


A young boy with dark hair is sitting on a grey carpeted floor in a library. He is wearing a dark jacket over a grey sweater and dark pants. He is hugging his knees with his arms and looking down with a sad or thoughtful expression. He is surrounded by tall bookshelves filled with books. The bookshelves are made of light-colored wood and the books have various colored spines. The lighting is soft and even.

Beyond Event Detection: A Practice-Based Study of AI-Supported Responses to Relational Bullying in Educational Contexts

CONTENTS

Abstract.....	1
1. Introduction	2
1.1 Bullying as a Persistent and Often Invisible Harm	2
1.2 AI Detection Systems and the Limits of Event-Based Recognition	2
1.3 Structural Violence, Trauma, and the Limits of Visibility	3
1.4 Research Aim and Design Direction	3
2. Literature Review	5
2.1 Mapping the Field: Technological, Sociological, and Psychological Perspectives	6
2.1.1 Technological Approaches: AI Detection Systems and Their Limitations	6
2.1.2 Sociological Perspectives: Structural Violence and Institutional Power	7
2.1.3 Psychological Perspectives: Trauma, Silence, and Tacit Experience	9
2.2 Identifying the Gaps: From Event-Based Detection to Structural Intervention	10
2.2.1 From Events to Patterns: The Recognition Gap	10
2.2.2 From Explicit Reporting to Silent Data: The Epistemic Gap	12
2.2.3 From Individual Intervention to Structural Intervention: The Design Gap	12
2.3 Theoretical Orientation: Tacit Knowledge and Reflective Design Practice	13
2.3.1 Tacit Knowledge and the Problem of Unspoken Experience	13
2.3.2 Reflective Practice and Design as Knowledge Production	14
2.4 From Critique to Construction	15
3. Research Methodology	16
3.1 A Practice-Based Approach to Structural Intervention	16

3.2 Conceptual Framework of the Guardian System	18
3.3 Research Design Overview	20
3.4 Co-Design as Situated Inquiry	21
3.5 Prototype Development and Iterative Exploration	23
3.6 Situated Evaluation: From Effectiveness to Interpretation	24
3.7 Practice as Data	25
3.8 Data Analysis and Reflexive Interpretation	26
3.9 Limitations of the Approach	27
3.10 Methodological Summary	28
4. Ethical Considerations	29
4.1 Ethics Beyond Procedural Compliance	29
4.2 Researcher Positionality and Reflexivity	29
4.3 Participants, Consent, and Emotional Safety	30
4.4 Data Interpretation and Behavioural Representation	31
4.5 Surveillance, Care, and Ongoing Ethical Reflection	31
5. Reflection and Conclusion	33
5.1 Revisiting the Research Problem	33
5.2 Design Practice as Inquiry	33
5.3 Reflective Making and Tacit Experience	34
5.4 Tensions Within Visibility and Interpretation	34
5.5 Methodological and Practical Limitations	35
5.6 Ethical Implications and Future Research	35
5.7 Final Reflection and Contribution	36
6. References	37

Abstract

Current AI anti-bullying systems mainly focus on detecting explicit incidents such as abusive language, physical aggression, or direct victim reports. However, many forms of bullying emerge through long-term exclusion, silence, and relational patterns that are difficult to identify through event-based detection alone. This research explores how practice-based design methods might support alternative approaches to understanding these less visible forms of harm.

Drawing on theories of structural violence, trauma, tacit knowledge, and reflective practice, the study argues that existing intervention models often overlook the ambiguity and contextual nature of bullying experiences. In response, the project develops a speculative Guardian System as a design research inquiry rather than a finished technological solution.

Using co-design workshops, iterative prototyping, and situated evaluation, the research investigates how relational patterns and behavioural changes might become partially visible without relying solely on explicit reporting or automated judgement. The study ultimately reframes AI anti-bullying intervention from event detection toward a more reflective and structural understanding of care, interpretation, and social relations.

Keywords:

AI-supported intervention, school bullying, structural violence, relational harm, tacit knowledge, practice-based research

1. Introduction

1.1 Bullying as a Persistent and Often Invisible Harm

School bullying remains a persistent issue across educational contexts worldwide. According to UNESCO (2019), nearly one-third of students globally have experienced some form of bullying. While public discussions often focus on physical violence or direct verbal harassment, many harmful experiences take more subtle forms, including social exclusion, relational manipulation, and ongoing humiliation in online spaces. Research on relational bullying suggests that these forms of harm are often difficult for teachers, parents, and school systems to recognise, precisely because they emerge gradually through everyday interactions rather than through singular visible incidents (Crick and Grotpeter, 1995). At the same time, many students do not actively seek help. Existing surveys indicate that victims may remain silent due to fear of retaliation, concern about not being believed, or difficulty explaining what they are experiencing (Anti-Bullying Alliance, 2023). As a result, bullying frequently exists not as a clearly identifiable event, but as a pattern of ongoing social experiences that can remain unnoticed for long periods of time.

1.2 AI Detection Systems and the Limits of Event-Based Recognition

As AI systems become increasingly embedded within educational platforms, communication technologies, and student support services, questions about how these systems recognise and respond to harm have become more urgent. Recent research has explored automated approaches to bullying detection through text analysis, sentiment recognition, and behavioural monitoring (Dinakar et al., 2021; Nobata et al., 2016). These approaches have shown some success in identifying explicit aggression and high-risk language. However, their underlying logic still largely centres on event detection, relying on keywords, emotional intensity, or identifiable behavioural patterns to determine risk. This creates a significant limitation when harm is indirect,

relational, or difficult to articulate. Existing research has also shown that automated moderation and detection systems do not operate independently from their surrounding environments, but are shaped by school cultures, platform governance, and broader social relationships (Chandrasekharan et al., 2017). For design research, this raises a broader question about how supportive technologies might engage with forms of harm that are gradual, relational, and not always immediately visible.

1.3 Structural Violence, Trauma, and the Limits of Visibility

Research from sociology and psychology further suggests that bullying cannot simply be understood as isolated conflict between individuals. Structural violence theory argues that certain forms of harm become embedded within institutional arrangements and everyday social structures (Galtung, 1969). Related work by Pierre Bourdieu (1986) further demonstrates how unequal social relations may continue to be reproduced through ordinary practices and cultural norms. Meanwhile, trauma research suggests that individuals experiencing ongoing harm may not always be able to clearly articulate their experiences through language (Herman, 1998; van der Kolk, 2014). Some experiences may instead appear through silence, withdrawal, behavioural change, or emotional instability. This suggests a limitation in AI systems that rely heavily on explicit reporting or observable incidents, as such systems may only recognise the most visible parts of a much broader problem.

