

Linköping Studies in Science and Technology  
Dissertation No. 2542

# Team Intentionality by Design

A Framework for Joint Human-Technology Activity

**Jimmy Hammarbäck**



Linköping Studies in Science and Technology  
Dissertations No. 2542

# Team Intentionality by Design

A Framework for Joint Human–Technology Activity

Jimmy Hammarbäck



Department of Science and Technology, Division of Media Information  
and Technology  
Faculty of Science and Engineering  
Linköpings universitet, SE-581 83 Linköping, Sweden  
Linköping 2026

Supervisor:  
Jonas Lundberg  
Professor  
Department of Science and Technology

Co-supervisor:  
Björn Johansson  
Professor  
Department of Computer and Information Science

Co-supervisor:  
Jens Alfredson  
PhD  
Saab AB

Faculty opponent:  
Johan Bergström  
Professor  
Lund University

© Jimmy Hammarbäck, 2026

 Except where otherwise noted, this work is licensed under a Creative Commons Attribution 4.0 International license. To view a copy of this license, visit:

<https://creativecommons.org/licenses/by/4.0/>

Printed in Sweden by LiU-tryck, 2026

ISBN 978-91-8118-658-1 (print)  
ISBN 978-91-8118-659-8 (PDF)  
<https://doi.org/10.3384/9789181186598>  
ISSN 0345-7524

# Abstract

Advances in artificial intelligence and autonomy are reshaping how humans and technology work together. Increasingly, technology is understood not only as a tool, but as a teammate, reflecting an emerging hybrid teaming paradigm. This transition is exemplified in military aviation through Manned–Unmanned Teaming, where fighter pilots may operate alongside unmanned aircraft in joint activities. Designing for such hybrid team activities, however, is a wicked problem. It demands the integration of the human, the technology, and their relation around the phenomenon of team intentionality: what is going on, and what are we going to do about it. Manned–Unmanned Teaming sharpens this difficulty under conditions of time pressure, task difficulty, and risk exposure. However, the conceptual basis for designing for team intentionality remains fragmented across disciplines. This makes it difficult for researchers and practitioners to determine which intentional aspects of hybrid teaming to design for, and how. To this end, this thesis develops a practical framework of team intentionality. Drawing on team science and intentionality research, the framework is grounded in a synthesis of established theory and empirical insight. Four included papers address how intent can be understood, modelled, and designed for in hybrid teaming. They substantiate the framework theoretically, methodologically, and empirically, situating this in the Manned–Unmanned Teaming context under complex and dynamic conditions. The framework establishes concepts and relations around intentional states and processes at different levels of agency. It serves as an analytical and evaluative lens that connects the hybrid teaming paradigm to the design process. Applied across that process, it foregrounds the understanding and sharing of intent as essential to such design, identifying the aspects that members must align. Theoretically, it introduces team intentionality as a distinct construct within the hybrid teaming paradigm. In practice, it helps determine which intentional aspects to design for and how, providing a shared vocabulary across stakeholders and stages of the design process.



# Sammanfattning

Framsteg inom artificiell intelligens och autonomi förändrar hur människor och teknik arbetar tillsammans. I allt högre grad förstås teknik inte bara som ett verktyg, utan som en medaktör, vilket speglar ett framväxande paradigm kring hybrida team. Denna övergång exemplifieras inom militär flygverksamhet genom bemannat–obemannat teamarbete, där stridspiloter förväntas kunna arbeta tillsammans med obemannade flygfarkoster i gemensamma aktiviteter. Att designa för sådana hybrida teamaktiviteter är dock ett svårlöst problem. Det kräver integration av människan, tekniken och deras relation kring fenomenet teamintentionalitet: vad är det som pågår, och vad ska vi göra åt det. Bemannat–obemannat teamarbete skärper denna svårighet under tidspress, hög uppgiftskomplexitet och riskfyllda förhållanden. Emellertid är den begreppsliga grunden för att designa för teamintentionalitet splittrad mellan olika forskningsfält. Det gör det svårt för forskare och praktiker att avgöra vilka intentionella aspekter av hybrida team som bör designas för, och hur. I detta syfte utvecklar denna avhandling ett praktiskt ramverk för teamintentionalitet. Med utgångspunkt i forskning om team och intentionalitet grundas ramverket i en syntes av etablerad teori och empirisk kunskap. Fyra inkluderade artiklar behandlar hur intentioner kan förstås, modelleras och designas för i hybrida team. De understödjer ramverket teoretiskt, metodologiskt och empiriskt, och situerar detta i bemannat–obemannat teamarbete under komplexa och dynamiska förhållanden. Ramverket etablerar begrepp och relationer kring intentionella tillstånd och processer på olika nivåer av agens. Det fungerar som en analytisk och utvärderande lins som kopplar paradigmet kring hybrida team till designprocessen. Tillämpat genom hela processen lyfter det fram förståelsen och delningen av intentioner som väsentliga för sådan design, och identifierar vilka aspekter medlemmarna behöver samordna. Teoretiskt introducerar det teamintentionalitet som ett eget begrepp inom paradigmet kring hybrida team. Praktiskt hjälper det till att avgöra vilka intentionella aspekter som bör designas för, och hur, och tillhandahåller ett gemensamt vokabulär för olika aktörer och faser i designprocessen.



# Preface

Once upon a time, I meant to write a book. It had two sides: one story told from two points of view, each seeing part of it, neither seeing the whole. I never got past meaning to. This is not that book, though it is, I have come to see, the same problem asked in earnest.

The story that actually happened starts later. Not long after the Big Bang, give or take, I was working in industry, with no plan to leave. Then a book on research philosophy fell into my hands, and I made the mistake of finding it fun. Fun turned into study, philosophy and then psychology, alongside the job, for no better reason than that each answer kept opening another question. Then cognitive science. Then a doctoral post. You know what they say about curiosity and boxes.

None held a cat, alive or dead, and none was quite Pandora's. Each just showed me something worth a look, and then, always, one more worth a look after that. That is the whole of my method, if it deserves the name: the next thing looked interesting, so I opened it, and never turned round to see what was piling up behind me. Individual intent looked worth a look, so I opened it. Inside was a team, which turned out to have an intentionality of its own, not just intent added up, so I opened that too. Each step was sensible. It was only ever the sum that got out of hand, and by the time the heap behind me had quietly become a neighbourhood, with cross-streets, a drainage problem, and its own weather, I was already reaching for the next box.

Which is where the doctorate had its way with me, and not gently. Opening things, it turns out, is the smallest part of the work, and the one part I turned up already able to do. Everything else simply required me to become. A presenter, which I dreaded. An editor of my own worst sentences, a keeper of references, a maker of figures, a defender of conclusions I had only lately reached, a reopener of everything I had thought settled, on the advice of Reviewer 2, and a good many other trades I do not recall applying for. Curiosity got the lid off. It was no help whatever with anything inside.

For a long time I kept track of all this the way I always had, a note here and there for whatever seemed worth keeping. One to a box, that served me well enough. It is no way to keep a place. A note can tell you a

place is there, but not where it stands against the others, which way its streets run, what it shares a wall with. For that you need a map, and a map you draw from inside it, one street at a time, never once seeing the whole while you are down there among it. So the map came on slowly, always a little behind the walking, never quite in time to save me the next wrong turn. What I can hand you is the map, not the walking. It is not much to look at. Some quarters are tidier than others, and a street or two plainly went in on a bad week. But it is honest, and it will get you about, which is more than it ever managed for me in the years I spent laying it.

That map is this thesis, more or less. Which is to say it is the book after all, or near enough: two points of view, one human and one not, each seeing part of a situation and neither the whole of it. That was never a flaw to be fixed. It is the reason there is a thesis at all, because two partial views do not add up on their own, and making them work together is the whole of the problem. Muddier than the book I meant to write, wearing far more theory than it set out to, but recognisably the same honest question. Not whose view is right, but how two of us, seeing differently, ever come to understand what is going on between us, and what to do about it.

Jimmy Hammarbäck  
*Linköping, August 2026*

# Acknowledgements

This thesis began long before the doctoral post, and so do the debts. I am grateful to everyone who was there along the way: those who inspired me to start, those who thought I could and said I should, and those who made it possible to keep going. Many who shaped this work will not find their names below, and I hope they know the omission is one of space, not of gratitude.

For making the work possible, I thank its funders, Vinnova and the Swedish Defence Materiel Administration. The work was carried out at Saab AB, where I was employed throughout, and which was home to this thesis in every practical sense. My thanks to Linköping University for the academic side, and for being the scholarly counterweight to industry. Much of what the thesis became came out of the tension between the two, and I am the better for having had a foot in each.

I owe most to my supervisors, each of whom taught me a different way to find my way. Jonas Lundberg, my main supervisor, taught me to look for the crux, to keep asking what a thing was really about. Jens Alfredson, at Saab, showed me the terrain, pointing across the ground toward what to head for and what to avoid. Björn Johansson caught the words I tried to use before I had reckoned with the baggage they carried. Between them, they let me walk the way I wanted, which was better than the way I would have found alone.

Last, and most, my family. Over these years we went from two to three, and much of what I was learning had nothing to do with research: how to be a father to Cornelia, and the keeper of an old house with opinions of its own, all while spending my working days on things that had not happened yet. It was more than I could have managed alone. What carried me through was them, close family and far, and Madeleine above all, who held the ordinary things together while I was off somewhere in the future. I could not have done it without them, and I did not.



# Contents

|                                                   |      |
|---------------------------------------------------|------|
| Abstract.....                                     | iii  |
| Sammanfattning.....                               | v    |
| Preface .....                                     | vii  |
| Acknowledgements .....                            | ix   |
| List of Figures.....                              | xv   |
| List of Tables .....                              | xvii |
| List of Papers.....                               | xix  |
| Thesis Outline .....                              | xxi  |
| Introduction .....                                | 1    |
| Background .....                                  | 1    |
| Purpose, Aim, and Research Questions .....        | 4    |
| Scope and Delimitations .....                     | 4    |
| Research Positing.....                            | 5    |
| Grounding Team Intentionality .....               | 7    |
| Teams and Teaming.....                            | 7    |
| <i>Defining Team Structures</i> .....             | 8    |
| <i>Teaming as an Activity</i> .....               | 10   |
| <i>The Evolution of Hybrid Teaming</i> .....      | 12   |
| <i>Foundations in Hybrid Teaming</i> .....        | 17   |
| Individual and Collective Intentionality .....    | 20   |
| <i>Epistemic Aboutness</i> .....                  | 21   |
| <i>Conative Directedness</i> .....                | 23   |
| <i>Shared and Collective Intentionality</i> ..... | 26   |
| Teams and Intentionality .....                    | 27   |
| <i>Structural Foundation</i> .....                | 28   |
| <i>Representational Foundations</i> .....         | 32   |
| <i>Mindreading</i> .....                          | 34   |
| Designing for Hybrid Teaming .....                | 39   |
| Design as Research and Practice .....             | 39   |
| <i>The Nature of Design</i> .....                 | 40   |
| <i>Design Process Stages</i> .....                | 40   |
| <i>Design Science Research</i> .....              | 42   |
| <i>The Integrative Design Process</i> .....       | 43   |
| Sharing Intentionality .....                      | 47   |

|                                              |    |
|----------------------------------------------|----|
| <i>Modelling Intentionality</i> .....        | 47 |
| <i>Transparency and Explainability</i> ..... | 49 |
| The Operational Domain .....                 | 51 |
| Command and Control .....                    | 51 |
| Manned–Unmanned Teaming.....                 | 52 |
| Methodology.....                             | 55 |
| Framework Development.....                   | 55 |
| Studies .....                                | 56 |
| <i>Research Setting</i> .....                | 56 |
| <i>Knowledge Acquisition</i> .....           | 59 |
| Summary of Papers .....                      | 67 |
| Outline .....                                | 67 |
| Paper I.....                                 | 68 |
| <i>Problem and Aim</i> .....                 | 68 |
| <i>Approach</i> .....                        | 69 |
| <i>Findings</i> .....                        | 70 |
| <i>Contributions</i> .....                   | 71 |
| Paper II .....                               | 72 |
| <i>Problem and Aim</i> .....                 | 72 |
| <i>Approach</i> .....                        | 73 |
| <i>Findings</i> .....                        | 73 |
| <i>Contributions</i> .....                   | 75 |
| Paper III.....                               | 76 |
| <i>Problem and Aim</i> .....                 | 76 |
| <i>Approach</i> .....                        | 76 |
| <i>Findings</i> .....                        | 77 |
| <i>Contributions</i> .....                   | 79 |
| Paper IV .....                               | 80 |
| <i>Problem and Aim</i> .....                 | 80 |
| <i>Approach</i> .....                        | 80 |
| <i>Findings</i> .....                        | 82 |
| <i>Contributions</i> .....                   | 83 |
| Team Intentionality by Design.....           | 85 |
| Framing Team Intentionality.....             | 85 |
| <i>The Hybrid Team</i> .....                 | 86 |
| <i>The Strata of Intentionality</i> .....    | 89 |

|                                           |     |
|-------------------------------------------|-----|
| <i>The Situation as Framed</i> .....      | 92  |
| <i>Agency as Regulating</i> .....         | 96  |
| <i>The Agentive Unit</i> .....            | 99  |
| <i>Autonomy as Relations</i> .....        | 101 |
| <i>Intent as Commitment</i> .....         | 104 |
| <i>Aligned Intentionality</i> .....       | 110 |
| <i>The Analytical Account</i> .....       | 113 |
| Realising Team Intentionality .....       | 113 |
| <i>Building a Toolkit</i> .....           | 114 |
| <i>Identifying What to Convey</i> .....   | 120 |
| <i>The Design Account</i> .....           | 128 |
| The Case of Manned–Unmanned Teaming ..... | 129 |
| Limitations and Future Work .....         | 131 |
| Conclusion .....                          | 133 |
| References .....                          | 135 |
| Appendix.....                             | 157 |



# List of Figures

Figure 1. Positioning among many contributing disciplines..... 5

Figure 2. Task interdependence across coordination, cooperation, and collaboration. .... 12

Figure 3. Publications on human–technology relations indexed in Scopus, 1975–2025..... 13

Figure 4. Proportional distribution of terminology used to describe the relationship between Proportional distribution of relational terminology across time intervals. .... 14

Figure 5. Analytical set relations between individually held representations in joint human–technology activity. ....33

Figure 6. Recursive mindreading between members. ....36

Figure 7. Illustration of the integrative design process. ....44

Figure 8. Simulation forms used across the studies, from pen-and-paper (top left) through animated (bottom left) to pilot-in-the-loop (right). ....59

Figure 9. Illustration of a score episode diagram. ....64

Figure 10. The four papers positioned across the integrative design process. ....68

Figure 11. Five structural facets of a hybrid team with example questions. ....86

Figure 12. Interaction mode profile example for shared reconnaissance under normal conditions.....89

Figure 13. The three strata of intentionality within agency. ....90

Figure 14. The frame projecting an ascribed situation against the actual. .... 93

Figure 15. Situations nested by granularity and temporality within a context. .... 95

Figure 16. Agency as a recursive loop of intentional processes coupled with the environment. ....96

Figure 17. An agentive unit configuration within an operational environment.....99

Figure 18. Autonomy boundaries for members  $\alpha$  and  $\beta$ , synthesised into unit  $\Gamma$ . .... 102

Figure 19. An illustrative cross-temporal intent structure..... 105

|                                                                                                                               |     |
|-------------------------------------------------------------------------------------------------------------------------------|-----|
| Figure 20. Paired commitments by status and admissibility. ....                                                               | 108 |
| Figure 21. Specification and enactment of a commitment over time. .                                                           | 109 |
| Figure 22. The alignment plane: awareness and intent positioned by<br>sharedness and world-fit.....                           | 111 |
| Figure 23. Reference points traced across the papers, converging on the<br>intent transparency instantiation in Paper IV..... | 115 |
| Figure 24. The situated intent model and intentional process model as<br>complementary representations.....                   | 116 |
| Figure 25. A situated intent model decomposed into belief and desire.<br>.....                                                | 118 |
| Figure 26. Relations within and between intentional process models.                                                           | 119 |
| Figure 27. Tracing explanatory questions within an intent structure<br>(illustrative).....                                    | 121 |
| Figure 28. Two members' intent structures under a shared team frame.<br>.....                                                 | 124 |
| Figure 29. Autonomy boundary models for members $\alpha$ and $\beta$ ,<br>synthesised into unit $\Gamma$ . ....               | 126 |

# List of Tables

Table 1. Example of team definitions. ....8

Table 2. Example definitions of human–technology relations..... 15

Table 3. Overview of how intent has been conceptualised across domains.  
.....24

Table 4. Comparison of design process stages..... 41

Table 5. Knowledge acquisition methods used in each paper, across the  
three stages of elicitation, analysis, and representation. ....60

Table 6. Aspects of team intentionality a model may represent. ....117

Table 7. Questions for tracing disclosure and clarification within an  
intent structure. .... 122

Table 8. Questions for analysis from the framing of team intentionality.  
.....157

Table 9. Questions for analysis from the realising of team intentionality.  
..... 160



# List of Papers

This compilation thesis is based on the following four papers, hereafter referred to as Papers I–IV.

