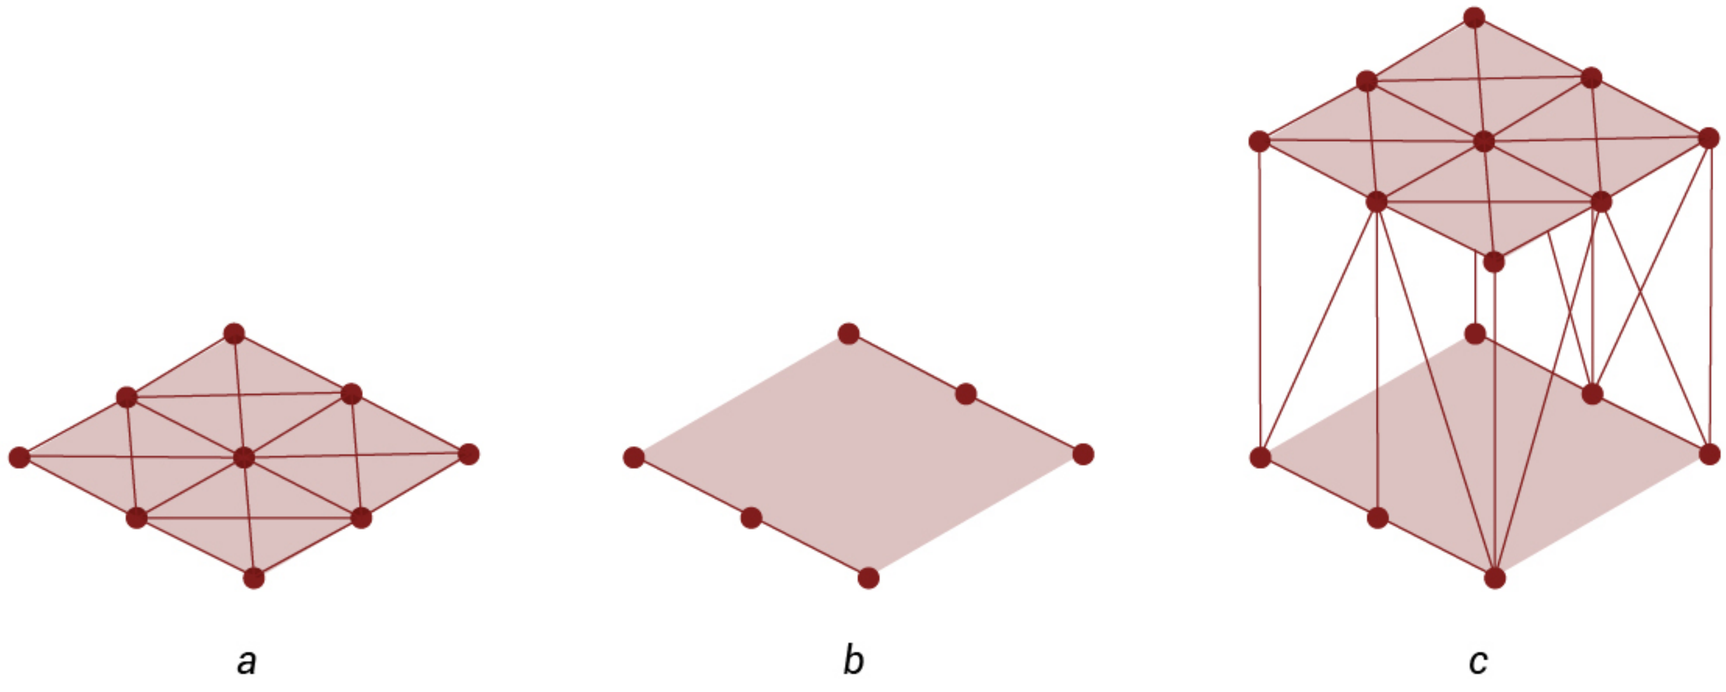

Figure 6.4 *Network redundancy as a tool to increase friction and optionality: a: the primary network without any obstruction, b: the fragmented network, c: by adding a bridge and additional nodes, the connection is restored [Author]*

6.1.4 Network Redundancy (Counter the Inaccessible Void)

The discovery of the Inaccessible Void—patches lost not through direct threat but through network fragmentation—reveals a critical architectural challenge: spaces become unusable when they are structurally isolated from viable circulation routes, even if they are not themselves surveilled. [Figure 6.4]

Design for Structural Resilience

The solution is network redundancy: ensuring that every patch has multiple viable, low-stress connections to main circulation arteries. This prevents the psychological erasure of a single edge from causing the structural isolation of an entire area.

Architectural tactics include:

- **Multiple access points:** Avoid designs where a space has only one entrance or connection
- **Parallel circulation:** Create alternative routes between key nodes so that the closure (physical or psychological) of one path does not sever access
- **Distributed connectivity:** Ensure that minor spaces connect to the network at multiple points, not just through a single checkpoint

Redundancy as Resilience

This principle extends beyond physical circulation. Social networks, informal communication channels, and multiple modes of visibility all serve as redundancy mechanisms that prevent the collapse of optionality when a single route is compromised. Architecture cannot guarantee use, but it can prevent isolation.

6.1.5 Material Complexity (Disrupt Automation)

Contemporary surveillance increasingly relies on algorithmic systems—facial recognition, movement tracking, automated classification. Each layer of material complexity increases the cost and difficulty of comprehensive automated monitoring.

Disrupting Automation via Texture and Ambiguity

Complex textures, dense vegetation, reflective surfaces, and irregular geometries introduce noise into surveillance systems, increasing error rates and forcing reliance on slower, more expensive human labor.

Design tactics include:

- **Visual complexity:** Hedges, trees, planted screens, and textured facades that fragment sightlines and complicate image recognition
- **Reflective and refractive surfaces:** Glass, water features, polished materials that distort visual data
- **Temporal change:** Seasonal foliage, user modifications (graffiti, posters, temporary structures), and movable elements that prevent the environment from stabilizing into a known baseline