1.4 Research Aim and Design Direction

Against this background, this research asks how AI systems might support responses to school bullying when harm is relational, difficult to articulate, and not always visible through discrete events. From the perspective of structural intervention, the project explores an alternative approach to AI-supported care within educational contexts. Through the design and iterative development of the Guardian System, the research investigates how systems might provide

support when victims are unable to actively seek help, how potential risks might be recognised without depending solely on explicit conflict detection, and how long-term structural patterns of harm might gradually become visible while still protecting individual privacy and experience.

2. Literature Review

Current AI anti-bullying systems primarily focus on detecting explicit incidents and victim self-reporting, while paying less attention to the relational and structural forms of harm that often characterise bullying in everyday life. Existing research has made important progress in automated moderation, cyberbullying detection, and behavioural recognition. However, many interventions continue to assume that bullying appears as identifiable events that can be directly observed, verbally reported, or computationally classified. Such assumptions become increasingly limited when bullying emerges through gradual exclusion, unequal social positioning, and long-term relational dynamics.

Rather than understanding bullying as isolated incidents, this review approaches it as a structural and relational condition shaped through everyday interaction, institutional norms, and unequal power relations. From this perspective, current technological approaches reveal important limitations. While sociological and psychological research has expanded understanding of structural violence, trauma, and victim silence, these insights have not yet been fully translated into AI-supported intervention systems.

This chapter critically reviews literature from three interconnected perspectives: technological approaches to bullying intervention, sociological theories of structural violence and institutional power, and psychological studies of trauma, silence, and tacit experience. Rather than simply summarising previous work, the review examines the assumptions and limitations underlying current intervention models in order to identify gaps that motivate the present research.

2.1 Mapping the Field: Technological, Sociological, and Psychological Perspectives

Existing anti-bullying research can broadly be divided into three overlapping areas. The first focuses on AI-driven detection systems designed to identify bullying-related behaviour. The second examines bullying through sociological frameworks such as structural violence, institutional power, and social reproduction. The third explores psychological perspectives on trauma, silence, and victim experience. Although these perspectives contribute valuable insights, they often remain insufficiently connected in practice. Technical systems frequently reduce bullying to observable events, while sociological and psychological research suggests that many harmful experiences are relational, contextual, and difficult to articulate.

2.1.1 Technological Approaches: AI Detection Systems and Their Limitations

Research on AI-supported anti-bullying intervention has largely focused on automated detection systems. Within natural language processing research, studies have explored semantic, contextual, and thematic approaches to identifying abusive language and cyberbullying content (Dinakar et al., 2021; Hee et al., 2018; Nobata et al., 2016; Smith et al., 2008). Alongside textual analysis, computer vision research has examined behavioural recognition systems capable of identifying physical aggression through gesture analysis, pose estimation, and facial recognition technologies (Omarov et al., 2022).

Although these systems demonstrate increasing technical sophistication, many remain dependent on predefined and observable indicators of harm, such as explicit insults, recognisable distress behaviour, or direct verbal threats. In experimental datasets, bullying is frequently represented through clearly identifiable actions and standardised behavioural patterns.

However, bullying in everyday educational environments is often far less visible. Harm may emerge through exclusion, humiliation disguised as humour, long-term social isolation, or repeated avoidance within peer relationships. Such forms of relational aggression rarely appear as isolated incidents and are difficult to reduce to explicit keywords or short-term behavioural events.

This reveals a broader tension between computational systems that depend on observable and labelable signals, and sociological understandings of bullying as gradual, relational, and structurally embedded. Existing AI systems remain largely optimised for moments of recognisable conflict rather than long-term social dynamics.

Another limitation concerns the assumption that detection itself constitutes intervention. Many systems follow a relatively linear logic in which harmful behaviour is identified, alerts are triggered, and institutional authorities are expected to respond. However, the consequences of technological intervention are often shaped by wider social and institutional conditions. Chandrasekharan et al. (2017), in their analysis of automated moderation systems on Reddit, demonstrate that the effects of moderation depend not only on technical accuracy but also on governance structures, community norms, and existing social relationships.

While AI research has become increasingly effective at recognising explicit aggression, considerably less attention has been paid to ambiguity, silence, and structural forms of exclusion.

2.1.2 Sociological Perspectives: Structural Violence and Institutional Power

Sociological research offers an alternative understanding of bullying by shifting attention away from individual incidents and toward broader systems of power,

inequality, and social reproduction. Galtung's (1969) theory of structural violence argues that harm can become embedded within institutional arrangements and social structures rather than appearing solely through direct aggression. Within educational environments, unequal expectations, institutional categorisation, and everyday social practices may gradually reinforce marginalisation and exclusion.

Bourdieu (1986) further demonstrates how educational institutions reproduce social hierarchies through forms of symbolic power that often appear neutral or legitimate. From this perspective, bullying cannot be understood solely as interpersonal hostility, but must also be considered in relation to broader cultural and institutional conditions that shape belonging and social legitimacy.

Research on cyberbullying similarly suggests that harmful behaviour is deeply connected to wider social environments rather than individual intention alone. Hinduja and Patchin (2020) argue that online hostility is shaped by social norms and cultural conditions that normalise exclusion and aggression. Foucault's (1977) analysis of discipline and surveillance further complicates technological intervention by demonstrating how observation and behavioural regulation function as forms of power.

These sociological perspectives significantly expand understanding of bullying beyond isolated conflict. At the same time, they reveal an important methodological difficulty. While structural theories explain how violence becomes normalised within institutions and everyday life, they offer limited guidance regarding how such conditions might be translated into practical technological intervention systems.

Espelage and Swearer (2003), drawing on ecological systems theory, argue that bullying must be understood through interactions between peers, schools,

communities, and wider social structures. Yet such relational conditions remain difficult to represent within conventional computational models. Existing intervention systems continue to prioritise identifiable acts of aggression rather than long-term relational dynamics.

2.1.3 Psychological Perspectives: Trauma, Silence, and Tacit Experience

Psychological research further complicates assumptions underlying many anti-bullying interventions by examining how trauma, memory, and emotional experience shape victims' ability to recognise and communicate harm.