## Paper I

Hammarbäck, J., Johansson, B. J. E., Alfredson, J., & Lundberg, J. (2024). Modelling situated intent for human-autonomy teaming: A human-centric approach. *Theoretical Issues in Ergonomics Science*, 25(6), 731–752. <https://doi.org/10.1080/1463922X.2024.2337683>

## Paper II

Hammarbäck, J., Alfredson, J., Johansson, B.J.E., & Lundberg, J. (2024). My synthetic wingman must understand me: Modelling intent for future manned–unmanned teaming. *Cognition, Technology & Work*, 26, 107–126 (2024). <https://doi.org/10.1007/s10111-023-00745-3>

## Paper III

Hammarbäck, J., Lundberg, J., Johansson, B. J. E., Alfredson, J. (Submitted). When Things Change: Temporal Dynamics of Intent Formation and Interpretation in Critical MUM-T Situations.

## Paper IV

Hammarbäck, J., Alfredson, J., & Lundberg, J., Johansson, B. J. E. (Submitted). Fast but fragile Intent Interpretations: Effects of Transparency on Awareness and Performance in MUM-T.

## Other Publications by the Author

Hammarbäck, J. (2022). *Modelling intent for Manned-Unmanned Teaming: Exploring human-centric approaches for future combat aircraft systems*. (Licentiate dissertation, Linköping University Electronic Press). <https://doi.org/10.3384/9789179295486>



# Thesis Outline

This thesis is structured as follows. The introduction frames the problem, states the purpose, aim, and research questions, sets the scope, and positions the research. Three chapters then establish the theoretical grounding: team intentionality and its foundations in hybrid teaming, the design perspective through which it is approached, and the operational domain of Manned–Unmanned Teaming. The methodology chapter describes the research approach and the knowledge acquisition underpinning the four papers. The papers are then summarised and positioned across the design process. The framework chapter develops team intentionality by design in two parts, framing and realising, and then reads the case of Manned–Unmanned Teaming through it. A separate chapter sets out the limitations and future work, and the conclusion returns to the research questions and to the question that motivated the thesis.



# Introduction

The introduction presents the problem this thesis sets out to solve: how to design for team intentionality in joint human–technology activity. As technology becomes a teammate rather than a tool, understanding and sharing intent grows mutual, and the conditions of joint activity make it difficult to establish. Situating this problem requires grounding in both the nature of the challenge and the state of the research that has approached it. The remainder is organised in four parts. Background grounds the challenge and surveys the traditions that have addressed it, identifying what remains missing: an integrative framework connecting them to design. Purpose, aim, and research questions set out how this thesis responds, developing a framework of team intentionality and asking what may be central to it and how it can inform design.. Scope and delimitations set the bounds within which its claims hold. Research positioning establishes the stance from which the work proceeds.

## Background

‘I do not understand you. Do you understand me?’ The question is an old human one. For millennia it has passed between people, never to the things they made. From the stone tool to the autonomous system, artefacts have been extensions of human agency, the functional forms through which existing situations are turned into preferred ones (Simon, 1996). In this ancient relationship, intent flowed in one direction: from the maker into the made. Advances in artificial intelligence and autonomy are reshaping that relationship. Technology is increasingly envisioned not as an instrument but as a teammate, with intent of its own to be read, and a need to read ours in return. For the first time, the question passes both ways across the line between the human and the technology as peers in joint activity.

Teams have long been the established way of organising complex work, often characterised by distinct membership, clear boundaries, shared purpose, collective accountability, and activity interdependence (Cohen, 1997; Hackman, 2011; Kozłowski & Ilgen, 2006; Salas et al., 1992). For most of this history, the members have been human and the

technology around them used as a tool. Over the past half-century, however, this perspective has steadily shifted, with technology increasingly being envisioned as an active member of the team. Within this joint human–technology activity, various forms of hybrid teaming have emerged as a new interaction paradigm and research field (Lyons et al., 2021; O’Neill et al., 2022). Such hybrid teaming spans multiple domains (O’Neill et al., 2022); for example industry (Baratta et al., 2023), healthcare (Strong et al., 2026), and defence (El Alami et al., 2023).

A particular case of hybrid teaming is found in military aviation, where fighter pilots are envisioned to operate alongside unmanned aircraft acting as ‘wingmen’ (Jordan, 2021; Stensrud et al., 2024; Taylor & Turpin, 2015). In such Manned–Unmanned Teaming, the pilot and wingman are expected to take on different roles and responsibilities in highly interdependent activities, complementing their respective strengths to achieve operational synergy (Department of Defense, 2018). Central to this vision is a wingman with a high degree of autonomy, acting within its role toward the mission end. Such agency makes understanding and sharing intent essential, in both directions: the pilot must understand the wingman, and the wingman the pilot. This points toward a need for team intentionality: a shared understanding of what the situation is about, and a coordinated focus on what they are trying to achieve.

Designing for such team intentionality is difficult even among humans, where it must be continually established and maintained. In designing future hybrid teaming, it becomes an envisioned world problem that is inherently wicked (Rittel & Webber, 1973; Woods & Dekker, 2000): the teaming does not yet exist, the target moves as it is designed, and the problem resists clear formulation and solution. In Manned–Unmanned Teaming, the problem is sharpened, where the pilot and wingman may operate under high time pressure, task difficulty, and risk exposure, with high stakes and often irreversible consequences. The team is also distributed in space, operating across distances that may be close or great, and cannot rely on communication to maintain team intentionality. They must instead rely on expectation and observation, each reading what the other is doing, why, and what will come next,

rather than communicating directly. Together, these conditions place the weight of teaming on intentionality itself: members must understand and share intent when the means of conveying it are least available. Though acute in Manned–Unmanned Teaming, this challenge runs through hybrid teaming more broadly, wherever humans and technology engage in joint activity within dynamic and complex environments.

Underlying this challenge is a theoretical and methodological difficulty: how to conceptualise team intentionality in joint human–technology activity, identify the aspects that warrant design attention, and translate those aspects into the design process. Several traditions have addressed pieces of this puzzle. For example, team science addresses what this collective unit is, what makes it perform, and how it can be developed (Kozlowski & Ilgen, 2006; Salas et al., 1992). Within philosophy, work on collective intentionality addresses phenomena such as how intentional states extend across members and how joint activity is possible (Bratman, 1992; Searle, 1983; Tuomela, 2007). In psychology, research on ‘mind-reading’ examines how humans attribute mental states to others, laying a foundation for understanding and sharing intent (Frith & Frith, 2006; Premack & Woodruff, 1978; Tomasello et al., 2005). In artificial intelligence, ongoing research aims to enable technology to emulate similar human-like reasoning (Albrecht & Stone, 2018; Tabrez et al., 2020; Van-Horenbeke & Peer, 2021). Within multidisciplinary hybrid teaming research, concepts such as bidirectional transparency and explainability seek to make members more interpretable, comprehensible, and predictable to one another (Bhaskara et al., 2020; Chen et al., 2018; Endsley, 2023). Integrative perspectives such as Human Factors provide practical guidance for analysing, designing, and evaluating complex socio-technical systems, including hybrid teams (Stanton et al., 2017; Wickens et al., 2014). Thus, while much related research exists, what is missing, as argued here, is a framework that brings these contributions together by connecting the conceptual basis for team intentionality to the design of joint human–technology activity.

## Purpose, Aim, and Research Questions

The overarching purpose of this thesis is to support the design of future joint human–technology activity in operational settings. To this end, the aim is to *develop a practical framework of team intentionality* that seeks to establish central aspects and offer guidance for their use across the design process. This aim is addressed through two research questions:

**RQ1.** *What concepts and relations may be central to analysing team intentionality in joint human–technology activity?*

**RQ2.** *How can these concepts and relations be used in the process of designing hybrid teaming?*

The first question supports the development of the framework itself, synthesising multiple perspectives into a structured account. It seeks to provide a unified perspective that makes central concepts and relations of team intentionality analysable and communicable for researchers and practitioners. The second question addresses how the framework can inform design, showing how its concepts and relations may be used across the design process. It seeks to provide guidance for identifying and analysing relevant aspects of team intentionality, so that the framework may serve as an analytical and evaluative lens on design situations. Together, these questions seek to support the overarching purpose by making team intentionality more tractable and traceable in the design of future joint human–technology activity.


## Scope and Delimitations

Three deliberate delimitations bound this work: what the framework covers, where it is grounded, and how it empirically accesses the intentionality phenomenon. First, the framework addresses central aspects of team intentionality rather than exhaustive coverage. The aspects it establishes, while argued to be central, form a bounded account, offering a structured starting point for analysis, design, and

evaluation of joint human–technology activity. Second, the domain of investigation and demonstration is Manned–Unmanned Teaming in military aviation. This domain is a salient and representative instance of hybrid teaming, and the applied setting within which this research was conducted. Its demonstration therefore draws on this domain alone. The framework is nonetheless intended to generalise beyond this setting, both to other forms of hybrid teaming and to joint human–technology activity more broadly. Third, the thesis is primarily grounded in intentionality as envisioned by domain experts rather than as observed in real settings. The studied domain is a future one, where the teaming does not yet exist operationally and intent is therefore accessible as espoused rather than observed in action. It addresses intentional states and processes as they are envisioned to unfold, even though the framework itself does not exclude study of enacted intent. Within these bounds, the thesis offers a principled and structured account of team intentionality for designing joint human–technology activity.

## Research Positing

Given the plurality of disciplines addressing the design of team intentionality in joint human–technology activity, this thesis argues that each provides a piece of the puzzle. However, none holds it all (Figure 1).



*Figure 1. Positioning among many contributing disciplines.*

The framework therefore draws selectively on established traditions, each contributing to a different part. Team science provides the structural basis for what constitutes a collective unit and its activities. Hybrid teaming research examines how such units relate to established team structures and what is expected of them, where intentionality within agency becomes central. Philosophy provides the conceptual basis of intentionality and agency, within the individual and across the collective. Psychology contributes the representational foundations of working together, including how members read one another's minds.

Hybrid teams are sociotechnical systems in which the human and the technology achieve together what neither could alone, and their reasoning must align to do so. Human Factors methods offer the practical orientation for studying such systems. In particular, methods from Cognitive Task Analysis and Cognitive Work Analysis focus on human reasoning and the constraints shaping activity, helping make team intentionality tractable and traceable. To apply the framework, design research and practice contributes an account of the design process and established approaches for supporting team intentionality. It also sets requirements for a rigorous contribution, so the resulting artefacts are both practically grounded and scientifically credible. Working at the intersection of these traditions, the framework developed in this thesis integrates their contributions to support the design of joint human–technology activity.

Throughout, the framework treats the technology as a member with representational and intentional capacity in the operative sense, as an entity whose states and processes can be analysed and designed for using the same concepts as the human member. It does not claim ontological equivalence between human and technological cognition.

# Grounding Team Intentionality

Understanding and designing for team intentionality in joint human–technology activity requires grounding in both the nature of teaming and the intentional character of the mental states that enable it. On that basis, the chapter builds toward team intentionality as a construct in its own right. It is organised in three sections. *Teams and Teaming* establishes the conceptual grounding for hybrid teaming: what teams are, how joint work is done, and the shift from tool to teammate. *Individual and Collective Intentionality* provides the philosophical and psychological foundations for mental states, directed agency, and shared commitment. *Teams and Intentionality* then synthesises these into a preliminary framework for human-technology teams. Together, the three build from established foundations toward a new account of team intentionality.

## Teams and Teaming

Research on traditional human teams has established strong theoretical foundations and practical guidance (Salas et al., 2024). Building on this foundation and recent advancements, new forms of hybrid teaming that integrate humans and technology have become a focal concern (Lyons et al., 2021; O'Neill et al., 2022). This transition has been described as a paradigm shift, in which technology is no longer perceived as a tool but as teammates. To situate this shift, this section establishes a broad conceptual foundation for what it means to work together. Defining team structures sets out what teams are as a specific form of collective. Teaming as an activity examines collaboration, cooperation, and coordination as distinct but interconnected dimensions of joint work. Building on these, *The Evolution of Hybrid Teaming* traces the shift from tool to teammate, and *Foundations in Hybrid Teaming* outlines the research topics this shift has generated. Together, they establish the conceptual grounding for hybrid teaming on which the chapter builds team intentionality.

## Defining Team Structures

To understand how hybrid human-technology teams relate to established forms of working together, there is a need to first clarify what is meant by a team. Team science offers several definitions, typically described as comprising two or more members with distinct roles who interact socially and are interdependent in pursuing a shared goal, while operating within a larger organisational entity (Benishek & Lazzara, 2019). To further examine team structures, Table 1 provides a selection of definitions.

*Table 1. Example of team definitions.*

| Source                           | Team definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| (Salas et al., 1992, p. 4)       | ‘A distinguishable set of two or more people who interact, dynamically, interdependently, and adaptively toward a common and valued goal/objective/mission, and who have each been assigned specific roles or functions to perform, and who have a limited life-span of membership’                                                                                                                                                                                                                             |
| (Kozlowski & Ilgen, 2006, p. 79) | ‘(a) two or more individuals <sup>3</sup> who (b) socially interact (face-to-face or, increasingly, virtually); (c) possess one or more common goals; (d) are brought together to perform organizationally relevant tasks; (e) exhibit interdependencies with respect to workflow, goals, and outcomes; (f) have different roles and responsibilities; and (g) are together embedded in an encompassing organizational system, with boundaries and linkages to the broader system context and task environment’ |
| (Cohen, 1997, p. 241)            | ‘A collection of individuals who are interdependent in their tasks, who share responsibility for outcomes, who see themselves and who are seen by others as an intact social entity embedded in one or more larger social systems, and who manage their relationships across organizational boundaries’                                                                                                                                                                                                         |

---

(Hackman,  
2011, p. 51)

‘Intact social systems whose members work together to achieve a common purpose. They have clear boundaries that distinguish members from nonmembers. They work interdependently to generate a product for which members have collective, rather than individual, accountability. And they have at least moderate stability, which gives members time to learn how to work well together’

---

By examining these definitions, five structural aspects emerge to help distinguish teams from other forms of groups or collectives: membership, boundaries, purpose, accountability, and interdependence.

1. **Membership.** The definitions emphasise that a team is a recognisable unit; a distinguishable set of people who socially interact (Kozlowski & Ilgen, 2006; Salas et al., 1992). It is perceived as an intact social entity, identifiable by both members and others, with clear delineations between members and non-members (Cohen, 1997; Hackman, 2011).
2. **Boundaries.** The definitions imply that teams are embedded in a broader social system or task environment, with boundary interfaces and linkages to the wider context (Cohen, 1997; Kozlowski & Ilgen, 2006). They also assume temporal boundedness—either through limited membership life-span or at least moderate stability over time (Hackman, 2011; Salas et al., 1992).
3. **Purpose.** The definitions converge on teams as purpose-oriented units: members pursue a common and valued goal, objective, or mission (Salas et al., 1992); and work together to achieve a common purpose or outcome (Cohen, 1997; Hackman, 2011). Kozlowski and Ilgen (2006) similarly emphasise one or more common goals.

4. **Accountability.** The definitions imply differentiation and accountability: members are assigned specific roles or functions (Salas et al., 1992); hold different roles and responsibilities (Kozlowski & Ilgen, 2006). Outcomes are treated as collectively owned through shared responsibility and accountability (Cohen, 1997; Hackman, 2011).
5. **Interdependence.** Across the definitions, interdependence is a core distinguishing feature: teams are described as interacting interdependently toward a mission (Salas et al., 1992); exhibiting interdependencies in workflow, goals, and outcomes (Kozlowski & Ilgen, 2006); and being interdependent in their tasks (Cohen, 1997). Hackman (2011) reinforces this by treating interdependent work as central to producing a shared output.

Together, these facets do not draw sharp lines between collectives, groups, and teams. Rather, membership, boundaries, purpose, accountability, and interdependence are structural properties that, present together, identify a team as a particular form of collective, more clearly bounded and tightly coupled than looser groups where these facets are weaker or less defined.