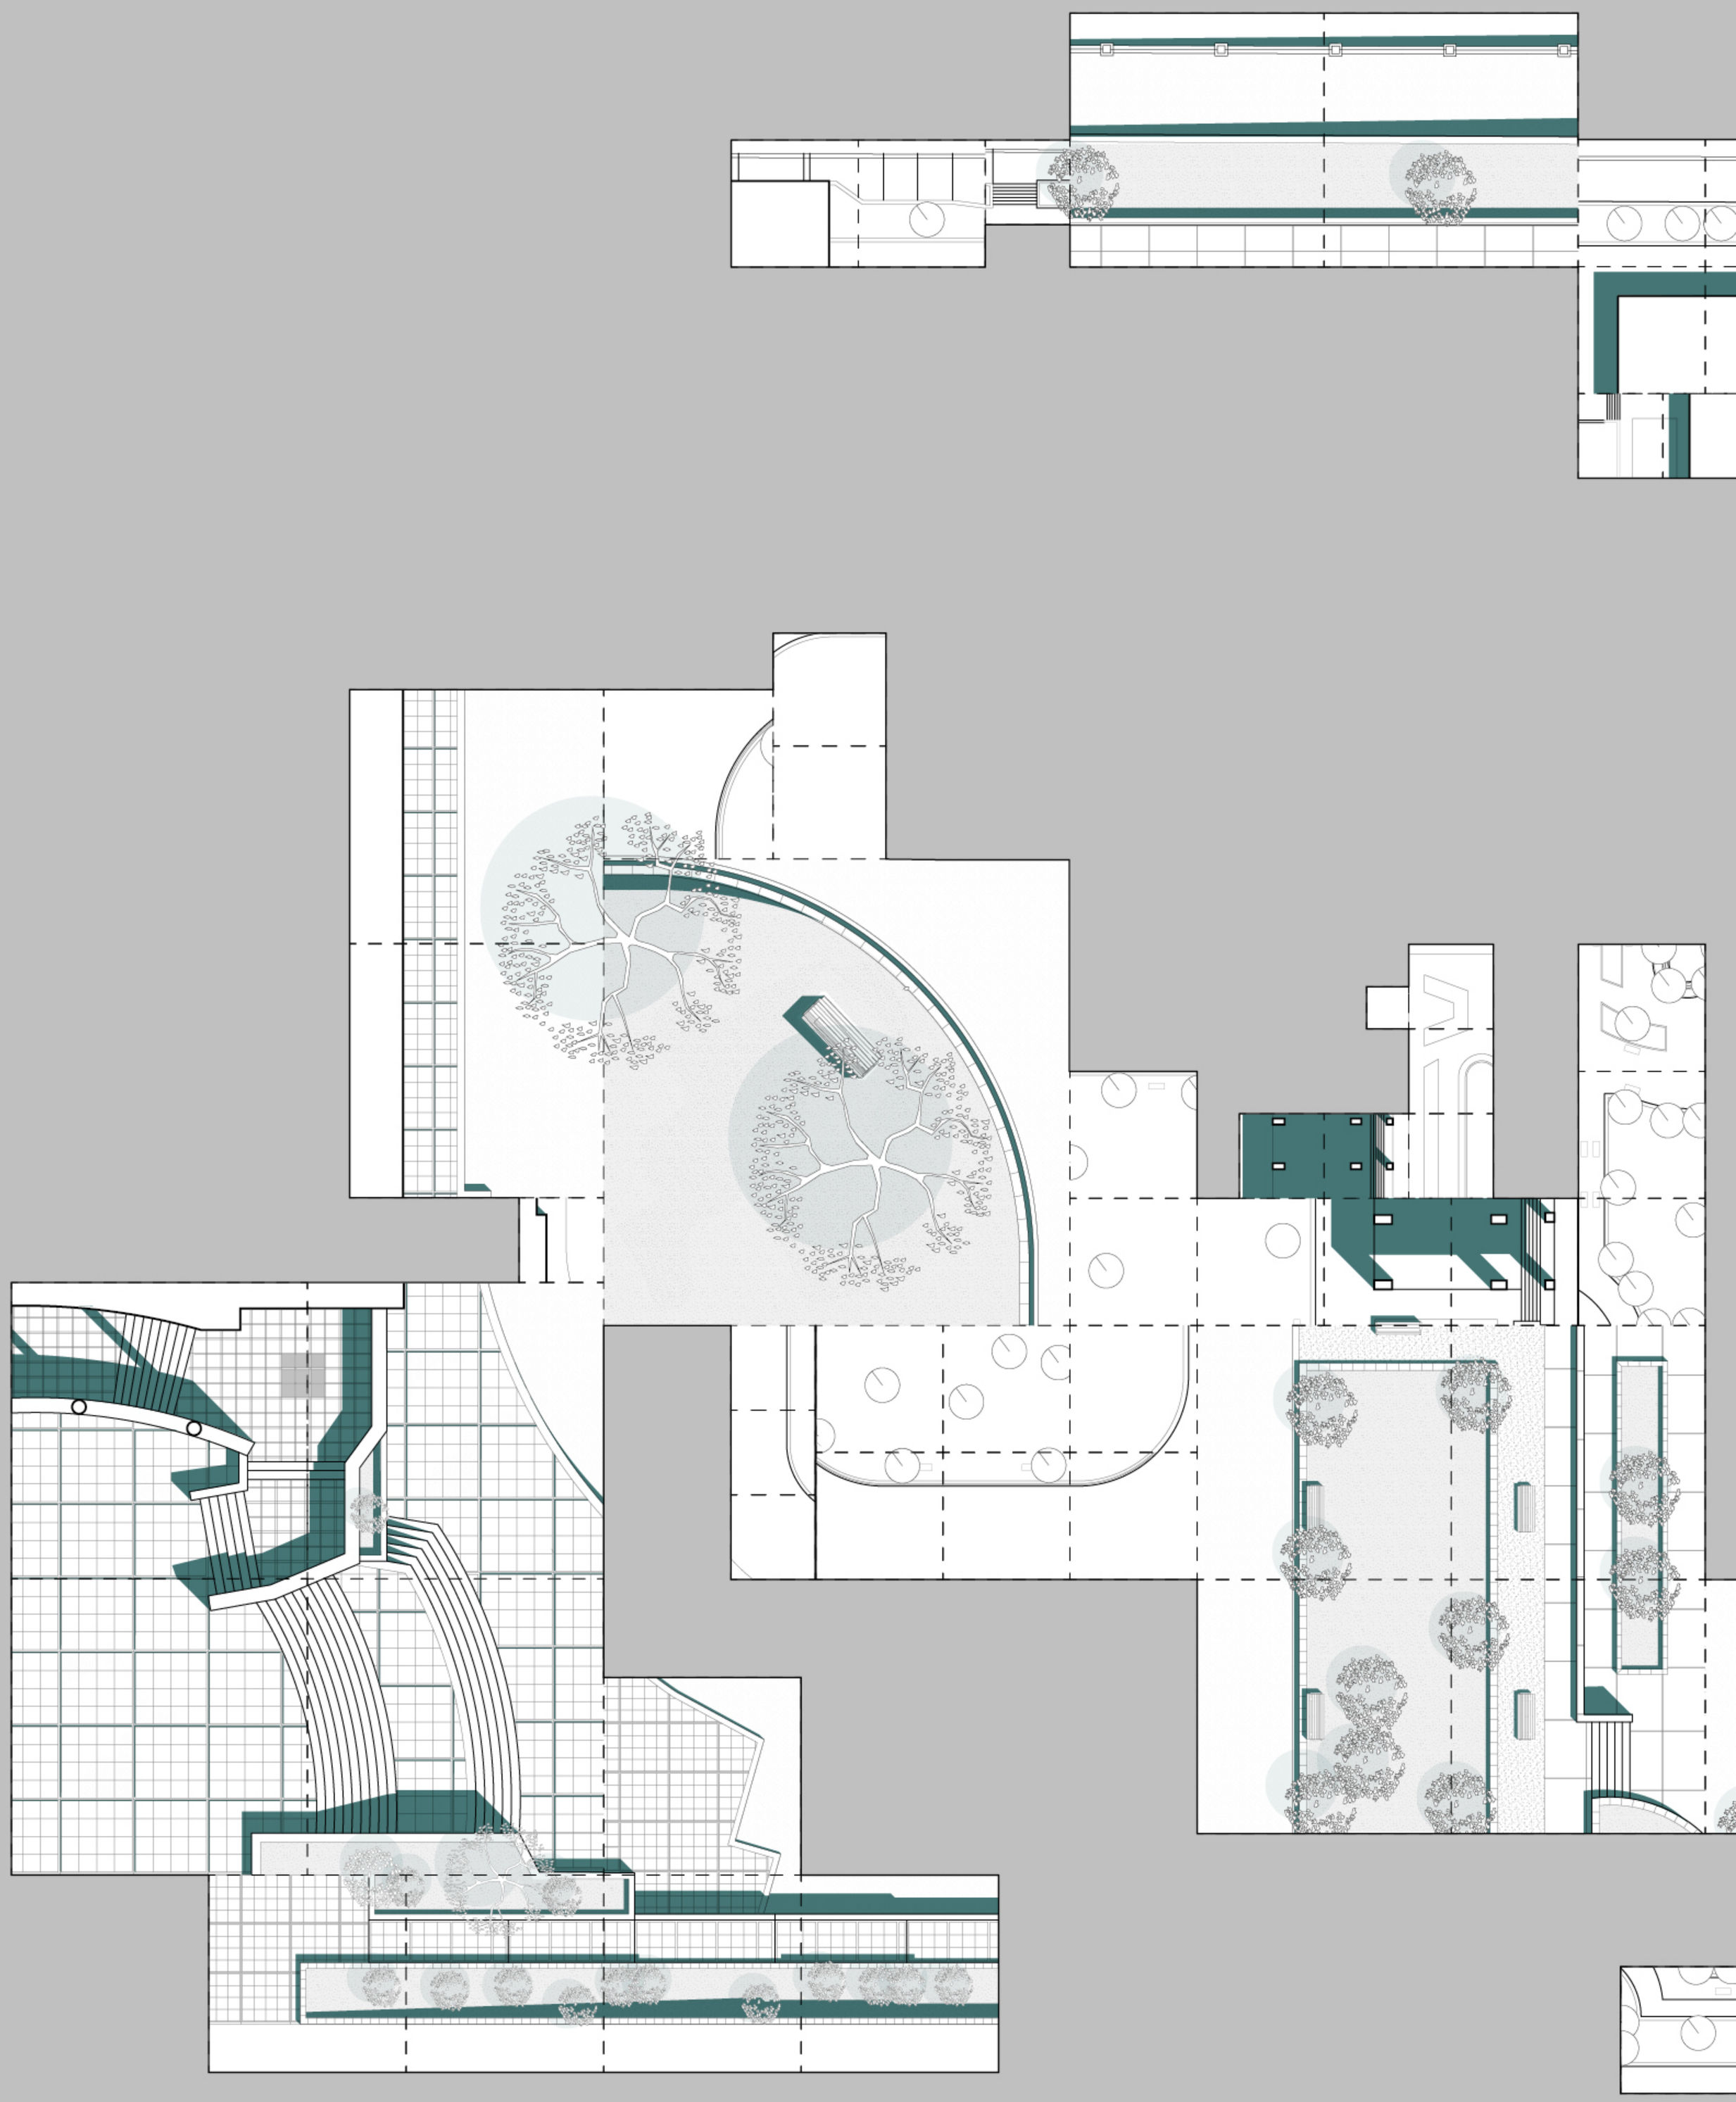

Figure 6.5 *The remembered landscape in varying scales based on intensity of detail memory. The visualization is based on the second scenario, mildly brave explorer [Author]*

| Micro-topography: Steps, levels, slopes, and grade changes that fragment sightlines and create multiple, overlapping visual planes

The Automation Tax

The goal is not to make surveillance impossible, but to make it expensive. Algorithmic systems rely on standardized, predictable environments. Introducing variability and complexity imposes an automation tax, surveillance remains possible, but at significantly higher cost in processing power, human oversight, and system maintenance.

In economic terms, this shifts the cost-benefit analysis away from total coverage and toward selective enforcement.

6.1.6 Revisiting the Corners

Another category of design strategies is revitalizing the forgotten corners, and re-designing the safe spots to create a space for micro-adaptive reuses. In other words, with the increase of familiarity (MD) and clustering, we can assume user's access at a more detailed, higher scale. While the main buildings are memorized only enough to act as landmarks for navigation, the safe spots become host of a vast variety of behaviors: sitting on the steps or the window sill, while using a niche in the wall as a table and the balcony above the head as a roof.

Figure 6.5 shows how this variation of resolution creates a gradient of detail and scale in the user's experience, when a bench in the central garden is twice the size of the main corridor. The concept can be expanded to detail design ideas, how a portion of facade can act as a basis for these user-centric micro adaptive reuses, enhancing the experience and optionality.

6.2 Synthesis: Architecture as Accumulation of Friction

None of these strategies, deployed in isolation, will prevent a determined state from establishing control. The point is not to make surveillance impossible, but to make it costly, visible, and partial.

The cumulative effect of these design decisions is friction: each strategy increases the labor, money, legitimacy burden, and visibility cost required to achieve comprehensive surveillance. Invert the gaze, and enforcement actions occur in public view. Zone efficiently, and most space remains undefined. Design for ambiguity, and control requires constant reclassification. Build in redundancy, and the network resists fragmentation. Add material complexity, and automation fails.

Together, these strategies do not eliminate the possibility of surveillance. They change its character—from ambient, comprehensive, and invisible to selective, effortful, and contested.

This is the architect's role: not to eliminate security, but to prevent the collapse of optionality that transforms space into a single-value optimization problem. The goal is to preserve the cognitive freedom to choose, to wander, to appropriate, and to resist, even, and especially, under conditions of observation.

Figure 6.6 *Detailed Plan of the Faculty of Fine Arts and identification of some of the proposed strategies used in its design: closed gardens and open corridor and placement of the entrances showcases efficient zoning while peripheral areas design and preserving the corners shows the "revisiting the corners", the difference of elevation levels, terraces and fences contributes to mutual visibility, and the placement of assets such as gardens, extra passages, and open gardens introduce functional ambiguity and network redundancy. Finally, the use of vegetation at intersections to break the sightline is a direct contribution to material complexity. [Author]*

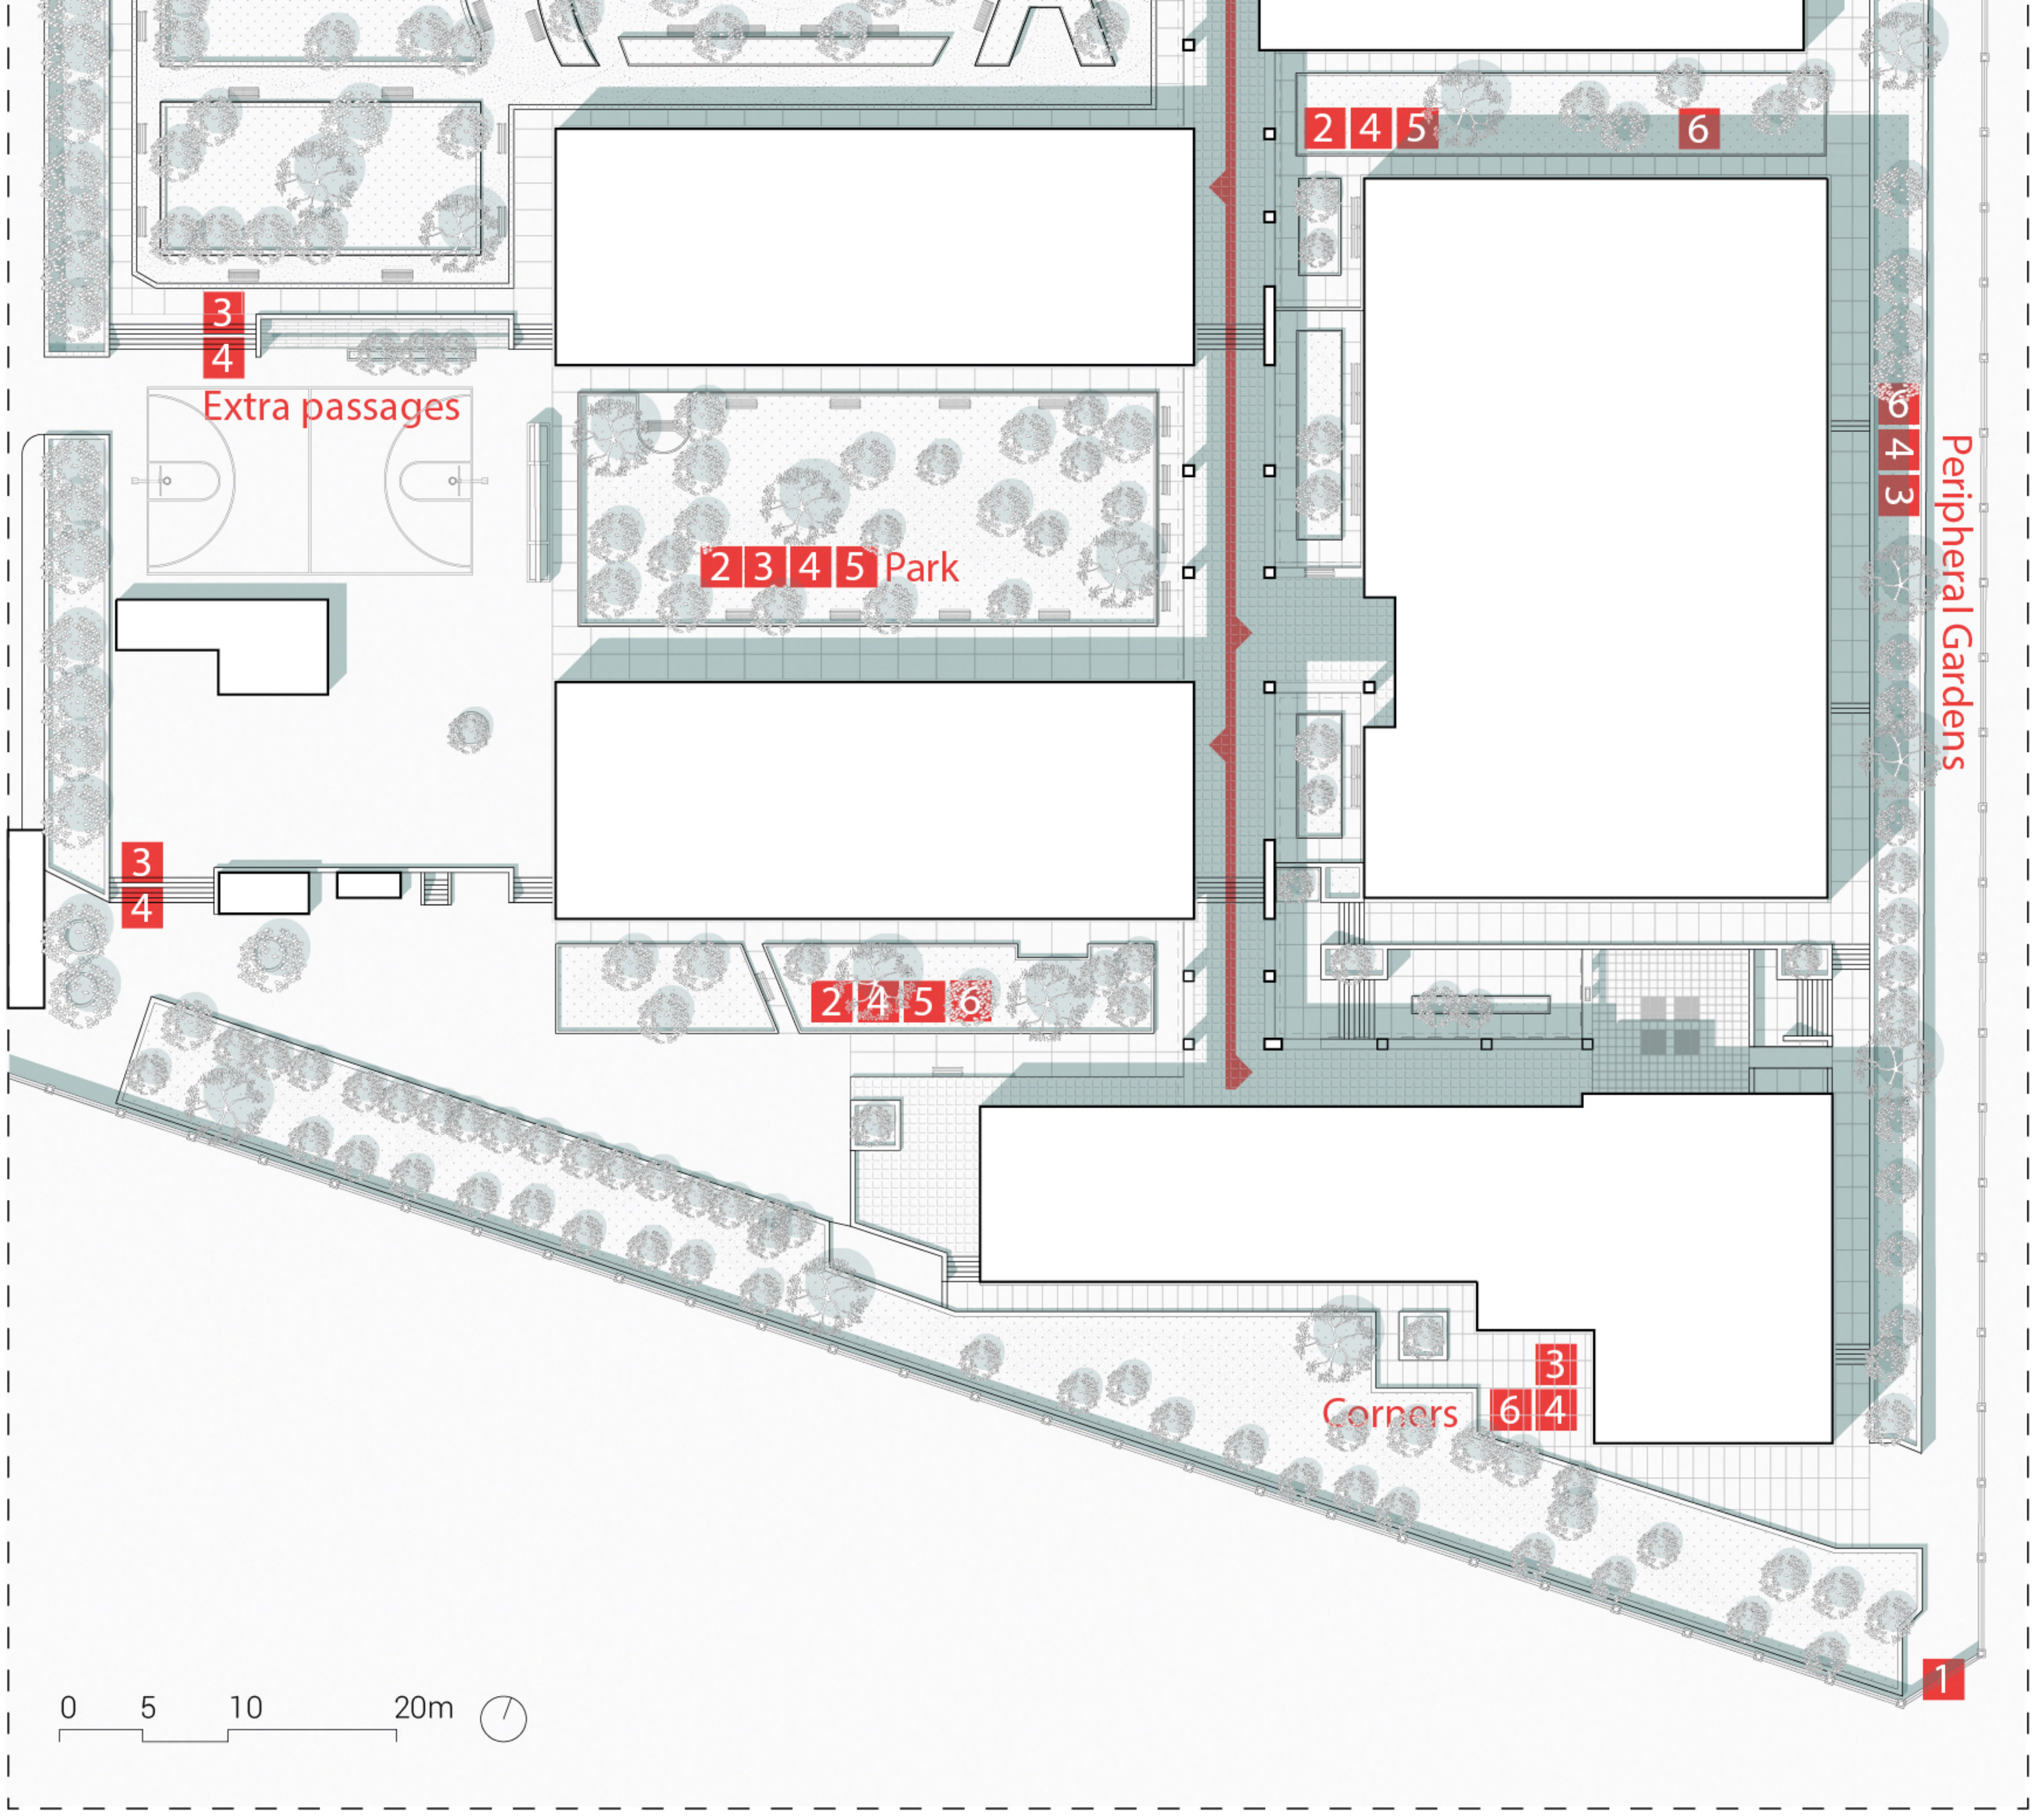
2 4 5
6
3
4
Extra passages
6 4 3
Peripheral Gardens
2 3 4 5 Park
3
4
2 4 5 6
3
Corners
6 4
1
0 5 10 20m