Herman (1998) argues that traumatic experiences are often characterised by difficulty of narration. Individuals experiencing prolonged humiliation, fear, or instability may struggle to translate their experiences into coherent verbal accounts. Similar arguments are developed in van der Kolk's (2014) research on trauma and embodied memory, which suggests that traumatic experiences may appear through fragmented memory, emotional response, or behavioural change rather than linear narrative.

These perspectives challenge a central assumption within many intervention systems: that victims are able to clearly identify, describe, and report harmful experiences. Pennebaker (1997) demonstrates that emotional disclosure is shaped by complex psychological conditions, while Fuggle, Dunsmuir and Curry (2013) show that negative experiences may produce overgeneralised memory patterns rather than precise recollection of events.

Morgan and Janoff-Bulman (1994) further argue that sustained victimisation may gradually reshape individuals' assumptions about themselves and the world, reducing the likelihood that they will seek support. Silence, therefore, should not simply be interpreted as absence of harm or unwillingness to

communicate. It may instead reflect trauma, emotional exhaustion, or internalised social pressure.

At this point, discussions of tacit knowledge become particularly relevant. Polanyi (1966) argues that individuals often know more than they can explicitly express. Experiences may remain embodied, emotional, and situational without becoming fully articulated through language. Within bullying contexts, withdrawal from social interaction, silence in particular environments, or long-term marginalisation within peer networks may all function as tacit indicators of distress.

Compared with these complex and partially unspoken experiences, many technological systems still rely heavily on explicit linguistic signals and stable behavioural categories. As a result, forms of suffering that remain ambiguous or difficult to articulate frequently fall outside computational recognition.

2.2 Identifying the Gaps: From Event-Based Detection to Structural Intervention

The literature reviewed above reveals several important disconnections between technological systems, sociological theory, and psychological understandings of victim experience. Although each perspective contributes valuable insights, important gaps remain regarding how hidden and structural forms of bullying might be recognised and addressed in practice.

2.2.1 From Events to Patterns: The Recognition Gap

Most existing technological systems treat bullying as a series of isolated incidents. Detection models are commonly trained to identify discrete events such as verbal insults, physical aggression, or threatening messages. However, sociological research suggests that harm often emerges gradually

through repeated exclusion, unequal participation, and shifting social positioning rather than singular events.

Such forms of relational aggression are difficult to translate into conventional datasets because their significance emerges cumulatively over time. Existing event-based models therefore remain poorly equipped to recognise long-term social dynamics.

This problem also connects to broader discussions of situated interaction within HCI research. Dourish (2006) argues that the meaning of behaviour cannot be separated from the social and temporal contexts in which it occurs. Future anti-bullying systems may therefore need to move beyond trigger-based detection and instead engage with relational patterns, changes in interaction frequency, and sustained forms of social isolation.

Furthermore, relational forms of bullying are often shaped through temporality rather than singular intensity. Harm may accumulate gradually through repeated micro-interactions that appear insignificant when observed independently but become psychologically meaningful over time. Small acts of exclusion, unequal participation, avoidance, or subtle humiliation may individually seem ambiguous, yet collectively contribute to sustained experiences of marginalisation.

This creates important challenges for computational systems that rely on short-term behavioural interpretation. Event-based detection models typically prioritise moments of visible escalation, while relational bullying frequently operates through repetition, atmosphere, and social positioning.

Understanding these dynamics may therefore require approaches capable of examining interaction histories and longer-term social patterns rather than isolated behavioural incidents alone.

2.2.2 From Explicit Reporting to Silent Data: The Epistemic Gap

A second limitation concerns the relationship between visibility, expression, and recognition. Most intervention systems rely heavily on explicit reporting or predefined behavioural indicators in order to identify risk. However, psychological research demonstrates that victims experiencing sustained harm may struggle to articulate their experiences due to trauma, normalisation, or fragmented memory.

As a result, some of the most significant experiences of harm may never enter institutional systems of recognition. The issue, therefore, is not only whether information is expressed, but also which forms of experience are considered legitimate and recognisable forms of data.

Polanyi's (1966) discussion of tacit knowledge becomes especially relevant here. Human experience often exceeds what can be formally articulated. Long-term withdrawal from group interaction, repeated avoidance behaviour, or persistent marginalisation may all function as subtle indicators of vulnerability.

Such traces should not be interpreted as objective proof, but as contextual signals that invite further attention and care. This gap therefore involves not only technological limitations but also broader epistemological assumptions regarding what counts as evidence and visibility.

2.2.3 From Individual Intervention to Structural Intervention: The Design Gap

A third limitation concerns intervention logic itself. Existing systems are often organised around individualised intervention pathways in which harmful behaviour is identified, responsibility is assigned, and institutional authorities are expected to respond.

While such approaches may address visible aggression, sociological research suggests that they are less capable of engaging with the broader structural and relational conditions that sustain exclusion over time. Chandrasekharan et al. (2017) further demonstrate that moderation systems may unintentionally reinforce exclusionary dynamics depending on how interventions are implemented.

This suggests the need to rethink the goals of technological intervention. Rather than focusing exclusively on identifying individual acts of aggression, systems might instead support awareness of long-term relational conditions, unequal participation patterns, and hidden forms of exclusion.

From this perspective, technological systems become more than detection tools. They may function as mediating structures that support reflection, discussion, and recognition of structural inequality within educational environments.

2.3 Theoretical Orientation: Tacit Knowledge and Reflective Design Practice

The limitations identified above suggest the need for alternative approaches capable of engaging with ambiguity, silence, and situated social experience. Two theoretical perspectives are particularly important for the present study: tacit knowledge and reflective design practice.

2.3.1 Tacit Knowledge and the Problem of Unspoken Experience

Polanyi (1966), in *The Tacit Dimension*, argues that individuals often know more than they can explicitly express. Human understanding is not limited to formal language or propositional knowledge. Experiences may remain embodied, emotional, and situational without becoming fully articulated.

This perspective offers an important way of understanding victim silence within bullying contexts. Victims may recognise discomfort, humiliation, or exclusion through bodily and emotional experience without necessarily being able to describe these conditions coherently.

This has important implications for research methodology and intervention design. Approaches based exclusively on interviews, surveys, or explicit reporting may struggle to capture tacit and relational dimensions of experience. Alternative approaches may therefore need to engage with behavioural traces, interaction patterns, and forms of situated participation.