### Teaming as an Activity

The noun team is closely related to the verb teaming, which shifts attention from the entity to the joint activity (Edmondson, 2012). Where classic team models emphasise designable structures within a unit, teaming emphasises the mindset and practices that enable teamwork when stable membership and boundaries cannot be assumed, as in high-tempo operational settings where teams form, adapt, and dissolve around changing demands. Teamwork, whether in stable teams or through occasional teaming, is often discussed in relation to terms such as collaboration, cooperation, and coordination. These concepts have been used to describe interaction in a variety of ways and are sometimes treated synonymously (Bedwell et al., 2012; Castañer & Oliveira, 2020). For clarity, distinctions are made here by focusing on different aspects of team interactions.

- **Collaboration** is defined as an evolving process in which two or more social entities actively and reciprocally participate toward at least one shared goal (Bedwell et al., 2012). It is often understood as the general, overarching notion of joint work, of which teams and teaming are an instantiation, and within which cooperation and coordination fall as narrower dimensions (Bedwell et al., 2012; Castañer & Oliveira, 2020).
- **Cooperation** occurs when an individual pays a cost for another individual to receive a benefit (Nowak, 2006); in collaborative settings, this is commonly interpreted as an attitudinal predisposition to prioritise collective goals over individual ones (Bedwell et al., 2012). For teams and teaming, this underscores willingness to incur cost for others as a key condition for joint work under uncertainty.
- **Coordination** has been defined in terms of managing interdependencies between activities (Malone & Crowston, 1994). For teams and teaming, it focuses attention on what these interdependencies in activities are. For example, different forms of such interdependencies have been identified, including pooled, sequential, reciprocal, and team interdependencies, which reflect increasing structural relationships among tasks (Saavedra et al., 1993).

Together, these conceptualisations frame collaboration as a joint process where participants are mutually engaged toward shared goals, cooperation as the attitudinal basis for prioritising collective ends, and coordination as the management of interdependencies. For teams and teaming, this distinction carries prescriptive implications, as effective teamwork depends on the simultaneous integration of all three modes. Building on the progression from pooled to sequential, reciprocal, and team interdependencies (Saavedra et al., 1993), Figure 2 relates forms of task structure to their inherent interaction requirements. As the activity moves toward tightly coupled, simultaneous, and multidirectional flows, the demands of all three modes rise together. Teaming thus occupies the high end of this continuum, where the complexity of these collective demands may necessitate their most integrated forms.




Figure 2. Task interdependence across coordination, cooperation, and collaboration.

Naturally, as noted by Salas (1992), what constitutes a ‘team’ may lie on a continuum. Cooke and colleagues (2024) similarly propose that ‘teamness’ can be positioned on a multidimensional continuum, including dimensions such as coordination, to examine how team-like a unit appears in a specific situation. As such, structural aspects of teams may allow a degree of fluidity across a variety of forms, as exemplified by human–animal collaboration (Quan et al., 2024). Such hybrid forms have been proposed as useful references for designing future human–technology teams (Cooke et al., 2024; Lum & Phillips, 2024; Phillips et al., 2015).

### The Evolution of Hybrid Teaming

Driven by increasing domain demands and enabling advancements, forms in which human and technology are working together as teams have emerged. This has been recognised as a new interaction paradigm, often emphasising that technology is no longer perceived as a tool but as a teammate (Lyons et al., 2021; O’Neill et al., 2022; Rix, 2022). Although the term ‘hybrid teaming’ is not established, it is deliberately used here as a broad unifying class that encompasses a variety of vernaculars such as Human-Agent Teaming (Chung et al., 2026), Human-AI Teaming (Berretta et al., 2023), Human-Autonomy Teaming (O’Neill et al., 2022), Human-Robot Team (Wolf & Stock-Homburg, 2023), or Human-Machine Team (Stowers et al., 2021).

While seen as a paradigm shift, the idea of human and technology working together has a long history. For example, human–technology relations have been described using a range of constructs, including systems, cooperation, coordination, collaboration, and teaming. To provide an overview of this literature, a Scopus search identified 20,313 records published between 1975 and 2025.<sup>1</sup> As illustrated in Figure 3, publications in this area remained relatively limited for several decades before increasing from the early 2000s and accelerating rapidly after the mid-2010s.




Figure 3. Publications on human–technology relations indexed in Scopus, 1975–2025.

Although no consistent terminology exists to distinguish the various vernaculars used in this literature (Sun & Qiu, 2025), it is here assumed that terms reflect the disciplinary lens through which the relation between human and technology has been examined over time. For example, earlier work predominantly relied on system-oriented

---

<sup>1</sup> Scopus search query: TITLE-ABS-KEY(("human-robot" OR "human-agent" OR "human-machine" OR "human-ai" OR "human-autonomy" OR "human-automation" OR "human-computer")) PRE/1 (system\* OR collaborat\* OR cooperat\* OR coordinat\* OR team\*). Results were limited to English-language journal articles, conference papers, reviews, and conference reviews, up to and including 2025

descriptions, whereas later literature increasingly adopted terms such as coordination, cooperation, collaboration, and, more recently, teaming. This variation in terminology is illustrated in Figure 4, showing a gradual shift in how authors have framed human and technology working together. As illustrated, system-oriented descriptions dominated until the early 2000s. From the mid-2010s onward, collaboration and teaming terms grew rapidly, reflecting the shift toward treating technology as a partner rather than a component.




Figure 4. Proportional distribution of terminology used to describe the relationship between human and technology across time intervals.

Further analysis of these trends shows that early descriptions of man-machine systems treated humans and mechanical elements as parts of a single system, focusing on their interactions and describing their relationship in terms of input and output reciprocity (McRuer & Krendel, 1962). Within this tradition, Sheridan (2002) argued that human-automation system design should move beyond allocating separate functions and instead aim for a symbiotic integration of complementary strengths, with cooperation as the goal. Concepts of human and technology collaborating date back about 50 years, and research in this area has grown substantially over time—almost to the

degree that it has replaced earlier system-oriented vernaculars. Here, Human-Machine Collaboration is often understood as an all-encompassing term that includes cooperation and coordination across multiple domains (Li et al., 2023; Sun & Qiu, 2025). At its core, the discipline seeks to support human tasks in joint activities. As with the notion that collaboration may, at times, subsume other concepts (Bedwell et al., 2012), such as cooperation (Sun & Qiu, 2025), it becomes difficult to separate the characteristics of the studied relationship. However, the trend also indicates the emergence of hybrid teaming concepts from the 1990s onward. Table 2 shows a selection of definitions that illustrate a shift from a system-oriented perspective to collaboration and, subsequently, teaming. While each vernacular may have its own time and place, they all share the idea that human and technology are engaged as partners in task performance.

*Table 2. Example definitions of human–technology relations.*

| Label                     | Definition                                                                                                                                                                                                                                                                   |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Man-Machine system        | ‘[A] system comprising a machine and an actively participating individual man’ (McRuer & Krendel, 1962, p. 1117)                                                                                                                                                             |
| Joint Cognitive system    | ‘[A] system that can modify its behaviour on the basis of experience so as to achieve specific anti-entropic ends’ (Hollnagel & Woods, 2005, p. 22).                                                                                                                         |
| Human–Machine Cooperation | ‘[A] trait concerning the degree to which a machine is generally agreeable in its relations, behavior and interaction with humans, or better: complementary to human needs, as opposed to competitive, aggressively self-centred or hostile’ (Flemisch et al., 2019, p. 558) |

---

|                             |                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Human–Machine Collaboration | ‘[H]uman and intelligent machines working together to accomplish work. There is a clear division of labor, with intelligent machines doing the highly repetitive and mundane work and humans tending to the complex and creative work. Both work together to make a significant impact on employees, organizations, and society. Intelligent machines refer to AI, robotics and algorithms’ (Li et al., 2023, p. 2922) |
| Human-Robot Collaboration   | ‘[E]nvironments in which human (human operator) and robot come into contact forming a tightly coupled dynamic system to perform a task’ (Baratta et al., 2023, p. 1890)                                                                                                                                                                                                                                                |
| Human–Robot team            | ‘[H]umans and robots, who perform joint tasks, share common goals, interact socially and exhibit task interdependencies’ (Wolf & Stock-Homburg, 2020, p. 247)                                                                                                                                                                                                                                                          |
| Human–Agent Team            | ‘[O]ne or more humans and one or more [intelligent agents] who socially interact (face-to-face or virtually); possess one or more common goals; exhibits interdependencies with respect to workflow, goals, and outcomes; and has distinct roles’ (Iftikhar et al., 2024, p. 10140)                                                                                                                                    |
| Human–Autonomy Teaming      | ‘[I]nterdependence in activity and outcomes involving one or more humans and one or more autonomous agents, wherein each human and autonomous agent is recognized as a unique team member occupying a distinct role on the team, and in which the members strive to achieve a common goal as a collective’ (O’Neill et al., 2022, p. 8)                                                                                |

---

---

|                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Human–AI teaming | ‘[A] process between one or more human(s) and one or more (partially) autonomous AI system(s) acting as team members with unique and complementary capabilities, who work interdependently toward a common goal. The team members’ roles are dynamically adapting throughout the collaboration, requiring coordination and mutual communication to meet each other’s and the task’s requirements. For this, a mutual sharing of intents, shared situational awareness and developing shared mental models are necessary, as well as trust within the team’ (Berretta et al., 2023, p. 23) |
|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

---

As the table shows, the definitions reflect a change. Early views considered humans and technology as components of a system. Later perspectives emphasised technology as supporting and complementing human tasks and needs. More recent definitions conceptualise humans and technology as tightly coupled in tasks and as collaborating, cooperating, and coordinating team members. Thus, when considering team definitions (see Defining Team Structures), there appears to be little difference from traditional human teams, with distinctions instead based on perceived autonomy and humanness of the technology (Harris-Watson & Larson, 2026).

Taken together, the rapid growth of research and the shift in conceptual framing reflect the recognition of hybrid teaming as an emerging paradigm. Indeed, while still a relatively new area of research, it has been suggested that it should be recognised as a field in its own right, while continuing to draw on related disciplines (McNeese et al., 2023).

### Foundations in Hybrid Teaming

Research on hybrid teams and teaming is multidisciplinary, drawing on human sciences such as psychology and team science, engineering disciplines such as artificial intelligence and robotics, and integrative fields such as human factors and human–computer interaction (Li et al., 2023; Lyons et al., 2021; O’Neill et al., 2023; Patel et al., 2025). Building on these foundations, there is a need for a unified perspective that

addresses the human, the technology, and their relation as an integrative whole (Berretta et al., 2023; Iftikhar et al., 2024; McNeese et al., 2023; Singh et al., 2026). Of the many structures and processes such a perspective would span, this overview takes up those foundational to team intentionality: agency, shared and team cognition, and the understanding and sharing of intent.

### *Agency and Autonomy*

Agency is central to hybrid teams and teaming (Berretta et al., 2023; Iftikhar et al., 2024; Lyons et al., 2021; O'Neill et al., 2022). Here, agency refers to a system's emergent property of autonomously regulating its coupling with the environment according to criteria established in its viability conditions (Barandiaran et al., 2009). As such, agential regulation is guided by standards that provide the evaluative basis for adaptivity. Such adaptivity is often described in terms of autonomy. In hybrid teaming, autonomy is commonly operationalised as hierarchical levels, often describing the allocation of roles or functions to the human or technology along a continuum from manual to autonomous performance (Lundberg et al., 2026; Vagia et al., 2016). More broadly, autonomy is a relational notion: an agent is autonomous in terms a goal in its dependence on other agents and environment (Castelfranchi, 2000; Kaber, 2018). On this view, autonomy is the complement of dependence, grounded in a member's repertoire for managing the interdependences through which its agency regulates its coupling with other members and the environment (Barandiaran et al., 2009; Castelfranchi, 2000). Autonomy may further be characterised by two dimensions: self-sufficiency, a practical capability to act without outside assistance; and self-directedness, a deontic authority to form and pursue goals of its own (Bradshaw et al., 2004; Castelfranchi, 2000; Kaber, 2018). Because it is relational, autonomy is variable, as interdependencies may change over time, shaping who, what, when, and how members can and may perform activities as situations unfold (Lundberg et al., 2026). Taken together, autonomy describes the extent to which a member governs its own agency within the team.

At the individual level, research consistently indicates that agency is a prerequisite for hybrid teaming, serving distinct functions for human

and technology. For the human, it centres on maintaining a sense of one's own agency—being an active participant in control of actions and their consequences (Berberian, 2019)—alongside perceiving the technology as an agent with utility in its distinct team role (Baltrusch et al., 2022; Wynne & Lyons, 2018). For the technology, such perceived agency and utility, together with sufficient autonomy, are considered necessary to move it from tool to teammate status (Harris-Watson & Larson, 2026; O'Neill et al., 2022; Schelble et al., 2020). Notably, such agency concerns what is perceived or attributed, rather than possessed. Together, the agency of individual members form team agency through a network of interdependent processes (Barandiaran et al., 2009), providing synergistic effects that seek to combine and complement the capabilities of human and technology (Berretta et al., 2023; Iftikhar et al., 2024).

### *Shared and Team Cognition*

When human and technology work together as group or team members, similar and accurate understanding of the task and mission environment is necessary (Andrews et al., 2023; Cooke et al., 2024; Mohammed et al., 2021; Ososky et al., 2012). Such collective understanding is often studied within the shared and team cognition frameworks, including constructs such as shared or team mental models alongside shared and team situation awareness. Generally, mental models are commonly described as internal knowledge structures that individuals use to describe, explain, and predict real or imagined states of affairs (Andrews et al., 2023; Johnson-Laird, 1983; Rouse & Morris, 1986). These models are limited and simplified representations that evolve over time through experience. Situation awareness, by contrast, is a transient state of knowledge that arises from these mental models together with perception of the environment, the two evolving reciprocally (Andrews et al., 2023; Endsley, 2015, 2023). Following Endsley (1995), situation awareness can be structured across three levels: perceiving environmental elements, comprehending their significance, and projecting their future status.

As part of this representational capacity, understanding and sharing intent is recognised as critical in all forms of teaming. It has long

been regarded as core to adaptive coordination in complex and dynamic situations (Klein, 1994, 1999; Pigeau & McCann, 2006), and increasingly so in hybrid teaming contexts (Berretta et al., 2023; Lyons et al., 2021; O'Neill et al., 2022; Schneider et al., 2022; Singh et al., 2026; Tsamados et al., 2025). In effect, understanding and sharing intent establish sufficient grounding for proactive coordination based on expectations, rather than reactive interpretations. Furthermore, when intent remains sufficiently stable, shared expectations allow members to distribute attention and planning over time. This supports workload management across the team (Howard & Cambria, 2013; Lyons et al., 2021; Schneider et al., 2022; Singh et al., 2026; Tsamados et al., 2025). Shared expectations also support trust calibration, helping members align interpretations and mitigate misperceptions of each other's behaviour (Chen et al., 2018; de Visser et al., 2020; Schaefer et al., 2017).

Together, agency, shared and team representations, and the capacity to understand and share intent constitute interconnected concerns. Through these, the hybrid teaming paradigm conceptualises effective joint action. These concerns share a common underlying character: they involve mental states and actions that are about and directed toward something in the world and toward one another. To this end, the following chapters turn to intentionality, first in its individual form, before extending to collective and team forms in the hybrid teaming context.

## **Individual and Collective Intentionality**

Understanding how team members come to be coordinated around shared ends requires an account of the states and processes underlying their joint activity. Intentionality, the property by which individuals and collectives are about and directed toward the world, offers such an account. Different traditions have emphasised different aspects: epistemic aboutness and conative directedness. Aboutness concerns the states that give rise to meaning, such as knowledge, assumption, and belief. Directedness concerns those that motivate and orient, such as wish, desire, and intent. The two cannot be meaningfully separated, since what a state is directed toward shapes what it is about, and what it is

about shapes what it can be directed toward. This section develops an account of intentionality, from the individual to the collective. Epistemic Aboutness examines how intentional states represent and give meaning to situations through frames and framing. Conative Directedness examines how they motivate and orient agents toward what is sought. Shared and Collective Intentionality then extends these to states held and directed across members of a collective. Together, they ground how intentionality arises and extends from the individual to the collective.

### Epistemic Aboutness

The contemporary use of intentionality as mental states being about or directed toward something is often traced to Brentano (1995). In his treatment of the psychology of perception, he described the ‘intentional (or mental) inexistence of an object’, a reference to a content and a direction toward an object. This account highlights that intentional states are always directed toward something and are about that thing. Intentionality is a way of structuring experience, perceiving-as, judging-that, or intending-to, rather than a feeling or sensation. As such, it is inherently representational and relational, independent of whether the object of perception is real or imagined.