1 Mutual Visibility
2 Efficient Zoning
3 Functional Ambiguity
4 Network Redundancy
5 Material Complexity
6 Revisiting The Corners

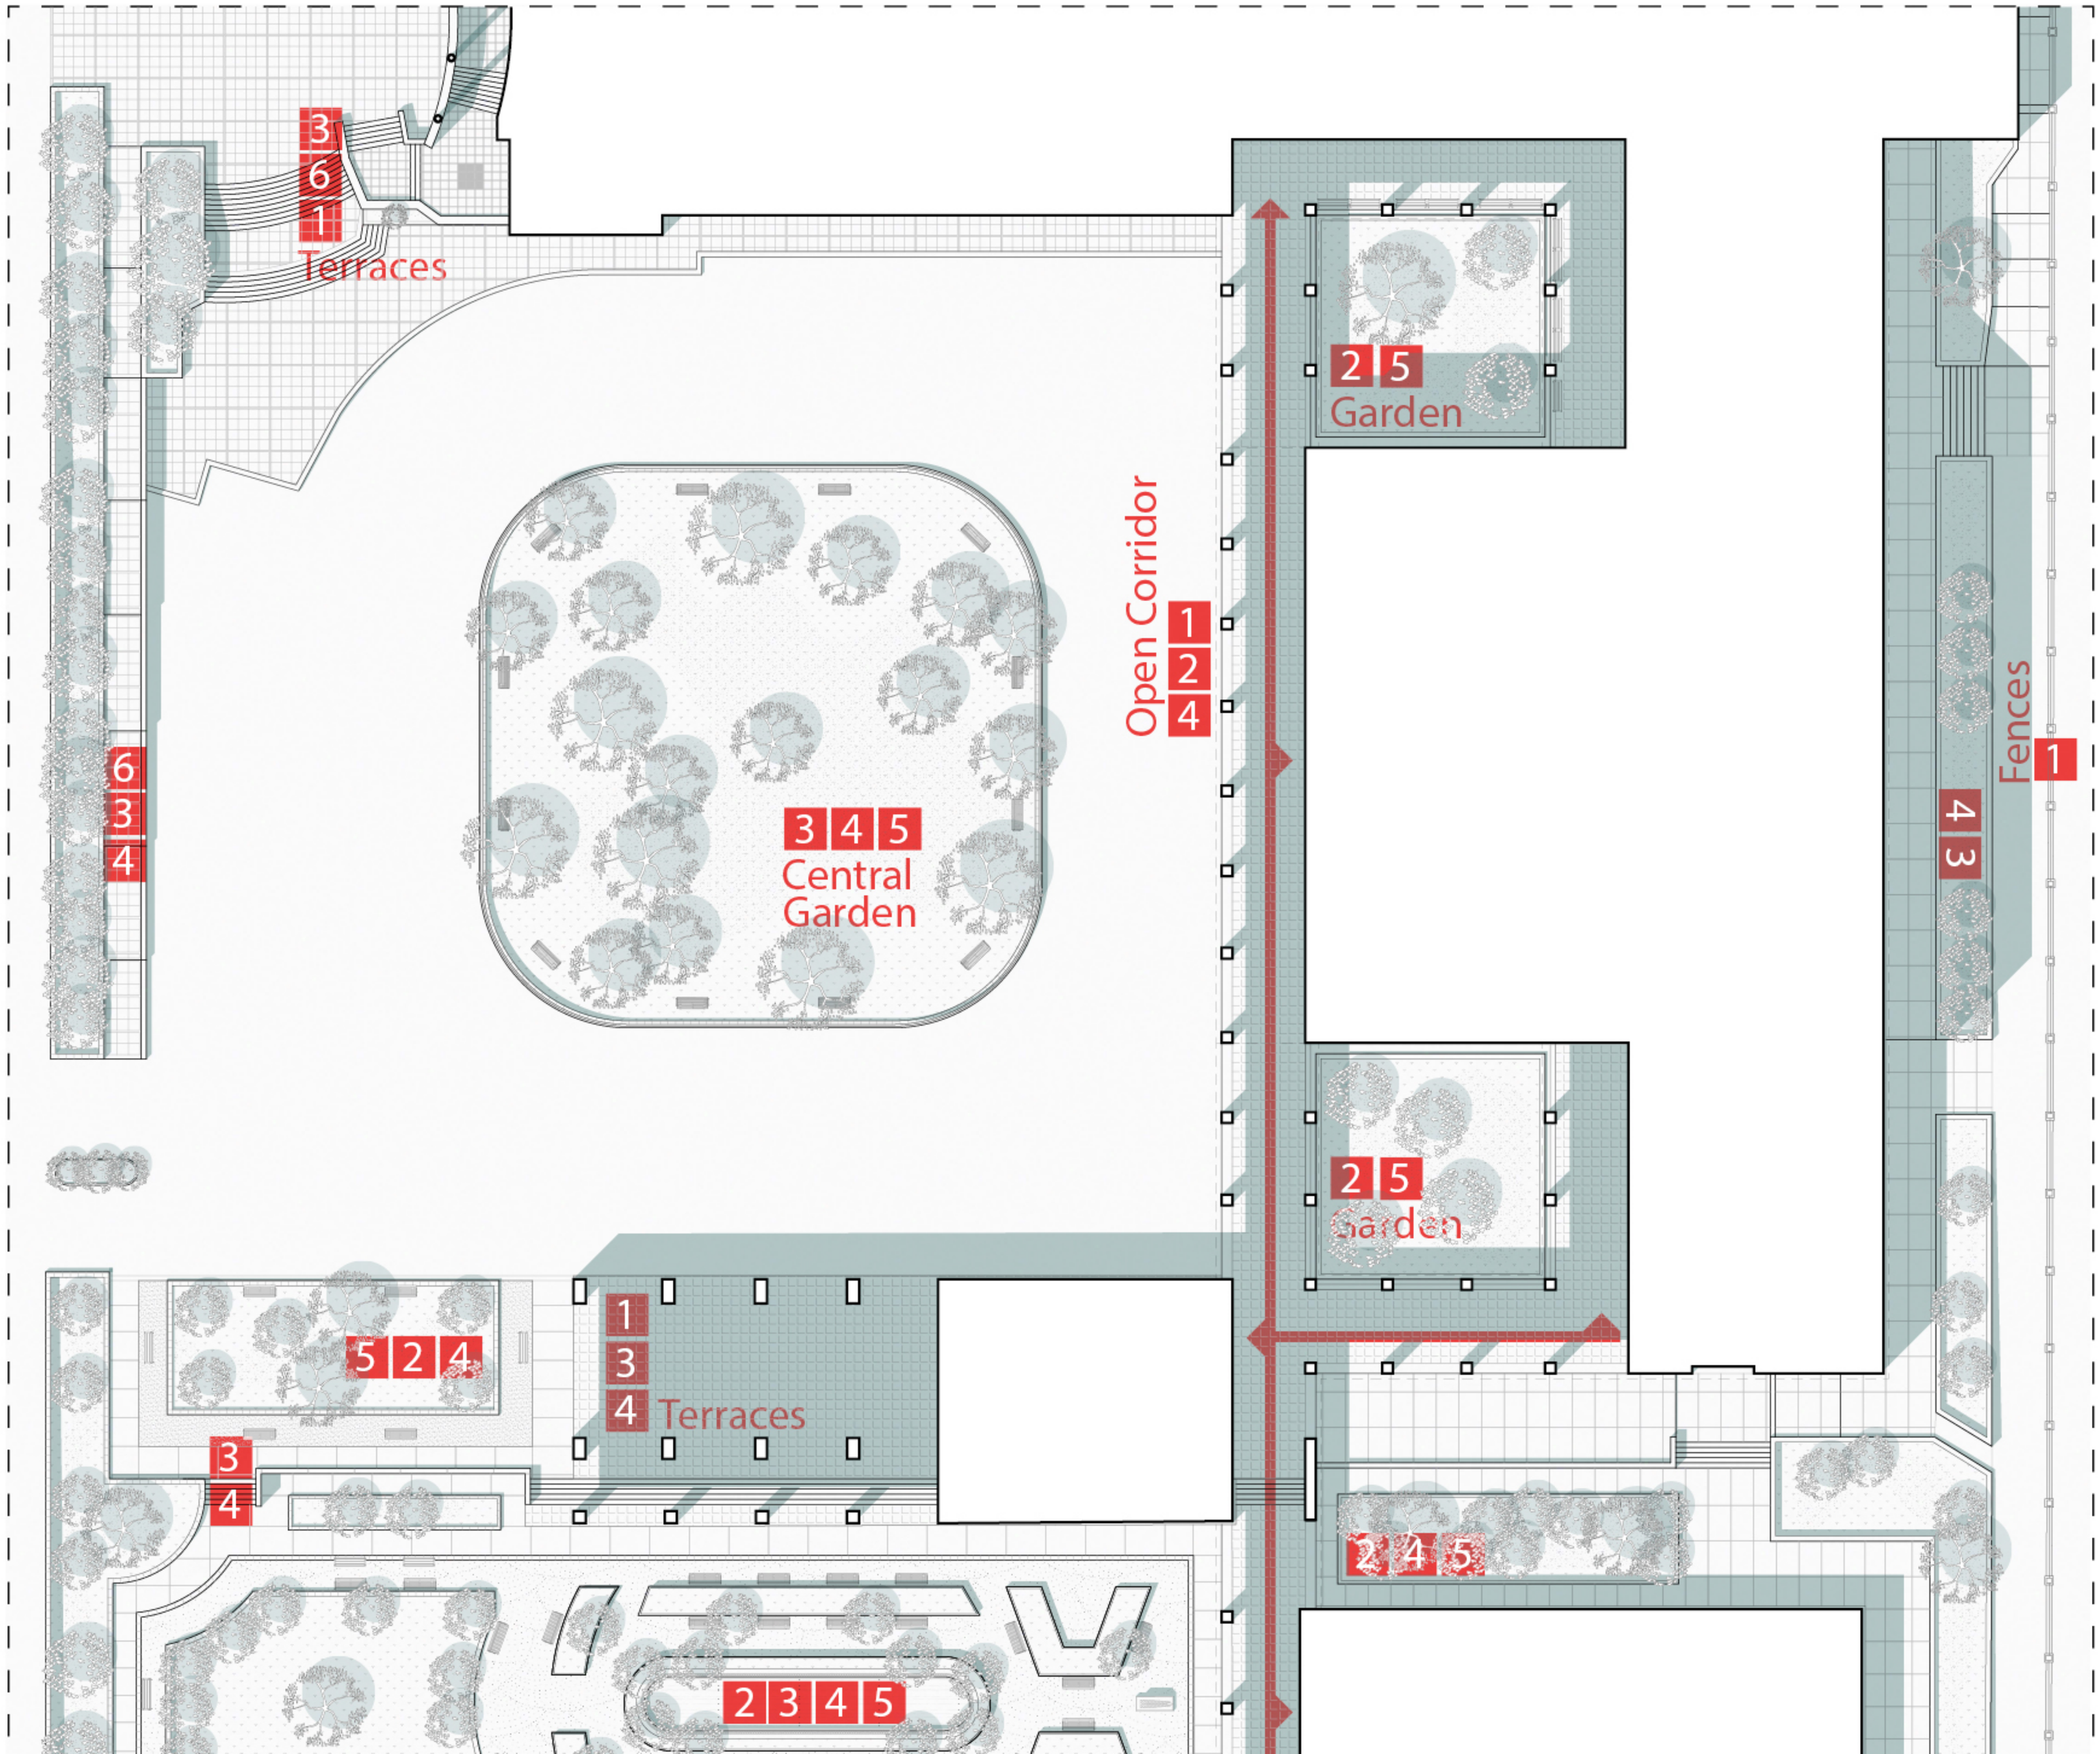

07 | Conclusion

The Fragmented Narrative
Summary of Findings

This thesis began with a question born from lived experience: "What is our city?" For women navigating Iran's surveilled urban landscape, the answer is not a matter of geometry or infrastructure, but of memory, fear, and the slow erasure of possibility. The city is not the physical grid of streets and buildings, but the fragmented cognitive map that remains after surveillance has removed choice, one path at a time.

Through the development and analysis of an agent-based model simulating navigation under systematic threat, this research has demonstrated that surveillance does not merely restrict movement—it fundamentally restructures spatial cognition. The accessible world collapses by at least 50% across all stress scenarios, not through physical barriers, but through the psychological erasure of viable options. The agent-based model, calibrated to the Faculty of Fine Arts at the University of Tehran, revealed patterns that align disturbingly well with the safe zones identified through auto-ethnographic mapping, confirming that the mechanisms of perception, memory, and learned avoidance captured in the simulation reflect real adaptive behavior.

7.1 The Mechanisms of Fragmentation

The research has identified three interconnected mechanisms through which surveillance fragments spatial experience:

1. The Tyranny of Value: Under persistent threat, the cognitive map does not simply shrink, it becomes dominated by a single dimension: avoidance. Navigation decisions that should balance multiple values (efficiency, social connection, aesthetic pleasure, curiosity) collapse into binary threat assessment. The environment becomes one-dimensional, and with it, the individual's capacity for choice.

The simulation quantified this collapse. Average degree increased by 127% in the mildly brave scenario, indicating extreme clustering and path repetition within a severely constrained territory. The agent continued moving, but all movement occurred within a residual safe zone, navigating the same familiar patches in repetitive circuits. This is not adaptation—it is cognitive imprisonment.