2.3.2 Reflective Practice and Design as Knowledge Production

Schön's (1983) theory of reflective practice challenges the assumption that complex social problems can be addressed through fixed technical solutions alone. Instead, knowledge emerges through ongoing cycles of reflection, experimentation, and situated engagement. Dourish (2006) similarly argues that interaction cannot be separated from its social and contextual conditions.

Within this research, design practice functions not simply as implementation, but as a mode of inquiry. Iterative making can reveal assumptions that may remain invisible within purely theoretical discussion. Prototype development does not simply test predetermined hypotheses, but may also expose tensions between institutional expectations, technological interpretation, and lived experience. Reflective design practice therefore becomes particularly valuable when addressing socially complex and emotionally ambiguous conditions such as bullying, where meanings and experiences often remain unstable across different contexts. Prototype development, participatory workshops, and iterative evaluation become ways of exploring how structural forms of harm may become perceptible within educational environments.

From this perspective, the proposed Guardian System is not understood as a final technological solution, but as an exploratory and reflective intervention through which assumptions about visibility, care, and relational harm can be critically examined.

2.4 From Critique to Construction

This chapter has critically reviewed technological, sociological, and psychological literature in order to examine the limitations of current AI anti-bullying interventions. The review identified important gaps between event-based detection systems and the relational, structural, and tacit dimensions of bullying experience.

Existing systems remain heavily dependent on explicit evidence, observable behaviour, and individualised intervention logic, while forms of structural violence and silent suffering often remain difficult to recognise computationally. In response, this research adopts a design-practice-based approach informed by tacit knowledge and reflective practice. Through the iterative development of the Guardian System, the study explores how AI-supported interventions might move beyond isolated event detection toward more relational and structural forms of support.

The following chapter outlines the methodological framework of the project, including participatory workshops, iterative prototyping, and approaches for examining the visibility of structural harm within educational settings.

3. Research Methodology

3.1 A Practice-Based Approach to Structural Intervention

This research adopts a practice-based research approach in order to explore forms of bullying that are often difficult to identify through existing intervention frameworks. As discussed in the previous chapter, current AI anti-bullying systems largely focus on event detection and explicit reporting, while paying less attention to long-term relational patterns, tacit experiences, and structural forms of harm. The three gaps identified in the literature review — the gap between events and patterns, between explicit reporting and silent data, and between individual intervention and structural intervention — suggest that the problem addressed in this study cannot be approached solely through conventional empirical methods.

Unlike problems that can be reduced to stable variables and measured through predefined indicators, structural violence and tacit harm are often ambiguous, context-dependent, and relational. In many cases, the issue is not simply whether bullying occurs, but how certain forms of exclusion gradually become normalised within everyday social interactions. Similarly, silence cannot always be interpreted as the absence of experience. Psychological research suggests that individuals who experience long-term humiliation, fear, or exclusion may struggle to describe their experiences directly, even when those experiences strongly shape their behaviour and social positioning.

Because of this, the study does not begin from the assumption that the phenomenon already exists in a clearly observable form waiting to be measured. Instead, the research problem itself is expected to evolve throughout the inquiry process. Certain behaviours may initially appear insignificant but later emerge as meaningful through repeated observation, participant discussion, or prototype interaction. In this sense, the research is

not only concerned with finding answers, but also with developing ways of perceiving and articulating problems that are not yet fully visible.

For this reason, practice-based inquiry is adopted not simply as a methodological preference, but as an epistemological necessity. Traditional methods such as surveys or structured interviews rely heavily on participants' ability to retrospectively explain and categorise their experiences. However, the kinds of experiences explored in this study are often fragmented, tacit, relational, or emotionally difficult to verbalise. Likewise, purely quantitative approaches tend to require clearly definable behavioural indicators, while structural exclusion frequently operates through subtle repetitions and social patterns that resist discrete classification. As a result, neither direct self-report nor predefined computational categories are sufficient on their own.

Within a practice-based framework, design is treated not only as a process of producing solutions, but also as a mode of inquiry. Knowledge emerges through iterative engagement with materials, participants, situations, and unexpected reactions. This position aligns with Schön's concept of reflection-in-action, in which understanding develops through cycles of making, testing, observing, and revising rather than through the direct application of fixed theory. The purpose of the research is therefore not to validate a single technological system, but to use the process of designing the Guardian System as a way of investigating how structural harm might become partially perceptible and discussable.

Importantly, the study does not assume that technology can objectively identify victimisation. Instead, it approaches behavioural traces, relational patterns, and interaction changes as ambiguous signals that require interpretation within specific contexts. This distinction is important because attempts to detect silence may themselves risk reproducing forms of surveillance,

overinterpretation, or institutional control. Consequently, the research does not aim to create a definitive detection model. Rather, it explores whether design interventions can support more careful and situated forms of awareness around relational harm.

3.2 Conceptual Framework of the Guardian System

Recent discussions surrounding AI-supported anti-bullying systems have largely focused on event detection, automated reporting, and behavioural monitoring. However, many forms of bullying do not emerge as isolated incidents, but rather through subtle relational patterns that develop gradually over time. Experiences such as exclusion, silence, withdrawal, or long-term social marginalisation are often difficult to capture through conventional detection-based systems. This limitation raises a broader methodological and design question: how might AI systems engage with forms of harm that are structural, relational, and not always directly expressible?

To explore this question, the present study develops the Guardian System as a speculative and reflective design framework. Rather than positioning AI as a tool for automated judgement, the project investigates how design might support more reflective forms of interpretation and care within school environments. In this context, the system is approached not as a finished technological solution, but as a research artefact through which questions surrounding visibility, silence, and structural intervention can be explored.

The conceptual framework presented below illustrates the overall logic guiding the research design. The framework proposes a shift away from event-based detection models toward a more relational and contextual understanding of school interactions. Instead of focusing solely on identifiable incidents, the system attempts to attend to long-term interaction patterns, behavioural traces, and subtle changes within social relationships. Importantly, the framework also

emphasises ambiguity and human interpretation, recognising that experiences of harm cannot always be reduced to fixed categories or automatically classified signals.




Figure 1. Conceptual framework of the Guardian System, illustrating a shift from event-based AI detection toward relational, reflective, and structurally informed approaches to support.