Building on this characterisation, Searle (1983) treats intentionality as a property of mental states and events by which they are directed at, about, or of objects and states of affairs in the world. These intentional states are structured as a psychological mode (the manner of the state, such as believing or desiring) paired with a representative content (what the state is about, such as a specific object or event). Their content is grounded in a Network of intentional states, including both epistemic (e.g., beliefs, assumptions, knowledge) and conative (e.g., desire, preference, wish), that define their conditions of satisfaction under which they are met. The functioning of this Network, and the conditions of satisfaction it defines, presupposes the Background: a precondition of capacities, abilities, skills, stances, practices, and know-how on which intentional states depend.

Considering intentionality as the mind being directed at, about, or of objects and states of affairs, frame theory provides both a structural account of representative content and of sensemaking processes. Frames

have been described as structural networks of nodes representing stereotypical situations (Minsky, 1974, 1975). The top layer of nodes captures features that must hold for the situation to be recognised as such. At the same time, lower-level terminals function as conditional slots to be filled with relevant data from memory or the environment. From the perspective of intentionality (Searle, 1983), this network provides an internal structure setting standards for what must be true for conditions of satisfaction to be met. That is, for the psychological mode of belief, conditions of satisfaction are met if and only if the representative content matches the actual state of affairs. Here, a frame forms that content, and the belief is satisfied only when the situation as framed matches the world.

Functionally, the frame structure also acts as a cognitive filter to manage information overflow (Klein et al., 2006, 2007). Because the frame structure, specifically its slots and relations, determines which aspects of the context are noticed, it selectively guides attention toward information that is relevant for the particular situation. In this manner, frames delineate the situation from its broader context by surfacing certain aspects while submerging others, thereby structuring attention and relevance. From the perspective of intentionality (Searle, 1983), this bounds which aspects of the world the mind is directed toward. Rather than attempting an exhaustive representation of reality, frames function to construct specific and plausible beliefs by directing attention only to the information required to satisfy their internal structural requirements.

With these representational capacities, frames also function as explanatory structures that ascribe meaning through the active process of sensemaking (Klein et al., 2006, 2007). This process is triggered when data serve as anchors, eliciting an initial frame from memory that matches the fixed structural features of a stereotypical situation, such as the top-level nodes proposed by Minsky (1974, 1975). Once a situation has been recognised, sensemaking continues as frames are elaborated and preserved by adding slots and relations, or by filling slots to make sense of events. Frames may also be questioned by detecting and tracking mismatches to assess their plausibility, which may lead to reframing: searching for and comparing alternative frames and, eventually,

adopting a more plausible one (Klein et al., 2006, 2007). From the perspective of intentionality (Searle, 1983), this corresponds to situations in which conditions of satisfaction are not met, prompting reassessment and potential reframing to align beliefs with the world. These notions highlight framing as a continuous process of aligning data and frames to construct coherent and plausible beliefs directed at, about, and of specific situations.

Beyond these representative and sensemaking capacities, frames connect to related constructs such as schema, scripts, mental models, and situation awareness (Klein et al., 2007; Lundberg, 2015). Klein and colleagues (2007) propose frames as an overarching construct that synthesises schema, scripts, and mental models within sensemaking, the process by which situation awareness manifests. This process is dynamic: framing is anchored in the present but also has a temporal dimension. For example, Lundberg's (2015) model suggests that situation awareness is guided by plans and event developments, yielding 'snapshots' that support diagnosis and prognosis of past, present, and projected states over a time horizon. From the perspective of intentionality (Searle, 1983), diagnosing and projecting across these states draws on the Background, the capacity to shift stances and revise the Network by instantiating and synthesising relevant mental models, schemas, and scripts.

Taken together, these accounts suggest that the epistemic aspect of intentionality is not a passive representation of the world but an active, selective, and dynamic relation to it. The mind is directed at, about, and of a situation, constructing a plausible rather than exhaustive picture. Grounded in the Network and enabled by the Background, this picture is continually realigned as events unfold.

### Conative Directedness

Where the epistemic aspect of intentionality concerns the meaning and content of states, the conative aspect accounts for their teleological orientation. To simplify the distinction, the following uses the term 'intent' to characterise the conative aspect of intentionality, where states and actions are deliberately directed toward ends under the guidance of reasoning and behavioural control. While intentionality as aboutness

and directedness is relatively stable across the literature, intent has been conceptualised in more diverse ways. Researchers often fail to distinguish between intent-as-expectations, intent-as-plans, and intent-as-wants (Söderlund, 2002). Intent has also been recognised as a motivation (Freedman & Zilberstein, 2019), a goal (Van-Horenbeke & Peer, 2021), and as a broader concept that connects means with ends (Tubbs & Ekeberg, 1991). Table 3 presents a selection of descriptions of intent across domains to illustrate the diversity.

*Table 3. Overview of how intent has been conceptualised across domains.*

| <b>Domain</b> | <b>Characterisation of intent</b>                                                                                                                                                                                                                                                                                                                                      |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Philosophy    | Intentional action is a description of behaviour identified by the question ‘Why?’, which connects physical acts to a logical chain of reasons directed toward an end (Anscombe, 2000)                                                                                                                                                                                 |
| Philosophy    | Intent is a primary reason—a pair consisting of a pro-attitude and a belief—that functions as the rational cause of an action under a specific description (Davidson, 1963)                                                                                                                                                                                            |
| Philosophy    | Intent is a causally self-referential mental state that unifies mental effort and bodily movement into a single unit to satisfy conditions (Searle, 1983)                                                                                                                                                                                                              |
| Philosophy    | Intent is characterised as a conduct-controlling pro-attitude whereby commitments to choices function as the building blocks of hierarchical plans to coordinate behaviour over time. Beyond initiating planned action when the time is right, reasoning around settled decisions functions to resist reconsiderations and filter inadmissible choices (Bratman, 1987) |
| Psychology    | Intent is a motivational state representing an individual’s willingness to exert effort; it is the immediate precursor to behaviour determined by attitudes, norms, and perceived control (Ajzen, 1991)                                                                                                                                                                |

---

|                   |                                                                                                                                                                                                                                                                             |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Psychology        | Intent is separated into goal intentions (commitment to an end) and implementation intentions (if-then plans) that delegate behavioural control to anticipated situational cues (Gollwitzer, 1993)                                                                          |
| Cognitive science | Intent is a three-tiered dynamic cascade consisting of distal (future-directed/conceptual), proximal (situated/indexical), and motor (sensor-motor) levels that progressively specify and monitor an action to satisfy various forms of agency and control (Pacherie, 2008) |
| Military          | Intent is a purpose or aim with all its associated connotations, functioning to bound behaviour within an acceptable solution space to enable coordinated action under various circumstances (Pigeau & McCann, 2006)                                                        |
| Hybrid teaming    | Intent is a mental model comprised of goals with associated qualities and priorities, hard and soft execution constraints, plans in terms of sequencing and timing (Schneider & Miller, 2018)                                                                               |

---

Given this diversity, it is unsurprising that theoretical and empirical ambiguities persist. This plurality need not be interpreted as conceptual inconsistency, however. Rather, the accounts selectively articulate different structural dimensions of intent, and read together they form a progression. Intent renders behaviour intelligible as authored action (Anscombe, 2000). It is formalised as the primary reason, a pro-attitude paired with a belief, that justifies and causes action (Davidson, 1963). Searle (1983) adds a further distinction: intent is not only a cause of action but causally self-referential, set apart from other pro-attitudes by referring to itself in its own execution. Over time, commitments to intended ends embed specific means into hierarchical plans, providing rational inertia that filters inadmissible options, supports coordination, and specifies partial plans as they fill into subplans toward action (Bratman, 1987). This control of conduct sharpens as action nears: goal intentions become implementation intentions triggered by situational cues (Gollwitzer, 1993), through the willingness to exert effort at the threshold (Ajzen, 1991). At the motor

end, intent becomes present-directed and self-referential: the commitment brings about its own conditions of satisfaction, no longer specifying toward action but regulating it as it unfolds (Pacherie, 2008; Searle, 1983). At the level of joint activity, the totality of intent, an aim or purpose with all its associated connotations, bounds behaviour within an acceptable solution space (Pigeau & McCann, 2006). This space may be operationalisable through goals, constraints, and plans (Schneider & Miller, 2018).

Taken together, these accounts suggest that intent is not a single, uniform construct but a multidimensional one. It spans from making individual actions intelligible to structuring agency over time. By foregrounding different elements, such as pro-attitudes, hierarchical plans, implementation rules, and structural constraints, they collectively outline the representative content through which agents guide their behaviour in changing circumstances.

### Shared and Collective Intentionality

The transition from individual to shared and collective intentionality marks a shift in the unit of agency from the singular to the plural subject. Bratman (1992) lays the groundwork for this transition by characterising Shared Cooperative Activity through three features: mutual responsiveness in aligning plans and coordinating actions, commitment to the joint activity, and commitment to mutual support in each other's roles. Coordination, on this account, arises through the meshing of plans, as participants align their individual strategies to be compatible and non-conflicting in pursuit of a shared goal. Building on this idea, Tomasello and colleagues (2005) model intentional action as a cybernetic control system: a self-regulating loop of an internal goal, an action plan, and continuous perceptual monitoring. In shared intentionality, humans are motivated to share these mental states for collaboration. This transforms the individual monitoring loop into joint attention, establishing the common ground necessary to coordinate perceptions and track each other's intent in real time. Crucially, cognitive representations of self and other provide a holistic bird's-eye view of the social engagement. Holding both in a single format, these structures allow participants to understand how their different but complementary roles intersect, coordinate

specific action plans, and provide mutual support as the interaction unfolds.

As agency extends toward the group level, Tomasello (2014) identifies a transition from shared to collective intentionality, defined by group-mindedness: individuals coordinate with any participant in the collective through shared norms, evaluating and regulating their own performance against collective standards rather than tracking a specific partner. Tuomela (2007) describes such agency through the distinction between the ‘I-mode’ and the ‘we-mode’. In the I-mode, a participant acts as a private agent for individualistic reasons. In the we-mode, the group ethos, the collectively accepted goals, values, and beliefs, serves as the reason for action. The transition is anchored by a collective condition, specifying that a goal is satisfied for one member if and only if it is satisfied for all. This is a social extension of Searle’s (1983) conditions of satisfaction into a shared standard for the ‘we-together’ unit.

Together, these accounts show that collective agency is not merely coordinated individual action but a structural reorganisation of agency itself: alongside organising perceptions and meshing plans through a bird’s-eye view, collectives regulate behaviour in the we-mode against higher-order standards.

## **Teams and Intentionality**

Team intentionality concerns the emergence and evolution of intentional states, encompassing members’ epistemic aboutness (e.g., awareness) and conative directedness (e.g., intent). It seeks to address the ‘lifespan’ of intentional states that enable individual and collective agency in team constellations. At the core of team intentionality is the capacity to understand and share intent, fostered by the continuous alignment of mental states and behaviours. This section develops a conceptualisation of team intentionality by drawing selectively on team science and intentionality research, focusing on the human–technology constellation. It is organised around three interconnected perspectives: the structural foundation sets out what team intentionality involves through the characteristics and relations that define a team; the representational foundations examine the models and knowledge that

enable team agency; and mindreading introduces the meta-representational capacity through which members access and reason about each other's intentional states. Together, these perspectives build toward a preliminary framework of team intentionality.

## Structural Foundation

While the concept of team intentionality is not established, related work suggests that insights into collective intentionality can be gained through team agency (Gold & Sugden, 2007). As teams constitute a specific form of group or collective (Bradshaw et al., 2004; Cannon-Bowers et al., 1993; Wynne & Lyons, 2018), this warrants a more focused delineation. Conversely, while the need for teams to understand and share intent is well established (Klein, 1999; Lyons et al., 2021; Pigeau & McCann, 2006), the research remains fragmented or implicit; shared and team cognition constructs, for example, are inherently representational, but their link to collective agency is less elaborated. Positioning team intentionality at the intersection of team science and intentionality research, the characterisation draws on team structures: perceivable membership, defined boundaries, shared purpose, established accountability, and activity interdependence (see Defining Team Structures). These serve as a basis for aligning converging ideas across the two fields.

**Membership** defines who is in or out of the team, thereby fixing the scope of whose intentionality counts in joint activity. In hybrid teaming, the criterion is well recognised; perceiving the technology as a team member is sometimes described as a prerequisite (Harris-Watson & Larson, 2026; O'Neill et al., 2022; Wynne & Lyons, 2018). Perceiving participants as members enables the adoption of a basic compact—a commitment that facilitates coordination when working toward shared goals in joint activity (Klein et al., 2004, 2005). Such commitment functions as a social glue, furthering the ethos by which members stand and fall together (Tuomela, 2007) and establishing a ‘sense of us’ that attracts members toward the team (Lakhmani et al., 2022). Recognising the team as a distinct unit also grounds a perspectival shift. This has been described variously as adopting a ‘we-frame’, from asking ‘what should I do?’ to ‘what should we do?’ (Gold & Sugden, 2007); transitioning from

an ‘I-perspective’ to a ‘we-perspective’ (Tuomela, 2007); shifting mindset for teaming in practice (Edmondson, 2012); or taking a ‘bird’s-eye view’ as a shared point of reference (Tomasello et al., 2005). These related notions describe the same capacity: changing stance across levels of agency. Thus, from a team intentionality perspective, membership not only defines who counts in joint activity but also binds members through shared commitments, enabling the shift from individual to collective agency.

**Boundaries** position the team within a social system, with relational interfaces to other entities. Within a social ontology, an organisation or institution may authorise operative members to act on its behalf. Its ethos, comprising goals, standards, and beliefs, situates and regulates that agency within a broader task context (Tuomela, 2007). Such declarations set the conditions of satisfaction for task and team criteria (Searle, 1983), typically bounding a solution space for acceptable behaviours and outcomes (Pigeau & McCann, 2006). In hybrid teams, these bounds can be viewed through the descriptive ability (self-sufficiency) and prescriptive authority (self-directedness) of the individual member or the team as a whole (Barandiaran et al., 2009; Bradshaw et al., 2004; Kaber, 2018): the former assessed through, for example, knowledge, skills, competencies, and characteristics (Bradshaw et al., 2004; McNeese et al., 2019); the latter defined by policies specifying what should be done, including the obligations and requirements governing action. Thus, from a team intentionality perspective, the shared ethos provides the purpose and norms by which members regulate and evaluate (Tomasello et al., 2005; Tsamados et al., 2025; Tuomela, 2007), supplying both reasons and means for implicit coordination.

**Purpose** is a central tenet for motivating and committing members toward shared and valued objectives or missions. It establishes a teleological anchor for the team’s existence, setting the direction and aboutness of individual and collective agency. In this respect, the ethos and operational standards may enable the adoption of a ‘team frame’ when striving to accomplish authorised and delegated goals (Gold & Sugden, 2007; Tomasello et al., 2005; Tuomela, 2007). Such a frame is necessary for understanding and sharing intent (Tomasello et al., 2005),

which is foundational for adaptive coordination (Klein, 1999; Pigeau & McCann, 2006). In hybrid teams, this broader notion of intent, along with the team perspective, may be established through the basic compact, in which members are committed to continuous alignment to maintain common ground and remain mutually directable and predictable as conditions and priorities change (Klein et al., 2004, 2005). Thus, from a team intentionality perspective, purpose not only motivates pursuit of shared objectives but also entails the continuous alignment required for coordinated activity.

**Accountability** for task and team outcomes is often structured through assigned roles or functions, with associated responsibilities. The directedness and aboutness of members' intentionality are oriented through these roles to support the delineation of tasks, while commitment to collective responsibility establishes accountability for outcomes within team agency (Gold & Sugden, 2007; Tuomela, 2007). This implies a perspectival shift between the I-mode and we-mode, to grasp both the responsibilities distinct to each member and the shared responsibilities that collective accountability requires. Such shifts necessitate a team-level view, fostering commitment to one's own and others' contributions across mental states, actions, and outcomes (Bratman, 1992; Dennett, 1987; Tomasello et al., 2005). In hybrid teams, the need for distinct roles for human and technology is well recognised (Berretta et al., 2023; Harris-Watson & Larson, 2026; O'Neill et al., 2022), often envisioned through dynamic task or function allocation that assigns responsibility and authority to the most appropriate member in the situation (Ososky et al., 2012; Roth et al., 2019). Notably, this distribution goes beyond mere allocation, as perceived interdependence shifts from task-sharing to consequence-sharing (Lyons et al., 2021; Wynne & Lyons, 2018), implying a 'bird's-eye view' of mutual support in which roles may be interchangeable or assistance provided (Tomasello et al., 2005). Thus, from a team intentionality perspective, accountability highlights fluency in role and responsibility allocation under shared accountability, where the adopted ethos grounds the expectations for mutual support.