2. The Inaccessible Void: Perhaps the most troubling finding is the emergence of spaces lost not through direct threat, but through network fragmentation. The Inaccessible Void comprises patches that remain physically passable and psychologically acceptable (avoidance memory below threshold), yet become unreachable because the surrounding network edges have been erased.

In the avoidance threshold 0.4 scenario, 10,965 patches were pruned from the navigation graph, yet fewer than this number directly exceeded the avoidance threshold. The difference, thousands of patches, represents spaces rendered unusable not by surveillance itself, but by the structural contamination of the network. Surveillance spreads like poison through the topology of navigation, isolating spaces that were never themselves threatened.

This demonstrates that the loss of spatial access is not proportional to the extent of surveillance coverage. A relatively limited surveillance system can fragment the cognitive map far beyond its immediate reach, multiplying its effect through network dynamics.

3. The Breaking Point: From Constraint to Paralysis: The relationship between risk tolerance and spatial behavior proved non-linear. The most severe fragmentation did not occur in the most risk-averse agent, but in the mildly brave scenario—agents courageous enough to continue moving but cautious enough to avoid large portions of space. This produces the worst of both conditions: continued exposure to threat without the protective withdrawal of the most risk-averse agents, and severe territorial constraint without the expanded access of the brave.

The risk-avoider scenario, however, revealed a different form of collapse: cognitive paralysis. Total movements dropped to 67% of baseline, with the remaining 33% representing not detours or alternative routes, but freezing

and withdrawal. At this point, the agent loses not just territory, but the will to navigate at all.

Architecturally, this means that spaces under heavy surveillance may become effectively unusable even when technically accessible. The agent knows where it could go but cannot bring itself to move. The space exists, but the individual has disappeared from it.

7.2 From Simulation to Lived Reality

The alignment between simulation results and the reported safe zones from peer mapping exercises confirms that these are not merely computational artifacts, but mechanisms operating in lived spatial experience. The scattered clusters identified by the agent correspond to the actual gathering spots used by students—spaces chosen not for their functional superiority, but for their structural position within a fragmented network.

This suggests that, consciously or unconsciously, we were following the same navigational logic captured in the model. Surveillance reduced our cognitive maps to a single dimension: the level of threat. Individuality, mood, social desire, aesthetic preference—all subordinated to the calculus of avoidance. And as chunks of the urban environment were removed from individual memory, we disappeared from those parts of the city.

The costs of this disappearance extend beyond navigation. The simulation cannot capture the psychological weight of constant vigilance, the exhaustion of risk assessment performed thousands of times per day, the erosion of spontaneity and curiosity, or the slow internalization of alienation. These are the residues of the surveilled experience—unconscious impulses that persist even after leaving the environment, the feeling of one's body superseding one's individuality.

7.2.1 The Paradox of Constrained Inhabitation

Yet the conclusion cannot be one of complete erasure. While unjustly deprived of navigational options and urban space, the constrained inhabitation leads to clusters of patches with high detail-memory. As familiarity with remaining spaces increases, small-scale decisions appear: painting graffiti, sitting on window sills, appropriating neglected corners, redefining spaces through temporary installations and social practices.

Spatial oppression may paradoxically intensify certain forms of place-making within the residual territory. This does not compensate for the violation of the most basic rights, to body, to place, to expression, but it refuses the narrative of total defeat. The individual retains agency within constraint, and the cognitive map, though fragmented, remains a site of creativity and resistance.

This tension is central. The simulation demonstrates quantifiable collapse, but lived experience contains forms of adaptation that cannot be reduced to graph coverage statistics. The question is not whether individuals survive, clearly, we do, but what survival costs, what is lost in the process, and what forms of life remain possible within the residual spaces.

7.2.2 Architecture as Enabler of Choice

The findings of this research point toward a fundamental re-framing of architectural responsibility. The architect and urban designer function as enablers of choice, and every design decision—every wall, every sightline, every circulation pattern—creates and removes options from the possibility space.

The design implications outlined in Chapter 6 propose strategies to increase friction: the cost in labor, money, legitimacy, and visibility required to achieve total surveillance. These strategies, mutual visibility, efficient zoning, functional ambiguity, network redundancy, material complexity, do not eliminate the possibility of surveillance. They change its character, making control costly, visible, and partial rather than ambient, comprehensive, and invisible.

The goal is to defend cognitive optionality: ensuring that individuals retain the capacity to make choices based on multiple, competing values rather than fear alone. This is not naïve resistance. It accepts that security exists while refusing to make oppressive control *easy*. It acknowledges that architecture cannot prevent a determined state power from establishing surveillance, but it can prevent the designer from becoming complicit in making that surveillance effortless.

The Faculty of Fine Arts serves as an example. Its scattered pavilion structure, porous boundaries, and distributed circulation create a spatial environment where comprehensive surveillance requires disproportionate effort. There is no single panopticon point. Control remains possible, but it must be actively maintained across multiple, non-overlapping zones—a task that is economically and logistically prohibitive compared to a centralized, legible plan.

This is the lesson: architecture can resist not by denying the reality of surveillance, but by refusing to optimize for it.

7.3 Limitations and Future Directions

This research represents an initial exploration, and its limitations suggest pathways for future work: The simulation captures individual navigation under threat, but excludes collective dynamics and social learning. Future iterations should model multiple agents, social communication networks, and the ways groups coordinate to navigate hostile environments.

The threat patterns are static and pre-defined, while real surveillance adapts

and learns. Introducing adaptive threat agents that respond to user behavior would model escalation dynamics and the arms race between surveillance and resistance.

The case study is bounded—a single institutional space in a specific political context. Comparative studies across different environments (urban public space, commercial districts, residential neighborhoods) and different populations (varying by gender, class, ethnicity, political affiliation) would test the generalizability of the fragmentation mechanisms.

The temporal scope is limited to 2016-2021, and conditions have intensified significantly since. Longitudinal studies tracking the evolution of spatial memory under changing surveillance intensity would capture adaptation over time.

Finally, the framework developed here—understanding environment as a terrain of uncertainty storing ever-changing information relative to foci of attention—is expandable beyond surveillance contexts. The mechanisms of perception, memory decay, and value-driven navigation apply to non-human agents, robot sensing, autonomous systems, and any situation where navigation occurs under conditions of incomplete information and dynamic risk.

7.4 The Right to the City

The right to the city, as Henri Lefebvre conceived it, is not merely the right to access urban space, but the right to participate in its production, to inhabit it freely, to appropriate and transform it through everyday practice. Surveillance systematically denies this right by fragmenting the cognitive map, erasing options, and reducing navigation to threat-avoidance.

The women navigating Iran's surveilled landscape do not have the privilege of unconscious movement, of wandering without calculation, of experiencing the city as anything other than a terrain of threat. The cognitive map becomes a record of erasure, spaces avoided, paths abandoned, territories rendered inaccessible not through physical barriers, but through accumulated fear. It leads not only to an individual deprivation of the right to the city and impediment of intergenerational transmission of spatial practices, but as an erasure of shared place-based memory that constitutes intangible heritage.

This thesis has attempted to make this erasure visible, to quantify the collapse of optionality, and to trace the mechanisms through which surveillance transforms accessible space into cognitive void. It has demonstrated that the loss is not merely personal or psychological, but structural—a fundamental reorganization of spatial possibility that operates through the topology of networks and the dynamics of memory.

But visibility alone is not enough. The design implications proposed here outline strategies for resistance, but their effectiveness depends on political will, social solidarity, and the collective refusal to accept fragmentation as inevitable. Architecture can increase friction, but it cannot replace the broader struggle for bodily autonomy, political freedom, and the right to exist without surveillance.

The city we deserve is not the fragmented terrain of avoidance, but a space where navigation is guided by curiosity, connection, pleasure, and purpose. Where the cognitive map expands through exploration rather than contracting through fear. Where the individual does not disappear from space, but claims it.

This is the city we fight for. This research is one tool in that fight—a demonstration that what appears natural and inevitable is, in fact, constructed and contestable. The fragmentation of spatial memory under surveillance is not a fixed condition, but a dynamic process that can be resisted, interrupted, and transformed.

Figure 7.1 *A young man breaking the security cameras used to identify and eventual arrest of protesters, during the national protests of Jan. 2026, Image enhanced slightly using Google Gemini to increase the resolution by the author [VahidOnline, Jan. 2026]*

0B | Bibliography

0A|Bibliography

1. Ingold, Tim. "The Temporality of the Landscape." World Archaeology 25, no. 2 (1993): 152–74. http://www.jstor.org/stable/124811

2. Latour, Bruno. Reassembling the Social: An Introduction to Actor-Network-Theory. New York: Oxford University Press, 2005.

3. Mallot, Hanspeter A. From Geometry to Behavior: An Introduction to Spatial Cognition. Cambridge, MA: MIT Press, 2024.https://doi.org/10.7551/mitpress/9621.001.0001.

4. Tolman, Edward C. "Cognitive Maps in Rats and Men." Psychological Review 55, no. 4 (1948): 189–208. https://doi.org/10.1037/h0061626.

5. Gibson, James J. The Ecological Approach to Visual Perception. Boston: Houghton, Mifflin and Company, 1979.

6. Wilensky, Uri. NetLogo. Evanston, IL: Center for Connected Learning and Computer-Based Modeling, Northwestern University, 1999. http://ccl.northwestern.edu/netlogo/

7. Yang, Q., R. Dubey, and S. Kalantari. "An Evidence-based Cognitive Model of Uncertainty during Indoor Multi-level Human Wayfinding." bioRxiv (2023). https://doi.org/10.1101/2022.07.27.501728

8. Staddon, J. E. R., I. M. Chelaru, and J. J. Higa. "Habituation, Memory and the Brain: The Dynamics of Interval Timing." Behavioural Processes 57, nos. 2–3 (2002): 71–88. https://doi.org/10.1016/S0376-6357(02)00006-2 .

9. Sohn, Samuel, Serena DeStefani, and Mubbasir Kapadia. "Dynamic Cognitive Maps for Agent Landmark Navigation in Unseen Environments." In Motion, Interaction and Games (MIG 2018), 1–6. https://doi.org/10.1145/3274247.3274518.

10. Grimm, Volker, Railsback, Steven F., Vincenot, Christian E., Berger, Uta, Gallagher, Cara, DeAngelis, Donald L., Edmonds, Bruce, Ge, Jiaqi, Giske, Jarl, Groeneveld, Jürgen, Johnston, Alice S.A., Milles, Alexander, Nabe-Nielsen, Jacob, Polhill, J. Gareth, Radchuk, Viktoriia, Rohwäder, Marie-Sophie, Stillman, Richard A., Thiele, Jan C. and Ayllón, Daniel (2020) 'The ODD Protocol for Describing Agent-Based and Other Simulation Models: A Second Update to Improve Clarity, Replication, and Structural Realism' Journal of Artificial Societies and Social Simulation 23 (2) 7 <http://jasss.soc.surrey.ac.uk/23/2/7.html>. doi: 10.18564/jasss.4259

11. Mojab, Shahrzad. "State-University Power Struggle at Times of Revolution and War in Iran." International Higher Education 36 (2004): 11–13.

12. Foucault, Michel. Discipline and Punish: The Birth of the Prison. New York: Pantheon Books, 1977.

13. Rivetti, Paola. "What Happens In-Between Mobilizations? Building and Organizing Contentious Politics at the University of Tehran (2007–2017)." Partecipazione e Conflitto 13, no. 1 (2020): 586–606. ISSN 2035-6609.

14. El-Hage, Toni. 2002. "Religio-Patriarchy and the Gendered Risk : The Regulation of Iranian Femininity in Public Spaces through the Veil." Retrospective Theses and Dissertations, 1919-2007. T, University of British Columbia. doi:http://dx.doi.org/10.14288/1.0090401.

15. Zarbighalehhammami, Shakib, and Gamze Aksan. "Analysis of Hijab Ban Practices in Iran after 2022 in Terms of Panopticon Theory: A Case Study." Selçuk Üniversitesi Hukuk Fakültesi Dergisi 33 (2025): 745–775. https://doi.org/10.15337/suhfd.1642295.

16. Abdmolaei, Shirin. "(Re)Fashioning Resistance: Women, Dress and Sexuality in Iran." Anthropology of the Middle East 9 (2014): 38–55. https://doi.org/10.3167/ame.2014.090204

17. Archive of maps of Tehran, Municipality of Tehran.

18. University of Tehran. Accessed 15/09/2025. https://ut.ac.ir/en/

19. CAOI (Architecture of Iran). Accessed 15/09/2025. https://www.caoi.ir/en/projects/item/1123

20. Mehrabi, Ehsan. "From Sarallah and Jundallah to the Morality Police: The Means of 40 Years of Mandatory Hijab Enforcement in Iran." BBC Persian, 2022. https://www.bbc.com/persian/articles/c728x75gpz0o

21. Radio Farda. "Iran Bill on Hijab Law." Accessed 21/09/2025. https://www.radiofarda.com/a/iran-bill-on-hijab-law/32523675.html

22. Lefebvre, Henri. The Right to the City. In Writings on Cities, selected, translated, and introduced by Eleonore Kofman and Elizabeth Lebas, 63–184. Oxford: Blackwell, 1996.

0G | Glossary

Avoidance Memory (Ma)
Long-term memory layer storing learned negative utility from threat exposure. The agent may forget visual details but retain avoidance.

Avoidance Threshold (Ta)
Maximum perceived stress an agent tolerates before rejecting a space. The decisive variable determining fragmentation severity.

Breaking Point
Moment when the learning curve flattens and safe areas become fixed. Marks transition from exploration to constrained circulation.

Cognitive Map
Flexible memory structure built from associations between states, actions, and outcomes. Represents knowledge of possibilities, not physical layout.

Cognitive Paralysis
Extreme constraint where agents freeze rather than navigate. In risk-avoider scenarios, 33% of movements became withdrawal.

Detail Memory (Md)
Long-term memory layer storing neutral spatial knowledge—visual arrangement and geometric configuration.

Discretionary Spatial Behavior
Wandering without prescribed destination. Navigation type studied to isolate stress-driven choice from obligatory movement.

Foci of Uncertainty
Active threatening elements within the ToU. In this thesis: surveillance cameras and patrols propagating stress.

Friction
Cumulative cost in labor, money, legitimacy, and visibility required to achieve comprehensive surveillance. Architecture increases friction to prevent effortless control.

Graph Coverage
Percentage of explored area relative to total accessible area. Primary metric for fragmentation—collapsed 50-75% under surveillance.

Inaccessible Void

Spaces physically passable and psychologically acceptable but structurally unreachable because surrounding edges are erased. Surveillance contaminates beyond its immediate reach.

Peripersonal Space / Perception Zone

Immediate surrounding of the agent—an active representation continuously updated with movement. Operationalized as limited, heading-dependent information intake.

Psychological Erasure

Physically accessible space becoming cognitively unavailable through accumulated threat. Removal of viable options from the cognitive map when avoidance exceeds threshold.

Reinforcement Learning (RL)

Assignment of value to state-action pairs through rewards/penalties. Under surveillance, persistent threat causes psychological erasure of exposed paths.