As shown in Figure 1, the proposed system does not aim to replace human judgement or automate disciplinary intervention. Instead, it functions as a reflective layer that may help teachers, counsellors, or support staff notice emerging relational imbalances that might otherwise remain invisible within everyday school life. This distinction is central to the methodological position of the study. The project is not concerned with building a predictive surveillance system, but with exploring how design practice might make certain structural conditions more discussable, interpretable, and collectively visible.

The framework therefore serves not merely as an illustration of the proposed system, but also as a visual articulation of the study's theoretical position. It brings together several themes developed throughout the literature review, including structural violence, tacit knowledge, trauma-related silence, and reflective practice, while also providing a conceptual bridge into the methodological framework discussed in the following sections.

3.3 Research Design Overview

The research process is organised around three interconnected and iterative stages: problem reframing, prototype exploration, and situated evaluation. These stages are not intended to function as a strictly linear sequence. Instead, findings from one stage may repeatedly reshape assumptions developed in another. New ethical concerns, participant interpretations, or design failures may lead to revisions of earlier concepts and research directions.

The first stage focuses on reframing the problem space. Building on the literature review, the research examines how silence, relational exclusion, and structural violence appear within everyday school interactions. At this stage, the intention is not to produce a definitive categorisation of bullying behaviours, but to identify recurring tensions, ambiguities, and overlooked experiences that may not be adequately addressed through existing intervention models. Early exploratory discussions and contextual observations help generate initial design questions regarding how patterns of exclusion might become perceptible without reducing them to simplified indicators.

The second stage involves the development of a series of low-fidelity and medium-fidelity prototypes. These prototypes are not treated as final products or experimental instruments designed to verify fixed hypotheses. Instead, they function as exploratory artefacts that make particular assumptions visible and open to critique. Some prototypes focus on visualising long-term interaction

patterns, while others experiment with indirect ways of representing silence, withdrawal, or relational imbalance. The purpose of prototyping is therefore not only technical development, but also conceptual investigation. Through the process of making and revising, the research examines which forms of representation become meaningful, problematic, uncomfortable, or ethically questionable.

The third stage focuses on situated evaluation through participatory discussion, scenario-based interaction, and reflective interpretation. Rather than measuring success through detection accuracy alone, the study investigates how participants interpret the system, how meanings shift during interaction, and what kinds of social effects the system might produce. Particular attention is paid to moments of misunderstanding, discomfort, resistance, or reinterpretation, as these moments often reveal underlying assumptions embedded within the design itself.

Across all three stages, the research remains open to uncertainty and revision. The study does not assume that structural harm can be fully captured through technological systems. Instead, the design process is used to explore the limits, tensions, and possibilities involved in attempting to engage with such experiences.

3.4 Co-Design as Situated Inquiry

Co-design methods are introduced in order to better understand how participants perceive exclusion, discomfort, silence, and relational tension within everyday school life. This approach is particularly important because the research does not assume that participants can always directly articulate experiences of bullying through straightforward questioning.

Participants will primarily include university students and individuals with relevant school experiences. Given the sensitivity of the topic, workshops will avoid directly asking participants whether they have been bullied. Instead, indirect and situation-based activities will be used to encourage discussion around everyday experiences of discomfort, avoidance, exclusion, or silence.

For example, participants may be invited to describe situations in which they chose not to speak, moments where they felt socially invisible, or circumstances where group interactions became uncomfortable without explicit conflict occurring.

This indirect approach is methodologically important because it acknowledges the limits of explicit reporting identified in the literature review. Rather than treating silence as missing data, the workshops attempt to create conditions in which tacit or emotionally difficult experiences may emerge indirectly through discussion, storytelling, role-play, speculative scenarios, or interaction mapping.

Co-design is also used to investigate how participants interpret potential technological interventions. Participants may discuss different forms of data representation, system feedback, or interaction visualisation. These discussions are intended not only to improve usability, but also to examine how participants negotiate questions of trust, discomfort, visibility, and surveillance. In some cases, participants may reject certain forms of system intervention entirely, and such rejection is treated as valuable research insight rather than design failure.

Importantly, co-design is not understood as a neutral method for extracting user requirements. Participants do not simply provide information that is then translated into design solutions. Instead, the workshops function as spaces

where meanings are collaboratively negotiated, challenged, and reformulated. This aligns with the broader practice-based position of the study, in which understanding is generated through interaction rather than discovered in advance.

3.5 Prototype Development and Iterative Exploration

The prototype phase explores three interconnected directions: relational pattern visualisation, silent behavioural signals, and structural representation. The first direction focuses on relational patterns rather than isolated incidents.

Existing anti-bullying systems often rely on identifying discrete events such as insults, threats, or physical aggression. However, as discussed in the literature review, structural forms of exclusion frequently emerge through repetition, accumulation, and long-term relational positioning rather than through singular events. In response, the prototypes experiment with visualising interaction frequency, participation distribution, and changes in social positioning over time. For example, interaction networks may indicate which individuals consistently remain peripheral within group communication structures.

The second direction explores how systems might respond to silence without requiring explicit self-report. Instead of recording conversational content directly, some prototypes focus on behavioural changes such as declining participation, reduced response frequency, or prolonged withdrawal from interaction spaces. These signals are not interpreted as objective indicators of victimisation. Rather, they are approached as contextual traces that may prompt reflection or discussion.

The third direction concerns structural representation. Instead of centring intervention solely on individuals, the prototypes explore ways of presenting group-level dynamics and collective interaction trends. This may include visual

comparisons between participation distributions across groups or representations of recurring asymmetries within communication patterns. The intention is not to identify problem individuals, but to make broader relational structures more visible.

Each iteration is expected to generate new tensions and uncertainties. Some forms of visualisation may unintentionally stigmatise individuals. Certain feedback mechanisms may create anxiety rather than support. Participants may interpret system outputs in ways that differ significantly from the researcher's assumptions. These moments are treated as central to the research process because they reveal how technological interpretation itself can become socially consequential.

As a result, iteration is not framed as a process of gradually optimising a stable system. Instead, it functions as a continuous re-examination of the assumptions embedded within the design.

3.6 Situated Evaluation: From Effectiveness to Interpretation

Conventional AI system evaluation often prioritises performance metrics such as classification accuracy, detection rate, or predictive reliability. While these measures may be useful within controlled computational settings, they are insufficient for addressing the concerns explored in this study.

This research therefore adopts a situated evaluation approach that focuses less on whether the system is correct and more on how the system is interpreted, negotiated, and socially experienced within specific contexts.