**Interdependence** is a hallmark characteristic of teams, ranging from loosely coupled tasks to highly coordinated activities requiring

continuous alignment. In hybrid teams, interdependencies in joint activity are widely recognised as the feature that sets such units apart from other groups or collectives (Bradshaw et al., 2004; Cannon-Bowers et al., 1993; Wynne & Lyons, 2018). These interdependencies can be understood as required dynamic couplings (Cuevas et al., 2007), through which human and technology provide synergistic capabilities to achieve outcomes no member could accomplish alone (Berretta et al., 2023). While less explicitly developed in the intentionality literature, managing interdependencies is implied in accounts of group and collective agency: members act interdependently to transform inputs into outcomes through cognitive, verbal, and behavioural activities directed toward structuring tasks for shared goals (Gold & Sugden, 2007). In this view, such management relies on a commitment to interlock a web of intent by meshing plans for coordinated action (Bratman, 1992). Complementing this individual mode of reasoning, the team perspective enables members to manage interdependencies through commitment to the collective ethos and operational standards (Tomasello et al., 2005; Tuomela, 2007). Together, the two modes support managing individual dependencies and team interdependencies, enabling coordinated action and team agency (Gold & Sugden, 2007). Understanding the forms these dependencies take at both levels shapes situational requirements and opportunities for effective taskwork and teamwork (Johnson et al., 2014). Thus, from a team intentionality perspective, furthering the collective ethos entails not only managing required dependencies and interdependencies but also creating opportunities for mutual support.

From this initial synthesis, several converging ideas emerge. Team science provides structural characteristics—membership, boundaries, purpose, accountability, and interdependence—while intentionality research highlights commitments and different units of agency. Positioned within a social ontology, members may need to shift and mix perspectives to understand and share intent within the team constellation and across its boundaries. Importantly, these structural characteristics and perspectival capacities do not operate in a vacuum: they depend on underlying representational structures through which intentional states acquire their content and direction. To adopt a team frame, manage interdependencies, or align toward a shared purpose,

members must draw on representations of the task, the situation, each other, and the team as a whole.

## Representational Foundations

Team intentionality rests on underlying representational structures through which members' intentional states acquire aboutness and directedness. These foundations are primarily epistemic, providing the knowledge necessary for understanding the situation, each other, and the task. They take plural forms—mental models alongside schemas and scripts organising typified situations and courses of action—since no single model captures such complexity. Team mental models involve knowledge of the task, equipment, members, and interactions: strategies, procedures, and contingencies; system functions and limitations; members' skills and abilities; and roles, responsibilities, and interdependencies (Cannon-Bowers et al., 1993). The technology draws on corresponding models of its own, supporting prediction and interaction with humans, such as task processes and strategies, teammate descriptors (skills, preferences, and dispositions), and rich contextual knowledge (Bendell et al., 2024; Tabrez et al., 2020; Williams et al., 2022). Hybrid teaming research further suggests, for example, intentional, task, analytical, environmental, teamwork, and human state models for shared awareness and intent (Lyons, 2013; Lyons & Havig, 2014). Team intentionality thus rests on plural, potentially simultaneous, active representations across members.

These representations are distributed across the team, no single member holding all of them, each pointing to a different set of elements. Conceptualising each member's representation as a set of elements, the relationship between two members can broadly be described through basic set operations. When the sets are disjoint, members share no elements and operate independently. When they overlap, the relative complement refers to the non-shared portion, the role-specific representations distinct to each member. In contrast, the intersection, refers to the shared portion, capturing the commonality needed to coordinate overlapping responsibilities. The union combines both across members, covering all responsibilities required of the team. Figure 5 illustrates these relationships in joint activity, from disjoint sets at the

independent level, through the relative complement and intersection, to their union at the team level. For example, two members holding  $\{x, y\}$  and  $\{z, w\}$  are disjoint; if instead they hold  $\{x, y\}$  and  $\{y, z\}$ , the relative complements are  $\{x\}$  and  $\{z\}$ , the intersection is  $\{y\}$ , and the union is  $\{x, y, z\}$ . The relationship applies to representational constructs such as mental models, situation awareness, and intentional states. Note, each representation remains individually held, not a pooled cognitive state.




Figure 5. Analytical set relations between individually held representations in joint human-technology activity.

These representational structures are not given; they are actively developed and maintained. For example, drawing on Pigeau and McCann (2000, 2006), understanding and sharing intent involves four mechanisms: communication and socialisation, through which intent is conveyed, and externalisation and internalisation, through which it is converted between forms. Communication concerns the explicit expression of intent through publicly shared media, such as utterances, gestures, text, and symbols; socialisation, the implicit conveyance of intent through tacit understandings developed through training, interaction, and shared experience. Externalisation and internalisation bridge these: externalisation surfaces implicit intent into explicit form; internalisation incorporates explicit intent into tacit understanding. In hybrid teams, bidirectional communication and socialisation have similarly been recognised as essential for establishing and maintaining appropriate representations, for example through conveying intent and through training (Andrews et al., 2023; Chen et al., 2018; Lyons et al., 2021; Schelble et al., 2020). Together, these mechanisms support the development and alignment of the representations that mentalistic attribution depends on.

Framed around team intentionality, these mechanisms operate across levels of representation. Communication and socialisation align members' emergent intentional states and develop the shared Network from which they arise, while externalisation and internalisation bridge the explicit and implicit within these processes. Over repeated interaction, a common Background of shared capacities, skills, practices, and ethos takes shape, enabling those states to arise. Together, these structures and mechanisms constitute the foundations from which team intentionality emerges. Accessing and reasoning about each other's states in the flow of activity requires a further perspective: mindreading, the meta-representational capacity through which members represent and reason about the intentional states of others.

### Mindreading

Mindreading—the capacity to attribute mental states to others—is a naturally evolved human ability and a prerequisite for the collaborative process itself (Tomasello et al., 2005). As such, it underpins team intentionality: it supports the attitudinal predisposition toward collective goals that enables cooperation (Han et al., 2011, 2015), and enables the management of interdependencies that coordination requires (Bratman, 1987; Klein et al., 2005; Pezzulo, 2011; Pigeau & McCann, 2006).


The capacity also affords deliberative efficiency. Shared representations reduce the need for recursive mindreading (Pezzulo, 2011), particularly under high-confidence intent recognition (Han et al., 2015). Recognised intent further shifts the task to tracking and projecting its status rather than reacting to situational data (Howard & Cambria, 2013), as when chess players' recognised patterns constrain plausible continuations, focusing attention on likely moves and deviations (Dennett, 1987). In this sense, individuals adopt an intentional stance: a cognitively efficient strategy for interpreting, explaining, and predicting behaviour. The strategy is indicated in cases ranging from perceiving agency in animated geometrical shapes (Heider & Simmel, 1944) to interacting with humanoid robots (Perez-Osorio & Wykowska, 2020; Thellman et al., 2017, 2022). While idealised and not

always accurate, the stance remains a pragmatic shortcut for making sense of complex behaviour (Dennett, 1987).

To make such attributions, the observer draws on an internal interpretive framework, often referred to as folk psychology, common-sense psychology, or naïve psychology (Dennett, 1987; Malle, 2004; Thellman et al., 2022). These frameworks are instantiated through cognitive structures, such as mental models, schemas, and scripts, that organise the representational content used in mentalistic reasoning. Drawing on frames as previously established (see Epistemic Aboutness), making sense of another agent involves framing their situation: what they perceive, what it means to them, and how they anticipate it will develop. Frames thereby structure the representational content and inferential processes through which others' mental states and behaviours are made meaningful. Importantly, the process is inherently meta-representational: the observer constructs a mental representation of another's internal mental states, effectively a model of a model.

Several perspectives account for how such meta-representational processes support the interpretation, explanation, and prediction of behaviour. Attribution Theory concerns how individuals perceive the causes of behaviour, ascribing it to internal or external factors (Heider, 1958; Kelley & Michela, 1980). Individuals are further predisposed to distinguish intentional from unintentional acts (Malle, 1999; Malle & Knobe, 1997). Intentional behaviour is judged through desires, beliefs, decisions, skill, and awareness, and intent is inferred by whether the actor 'can' perform the act and is deliberately 'trying' to do so (Jones & Davis, 1965). Notably, such attributions need not rest on observation alone: verbal accounts may function as sole or complementary sources, simplifying complex inferences (Heider, 1958). Relatedly, Theory of Mind denotes the system of mental states imputed to others in support of prediction (Premack & Woodruff, 1978). Mentalising is the active process through which it is applied, integrating observed cues, world knowledge, and perspective-taking (Frith & Frith, 2006). In essence, the process judges plausible states by comparing what oneself would experience in the same situation. These capacities likewise extend beyond observation: effective communication requires mutual Theory of Mind (Frith & Frith, 2006), at times recursive, as in 'I believe that you

believe that I believe...' (Tabrez et al., 2020). Such model of a model is illustrated in Figure 6.



*Figure 6. Recursive mindreading between members.*

While the human Theory of Mind of the technology is essential for explaining and predicting its behaviour (Thellman et al., 2022), incorporating an artificial Theory of Mind of the human is equally important for fluent interaction (Bendell et al., 2024; Tabrez et al., 2020; Williams et al., 2022). Further, drawing on collective intentionality, higher-order standards may be necessary for understanding and sharing intent from a bird's-eye view (Tomasello et al., 2005), and the recognition of norms and conventions similarly supports the understandability and predictability of behaviour in hybrid teams (Tsamados et al., 2025). Hybrid teaming thus establishes a bidirectional need for mindreading capacity, extending across human and technology alike.

In summary, adopting an intentional stance constitutes a worthwhile strategy, enabling mental state attribution through Theory of Mind and mentalising capacities (Harris-Watson & Larson, 2026; Matta, 2026; Thellman et al., 2022; Williams et al., 2022). Such capacities draw on multiple meta-representations, wherein the team ethos and perspectival shifts become central for individual and team agency.

Tomasello and colleagues (2005) describe this as dialogic representations encompassing first-, second-, and third-person perspectives—personal states, team members' states, and a holistic bird's-eye view of the social engagement. For team intentionality, this bird's-eye view extends beyond the dyad to the team as a whole, reflecting the transition toward collective intentionality and the we-mode (Tuomela, 2007). Mindreading in team contexts is therefore not a singular capacity but a multi-perspectival and meta-representational activity, one that must account for intentional states across individual members and the team as a whole.



# Designing for Hybrid Teaming

This chapter concerns how the design of hybrid teaming may be approached, from design as a general practice to the particular means it offers for team intentionality. Designing for hybrid teaming is a wicked problem, so it calls for both a grounded account of what design is and how it proceeds, and the practical means through which team intentionality can be modelled and made transparent. The one situates the work; the other equips it. This chapter is organised in two sections. Design as Research and Practice sets out the nature of design, its stages, and its grounding as research. Modelling and Transparency turns to the practical means for designing team intentionality, through modelling and through transparency between members. Together, they aim to connect the design process to the means by which team intentionality is put into practice.

## Design as Research and Practice

Designing for team intentionality requires more than theoretical insight. It requires a practical and scientifically grounded capacity to shape human–technology constellations toward preferred situations. Design, understood broadly, is the human capacity to envision and realise such transitions, moving from present states to preferred ones through deliberate choices and purposeful artefacts (Heskett, 2005; Simon, 1996). In complex socio-technical contexts, such as hybrid teaming, this capacity must be both procedurally guided, providing a structured path from problem to solution, and scientifically rigorous, ensuring that the resulting artefacts are grounded in and contribute to existing knowledge. Neither alone is sufficient: procedural guidance without scientific grounding risks producing solutions that cannot be generalised or validated, while scientific rigour without procedural clarity risks producing knowledge that cannot be operationalised in practice. To establish this dual foundation, this section examines design from three perspectives. It begins with the nature of design as a purposeful activity, before examining design process models that provide structured guidance from abstract problems to concrete solutions. It then

introduces Design Science Research as the paradigm that grounds design activity in scientific rigour, bridging the application environment and the knowledge base. Building on these, the section aligns the design process with Design Science Research into an integrative model that provides both procedural and scientific grounding.

### The Nature of Design

Design can be understood in several ways: as a field, an activity, a plan, or an outcome (Heskett, 2005). Common across these senses is a purposeful and directed activity: ‘everyone designs who devises courses of action aimed at changing existing situations into preferred ones’ (Simon, 1996, p. 111). Where the natural sciences concern how things are, design concerns how things ought to be, devising artefacts to attain goals (Simon, 1996). Alexander (1964) describes such state discrepancies through the relational fitness between form, the structure of the artefact, and context, the demands and constraints acting upon it. Incongruences, or misfits, motivate error reduction until a satisfied (good enough) level of fit is achieved (Alexander, 1964; Simon, 1996). Structuring such transformational pathways, from abstract problems to concrete solutions, is the concern of design process models.

### Design Process Stages

Design process models seek to facilitate effective design outcomes by outlining a structured and systematic process (Wynn & Clarkson, 2018). While numerous models exist across domains, they show strong consensus at an abstract level: a cross-disciplinary review of 124 process models demonstrates agreement around core stages (Gericke & Blessing, 2012), consolidated into four core stages recognised across academia and practice, with a fifth (implementation) treated as optional (Bobbe et al., 2016). Table 4 presents a selection of such models and their stages.

Table 4. Comparison of design process stages.

| Source                     | Design stages                                                                                                                                                                                                                |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| (Howard et al., 2008)      | Establish need → Analysis → Conceptual → Embodiment<br>→ Detailed → Implementation Creative process: Analysis<br>→ Generation → Evaluation →<br>Communication/Implementation (operating iteratively<br>within design stages) |
| (Gericke & Blessing, 2012) | Establish a need → Analysis of task → Conceptual design<br>→ Embodiment design → Detailed design →<br>Implementation → Use → Closeout                                                                                        |
| (Bobbe et al., 2016)       | Analyse → Define → Design → Finalize → Implement                                                                                                                                                                             |
| (Costa et al., 2015)       | Front-end/Product Planning → Requirement Definition<br>→ Conceptual Design → Embodiment Design → Detailed<br>Design → Production Preparation → Commercialization<br>→ Use → End-of-life                                      |

Across terminology and granularity, the models converge on design progressing through apparent stages. A need is recognised and analysed, a solution conceptualised and actualised, and the result finalised as an artefact released into the world. Consensus is strongest in the early stages, reflecting the importance of problem delineation. All models also recognise the iterative nature of design, with cycles within and between stages. Notably, the same logic recurs across levels of abstraction. A micro-level creative process operates iteratively within the macro-level design stages (Howard et al., 2008), and this creative loop largely mirrors the core stages themselves (Bobbe et al., 2016). While these process models structure the realisation of design artefacts, designing for team intentionality also requires grounding in relevant scientific knowledge. This concern is addressed by the design science paradigm.

## Design Science Research

Drawing on Simon (1996), Hevner and colleagues (2004) ground Design Science Research in the sciences of the artificial, framing it as a problem-solving paradigm that creates artefacts to address wicked problems. Such problems are characterised by shifting requirements and ambiguous contextual constraints. Design Science Research is positioned between the environment and the knowledge base: the environment provides the needs that define the problem space, and the knowledge base provides the applicable foundations to address them. Research activity centres on iteratively building and evaluating artefacts (Hevner, 2007). A purposeful artefact is expected to address a relevant problem, be rigorously evaluated for utility and effectiveness, and make verifiable contributions to the knowledge base (Hevner et al., 2004). Peffers and colleagues (2007) operationalise this into a research methodology spanning problem identification, objective definition, design and development, demonstration, evaluation, and communication. Although presented sequentially, the methodology is iterative and flexible in entry point and ordering.

Central to Design Science Research is what counts as a research artefact and how it should be judged. March and Smith (1995) identify four types: concepts (or constructs), which provide the theoretical vocabulary for describing a problem; models, which represent relationships between constructs to structure problems and solutions; methods, which prescribe processes and guidelines for solving problems; and instantiations, working implementations that demonstrate the former in a real context. Each type is evaluated on grounds proper to it: constructs by conceptual clarity, models by purpose adequacy, methods by their operationality and generality, and instantiations by their effectiveness and contextual impact (March & Smith, 1995; Parker, 2020; Suddaby, 2010). How an artefact is evaluated further depends on its maturity: formatively and in artificial settings when early, summatively and in naturalistic ones when mature (Venable et al., 2016). Notably, existing artefacts may provide entry points for further inquiry (Baskerville et al., 2018), wherein the appropriate evaluation method may range from logical argument and illustrative scenario to prototype, case study, and experiment (Peffers et al., 2012). Because evaluation may

occur at different points, designing for wicked problems such as hybrid teaming need not wait for a complete solution: each artefact carries scientific and practical value at the micro level, and contributions accumulate across the full process at the macro level.