Stress Propagation

Mechanism assigning negative utility to space. Real Stress (Sr) is objective; Perceived Stress (Sp) is subjective, influenced by visibility and memory.

Terrain of Uncertainty (ToU)

Complete experiential environment integrating physical space, memory, perception, and active agents. Possesses agency through contained elements, shaping choice and constraining movement.

Tyranny of Value

Under persistent surveillance, threat-avoidance dominates all other navigation criteria (efficiency, pleasure, curiosity). The environment becomes one-dimensional, collapsing choice.

0A | Appendices

Appendix A
ODD Protocol

Navigation Under Surveillance Agent-Based Model

Model Name	*Spatial Cognition Fragmentation Under Surveillance*
Version	*1.0*
Authors	*Kosar Mohammadi*
Institution	*Politecnico di Torino*
Supervisor	*Prof. Tomà Berlanda*
Last Updated	*Jan 19 2026*
Protocol Compliance	*ODD+D (Grimm et al. 2020)*

A.1 PURPOSE AND PATTERNS

A .1.1 Purpose

This model simulates the fragmentation of spatial cognition under systematic surveillance. It operationalizes how persistent threat exposure psychologically erases physically accessible paths from an agent's cognitive map, quantifying the collapse of navigational optionality in controlled environments.

Central Research Question: How does surveillance fragment spatial perception and memory, and how does viable navigation collapse?

The model tests whether stress-driven reinforcement learning, combined with dual-layer memory decay, can reproduce the safe-zone clustering patterns observed in auto-ethnographic mapping of the Faculty of Fine Arts, University of Tehran (2016-2021).

A.1.2 Patterns to Reproduce

The model is calibrated against peer-mapped safe zones from the case study. Success criteria include:

1. Spatial clustering that matches reported gathering locations
2. 50-75% reduction in graph coverage under surveillance (hypothesis-driven)
3. Emergence of "breaking point"where exploration curve flattens
4. Identification of structurally isolated "Inaccessible Void" zones

A.2. ENTITIES, STATE VARIABLES, AND SCALES

A.2.1 Agents

Protagonist (Rabbit)

```
rabbits-own[
    p-error                 ; Perception error rate [0-0.1]
    goal-patch              ; Current navigation target
    stuck-counter           ; Paralysis detection
    last-position           ; For movement validation
    stuck-threshold         ; Paralysis trigger (25 ticks)
]
```

Key Individual Parameters:

Avoidance Threshold (TA): Individual risk tolerance [0.2-0.9], the critical parameter varied across scenarios

Perception Zone: 30m radius, 120° field-of-view cone

Error Rate (Pe): ±10% perception noise

Threat Agents (Patrols & Cameras)

```
patrols-own[                ; Stress propagation strength [0-1]
    ps-intensity            ; Coverage radius (30m)
    ps-radius               ; Field of view (180°)
    ps-fov                  ; Movement history
    last-direction
]
cameras-own[
    cs-intensity            ; Fixed at 1.0
    cs-radius               ; Coverage radius (17.5m)
]
```

Agent Characteristics:

Patrols: Mobile, follow pre-defined red-path routes, heading-dependent perception cone

Cameras: Static, 15 fixed positions, 360° coverage

A.2.2 Spatial Environment (Patches)

Graph Structure

```
patches-own[
    ; Physical Properties
    passable?                   ; Boolean: navigation graph membership
    see-through?                ; Boolean: visual occlusion (walls vs. fences)
    basecolor                   ; Environment encoding (white/black/gray/red)

    ; Memory System
    detail-memory               ; Spatial knowledge MD  [0,1]
    avoidance-memory            ; Threat association MA   [0,1]
    time-last                   ; Ticks since last perception
    visit-frequency             ; Cumulative visual encounters

    ; Real-Time Perception
    r-stress                    ; Objective stress SR   [0,1]
    p-stress                    ; Perceived stress SP   [0,1]
    knowledge-degree             ; Distance-based fidelity K   [0,1]

    ; Navigation Graph
    nav-edges                   ; List of connected patches
    edge-weights                ; Usage frequency per edge
    presence-f                  ; Movement count (not just visual)
    pruned?                     ; Boolean: excluded from viable network
]
```

Memory Interpretation:

detail-memory: Latent learning, stores neutral spatial knowledge

avoidance-memory: Reinforcement learning, stores accumulated threat association

A.2.3 Scales

Spatial: 117 × 210 patch grid (1m resolution) = 14,806 passable patches

Temporal: 10,000 ticks per simulation (≈1 semester of spatial learning)

Perception Update: Every tick

Memory Export: Every 100 ticks

Graph Pruning: Every 25 ticks

A.3 PROCESS OVERVIEW AND SCHEDULING

A.3.1 Execution Sequence (per tick)

SETUP (once):

1. Define map from color-coded image
2. Initialize patch variables

3. Create agents (1 rabbit, 2 patrols, 15 cameras)
4. Pre-calculate camera stress fields
5. Build initial navigation graph

GO (repeat 10,000 times):

1. update-stress-fields
2. update-perception
3. update-memory
4. move-patrols
5. move-rabbit
6. [every 25 ticks] prune-navigation-graph
7. [every 5 ticks] log trajectory
8. [every 100 ticks] export memory timeseries
9. tick

A.3.2 Order Justification

Stress propagation precedes perception: Environment state must stabilize before sensing

Memory update follows perception: New information integrated before decision-making

- **Agent movement occurs after environmental updates:** Reactive navigation model

- **Graph pruning is periodic:** Computational efficiency; analog to cognitive map consolidation

A.4 DESIGN CONCEPTS

A.4.1 Theoretical Foundation

Spatial Cognition Framework (Tolman 1948, Mallot 2024):

Cognitive Map: Flexible, declarative state-action network enabling multi-goal navigation

Latent Learning: Non-reward-driven exploration stores potential routes (Detail Memory)

Reinforcement Learning: Negative incentive (threat) assigns value to state-action pairs (Avoidance Memory)

A.4.2 Emergence

Primary Emergent Phenomena:

1. **Safe-Zone Clustering:** High-frequency nodes emerge without explicit goal assignment

2. Inaccessible Void: Patches become unreachable through network topology collapse, not direct avoidance
3. Breaking Point: Non-linear threshold where exploration ceases (TA-dependent)

A.4.3 Objectives

Protagonist seeks to maximize Attractiveness while wandering (no fixed destination):

$$A(i) = w_base \times MD(i) - w_avoid \times MA(i) - w_perc \times SP(i) + 1/(visit\text{-}frequency + 1)$$

; If MA(i) > TA, A(i) = -9999 (psychological erasure)
w_base = -0.2 (negative to encourage exploration over familiarity)
w_avoid = 1.0 (primary constraint)
w_perc = 1.0 (immediate threat response)

Tyranny of Value Hypothesis: As MA increases, the avoidance term dominates, causing other criteria (familiarity, curiosity) to collapse once MA > TA.

A.4.4 Learning

Memory Acquisition (Reinforcement + Latent)

; When protagonist occupies patch i with knowledge-degree K:

$$MD(i) = MD(i) + \alpha_detail \times K(i)^6$$
$$MA(i) = MA(i) + \alpha_avoid \times SP(i) \times K(i)$$

α_detail = 0.05 (slow latent learning)
α_avoid = 0.1 (fast threat conditioning)

Memory Decay (Power-law forgetting)

$$M = M0 \times (1 - (time\text{-}last)^{(-cd)})$$

Decay constants:
cd_detail = 6-8 (faster decay, requires rehearsal)
cd_avoid = 8-10 (slower decay, persistent fear)

Justification: Avoidance memory persists longer than neutral spatial detail, reflecting traumatic memory retention patterns (Staddon et al. 2002).

A.4.5 Sensing

Perception Zone (Peripersonal Space)

1. Line-of-Sight Occlusion

```
to-report los-clear? [x0 y0 x1 y1]
    let steps max(abs(x1-x0), abs(y1-y0))
    let sx (x1-x0)/steps
    let sy (y1-y0)/steps
```

```
    repeat steps [
        if not [see-through?] of patch(x, y) [ report false ]
    ]
    report true
end
```

2. Knowledge Fidelity

$$K(i) = (radius - distance) / radius$$

$$SP(i) = SR(i) \times K(i) \times (1 + random\text{-}float(2 \times Pe) - Pe)$$

; Pe = 0.1 (±10% perception noise)

A.4.6 Interaction

Direct Interaction: None (single-agent simulation)

Indirect Interaction:

Threat agents modify patch `r-stress`

Protagonist perceives `p-stress`

Updates `avoidance-memory`

Constrains future navigation

Navigation creates `nav-edges` and `edge-weights` (emergent graph)

A.4.7 Stochasticity

Perception Error	*SR(i) × K(i) × (1 ± random 10%)*
Movement Micro-Variation	*rt random 10 - 5` (±5° heading noise)*
Exploration Probability	*10% chance to select least-visited patch instead of max-attractiveness*

A.4.8 Observation

Spatial Outputs

Heatmaps	*detail-memory, avoidance-memory, presence-f, visit-frequency*
Navigation graph	*Node size based on memory, edge thickness based on weight*
Pruning visualization	*Red (avoided), blue (active), gray (inaccessible void)*

Quantitative Metrics

Collected every 100 ticks:	*Graph coverage = active-nodes / total-passable* *Total movements (paralysis indicator)* *Average degree = edges-used / nodes-navigated* *Nodes navigated vs. avoided vs. never-visited*

A.5 INITIALIZATION

A.5.1 Environment Setup

```
to define-map
  ; Load base map (117×210 grid)
  ; Color encoding:
  ;   white, passable? = true, see-through? = true
  ;   black, passable? = false, see-through? = false (buildings)
  ;   gray, passable? = false, see-through? = true (fences)
  ;   red, patrol path markers
  ask patches [
    set visit-frequency 0
    set time-last 0
    set detail-memory 0
    set avoidance-memory 0
    set r-stress 0
    set p-stress 0
    set nav-edges []
    set edge-weights []
    set pruned? false
  ]
end
```

A.5.2 Agent Placement

Protagonist: Random entrance (4 options: campus gates)

```
entrances: [[3, -2], [116, -2], [2, -68], [2, -125]]
```

Cameras (15 fixed positions):

```
[[89,-167], [70,-155], [30,-149], [70,-123], [30,-115],
 [70,-87], [51,-76], [91,-69], [70,-60], [70,-42],
 [70,-25], [91,-25], [91,-8], [70,-8], [30,-14]]
```

Patrols: Start on red-path patches, heading = 0

A.5.3 Initial State

All memory values = 0 (tabula rasa). Camera stress fields pre-calculated once during setup to avoid redundant computation.

A.6 INPUT DATA

Spatial Data Source: Simplified site plan of Faculty of Fine Arts, University of Tehran:

Buildings, fences, obstacles reduced to primary occluders

Patrol routes traced from observational mapping (2016-2021)

Camera positions verified through peer-mapping and photographic evidence

No time-series inputs (stateless environment).

Color-Coded Encoding:

White pixels : Passable, see-through

Black pixels : Buildings (non-passable, non-see-through)

Gray pixels : Fences (non-passable, see-through)

Red pixels : Patrol paths

A.7 SUBMODELS

A.7.1 Stress Propagation

Camera Stress (static, pre-computed)

```
ask cameras [
  ask patches in-cone cs-radius 360 [
    if los-clear?(camera-x, camera-y, patch-x, patch-y) [
      let threat-stress = (radius - distance) × intensity / radius
      set r-stress max(r-stress, threat-stress)
    ]
  ]
]
```

Patrol Stress (dynamic, updated each tick)

```
ask patrols [
  ask patches in-cone ps-radius ps-fov [
    if los-clear?(patrol-x, patrol-y, patch-x, patch-y) [
      let threat-stress = (radius - distance) × intensity / radius
      set r-stress max(r-stress, threat-stress)
    ]
  ]
]
```

Stress Function: Linear decay with distance, maximum at threat source, zero beyond radius. Multiple threats use max() to avoid double-counting.

A.7.2 Decision-Making (Nested Filtering)

The protagonist uses a single-phase nested filtering approach, progressively narrowing viable options:

```
to move-rabbit
  ask rabbits [
    ; 1. Get all visible patches in perception cone
    let visible-patches patches in-cone (perception-radius / 2) fov-angle with
[passable?]

    ifelse any? visible-patches [
      ; 2. Filter by line-of-sight (occlusion check)
      let seeable-patches visible-patches with [
        los-clear? ([pxcor] of myself) ([pycor] of myself) ([pxcor] of self) ([pycor] of self)
      ]

      ; 3. Filter by avoidance threshold
      let candidate-patches seeable-patches with [
        avoidance-memory < avoidance-threshold
```

```
]

; 4. Check if area is viable (>50% of visible patches are safe)
ifelse (count candidate-patches / count seeable-patches) > 0.5 [
  ; Area is safe enough - select goal patch

  ifelse random-float 1 < 0.9 [
    ; 90% Exploit: maximize attractiveness
    let max-attractiveness max [calculate-attractiveness myself] of candidate-
patches
    let best-patches candidate-patches with [
      calculate-attractiveness myself = max-attractiveness
    ]
    set goal-patch one-of best-patches
  ][
    ; 10% Explore: go to least visited
    set goal-patch min-one-of candidate-patches [detail-memory + visit-
frequency]
  ]

  ; Move toward goal with noise
  face goal-patch
  rt random 10 - 5  ; ±5° heading noise

  ; Try to move forward
  ifelse [passable?] of patch-ahead 1 and
      [avoidance-memory] of patch-ahead 1 < avoidance-threshold [
    fd 1
  ][
    ; Path blocked - small correction
    rt random 60 - 30
    if [passable?] of patch-ahead 1 and
       [avoidance-memory] of patch-ahead 1 < avoidance-threshold [
      fd 1
    ]
  ]
 ][
  ; Area too dangerous (<50% viable) - turn away
  rt 90
 ]
][
 ; No visible patches - turn to sample new direction
 rt 90
]
end
```

A.7.3 Paralysis Recovery

```
to execute-unstuck-maneuver
 ; Triggered when stuck-counter > 25
 let escape-patches patches in-radius perception-radius with [
  passable? and MA < (TA × 2)  ; Relax threshold temporarily
 ]
```

```
  let safest-patch min-one-of escape-patches [MA]

  ifelse safest-patch != nobody [
    face safest-patch
    repeat 5 [ if [passable?] of patch-ahead 1 [ fd 1 ] ]
  ][
    execute-panic-retreat  ; Last resort: teleport
  ]
end
```

Panic Retreat (Emergency Teleport):

```
to execute-panic-retreat
  print (word "PANIC RETREAT at: " pxcor " " pycor)

  let safe-patches patches with [
    passable? and MA < TA
  ]

  if any? safe-patches [
    move-to min-one-of safe-patches [MA]
  ]
end
```

Justification: Models cognitive paralysis resolution through temporary threshold relaxation or emergency relocation to known-safe zones.