Scenario-based sessions and semi-structured discussions will be used to examine participant responses to different forms of system output. Participants may be shown hypothetical interaction patterns or visualisations and invited to

discuss what these representations suggest, how they might respond, and whether such representations feel supportive, intrusive, ambiguous, or misleading.

Particular attention will be paid to ambiguity and disagreement. If participants interpret the same visualisation in conflicting ways, this is treated not as methodological failure but as evidence of the inherently interpretive nature of relational data. The study therefore avoids assuming that behavioural patterns possess fixed meanings independent of social context.

The evaluation also considers broader ethical implications. Attempts to infer vulnerability through behavioural traces may unintentionally reproduce forms of monitoring or institutional surveillance. There is also a risk that certain visualisations could reinforce stigma, encourage overinterpretation, or shift responsibility onto already vulnerable individuals. For this reason, the study prioritises aggregated and de-identified representations where possible, and participants are informed about how interaction data is collected, interpreted, and discussed throughout the research process.

3.7 Practice as Data

A central methodological position of this study is that design practice itself constitutes a form of research data. The prototypes produced during the project are not treated merely as outcomes or demonstrations, but as material traces of evolving assumptions, interpretations, and conceptual shifts.

Data is therefore generated throughout the design process rather than only at the point of evaluation. This includes workshop discussions, participant reactions, observational notes, prototype revisions, design sketches, reflective journals, records of design decisions, and moments where assumptions are challenged or reformulated.

Special attention is given to moments of friction, uncertainty, or reinterpretation. For example, if participants perceive a visualisation as judgmental rather than supportive, this reaction becomes analytically significant because it reveals tensions between care and surveillance within the design logic. Similarly, when certain assumptions repeatedly fail during testing, these failures are documented as evidence of the limits of particular approaches.

The study therefore analyses not only what is designed, but also how understandings change through the process of designing. Iterations are compared across stages in order to examine how concepts such as silence, relational harm, and structural intervention become reframed over time.

This position differs from approaches in which design simply implements pre-existing theory. Instead, the act of making becomes a way of thinking through uncertainty, revealing hidden assumptions, and generating situated forms of knowledge that may not emerge through detached observation alone.

3.8 Data Analysis and Reflexive Interpretation

The research primarily adopts an inductive and reflexive approach to analysis. Rather than beginning with fixed coding categories, themes are expected to emerge gradually through repeated engagement with workshop materials, prototype interactions, field notes, and reflective documentation.

Alongside thematic analysis, reflexive analysis is used to examine how the researcher's own assumptions shape the design process. This is particularly important because the study deals with vulnerable experiences and interpretive technologies. Decisions regarding what counts as withdrawal, silence, or risk are never entirely neutral, and the research acknowledges that design interventions inevitably carry implicit values and interpretations.

Reflective journals will therefore document moments of uncertainty, disagreement, ethical tension, and conceptual revision throughout the project. These reflections are analysed alongside participant responses rather than treated as separate from the research process.

The study does not aim to produce universally generalisable findings. Instead, it seeks to generate situated insights into how structural harm may become partially visible through design practice, while also recognising the limitations and ambiguities involved in such attempts.

3.9 Limitations of the Approach

Several limitations shape this research approach. First, the findings are context-dependent and may not be directly transferable across different educational or cultural settings. Experiences of bullying, silence, and exclusion are shaped by highly specific institutional and relational conditions, meaning that the prototypes developed in this study cannot be assumed to function universally.

Second, behavioural traces remain inherently ambiguous. Reduced participation or social withdrawal may relate to many factors unrelated to bullying, including personality, stress, fatigue, or broader social circumstances. The study therefore does not claim that behavioural patterns can objectively identify victimisation.

Third, workshop environments and speculative scenarios cannot fully reproduce the complexity of real school dynamics. Participants may respond differently in actual situations involving peer pressure, institutional authority, or emotional vulnerability.

Finally, the study recognises that attempts to make structural harm visible may themselves produce new forms of visibility and surveillance. The research therefore remains cautious about treating technological interpretation as neutral or universally beneficial. Rather than resolving these tensions completely, the project treats them as part of the inquiry itself.

3.10 Methodological Summary

This research adopts a practice-based and reflective methodology in order to investigate forms of bullying that are often structurally embedded, tacit, and difficult to articulate directly. Through co-design, iterative prototyping, and situated evaluation, the study attempts to move beyond event-based detection models toward a more relational understanding of harm.

Rather than seeking a definitive technological solution, the research uses design practice as a way of exploring how silence, exclusion, and structural patterns might become partially perceptible within specific contexts. In doing so, the project positions technology not simply as a detection tool, but as a medium through which social relationships, assumptions, and institutional structures may be questioned, negotiated, and reinterpreted.

The following chapter will further discuss the ethical framework and practical implementation of the Guardian System, including participant recruitment, workshop organisation, and prototype deployment strategies.

4. Ethical Considerations

4.1 Ethics Beyond Procedural Compliance

This research engages with experiences related to bullying, exclusion, silence, and vulnerability, and therefore raises a number of ethical questions that extend beyond procedural approval or data protection alone. Because the project involves both participatory design activities and the interpretation of behavioural patterns, ethics in this study is understood not simply as a formal requirement, but as an ongoing reflexive process embedded within the research itself. In particular, the study recognises that attempts to make structural harm visible may also unintentionally produce new forms of visibility, interpretation, and power.

4.2 Researcher Positionality and Reflexivity

As a researcher working within the fields of design and emerging technology, my position inevitably shapes how problems are framed, what kinds of behaviours become visible, and which forms of experience are prioritised within the design process. This project is informed by an awareness that technological systems are never neutral tools. Decisions regarding what counts as withdrawal, risk, or abnormal behaviour are shaped by social assumptions, cultural expectations, and design choices.

The research therefore does not claim to objectively identify victimisation or emotional states. Instead, behavioural traces and interaction patterns are approached as partial and context-dependent signals whose meanings remain open to interpretation. Reflexive documentation will be maintained throughout the project in order to critically examine how my own assumptions influence prototype development, participant interaction, and data interpretation. In addition, the study acknowledges that my perspective as a researcher may differ significantly from the lived experiences of participants. Co-design

activities are therefore used not only to gather feedback, but also to challenge the researcher's assumptions and reveal perspectives that may otherwise remain overlooked.