### The Integrative Design Process

In practice, the application of Design Science Research often emphasises theoretical output over procedural guidance, leaving the analytical and creative processes of design implicit. The focus falls on what to achieve rather than how to achieve it. Related paradigms such as Research through Design share the integrative ambition (Zimmerman et al., 2007). However, how each stage actively draws on both the environment and the knowledge base often remains implicit. Team intentionality in joint human–technology activity exemplifies the challenge. Designed artefacts must support intentional states that emerge and evolve within operational contexts; however, such problems exceed what any single designer can grasp, requiring specialist perspectives that span scientific and practical domains (Alexander, 1964). Researchers, domain experts, and practitioners may each ascribe different meanings to the problem and hold different orientations toward its solution. A conceptual design process, anchored between the environment and the knowledge base, provides a shared reference point through which these perspectives can be integrated and traced. This is here operationalised as five stages: Recognise and Finalise as entry and exit points, with Analyse, Conceptualise, and Actualise forming the central iterative core. In and between these stages, realisations function as mediating representations. These internalised or externalised forms direct the transformational flow. Figure 7 illustrates the integrative design process model, with each stage and the role of realisations elaborated below.




Figure 7. Illustration of the integrative design process.

**Recognise (Misfit).** The first stage concerns the felt difficulty or misfit between form and context that brings itself to attention (Alexander, 1964). It initiates the transition from a current state of affairs to a preferred one (Simon, 1996). Recognition may be triggered by the environment, where operational challenges or breakdowns make the misfit salient, or from the knowledge base, where the absence or inadequacy of existing artefacts confirms a genuine gap. In the process models and guidelines, this stage corresponds to problem identification and motivation (Peffer et al., 2007), establishing a need (Gericke & Blessing, 2012; Howard et al., 2008), and ensuring relevance (Hevner et al., 2004). The realisation at this stage is primarily internalised: mental representations of the problem, articulated as experienced breakdowns or documented gaps that direct the process forward. In hybrid teaming, for example, core constructs, methods, and models may be lacking or require adaptation, thereby motivating further analysis.

**Analyse (Problem structure).** The second stage traces misfits and links them to their functional requirements (Alexander, 1964), structurally representing the problem space (Simon, 1996). For team intentionality, this means mapping the intentional states involved to understand the problem in context, such as the causes and reasons for the emergence of misfits. Activities draw on both the environment and existing scientific knowledge (Hevner et al., 2004), and criteria are operationalised to clarify what constitutes a good-enough solution

(Peffer et al., 2007; Simon, 1996). Across the process models, the stage aligns with analysis and definition (Bobbe et al., 2016; Costa et al., 2015; Gericke & Blessing, 2012; Howard et al., 2008; Peffer et al., 2007). The realisation is an externalised form: a structured problem representation that transforms the initial felt difficulty into a formal account of the problem space, bounding and enabling the search for viable solutions.

**Conceptualise (Generate form).** The third stage synthesises solutions to reduce misfits (Alexander, 1964), guided by a systematic search of the problem space for satisficing candidates, artefacts judged good enough for the desired situation (Simon, 1996). For team intentionality, this means generating candidate forms that could support epistemic and conative alignment, for example by conceptualising artefacts that support the understanding and sharing of intent. Across the process models, the stage corresponds to conceptual design, where abstract principles are identified and progressively specified (Bobbe et al., 2016; Costa et al., 2015; Gericke & Blessing, 2012; Howard et al., 2008; Peffer et al., 2007). The realisation is externalised candidate forms: proposed solutions drawn from scientific knowledge and practical experience, specified sufficiently to be demonstrated and evaluated.

**Actualise (Demonstrate form).** The fourth stage develops candidate solutions into concrete forms, sufficiently actualised to generate meaningful experiences of the fit between form and context (Alexander, 1964). Through this, opportunities arise to learn about both the problem and the solution, specifically the relationship between the proposed artefact and its intended task environment (Simon, 1996). For team intentionality, this means demonstrating whether the artefact supports the intentional states it was designed for, such that alignment among members satisfies the operationalised criteria. Across the process models, the stage concerns the transition from candidate solutions to demonstrated forms (Bobbe et al., 2016; Costa et al., 2015; Gericke & Blessing, 2012; Howard et al., 2008; Peffer et al., 2007). The realisation is a demonstrated form, concrete enough to be engaged with and assessed against fit criteria, whether through a worked example, simulation, or operational demonstration.

**Finalise (Transfer form).** The fifth stage concludes the process when fit criteria have been sufficiently actualised (Alexander, 1964), or

the artefact at least satisfies the desired situation (Simon, 1996). For team intentionality, this marks the point at which the artefact is released to support hybrid teams. Across the process models, the stage delineates the design from its implementation and communication to relevant communities (Bobbe et al., 2016; Hevner et al., 2004; Peffers et al., 2007). While no design is truly finished, the stage marks a commitment to function independently of the design process. The realisation is the transferred form, such as constructs, models, methods, and instantiations released to their contexts, exhibiting both scientific grounding and practical utility.

Across the stages, realisations carry the design process forward. Realisation here carries a dual meaning. As an internalised form, a realisation is the designer's understanding of the problem and its resolution, transformed through successive representations until the fit criteria are met or revised (Burnette, 2002). As an externalised form, a realisation is the tangible output through which that understanding is concretised and shared, such as a documented misfit, a requirement list, a concept model, a prototype, or an evaluation report (Burnette, 2002). Externalised realisations thereby function as shared objects, providing a common reference point through which designers, stakeholders, and practitioners coordinate on the same problem. Together, these forms drive the process forwards and backwards, translating understanding between stages and transforming the abstract into the concrete within them (Alexander, 1964; Dorst & Cross, 2001; Simon, 1996).

As the engine of the process, realisations may also trigger new design cycles. The same logical structure recurs at different levels of abstraction, from the micro creative loop (Howard et al., 2008) to the macro stages (Bobbe et al., 2016). A realisation may thus initiate a new cycle at any level: the absence of adequate constructs for team intentionality, or insufficient models for representing intent in hybrid teaming, may function as a Recognise stage triggering a new process. Recognise and Finalise thereby serve as entry and exit points not only within the current process but across levels of design activity. Which point is engaged depends on the maturity of the problem and the state of scientific knowledge.

Together, the five stages provide a procedural structure anchored between the knowledge base and the environment. By integrating the iterative logic of design process models with the rigour and contribution requirements of Design Science Research, the model supports the development of artefacts that are both practically grounded and scientifically credible. Through realisations as shared referents, it further provides a common reference point through which researchers, domain experts, and practitioners can coordinate their distinct orientations toward the same design problem. For team intentionality, the resulting artefacts are not merely theoretical constructs but designed forms, grounded in the operational realities of hybrid teaming and carrying both practical utility and scientific grounding.

## **Sharing Intentionality**

Designing for team intentionality requires more than an account of what it is. That account must be modelled, so the implicit is made explicit, and made transparent, so members become interpretable to one another. This section sets out these two means. Modelling Intentionality concerns making the domain and its intentional states explicit through models, building a shared understanding across those who design. Transparency and Explainability concern making members interpretable to one another, so that awareness and intent may be shared between them.

### **Modelling Intentionality**

Designing for team intentionality first requires understanding the domain and the intentional states within it. This understanding is gained through knowledge acquisition, which can be understood as designing a representation of a phenomenon at the scope and depth its purpose requires, that mediates between the world and theory (Leu & Abbass, 2016; Morgan & Morrison, 1999). A model, in this sense, is always limited by what it represents, and capturing what matters serves better than striving for completeness (Box, 1976). To capture what matters is itself a modelling process, often described in three stages: knowledge elicitation, analysis, and representation (Leu & Abbass, 2016). Each stage is a source of insight, making an implicit representation explicit

rather than merely producing an output (Epstein, 2008). Across them, a model is iteratively built and manipulated, and each affords a different means of learning from it (Morgan & Morrison, 1999), similar to the realisations within the design process (see The Integrative Design Process).

No single stakeholder holds the whole picture, so a means is needed to make the implicit explicit and build a shared understanding across them (Oppl, 2017). A shared understanding may be reached through one integrated model for a unified perspective (Renger et al., 2008), or, where that is not feasible, several models that reduce complexity while maintaining comparability (Hochstein, 2017; Yoshida, 2021). As such, conceptual models, alongside, for example, conceptual definitions, operationalisations, and methods, may function as boundary objects: shared reference points that support shared understanding across participants and coordination in projects (Carlile, 2002; Star, 2010). In this manner, such objects may support the transformations between design stages, guarding against what might be lost in translation, so that the misfit recognised at the beginning is the one addressed in the end (see The Integrative Design Process).

For hybrid teaming, building such a shared reference means making the implicit assumptions members hold explicit. For example, team science, psychology, and artificial intelligence converge on broadly similar aspects (see Representational Foundations, Mindreading): the task (strategies, procedures, contingencies), the members (skills, preferences, dispositions), the environment (system functions, contextual knowledge), and their intentional states (belief, desire, intent). Capturing what matters therefore calls for diverse approaches: a plurality of models resting on a plurality of methods. Given the nature of hybrid teaming, potential candidates are drawn from Human Factors methods (Stanton et al., 2017; Wickens et al., 2014), where Cognitive Task Analysis and Cognitive Work Analysis approaches have been suggested (Heinze, 2004; Norling, 2012; Schneider et al., 2023). The two approaches differ in focus: Cognitive Task Analysis attends to the cognitive work of humans, the reasoning, knowledge, and strategies behind performance (Crandall et al., 2006), whereas Cognitive Work Analysis attends to the constraints that shape such performance

(Vicente, 1999). While different, the two are also understood as complementary (Hoffman & Lintern, 2006), particularly in hybrid teaming, where together they can address a broader range of work demands than either alone (Roth et al., 2019). Coactive design is a further candidate, attending to the interdependencies between members in joint human–technology activity, which afford both requirements and opportunities (Johnson et al., 2011, 2014).

Taken together, these accounts show that modelling team intentionality is not the pursuit of a single, complete model but a plural and shared activity. It draws on complementary methods to represent different aspects, held explicit across those who design them, so that the misfit recognised at the outset remains traceable to the artefact released.

### Transparency and Explainability

Transparency and explainability are central to joint human–technology activity, providing mechanisms that foster shared awareness and shared intent between members in hybrid teams (Alonso & De La Puente, 2018; Bhaskara et al., 2020; Rajabiyazdi & Jamieson, 2020; van de Merwe et al., 2024). Transparency and explainability are often used interchangeably, but research has distinguished them (Verhagen et al., 2021): transparency discloses relevant elements to make a member interpretable, and explainability clarifies those disclosures to make the member understandable.

Despite the central role of transparency, the field currently lacks a unified theoretical foundation (Bhaskara et al., 2020). However, existing frameworks draw on, for example, situation awareness (Chen et al., 2014, 2018; Kridalukmana et al., 2020; Sanneman & Shah, 2022), mental models (Cabour et al., 2023; Eiband et al., 2018; Merry et al., 2021), Theory of Mind (Anjomshoe et al., 2019; de Visser et al., 2020; Verhagen et al., 2021), trust (Chen et al., 2014, 2018; de Visser et al., 2020), and intent (Chen et al., 2014, 2018; Johnson et al., 2011; Lyons, 2013; Lyons & Havig, 2014). Of these, frameworks often combine more than one theoretical foundation to identify and organise information requirements for transparency. For example, the situation awareness-based agent transparency model (Chen et al., 2014), later revised to a dynamic and bidirectional model (Chen et al., 2018), combines Endsley's

(1995) levels of situation awareness, Bratman's (1987) theory of practical reasoning, and Lee and See's (2004) theory of human trust in automation. Similarly, de Visser and colleagues (2020) combine mental models and theory of mind, framing transparency and explainability as means through which members share information about themselves, in proactive preparation and reactive repair.

Where these models generally concern disclosure, Verhagen and colleagues (2021) locate explainability in the relations between the elements disclosed. Such elements may concern, for example, perceptions, decisions, and actions, and a relation may hold between any of them. In this manner, disclosing elements may make incomprehensible activity interpretable, and clarifying them through their relations to other elements may make it understandable. A similar structure has been proposed for control systems, where the design of transparency requires an operator to see not only the information at each level of cognitive control, but also the relations between those levels (Lundberg & Johansson, 2021; see Representation). Research has further proposed a number of questions that explanations should answer to support interpretability and understandability (Lim et al., 2009; Lim & Dey, 2009; Miller, 2019; Mohseni et al., 2021). These broadly concern how, why, why not, what if, how to, and what else (Mohseni et al., 2021).

Taken together, transparency and explainability may be best understood not as separate mechanisms but as two complementary focuses: one on the elements disclosed, the other on the relations between them. Each disclosure invites clarification, and each clarification rests on the disclosure. Their foundations (situation awareness, mental models, theory of mind, and intent) overlap substantially with those of team intentionality as presented here. The questions above are then the ones members ask about one another's epistemic aboutness and conative directedness. Described together, transparency and explainability are the means of answering them for alignment.

# The Operational Domain

This chapter establishes the operational domain the thesis is situated in: Manned–Unmanned Teaming in military aviation. It first situates the domain within command and control, then characterises Manned–Unmanned Teaming as a specific instance of hybrid teaming, and the challenges its actualisation must address.

## Command and Control

Military force serves ends that a society sets and its government authorises through law and order (Pigeau & McCann, 2000). A nation pursues those ends through command and control, defined in doctrine as the exercise of authority and direction by a designated commander (Alberts & Hayes, 2006). It rests on two concepts (Pigeau & McCann, 2000): command, where a person wills and takes responsibility for what a mission is to achieve, and control, the structures and processes that keep a mission within feasible bounds. Command becomes essential once a force is divided into interdependent, specialised parts, and once it acts in ill-defined situations with uncertain outcomes (van Creveld, 1985). Authority thereby passes down a chain, from the government to the individual who acts, delegated at each level as a decision right that settles who may decide what (Pigeau & McCann, 2000). Such a right is defined by its function rather than the entity that holds it (Alberts & Hayes, 2006). At the end of this chain, authority is delegated to an operative unit, a team whose members are authorised to act within the bounds set for it. In this manner, command and control positions a team within the wider system and establishes the solution space of its members (see Defining Team Structures, Structural Foundation).

As a commander cannot foresee every situation those in the chain will meet, direction is given as intent: the Commander’s Intent states what a mission is to achieve and why, letting subordinates decide how to act (Pigeau & McCann, 2006). Given as intent, it establishes common intent across the chain, and through its implicit part is meant to enable coordination with less reliance on explicit communication (see Mindreading). More importantly, it allows entities to act in congruency

across the command chain, even under unanticipated situations without waiting for direction (Klein, 1999; Pigeau & McCann, 2006). For the operational domain, enabling such adaptability is the central demand, and it is sharpened across social systems in joint human–technology activity. Manned–Unmanned Teaming, for example, operates under varying degrees of time pressure, task difficulty, and risk exposure that, at times, require operating within the solution space without explicit communication in the command chain.

## **Manned–Unmanned Teaming**

Broadly, Manned–Unmanned Teaming refers to the integration of crewed and uncrewed platforms across ground, sea, and air domains toward shared mission objectives (Kim & Choi, 2024; Taylor & Turpin, 2015). The pairing of these entities is envisioned to lower cost and reduce risk to humans, while improving situational awareness, reach, and mission effectiveness (Athanasiadou et al., 2026; Department of Defense, 2013; Kim & Choi, 2024), in operations such as intelligence, surveillance, and reconnaissance (Banas et al., 2020; Taylor & Turpin, 2015).

Across almost three decades (Taylor & Turpin, 2015), Manned–Unmanned Teaming has evolved from basic control to high-level command of unmanned assets (Jameson et al., 2005), toward an active and collaborating partner (Athanasiadou et al., 2026). As such, it has largely followed the trajectory of hybrid teaming, with its structural properties (see *The Evolution of Hybrid Teaming*, Structural Foundation). For example, in aviation, the uncrewed element is often described as a wingman, assuming roles and responsibilities to support the commanding pilot in interdependent activities (Department of Defense, 2018; Jordan, 2021). The pilot, in turn, authorises the wingman’s decision rights within a bounded solution space through the command chain (Heilemann et al., 2021; Jameson et al., 2005). In this manner, the command chain remains: the pilot holds certain tasks and decision rights, such as mission oversight, planning, and supervision (Jordan, 2021), while the wingman autonomously executes the tasks

delegated to it, such as detecting, identifying, tracking, and assessing targets (Department of Defense, 2018).