A.7.4 Navigation Graph Pruning

```
to prune-navigation-graph
  set detail-threshold 0.4

  ask patches with [passable?] [
    let should-prune? (MA > TA or MD < detail-threshold)

    ifelse should-prune? [
      ; Prune node
      if not pruned? [
        set pruned? true
        set nav-neighbors no-patches
        ; Remove self from neighbors' edge lists
        ask neighbors with [passable?] [
          set nav-neighbors nav-neighbors with [self != myself]
        ]
      ]
    ][
      ; Restore if conditions improve
      if pruned? [
        set pruned? false
        set nav-neighbors neighbors with [passable? and not pruned?]
```

```
      ; Re-add to neighbors
      ask neighbors with [passable? and not pruned?] [
        set nav-neighbors (patch-set nav-neighbors myself)
      ]
    ]
  ]
]
end
```

Pruning Logic:

Nodes with `MA > TA`: Psychologically erased (threat-driven)

Nodes with `MD < 0.4`: Forgotten (insufficient familiarity)

Creates emergent "Inaccessible Void" as network fragments

A.7.5 Memory Update Mechanism

```
to update-memory
  ask patches [
    ; Increment time since last visit
    set time-last time-last + 1

    ; Memory decay (differential rates)
    ifelse visit-frequency < 50 [
      set detail-memory detail-memory - detail-memory * ((time-last + 1) ^ (-6))
      set avoidance-memory avoidance-memory - avoidance-memory * ((time-last
+ 1) ^ (-8))
    ][
      ; Slower decay for frequently visited patches
      set detail-memory detail-memory - detail-memory * ((time-last + 1) ^ (-8))
      set avoidance-memory avoidance-memory - avoidance-memory * ((time-last
+ 1) ^ (-10))
    ]

    ; Clamp values [0,1]
    set detail-memory max list 0 (min list 1 detail-memory)
    set avoidance-memory max list 0 (min list 1 avoidance-memory)
  ]

  ; Update on visit
  ask rabbits [
    let rabbit-agent self
    ask patches in-cone perception-radius fov-angle [
      if passable? [
        if los-clear? ([pxcor] of rabbit-agent) ([pycor] of rabbit-agent)
                ([pxcor] of self) ([pycor] of self) [
          ; Reset time counter
          set time-last 0
          set visit-frequency visit-frequency + 1
; Update memories
          set detail-memory detail-memory + 0.05 * (knowledge-degree ^ 6)
```

```
        if p-stress > 0 [
          set avoidance-memory avoidance-memory + p-stress * 0.1 * knowledge-
degree
        ]

        ; Clamp
        set detail-memory max list 0 (min list 1 detail-memory)
        set avoidance-memory max list 0 (min list 1 avoidance-memory)
      ]
    ]
  ]
end
```

A.8 DESIGN DETAILS (ADDITIONS TO ODD+D)

A.8.1 Parameter Sensitivity

Tested Ranges (full analysis in Appendix B):

Parameter	Range Tested	Selected Value	Rationale
Detail decay constant	2-10	6-8	Balances learning/forgetting timescale
Avoidance decay constant	2-10	8-10	Persistent fear (slower than detail)
Learning rate (detail)	0.01-0.2	0.05	Plausible semester-scale acquisition
Learning rate (avoid)	0.01-0.2	0.1	Faster threat conditioning
Stuck threshold	15-50	25	Balance responsiveness/noise

Robust Outcomes:

Qualitative fragmentation patterns persist across ±30% parameter variation

Breaking point timing shifts but emergence is consistent

Absolute values subjective; relative scenario comparisons valid

A.8.2 Validation Strategy

Face Validity: Safe zones match peer-mapped locations (3 independent mappers, blind mapping exercise)

Structural Validity:

Graph coverage reduction (50-75%) aligns with reported spatial constraint

Inaccessible Void emergence not predicted a priori (novel finding)

Average degree increase confirms clustering hypothesis

Behavioral Validity:

Exploration-consolidation-paralysis sequence matches auto-ethnographic narrative

Movement reduction under high stress corresponds to reported withdrawal behavior

A.8.3 Scenario Design

Scenario	TA	Stress Intensity	Interpretation
Control	N/A	0.0	Baseline free exploration
Mildly Brave	0.4	1.0	Reluctant compliance
Risk-Avoider	0.2	1.0	Hyper-vigilance
Brave	0.9	1.0	High risk tolerance

Hypothesis: Non-linear relationship between TA and fragmentation severity, with peak constraint occurring in intermediate TA range.

A.8.4 Limitations

1. Single-agent simulation: Excludes collective dynamics, social learning, information sharing

2. Static threat patterns: Patrols follow fixed routes; no adaptive learning from threat agents

3. Simplified environment: Small-scale obstacles, visual details, temporal variations not fully represented

4. Homogeneous baseline: Real environment has pre-existing heterogeneity in social meaning, historical associations

5. Auto-ethnographic data: Limited temporal scope (2016-2021), single perspective requiring peer validation

A.8.5 Extensions for Future Research

Multi-Agent Extensions:

Add collective stress perception (social amplification)

Model information sharing networks (warning systems)

Test emergence of coordinated resistance patterns

Adaptive Threats:

Machine learning patrol routes based on protagonist movement

Dynamic camera repositioning

Escalating enforcement intensity

Counter-Surveillance:

Bidirectional stress (protagonist as discomfort source to threats)
Temporary spatial modifications (furniture, installations)
Territory marking and reclamation dynamics

A.9 REFERENCES

Primary Methodology

Wilensky, Uri. NetLogo. Evanston, IL: Center for Connected Learning and Computer-Based Modeling, Northwestern University, 1999. http://ccl.northwestern.edu/netlogo/

Grimm, V., Railsback, S. F., Vincenot, C. E., et al. (2020). "The ODD Protocol for Describing Agent-Based and Other Simulation Models: A Second Update to Improve Clarity, Replication, and Structural Realism." *Journal of Artificial Societies and Social Simulation* 23(2):7. doi:10.18564/jasss.4259

Theoretical Framework

Mallot, H. A. (2024). "From Geometry to Behavior: An Introduction to Spatial Cognition". Cambridge, MA: MIT Press.

Tolman, E. C. (1948). "Cognitive Maps in Rats and Men." Psychological Review 55(4):189-208. doi:10.1037/h0061626

Gibson, J. J. (1979). "The Ecological Approach to Visual Perception". Boston: Houghton Mifflin.

Staddon, J. E. R., Chelaru, I. M., & Higa, J. J. (2002). "Habituation, Memory and the Brain: The Dynamics of Interval Timing." Behavioral Processes 57(2-3):71-88. doi:10.1016/S0376-6357(02)00006-2

Contextual Background

Foucault, M. (1977). "Discipline and Punish: The Birth of the Prison." New York: Pantheon Books.

Ingold, T. (1993). "The Temporality of the Landscape." World Archaeology 25(2):152-174.

A.10 MODEL AVAILABILITY

Platform	*NetLogo 7.0.0*
Code Length	*+1,400 lines*
Reproducibility	*Full ODD is provided in thesis appendix*
Data Outputs	*CSV timeseries, spatial heatmaps, navigation graph exports*
Contact	*Kosar Mohammadi, Politecnico di Torino*
Supervisor	*Prof. Tomà Berlanda*
Doc. Version	*1.0*
Last Updated	*Jan 19 2026*
Protocol Compliance	*ODD+D (Grimm et al. 2020)*

Appendix B
Sensitivity Analysis

A.1 decay constants

All the parameters except for decay constants are kept the same used for Mildly Brave scenario. In the first version, cd-avoid was set to 4 and cd-detail to 2. As a result, the decay happens very fast, keeping the detail memory at very low levels, disrupting the process. While the learning still takes place, and the final results are following the same pattern of peripheral clustering, the intensity of memory increases very slowly and in the timeline of 1000 ticks, it is not efficient.

By increasing the parameters to 8 and 10, decay is much slower and almost not effective on the learning process. The agent makes fewer mistakes, and although the final results follow the same pattern, the speed and accuracy of finding safe spots is not practical.

Values 6 and 8 prove to be matching the effectiveness of decay in the simulation conditions, maintaining an acceptable level of learning in the timeline while replicating human behavior better. In the following figure, maps of navigation and memory of the three scenarios over 10000 ticks is provided.

Figure A.1 *The maps of navigation and memory under different values of decay constant: a. Cd-avoid = 8, cd-detail = 10. Although the same clustering is visible, the agent's behavior shows fast learning and lack of mistakes which makes the boundaries of safe spaces very clear, b. Cd-avoid = 6, cd-detail = 8, the agent shows occasional clustering in the middle, c. Cd-avoid = 2, cd-detail = 4, the agent's memory is still quite underdeveloped, however the same pattern is visible [Author]*