4.3 Participants, Consent, and Emotional Safety

Because the study involves discussions related to bullying and exclusion, participant wellbeing is treated as a central ethical concern. Recruitment will focus primarily on adult participants, including university students and individuals with relevant school experiences, in order to avoid placing unnecessary pressure on minors during the early stages of the research. Participation in workshops, interviews, and prototype testing will be entirely voluntary. Before taking part, participants will receive clear information regarding the aims of the project, the nature of the activities involved, how data will be used, and their right to withdraw from the study at any stage without explanation or penalty.

Given the emotional sensitivity of the topic, the research avoids directly asking participants to recount traumatic experiences in detail. Instead, workshops will rely on indirect and scenario-based activities, such as discussing hypothetical situations, interaction patterns, or feelings of discomfort within group environments. Participants will also be free to skip questions or pause participation at any point if they feel uncomfortable.

All recorded materials will be anonymised during transcription and analysis. Identifiable information will be removed wherever possible, and digital files will be securely stored using password-protected systems in accordance with university research ethics guidelines.

4.4 Data Interpretation and Behavioural Representation

The project also raises ethical questions regarding how behavioural data is interpreted and represented. Since the research explores interaction patterns, participation frequency, and forms of silence, there is a risk that participants could feel monitored, categorised, or misrepresented by the system.

To address this concern, the study avoids treating behavioural traces as objective evidence of bullying. Reduced interaction or social withdrawal may reflect many different personal or contextual factors, and the research explicitly recognises the ambiguity of such signals. Prototype outputs will therefore prioritise aggregated or relational representations rather than individualised judgement or automated labelling.

Similarly, visualisations produced during the project will be approached as interpretive tools rather than factual diagnoses. Participants will be encouraged to critically discuss and question the meanings suggested by the system rather than passively accept them as accurate representations of reality.

4.5 Surveillance, Care, and Ongoing Ethical Reflection

The research further recognises that technological systems designed for care can unintentionally reproduce surveillance logics. As a result, ethical reflection is treated not as a separate stage conducted after design decisions are made, but as a continuous part of the design process itself. Moments where participants experience discomfort, resistance, or concern toward the system will therefore be treated as important research findings rather than problems to be eliminated.

Ultimately, the study does not aim to create a fully reliable system for identifying hidden harm. Instead, it seeks to explore how design practices might support more careful, situated, and ethically aware ways of engaging with experiences that are often difficult to articulate or recognise.

5. Reflection and Conclusion

5.1 Revisiting the Research Problem

This research began with a concern that many current AI anti-bullying systems remain centred on identifiable events and explicit victim reporting, while many experiences of bullying are embedded within longer-term social patterns that are difficult to articulate or recognise. Across the previous chapters, the study argued that existing approaches often struggle to address three interconnected gaps: the gap between isolated incidents and relational patterns, the gap between explicit reports and silent experiences, and the gap between individual intervention and structural understanding. Through a practice-based research approach centred on the iterative development of the Guardian System, this project explored whether design practice could offer another way of approaching these problems.

5.2 Design Practice as Inquiry

Rather than attempting to build a definitive detection tool, the research treated the design process itself as a way of investigating how structural violence might become partially visible within everyday school interactions. In this sense, the project does not claim to solve bullying through technology. Instead, it suggests that AI systems may function less as instruments of automated judgement and more as reflective infrastructures that help reveal patterns which would otherwise remain unnoticed.

One important finding emerging from the research is that harmful experiences are often expressed indirectly rather than explicitly. During co-design workshops and prototype testing, participants rarely described bullying through direct labels such as violence or abuse. Instead, they spoke through situations, atmospheres, and recurring feelings: being excluded from group interactions, repeatedly ignored in collaborative tasks, or becoming hesitant to speak within

particular environments. These observations reinforced the argument developed earlier in the dissertation: silence should not simply be interpreted as absence, but as a meaningful condition shaped by emotional, social, and institutional contexts.

5.3 Reflective Making and Tacit Experience

The practice-based methodology was particularly valuable in making these dynamics visible. In several workshops, participants initially struggled to explain uncomfortable situations verbally, yet became more expressive when responding to speculative scenarios, visual diagrams, or prototype interactions. This suggests that design artefacts can operate not only as outputs, but also as mediating tools that allow tacit experiences to surface indirectly. In this respect, the research supports the idea that knowledge can emerge through reflective action rather than existing fully in advance of practice.

5.4 Tensions Within Visibility and Interpretation

At the same time, the iterative design process revealed important tensions and limitations. While the system attempted to visualise long-term interaction patterns rather than isolated incidents, some participants interpreted these visualisations as forms of judgement or ranking. Network diagrams intended to show relational imbalance occasionally produced anxiety among users who feared being labelled as isolated or problematic. This complicated the assumption that making patterns visible is necessarily beneficial.

The project also revealed limitations in the idea of silent data. Behavioural traces such as reduced participation or prolonged social withdrawal may indicate distress, but they do not necessarily correspond directly to bullying experiences. Silence can emerge for many reasons, including personality, emotional state, or cultural communication styles. As a result, the research gradually shifted away from the idea of automated interpretation toward the

idea of contextual support. Rather than treating behavioural signals as evidence, the later prototypes explored how systems might encourage reflection and dialogue without making definitive claims about individual experience.

5.5 Methodological and Practical Limitations

These tensions highlight both the value and the limitations of practice-based research. The methodology allowed the project to engage with ambiguity in ways that more rigid empirical approaches may struggle to accommodate. Because the prototypes evolved through iterative testing and reflection, the research was able to respond to unexpected reactions and revise its assumptions throughout the process. At the same time, the study remained limited in scale and did not involve long-term deployment within real school environments. Institutional factors such as school policy, teacher interpretation, and long-term behavioural effects therefore remain outside the scope of this project.

5.6 Ethical Implications and Future Research

There are also ethical questions that remain unresolved. Although the research attempted to minimise surveillance by avoiding direct content collection and prioritising aggregated interaction patterns, the boundary between supportive monitoring and institutional observation remains difficult to define. A system intended to identify vulnerability may also unintentionally reinforce behavioural regulation or social pressure. This suggests that future work should engage more directly with questions of power, consent, and institutional accountability.

Future research could therefore explore longer-term participatory collaboration with schools, counsellors, and students across different educational settings. In particular, further work is needed to examine how such systems are

interpreted within everyday practice, how students negotiate their presence, and how design might support dialogue without increasing surveillance.