Whilst the envisioned form of Manned–Unmanned Teaming persists, its actualisation remains a challenge. For example, the commanding pilot’s tasks extend to include supervision, adding workload (Andrews et al., 2020; Heilemann et al., 2021). Further, the wingman’s high autonomy may degrade the pilot’s awareness of it and leave its state and intent less interpretable and understandable, raise the workload of interacting with it, and miscalibrate trust (Chen & Barnes, 2014; Roth et al., 2020). This is sharpened when operating beyond line of sight, where the team is distributed and depends on a data link for command, control, and communication (Banas et al., 2020; Hobbs & Lyall, 2016; Kaliardos et al., 2024). As the data link can never be guaranteed, there are times when the pilot and wingman need to operate independently (Meyer & Schulte, 2020), making demands on their own authority and solution space (see Command and Control). When the link is lost, the pilot and the wingman can no longer convey intent explicitly and must infer it instead, making the alignment of intentionality the team depends on harder to sustain (see Command and Control, Mindreading).

Several responses have been suggested to address these challenges. To manage workload, the level of control can shift from low to high, adaptively changing roles and responsibilities from the pilot to the wingman (Chen & Barnes, 2014; Heilemann et al., 2021). For example, at a lower level, the pilot may communicate ‘plays’ to the wingman, template command structures that package tasking (Calhoun et al., 2018). Monitoring the pilot’s cognitive state has been proposed to support the communication of such structures (Theissing & Schulte, 2013), and to drive adaptive interfaces for commanding the wingman (Lim et al., 2021; Lim, Ramasamy, et al., 2018). Further, since the pilot and the wingman must each interpret the other to act together, bidirectional, or teamwork, transparency has been proposed as a means for them to share awareness and intent (Chen et al., 2018; Lyons, 2013; Lyons & Havig, 2014). A shared semantic space has further been suggested for efficient communication, spanning, for example, goals, plans, constraints, and the situation behind them (Gustavsson et al.,

2011; Schneider & Miller, 2018). Recognising the intent of the pilot, the wingman, and other aircraft has further been proposed to support coordination by focusing attention, managing distributed planning, distributing information proactively, detecting conflicts, and aiding error correction (Geddes, 1994, 1997; Geddes & Lizza, 2001).

Whether conveying or inferring intent, teaming rests on members understanding and sharing it; however, in Manned–Unmanned Teaming the means for the pilot and the wingman to do so are not yet in place. The challenges are well recognised, and responses have been proposed, but they remain proposals, not yet realised as ways to design for team intentionality. What is envisioned and what can be built thus emerge as a misfit, one that motivates a framework for realising team intentionality by design.

# Methodology

This chapter addresses how the framework was developed and how the studies that inform it were conducted. Developing a framework requires synthesising its parts from theory and evidence, and grounding that synthesis in empirical study. These are distinct activities at distinct levels. The chapter is organised in two sections. Framework Development sets out how the framework was built from the background and the papers. Studies describes how the four papers were conducted. Together, the two sections show both the framework and the empirical work from which it draws.

## Framework Development

Frameworks are important tools across nearly all scientific inquiries, often serving as bridges between abstract paradigms and their knowledge foundations (McMeekin et al., 2020; Partelow, 2023). For example, Partelow (2023) describes frameworks as ‘bridging tools’ that enable the synthesis and communication of ideas, thereby driving the paradigm forward. Such synthesis is typically grounded in the identification and integration of, for example, specific theories, models, and empirical observations, which together form a set of core concepts and general relationships within the framework. Likewise, a methodological framework is foremost founded in existing methods and guidelines alongside experiences and literature reviews (McMeekin et al., 2020). These foundations are adapted, integrated, and built upon to structure processes and procedures.

As a composite of these foundations, a framework is itself a design artefact. Peffers and colleagues (Peffers et al., 2012) classify it as a meta-model, an artefact that assembles constructs, a model of their relations, and a method into a single organising whole. So understood, the framework is answerable to the same grounds as the artefacts it assembles, under one shared question: does it work? Its constructs should be clear, its models adequate to their purpose, its methods operational and general in use, and any instantiation effective in context (see Design Science Research). Beyond these grounds, a framework is

positioned in a specific field, offering a perspective on what exists in the world (Partelow, 2023). The framework developed here seeks to integrate theoretical, methodological, and empirical accounts in the context of hybrid teaming, with future Manned–Unmanned Teaming operations as its studied domain.

The framework itself was developed at a different level from the studies that inform it. Its parts were built by synthesis: theoretical and empirical accounts into the concepts and relations that frame team intentionality, and methodological accounts into methods for putting them to use. The included papers are both the input to this synthesis and where the framework is demonstrated through logical argument, illustrative scenario, and case study (Peffer et al., 2012). Its grounding is therefore formative and in constructed settings, not yet tested in the field. How the studies were conducted is described in the next section, separately from the framework's development.

## **Studies**

This section describes how the four papers were conducted. They share a common method, presented in two subsections. Research Setting describes the shared context, through a scenario and its simulation. Knowledge Acquisition describes how knowledge was elicited, analysed, and represented. Together, they set out the common method on which the studies draw.

### **Research Setting**

This subsection describes the context the four papers share. All are grounded in one scenario and its simulation. It is presented in two parts. Scenario sets out the mission context common to all four papers. Simulation describes how that context was played out for participants. Together, they establish the common setting the studies share.

#### *Scenario*

Scenarios have evolved from frameworks for speculation and ‘imagining the unimaginable’ (Kahn & Wiener, 1967) into rigorous tools for thinking critically and creatively about probable, possible, and preferable futures

(Bishop et al., 2007; Börjeson et al., 2006; Ramirez et al., 2015). They bridge scientific and practical rationality, challenging assumptions, opening new lines of inquiry, and supporting navigation from present conditions toward potential futures while preparing for surprises. At their core, scenarios are hypothetical sequences of events that focus attention on causal processes and decision points (Kahn & Wiener, 1967). In design, they are often presented as stories that provide a research context (Carroll, 1999; Ramirez et al., 2015), in which vignettes serve as scenes within a staged sequence through, for example, papers, animations, or videos (Jenkins et al., 2010). Such a context includes a setting and agents pursuing goals, providing a common medium and language for participation and communication.

In this thesis, all papers are grounded in a shared scenario context depicting a future Manned–Unmanned Teaming operation. The scenario was developed with domain experts and outlines a general situation and Commander’s Intent for a reconnaissance mission with the potential to encounter unknown entities. Conducted in an uncertain tactical situation, such reconnaissance is noted to be operationally relevant and sufficiently complex mission type (Theissing & Schulte, 2013). Within this scenario, transfer of control and link loss were selected as the target situations for investigation. Both are operationally critical and demanding.

Transfer of control is a fundamental challenge across high-stakes domains, described under labels such as handover/takeover and handoff/handback (Desmedt et al., 2021; Grindley et al., 2025; Hobbs & Lyall, 2016). Broadly, it refers to a process in which authorities, responsibilities, and relevant information transfer between members, whether human or technology, to maintain continuity of task or system performance. Although often described as unidirectional, it necessarily involves multidirectional flows, as aspects of control may be shared and distributed across the larger system. In Manned–Unmanned Teaming, it is especially demanding, as the control stations are geographically separated. Identified human factors include operator communication, awareness of the unmanned aircraft, and workload operations, and technical factors such as control-station configuration and data link performance affect the synchrony and stability of the transfer (Hobbs &

Lyall, 2016; Kaliardos et al., 2024; Pankok & Bass, 2018). Papers I and II examined transfer of control as a target situation.

Data link management and performance (availability, continuity, integrity, and latency), more generally, are critical for maintaining operational safety and permanence in Manned–Unmanned Teaming missions (Kaliardos et al., 2024; Lim, Gardi, et al., 2018). Whilst data link availability and sustainability are critical, they can never be guaranteed due to tactical considerations and technical limitations (Meyer & Schulte, 2020; Okcu, 2016). Consequently, a fundamental reliance on unmanned autonomy and standardised operational procedures emerges to ensure predictable aircraft behaviour during unplanned communication disruption episodes (Kaliardos et al., 2024; Stansbury et al., 2009). These disruptions represent a period of heightened operational demand that may break the team configuration, and potentially teaming itself. Papers III and IV examined unplanned link loss, and Paper III also examined a gradually degrading link.

### *Simulation*

Simulation is understood here through its etymology, ‘to imitate’ or ‘to behave as if’ (Rybing, 2018; Vincenzi et al., 2024). In this broad sense, anything can simulate anything else, an activity, process, or system, provided the simulator behaves similarly enough to the target to be studied in its place. Thus, when simulating, the scenario is central: it supplies the structure of events that play out, and where humans participate, lets them engage with it as if it were real (Huffman et al., 2016). Such forms range from low-fidelity vignettes, presented on paper, video, or computer (Jenkins et al., 2010), to human-in-the-loop and live simulation (Hodson & Baldwin, 2009). Choosing among them involves trade-offs, such as cost-effectiveness, fidelity, participant availability, and safety (Frank et al., 2000; Hodson & Baldwin, 2009). Notably, higher realism is not always better. For studying intentionality, what matters is cognitive fidelity, the match between the reasoning a simulation engages and that of the real task (Rybing, 2018). A low-fidelity form can therefore focus on the task and leave out realistic detail that adds cognitive load without engaging the reasoning of interest (Feinstein & Cannon, 2002). Across Papers I–IV, simulations supported

both elicitation and evaluation, matched to the stage of the work. Pen-and-paper representations built shared understanding of situations in Papers I and II. Animated presentations provided situational cues and temporal reference points as the scenario unfolded in Paper III. A higher-fidelity, pilot-in-the-loop environment then supported evaluation in Paper IV (Figure 8).




Figure 8. Simulation forms used across the studies, from pen-and-paper (top left) through animated (bottom left) to pilot-in-the-loop (right).

## Knowledge Acquisition

This subsection provides an overview of the knowledge acquisition process, the considerations that shaped it, and how these map to each paper. As noted, the acquisition process often applies in three stages, elicitation, analysis, and representation, each drawing on its own methods to transform its input and pass it to the next (see The Integrative Design Process). Table 5 maps the methods used across Papers I–IV onto these three stages, each elaborated in the following subsections. Note. Not every paper spans all three: Paper IV applies earlier models rather than building its own.

*Table 5. Knowledge acquisition methods used in each paper, across the three stages of elicitation, analysis, and representation.*

| <b>Paper</b> | <b>Elicitation</b>              | <b>Analysis</b>                                | <b>Representation</b>              |
|--------------|---------------------------------|------------------------------------------------|------------------------------------|
| I            | literature review;<br>interview | deductive thematic                             | adapted Work Domain<br>Analysis    |
| II           | literature review;<br>interview | deductive thematic                             | adapted Work Domain<br>Analysis    |
| III          | interview                       | deductive thematic                             | adapted Joint Control<br>Framework |
| IV           | observation;<br>questionnaire   | structured protocol;<br>descriptive statistics |                                    |

### *Elicitation*

The elicitation phase concerns gathering information from relevant sources. Across Papers I–IV, this drew on literature reviews, interviews, questionnaires, and observations, each described below.

### Literature reviews

Literature reviews serve as a fundamental methodology for synthesising fragmented and interdisciplinary knowledge, providing a collective assessment of evidence (Partelow, 2023; Snyder, 2019). Integrative reviews are particularly effective for developing models and frameworks in emerging fields (Snyder, 2019; Whittemore & Knafl, 2005). The approach is noted to be unique in its inclusivity of knowledge sources and bases to create preliminary conceptualisations of new topics and complex phenomena. The knowledge base may span a variety of sources and types, ranging from empirical journal articles to theoretical conference papers and technical reports. Paper I and Paper II used integrative literature reviews to establish the foundations of intent models and modelling. Reflecting an integrative approach, levels of cognitive control from the Joint Control Framework (Lundberg & Johansson, 2021; see Representation) was used to identify and integrate core aspects and elements of intent across disciplines and existing models. Both Papers used these models in subsequent development and application.

## Interviews

Interviews elicit subjective accounts of interviewees' situated perceptions and meaning-making, often yielding insights difficult to obtain otherwise (Alshenqeeti, 2014; DiCicco-Bloom & Crabtree, 2006; Holstein & Gubrium, 1995). They are time-consuming, limited in scale, and susceptible to interviewer influence and bias. Their interactive nature, however, allows the co-construction of understanding through in situ clarification and correction. Papers I–III used semi-structured interviews with subject matter experts. Drawing on Human Factors methods, two protocols were developed for different elicitation strategies. Following Norling (2012), the first used established probes from Applied Cognitive Task Analysis (Militello & Hutton, 1998) and the Critical Decision Method (O'Hare et al., 1998), supporting general task understanding and detailed accounts of decision-making. Reflecting an approach advocated by Elliott and colleagues (2000), the second protocol aligned elicitation more closely with the aspects of interest. It drew on the levels of cognitive control (Lundberg & Johansson, 2021; see Representation) and strategy analysis (Hassall & Sanderson, 2014), with probes developed to elicit specific model aspects under time pressure, task difficulty, and risk exposure.

## Questionnaires

Self-reports such as questionnaires offer standardised, resource-effective instruments for collecting subjective data (Krosnick, 1999; Ranganathan & Caduff, 2023). Responses are open-ended written accounts or closed-ended categorical or scale-based data. Unlike interviews, questionnaires minimise researcher-participant interaction, making the instrument itself the medium of inquiry. This is both an advantage and a limitation: administration and analysis are more practical, but opportunities for clarification and correction are lost, and the subjectivity of self-report may warrant complementary interviews or observation (Baumeister et al., 2007; Krosnick, 1999). In Paper IV, a questionnaire captured participants' demographics and their judgements of awareness and performance. Rather than adopt an existing instrument, a new measure was developed for the

operationalised components of awareness and performance, on the grounds that awareness is multi-dimensional, so its components are more informative when assessed separately than as a global state (Moens et al., 2026). Here the components are the six levels of cognitive control (Lundberg & Johansson, 2021; see Representation).

### Observations

Observation offers insights that self-report cannot: interaction and non-verbal cues in a naturalistic setting (Crandall et al., 2006; Kawulich, 2005). What people say, do, and say they do differ (Balcom et al., 2021). Argyris and Schön (1974) frame this as espoused theory, what people describe or justify, against theory-in-use, what governs conduct. The theory-in-use can only be reached from behaviour by an inferential leap, so when report and conduct agree, the espoused theory can be taken as a guide to it; when they diverge, two theories are at work. Observation is not purely objective either. The researcher's stance shapes the data, from complete participant to complete observer, and methods run from open-ended description to systematic coding of specific activities (Balcom et al., 2021; Kawulich, 2005). Stance and method therefore imply a degree of model co-construction (Braun & Clarke, 2006; Ford et al., 1993; Kawulich, 2005). Combining self-report and observation through triangulation mitigates these distortions (Balcom et al., 2021; Kawulich, 2005). Besides questionnaires, Paper IV gathered observational data from video recordings to measure awareness and performance indicators.

### *Analysis*

Analysis structures elicited material into patterns, qualitatively or quantitatively depending on the data. Qualitative accounts were examined through thematic analysis, guided deductively by frameworks and concepts while open to inductive refinement (Azungah, 2018; Braun & Clarke, 2006). Quantitative data were summarised with descriptive statistics. Papers I–III used deductive thematic analysis, guided by levels of cognitive control (Lundberg & Johansson, 2021; see Representation) in Paper I to build a preliminary model, refined in Papers II and III. In Paper IV, observational data were organised by a structured protocol into

a record of what happened and when, and self-report data were summarised through descriptive statistics.

### *Representation*

Representation makes the findings from the analysis explicit and communicable, expressing them in a form that can be shared and reused. Human Factors methods provide means for this, focusing on the human, the technology, and the relations between them, to improve the design of how they work together (Stanton et al., 2017; Wickens et al., 2014). In this, the Cognitive Task Analysis and Cognitive Work Analysis traditions have been suggested (see Modelling Intentionality), the former concerning how people think and make decisions (Crandall et al., 2006), the latter the work domain constraints that shape how work can be done (Naikar, 2013; Vicente, 1999). This distinction carries into the models each yields (Vicente, 1999): descriptive (or normative) models of how work is (or should be) done, formative models of what the constraints make possible to do.