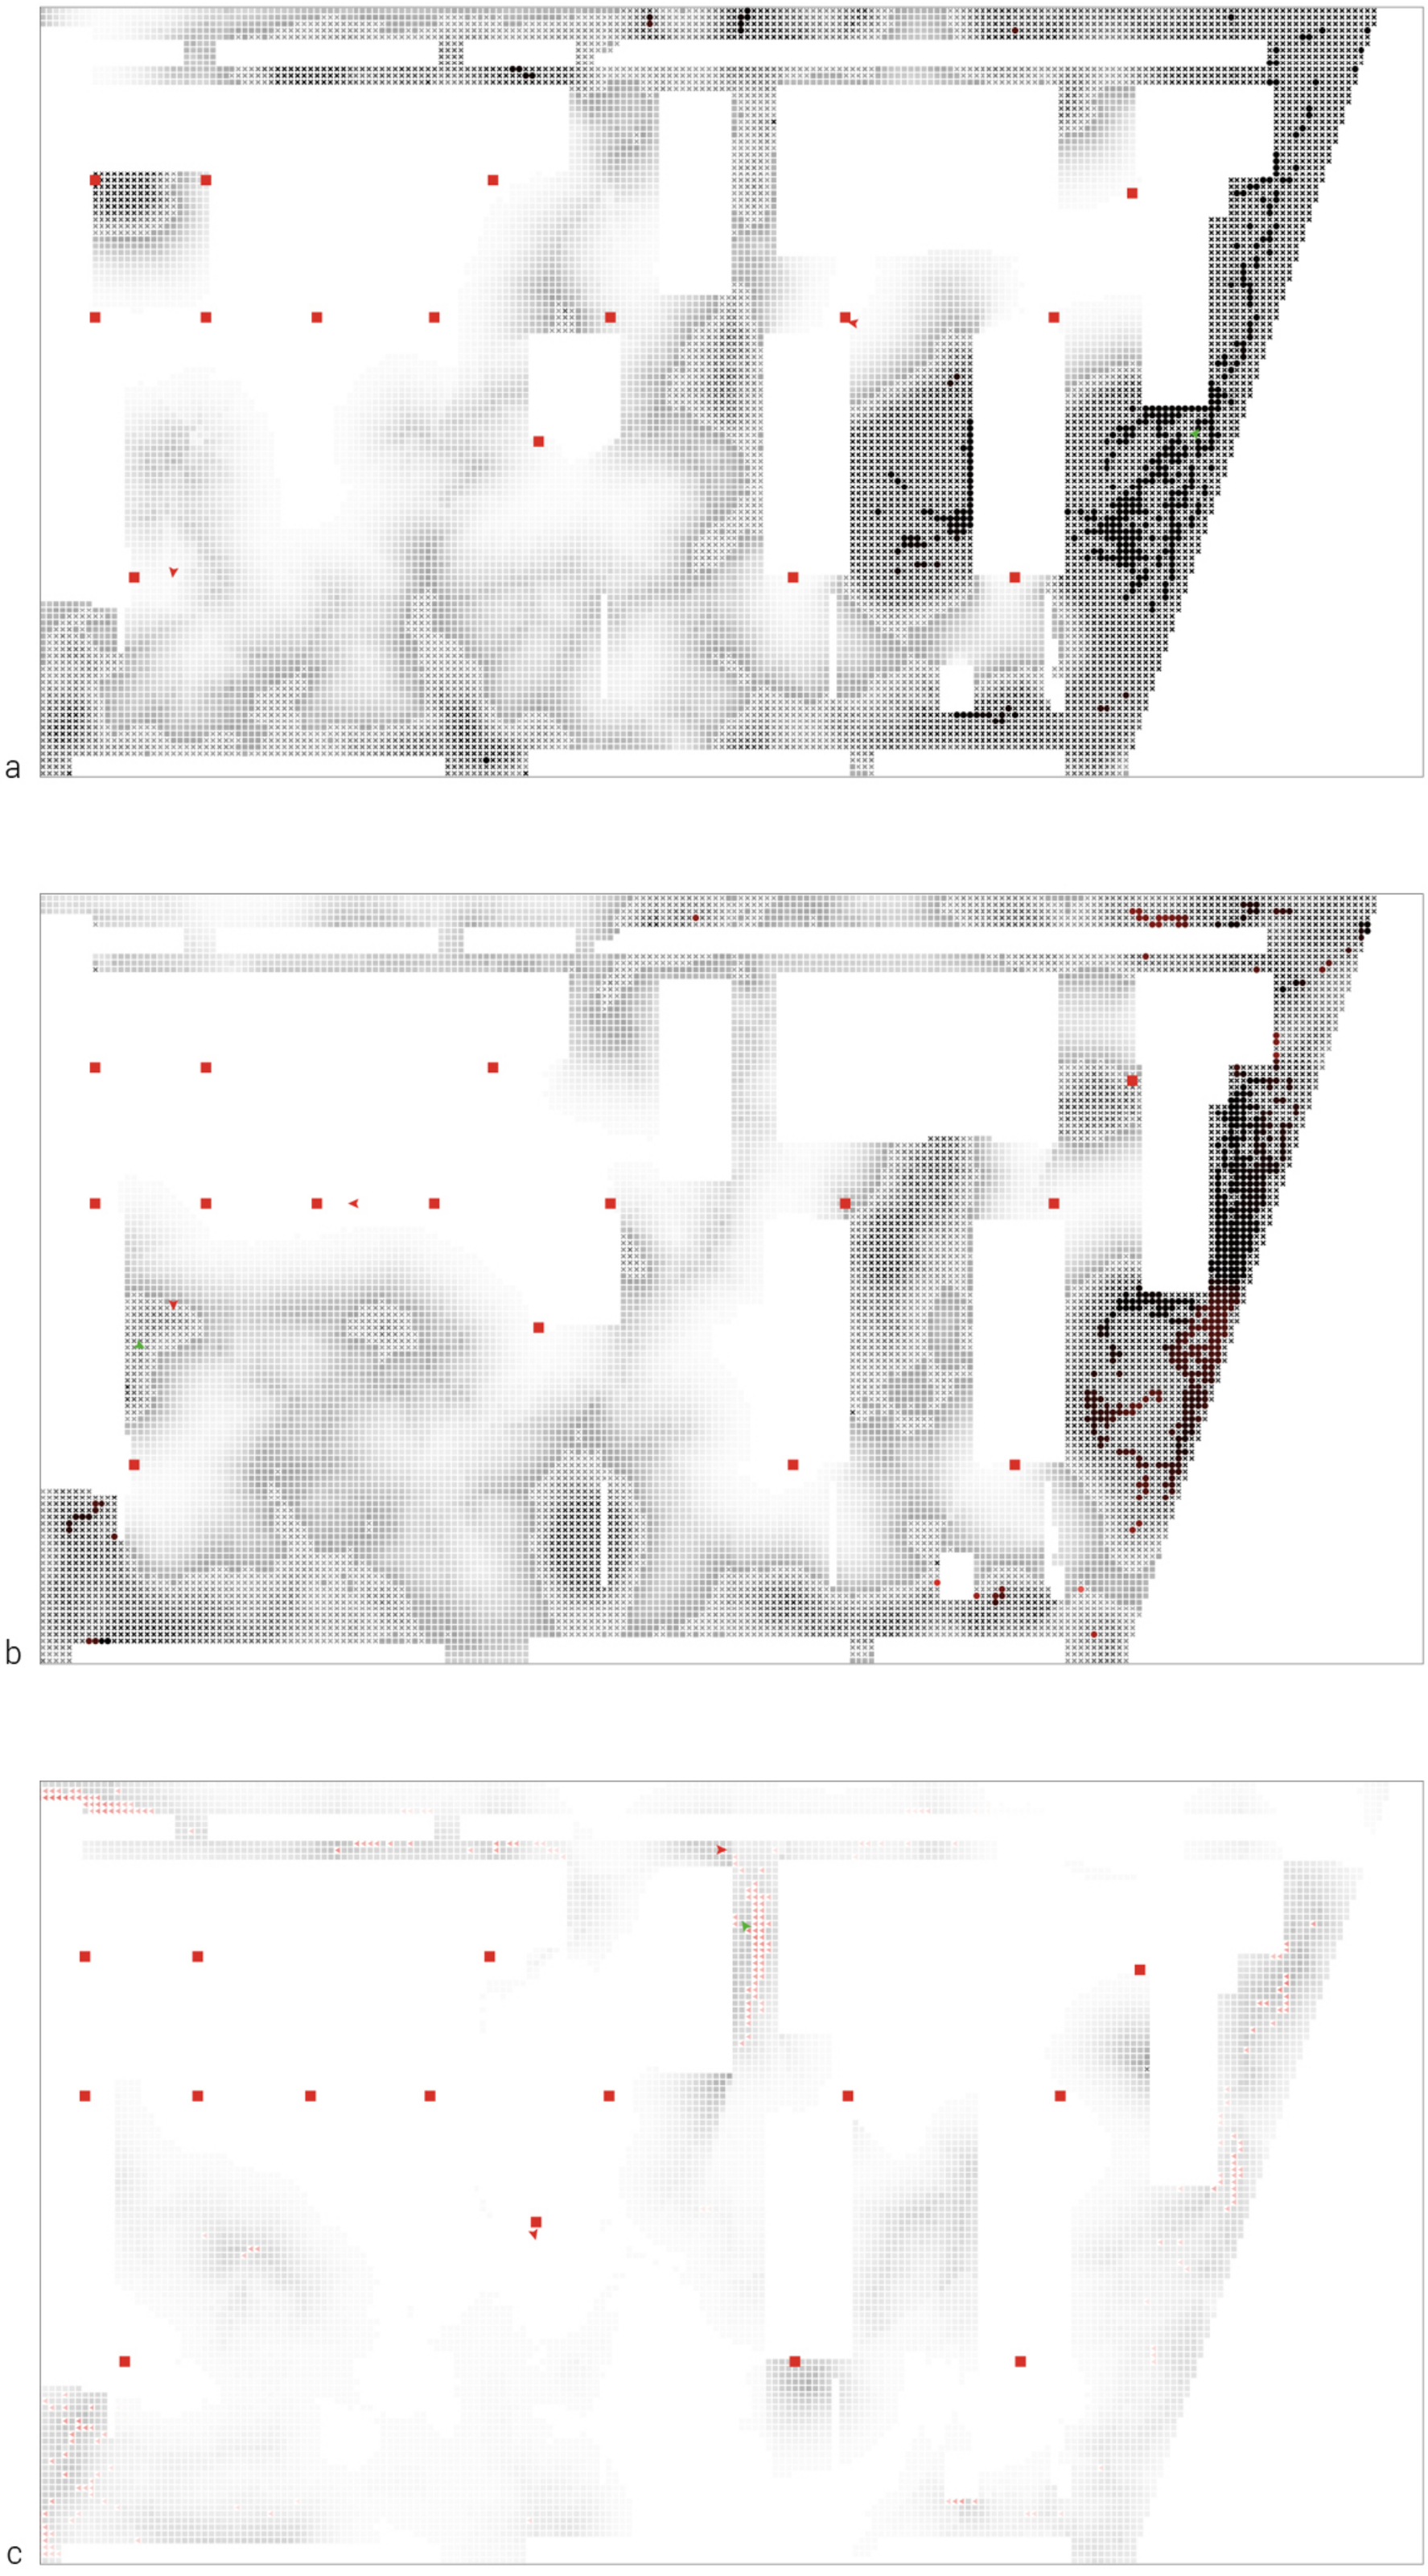

a

b

C

A.2 Learning rate

By setting learning rate for detail-memory and avoidance-memory to 0.01 and 0.02 respectively, learning is too slow and inefficient, however the results would not change, though the simulation time needed to reach the breaking point increases significantly. As shown in figure A.2.c, the accumulation of memory in 10000 ticks is much slower than the other versions.

By increasing the learning rates to 0.1, 0.2, learning happens faster than ordinary, making double visits of a patch almost non existent, although the results show the same pattern of avoiding the threats. This almost resembles what happened in figure A.1.a, when by setting the decay constants to 8 and 10, the mistakes decreased significantly, leading to less human-like behavior.

Finally, by setting the learning rate to 0.05 and 0.1 respectively, the accumulation of memory is somewhere in-between for the timeline of 10000 ticks. However, the results follow the same pattern with degrees of expected variation.

Figure A.2 *The maps of navigation and memory under different values of learning rate: a. Alpha-avoid = 0.2, alpha-detail = 0.1. Although the same clustering is visible, the agent's behavior shows fast learning and lack of mistakes and clear borders, b. Alpha-avoid =0.1, alpha-detail = 0.05, the agent shows occasional clustering in the middle, c. Alpha-avoid =0.01, alpha-detail = 0.02, the agent's memory is still quite underdeveloped, however the same pattern is visible [Author]*