5.7 Final Reflection and Contribution

Ultimately, this research does not conclude that technology can fully recognise hidden suffering. Instead, it argues that practice-based design can help expose the limits of current event-based approaches while opening alternative ways of understanding vulnerability, silence, and relational harm. In this sense, the contribution of the project lies less in proposing a finished technological solution than in reframing what anti-bullying technologies are designed to notice, interpret, and respond to in the first place.

References

1. Anti-Bullying Alliance (2023) *Bullying and wellbeing research findings*. London: Anti-Bullying Alliance. Available at: <https://anti-bullyingalliance.org.uk/> (Accessed: 29th April 2026).
2. Bourdieu, P. (1986) 'The forms of capital', in Richardson, J. (ed.) *Handbook of Theory and Research for the Sociology of Education*. New York: Greenwood Press, pp. 241–258. Available at: <https://web.stanford.edu/~eckert/PDF/Bourdieu1986.pdf>
3. Chandrasekharan, E., Pavalanathan, U., Srinivasan, A., Glynn, A., Eisenstein, J. and Gilbert, E. (2017) 'You can't stay here: The efficacy of Reddit's 2015 ban examined through hate speech', *Proceedings of the ACM on Human-Computer Interaction*, 1(CSCW), pp. 1–22. Available at: <https://doi.org/10.1145/3134666>
4. Crick, N.R. and Grotpeter, J.K. (1995) 'Relational aggression, gender, and social-psychological adjustment', *Child Development*, 66(3), pp. 710–722. Available at: <https://doi.org/10.2307/1131945>
5. Dinakar, K., Reichart, R. and Lieberman, H. (2021). Modeling the Detection of Textual Cyberbullying. *Proceedings of the International AAAI Conference on Web and Social Media*, 5(3), pp.11–17. Available at: <https://doi.org/10.1609/icwsm.v5i3.14209>
6. Dourish, P. (2006) 'Implications for design', in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. New York: ACM, pp. 541–550. Available at: <https://doi.org/10.1145/1124772.1124855>

7. Espelage, D.L. and Swearer, S.M. (2003) 'Research on school bullying and victimization: What have we learned and where do we go from here?', *School Psychology Review*, 32(3), pp. 365–383. Available at: <https://doi.org/10.1080/02796015.2003.12086206>
8. Foucault, M. (1977) *Discipline and Punish: The Birth of the Prison*. London: Penguin Books, pp. 291–296. Available at: https://monoskop.org/images/4/43/Foucault_Michel_Discipline_and_Punish_The_Birth_of_the_Prison_1977_1995.pdf
9. Fuggle, P., Dunsmuir, S. and Curry, V. (2013) 'CBT with children, young people and families', London: Sage, pp. 194–198. Available at: <https://doi.org/10.4135/9781473914858>
10. Galtung, J. (1969) 'Violence, peace, and peace research', *Journal of Peace Research*, 6(3), pp. 167–191. Available at: <https://doi.org/10.1177/002234336900600301>
11. Hee, C.V., Jacobs, G., Emmery, C., Desmet, B., Lefever, E., Verhoeven, B., De Pauw, G., Daelemans, W. and Hoste, V. (2018) 'Automatic detection of cyberbullying in social media text', *PLOS ONE*, 13(10), pp. 79–83. Available at: <https://doi.org/10.1371/journal.pone.0203794>
12. Herman, J.L. (1998). Recovery from psychological trauma. *Psychiatry and Clinical Neurosciences*, 52(S1), pp.S98–S103. Available at: <https://doi.org/10.1046/j.1440-1819.1998.0520s5s145.x>
13. Hinduja, S. and Patchin, J.W. (2020). Digital Dating Abuse Among a National Sample of U.S. Youth. *Journal of Interpersonal Violence*, 36(23-24), pp. 286–293. Available at: <https://doi.org/10.1177/0886260519897344>

14. Morgan, H.J. and Janoff-Bulman, R. (1994) 'Positive and negative responses to traumatic loss: An exploration of meaning and coping', *Journal of Social and Clinical Psychology*, 13(4), pp. 334–355. Available at: <https://doi.org/10.1521/jscp.1994.13.4.334>
15. Nobata, C., Tetreault, J., Thomas, A., Mehdad, Y. and Chang, Y. (2016) 'Abusive language detection in online user content', in *Proceedings of the 25th International Conference on World Wide Web*. Geneva: International World Wide Web Conferences Steering Committee, pp. 145–153. Available at: <https://doi.org/10.1145/2872427.2883062>
16. Omarov, B., Nurbekova, Z., Suliman, A. and others (2022) 'Violence detection using computer vision and deep learning: A systematic review', *PeerJ Computer Science*, 8, pp. 10–13. Available at: <https://doi.org/10.7717/peerj-cs.1096>
17. Pennebaker, J.W. (1997) *Opening Up: The Healing Power of Expressing Emotions*. New York: Guilford Press, pp. 41–43. Available at: <http://nwkpsych.rutgers.edu/~kharber/healthpsychology/READINGS%202024/Pennebaker.%20Opening%20Up.pdf>
18. Polanyi, M. (1966). The Logic of Tacit Inference. *Philosophy*, 41(155), pp.1–18. Available at: <https://doi.org/10.1017/s0031819100066110>
19. Schön, D.A. (1983) *The Reflective Practitioner: How Professionals Think in Action*. New York: Basic Books, pp. 1–8. Available at: <https://studenthublive.open.ac.uk/sites/studenthublive.open.ac.uk/files/reflective%20practitioner%20-%20schon.pdf>

20. Smith, P.K., Mahdavi, J., Carvalho, M., Fisher, S., Russell, S. and Tippett, N. (2008) 'Cyberbullying: Its nature and impact in secondary school pupils', *Journal of Child Psychology and Psychiatry*, 49(4), pp. 376–385. Available at: <https://doi.org/10.1111/j.1469-7610.2007.01846.x>

21. UNESCO (2019). *Behind the numbers: Ending school violence and bullying*. Available at: <https://www.unicef.org/media/66496/file/Behind-the-Numbers.pdf>

22. van der Kolk, B. (2014). *The body keeps the score : Brain, mind, and body in the healing of trauma*. New York: Viking, pp. 464. Available at: <https://ia601604.us.archive.org/35/items/the-body-keeps-the-score-pdf/The-Body-Keeps-the-Score-PDF.pdf>