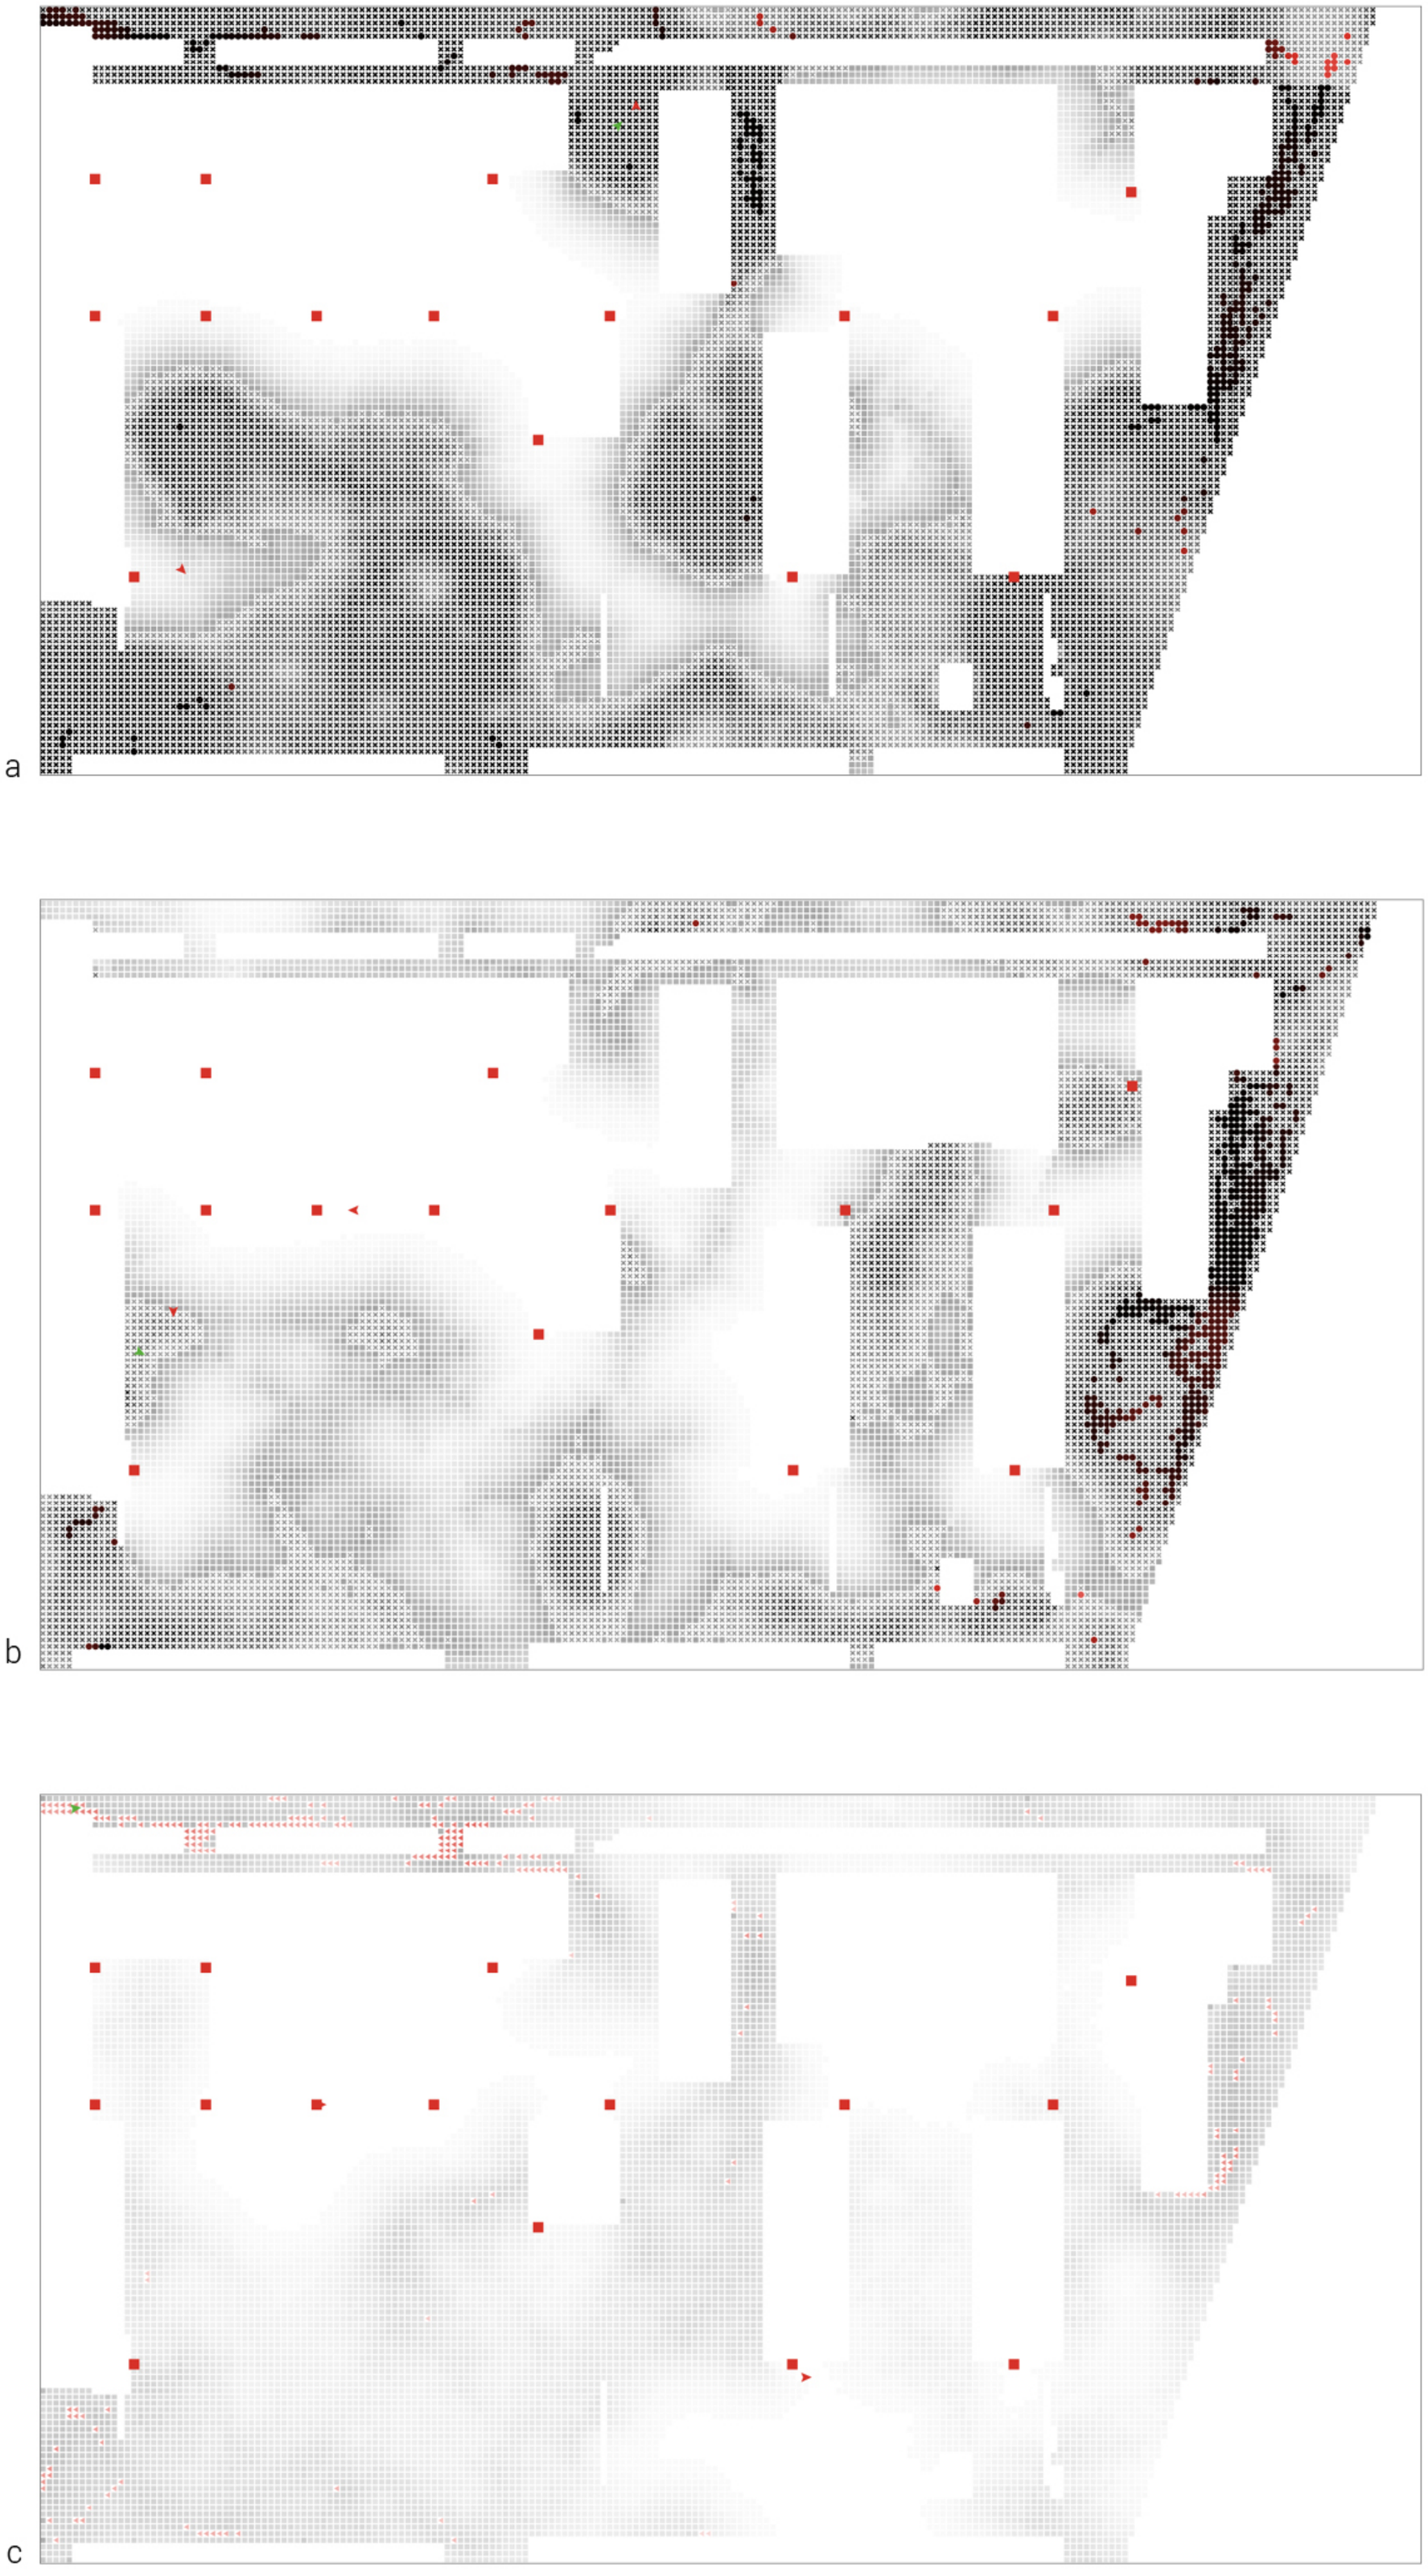

a

b

C

A.3 Stuck Threshold

By changing the value of stuck threshold three different scenarios were examined to calibrate the number. While the values are not much effective on the fragmentation results, in the scenario of stuck-threshold = 45, it can produce unwanted detail memory solely because the agent is trapped and has no options. The unwanted memory accumulation leads to appearance of clusters that in fact are not visited by the agent more than once.

By decreasing the value to 15, the frequency of unstuck maneuver increases, making abrupt movements. As I previously mentioned, while the general outcome remains largely the same, the abrupt jumps do not comply with human behavior and practicality of the mechanism.

Finally, by setting the value to 25, the necessary unstuck maneuver is preserved, while the frequency of it decreases.

Figure A.3 *The maps of navigation and memory under different values of Stuck Threshold: a. Threshold = 45, the overall pattern is not changes but the unwanted accumulation of memory when the agent is trapped adds unwanted noise, b. Threshold = 25, the noise is reduced while the unstuck function is not used frequently, c. Threshold = 15, the fragmentation remains the same, but the frequency of jumps increases unnecessarily [Author]*

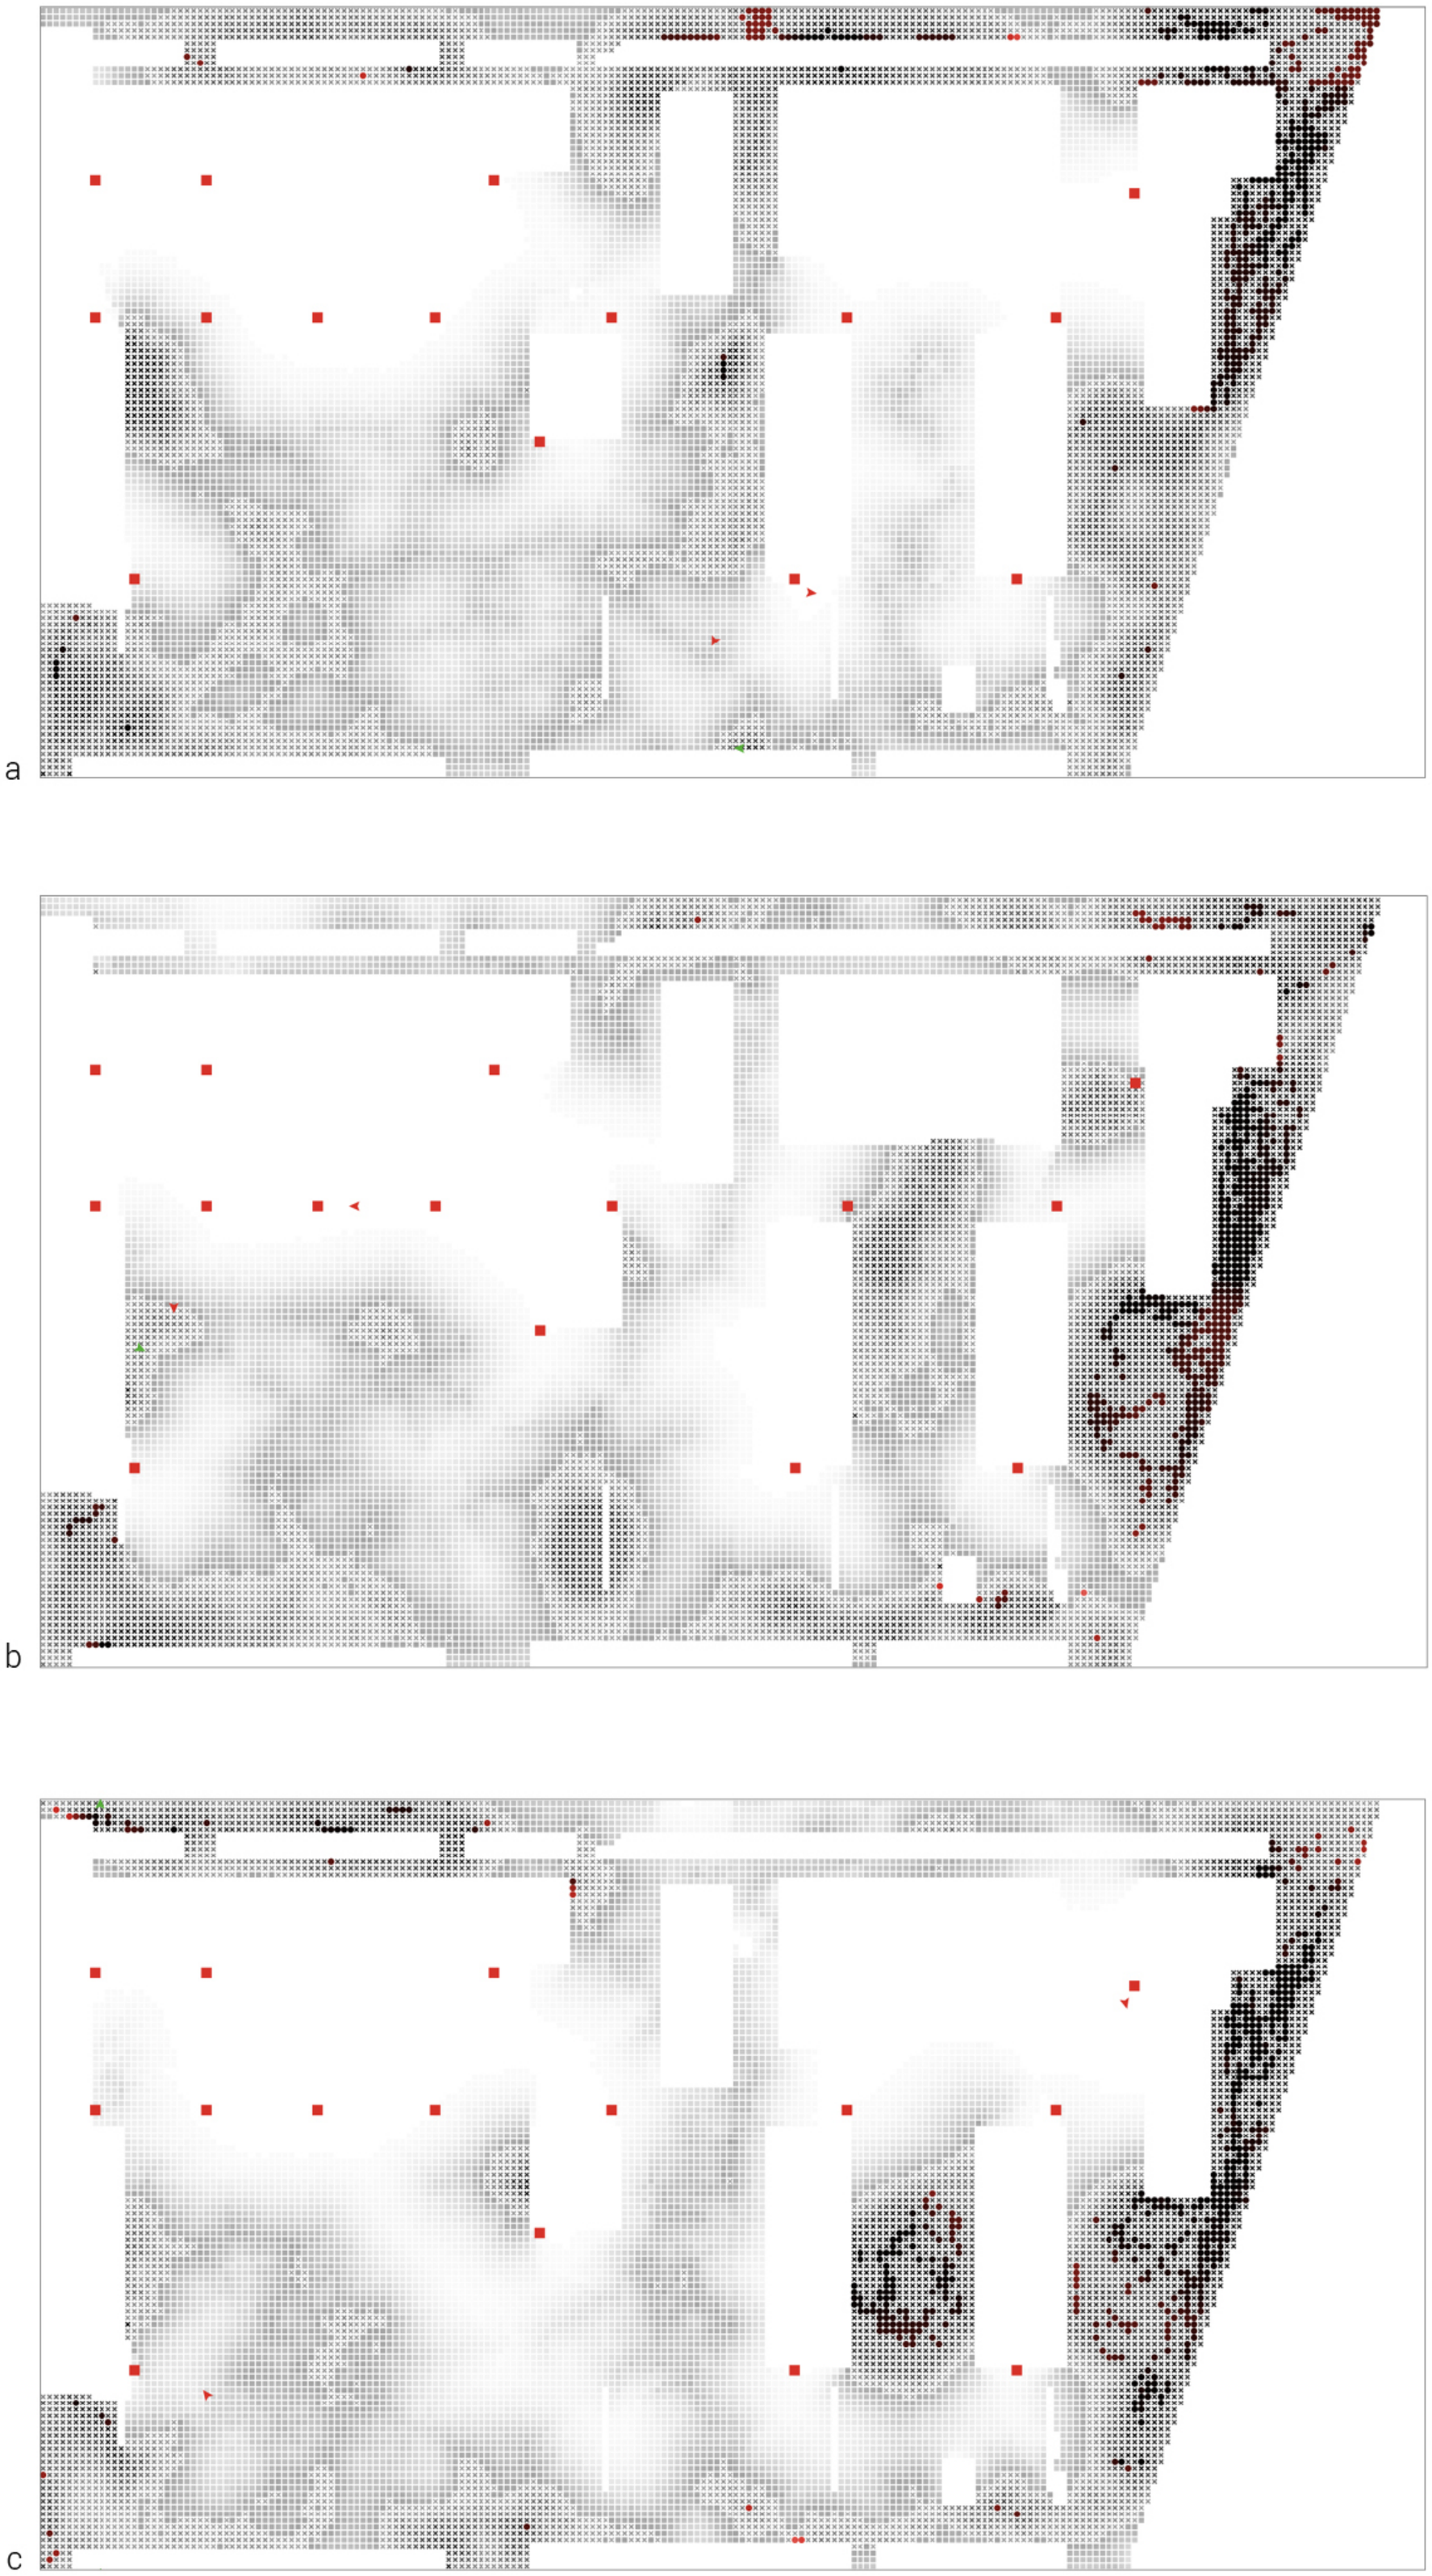

a

b

C

Politecnico di Torino | Feb. 2026