Methods from both traditions have been used to represent intentional states. From the Cognitive Task Analysis tradition, intentional states have been represented as causal cognitive maps of individually held or shared beliefs (Langfield-Smith, 1992), as task diagrams of goals and plans in which perceptual cues trigger courses of action (Norling, 2012), and as hierarchical graphs of goals, subgoals, decisions, and the information those decisions require (Schneider et al., 2023). From the Cognitive Work Analysis tradition, Work Domain Analysis has been used to represent intentional domains as abstraction hierarchies of the constraints that shape action in a specific situation (Hajdukiewicz et al., 1999), to model an operator's decision-making (Lui & Watson, 2002), and to make design intent explicit and traceable to support human problem-solving (Leveson, 2000). These forms draw on a common stock of concepts, beliefs, goals, plans, perceptual cues, decisions, and the constraints that shape action. The Joint Control Framework (Lundberg & Johansson, 2021) is one means of representing intentionality within this space, chosen here because it brings several of these concepts together within a single unified structure.

The Joint Control Framework grounds and harmonises models and constructs in Human Factors: the abstraction hierarchy and decision ladder (Rasmussen, 1986), the extended control model (Hollnagel & Woods, 2005), and situation awareness (Endsley, 1995). The framework describes joint control processes, the exchanges of control between human operators and automated or autonomous systems, across six levels of cognitive control and over time. The levels range from the abstract to the concrete: the frames that recognise a situation, the effects and goals pursued, the values and criteria weighed, the generic plans and patterns available, the implementations that assign resources, and the physical objects and actions. Its representation, the Score, places perception (P), decision (D), and action (A) points as control activities across these levels and over time, capturing an unfolding episode rather than a single moment (Figure 9).




Figure 9. Illustration of a score episode diagram.

Building on the Joint Control Framework, the levels of cognitive control were used across the papers to represent intent. In Papers I–II, they structured an adapted Work Domain Analysis (Naikar, 2013; Vicente, 1999), modelling situated intent as the choices a situation makes possible and those taken; that is, formative and descriptive content within one model. In Paper IV, the levels helped organise transparency information, instantiated in an HMI and subsequently evaluated against intent awareness. Paper III instead reconceptualised the Joint Control Framework in terms of intentional processes: where Papers I–II held choice and decision as possible and actual intent within one model, Paper III added choices as a distinct activity alongside perceptions, decisions, and actions, mapped across the levels of cognitive control and time in score episode diagrams.



# Summary of Papers

This chapter presents the four included papers, which address how intent can be understood, modelled, and designed for in hybrid teaming. It begins with an outline that positions their contributions across the design process and describes how those contributions are characterised. It then presents each paper in a consistent structure: the problem and aim it addresses, the approach it draws on, its findings, and its contributions.

## Outline

The papers included in this thesis can broadly be positioned across the integrated design process (see The Integrative Design Process), from recognised misfits through the iterative core to finalisation. From this perspective, the papers do not map onto single stages but span partly different and overlapping ones, with their emphasis shifting across the process. Realisations produced at one stage serve as inputs to the next, so that contributions accumulate and carry forward across the design process. For example, the limitations noted in Paper I motivate the temporal approach of Paper III, while the requirements of Paper II and the scenario of Paper III carry directly into the design and structure of Paper IV.

Paper I establishes the conceptual and methodological foundations for modelling intent; Paper II grounds these empirically and derives requirements for synthetic wingmen; Paper III extends them to the temporal dynamics of intent formation and interpretation; and Paper IV instantiates them in a human-in-the-loop study, exploring intent transparency in design. Together, the four papers move from defining intent as a modellable construct toward demonstrating how it can be designed for in hybrid teaming (see Figure 10). Contributions are characterised by two aspects. The first concerns the type of design artefact produced: concepts, models, methods, and instantiations, together with prescriptive design. These dimensions combine in each paper and are summarised in the contributions sections that follow.




Figure 10. The four papers positioned across the integrative design process.

## Paper I

In ‘Modelling Situated Intent for Human-Autonomy Teaming: A Human-Centric Approach,’ conceptual and methodological foundations for the framework is established by addressing what intent is and how it can be modelled.

### Problem and Aim

The paper identifies two significant issues when modelling intent in a hybrid teaming context. First, the concept of intent is ambiguous, making it difficult to determine what academics and practitioners alike are, or should be, modelling. This ambiguity may lead to misconceptions and misalignment within and across stages of the design process. To address this issue, defining and operationalising intent emerged as a requirement for achieving conceptual clarity and practical applicability. Second, there is a lack of methods for modelling intent that make its key aspects explicit and structurally intelligible. Such representations are essential to the design process, as they provide a shared reference point for analysis, conceptualisation, and evaluation across stages and individuals. To address this issue, there is a need to develop methods that are aligned with the definition and operationalisation of intent. Accordingly, the paper aims to define and operationalise intent, and to outline an approach for modelling situated intent. Through the demonstration of concepts, models, and methods in a Human-

Autonomy Teaming context, it provides a coherent and practical foundation for understanding and representing intent.

## Approach

To define and operationalise intent, a literature review identified, analysed, and synthesised how the concept has been treated across established frameworks and models, focusing on its nature, function, and representational content. To account for the nature and function of intent, Bratman's (1987) theory of practical reasoning was adopted. From this, intent was defined as 'choices characterised by commitment' with the functional role of coordinating action.' While this clarified the nature and function of intent, its representational content was left unclear. To address this misfit, the six levels of cognitive control were adopted from the Joint Control Framework (Lundberg & Johansson, 2021). Originally describing joint control between human and technology, the levels were here reconstituted as a structure for intent's representational content. Through these, the facets and elements proposed across domains were analysed and synthesised, ranging from situations and goals, through considerations and plans, to activities and objects. From this synthesis, a conceptual model of situated intent was proposed. Using prior definition and the conceptual model as a basis, intent was operationalised as: 'deliberated and connoted choices connected by context–situation and means–end reasoning within an intent space with six levels of cognitive control'.

Grounded in the operationalisation, Applied Cognitive Task Analysis, the Critical Decision Method, and Work Domain Analysis were adapted and integrated into the proposed approach, providing an account of both possible and actual intent in a specific situation (see Knowledge Acquisition). From the former two, the main usage was an integrated set of probes for eliciting human reasoning about possible and actual intent in target situations. Work Domain Analysis was adapted in two respects. First, the standard abstraction hierarchy was replaced with the six levels of cognitive control to model situated intent. Second, the representational structure was retained for its ability to hold formative and descriptive intent within the same model: what can be committed is stated as a choice, and what is committed as a decision. Together, the

adapted approach provided a unified representational form for analysing intent across its facets, from abstract situational framing to concrete objects, while making visible the structures of commitments formed through reasoning.

The conceptualisation of intent and the proposed approach were then explored, enabling demonstration. Seven subject matter experts participated in scenario-based interviews, using the integrated probes supported by low-fidelity simulations, set within the transfer of control situation of the shared reconnaissance scenario. Knowledge analysis used a theory-driven thematic analysis aligned to the proposed conceptualisation of intent, followed by knowledge representation using the outlined approach.

## Findings

Applying the approach, the definition of intent proved valuable in directing attention to commitments rather than other pro-attitudes such as desires or wishes. Treating intent as commitment also provided a theoretical rationale to constrain reasoning, effectively acting as a filter against inconsistent or conflicting choices. Further, as commitments serve as stable reference points, considerations and commitments may be distributed cross-temporally, making intent structurally coherent and consistent. Anchoring such structures to anticipated situations allows a shift towards enactment through recognition, guiding behaviour rather than requiring construction anew in the present. As such, conceptualising intent as a structure of commitments supports the management of cognitive (or computational) load through filtering, cross-temporal distribution, and recognition-based enactment.

The operationalisation was considered valuable as it presents intent as situated within a context and shaped by constraints, and therefore as multi-faceted. The context-situation separation allows a target situation to be treated as embedded within its wider context, whilst the levels of cognitive control differentiate the facets through which intent is formed. The connections drawn by context-situation and means-end reasoning further carry explanatory power, accounting for the reasoning behind decisions. As such, the operationalisation directs attention to the facets, elements, and potential structures of intent,

providing an analytical basis for understanding why it is formed, what it represents, and how it is (to be) enacted.

The methodological approach was considered applicable and practical, particularly because it accounted for both possibilities and actualities. This is valuable: design ought to account for what is or should be done in situations, while also supporting what can be done in unanticipated situations to remain adaptable. Further, the proposed approach helps explicate choices and decisions, making it easier to trace and explain the emergence of inconsistent or conflicting intent (see Paper II). For design, providing such traceability and explainability is essential to support coordination. As proposed, however, the approach did not address the strength and temporality of intent, motivating complementary approaches (see Paper III). In conclusion, whether used to design interfaces, semantic spaces, or synthetic agents in hybrid teams, the methodological approach was considered a valuable means to identify and convey the informational and functional requirements necessary to account for agency within the design process.

## Contributions

The paper mainly provides theoretical and methodological contributions toward defining, operationalising, and a method for modelling situated intent.

**Concepts.** Intent was both defined and operationalised, transforming it from an ambiguous concept into a structured design artefact. This was accompanied by a demonstration in a Manned–Unmanned Teaming context to provide practical guidance for hybrid teaming development. From the operationalisation, the concept of an intent space was also suggested, providing a means to represent both potential and actual intent. A further contribution is the concept of situations-of-situations, which recognises that intent may be nested, enabling the separation of mission context from specific target situations.

**Models.** The conceptual model of situated intent provided a link between theoretical, methodological, and empirical considerations, offering a practical tool to structure and analyse different aspects and

elements of intent. By applying it to the case, it was demonstrated as an initial scaffolding tool for the development of models of situated intent.

**Method.** A novel human-centric approach to represent intent by combining well-established Human Factors methods. This approach provides a unified representational form that emphasises both the formative (what can be considered) and descriptive (what is committed) nature of intent through context–situation and means–end reasoning. Through demonstration and evaluation in a Manned–Unmanned Teaming context, it provides a practical and structured account of how situated intent can be modelled to support coordination in hybrid teams.

## Paper II

In ‘My Synthetic Wingman Must Understand Me: Modelling Intent for Future Manned–Unmanned Teaming,’ the foundation of modelling situated intent is grounded empirically.

### Problem and Aim

The paper recognises a lack of practical approaches for modelling intent, and of models of intent that can inform design and development, in the context of Manned–Unmanned Teaming. Such approaches and models are critical for the design of future air capabilities. They can highlight which aspects must be understood and shared to maintain coordination between human pilots and synthetic wingmen. To address this issue, there is a need for applicable approaches to represent and analyse intent, and for corresponding models that are empirically grounded to inform design decisions. Building on the theoretical and methodological foundations described in Paper I, the present paper aims to complement these with practical and empirical insights into modelling situated intent. It further aims to identify requirements enabling synthetic wingmen to reason about their human partner’s intent. The paper contributes practical insights from applying the modelling approach, empirically grounded models of situated intent, and design considerations for the development of future Manned–Unmanned Teaming capabilities.

## Approach

To model situated intent, the same theoretical and methodological foundations as in Paper I were used, drawing on the same set of subject matter expert interviews. Paper II, however, places a stronger emphasis on what can be learned in practice from modelling and from the resulting models in a Manned–Unmanned Teaming scenario. To narrow the scope, the paper models pilots’ intent in a reconnaissance mission, focusing on transfer of control and link loss situations. To complement the interviews, a literature review was added, focusing on published work on domain models in the unmanned aircraft domain and enabling the identification and synthesis of relevant structures. Participants included four fighter pilots, two specialists with expertise in both manned and unmanned aircraft systems, and one unmanned aircraft system operator. The two sources were analysed separately through theory-driven thematic analysis, guided by the six levels of cognitive control. The results were then used to inform and complement one another, and synthesised into a conceptual model of the fighter pilot’s intent space in the mission context. From this synthesised model, the transfer of control and link loss models were constrained as subsets, representing and analysing pilots’ intent in the two situations separately.

## Findings

Learning through modelling and models provided valuable insights for the design of future Manned–Unmanned Teaming capabilities. From the applied approach, three empirically grounded models of situated intent were designed and analysed. The first model provides a novel conceptual representation of the Manned–Unmanned Teaming context, synthesising and extending earlier unmanned aircraft models through subject matter expert interviews. The elements identified across the six levels of cognitive control provide the necessary foundation for understanding aspects of intent in this domain. Using this contextual knowledge, intent in the transfer of control and link loss situations was represented and analysed to establish a better understanding of these recognised challenges.

The methodological approach was considered practical and applicable, particularly as the models of situated intent were judged to be reusable and scalable. Such reusability and scalability reflect the concept of situations-of-situations introduced in Paper I: the transfer of control and link loss models were embedded within the broader reconnaissance mission model. The chosen level of decomposition thereby controlled the granularity at which intent was represented. This is essential, as modelling efforts are typically resource-intensive and often suffer from scope. Existing representations thus serve as valuable scaffolding tools. In effect, the ability to re-utilise these models implies they may be applicable to similar situations. The design focus can then shift toward identifying misfits between the baseline model and new situational requirements. Further, the ability to represent and analyse choices considered and committed within a single model proved valuable, particularly for tracing where intent converges or diverges across agents. For example, the transfer of control analysis indicated a strong consensus, represented as a single intent structure, reflecting an anticipated situation in which participants could rely on shared plans and procedures. In contrast, the link loss analysis revealed three personas with distinct intent structures, reflecting an unanticipated situation. Small differences in the framing of the situation propagated across levels of cognitive control, resulting in inconsistent and conflicting intent. The ability to trace where such divergences originate, and how they propagate, strengthens the analysis for the design of future Manned–Unmanned Teaming capabilities.

To further support the design of such capabilities, several requirements were identified, focusing on the higher levels of cognitive control. In general, these draw on human abilities and tendencies, such as Theory of Mind, conceptualised as capabilities to be emulated in technology for more natural interaction. For example, unmanned aircraft acting as wingmen need to be able to reason about the mission context (e.g., Commander’s Intent and operational environment) and situational needs (e.g., priorities and trade-offs) to be sufficiently agentic. Further, the ability to represent frames and framings is considered valuable, as this provides a means to anticipate situational needs and recognise incongruences with the pilot’s framing of the

situation. Understanding own and human performance (e.g., mission and team status) and constraints (e.g., causal limitations and mission boundaries) was also considered important, providing the explainability required when shifting goals and courses of action. Finally, the importance of shared plans and procedures was emphasised, as they provide a stable point of reference for detecting incongruences. Rather than strictly recognising what the pilot is doing, this becomes a monitoring function: measuring the congruence between observation and expectation carries a comparatively low computational burden.

## Contributions

The paper provides methodological and empirical contributions toward grounding and demonstrating the modelling of situated intent in a Manned–Unmanned Teaming context.

**Models.** Grounded in a literature review and seven subject matter expert interviews, three models of situated intent were presented. Each offers novel insight into envisioned situations in the emerging research and development of Manned–Unmanned Teaming capabilities.

**Method.** Demonstrating the practical applicability of the approach proposed in Paper I, Paper II showed how existing models of situated intent can function as reusable and scalable scaffolding tools for modelling intent in similar situations. Specifically, existing models reduced the need to construct new ones, instead enabling the extension, constraining, and alignment of existing representations. Further, as anticipated in Paper I, the models enabled the identification and analysis of considered and committed choices, tracing the reasoning between them, and comparing intent structures across situations and agents.

**Design requirements.** Grounded in the learning from modelling intent and the resulting representations, several requirements for enabling unmanned aircraft to act as human-like wingmen were identified. In essence, they suggest that the technology needs to reason across the levels of cognitive control, particularly the higher ones, and to understand such reasoning in relation to its human partner. While the focus was on unmanned aircraft capabilities, supporting the human in a similar manner is warranted.

## **Paper III**

‘When Things Change: Temporal Dynamics of Intent Formation and Interpretation in Critical MUM-T Situations’ extends the prior foundation into the temporal domain.

### **Problem and Aim**

The paper recognises the need for methods to model the temporal aspects of intentional processes, and for representational forms that make these aspects analysable, in critical Manned–Unmanned Teaming episodes. In the design of future air capabilities, accounting for such temporal dynamics is essential for understanding the deliberative process whereby pilot intent emerges, evolves, and is to be enacted over time. Accordingly, the aim of the study is to model temporal aspects of intentional processes in a Manned–Unmanned Teaming scenario. Intent is addressed both as a formative process, in which intent is formed and transformed as situations evolve, and as an interpretative process, in which the intent of other agents is inferred. The investigation focuses on data link disruption and aircraft encounters. Through application and demonstration, the paper contributes an improved understanding of how temporal aspects of intentional processes can be modelled, and insights into the deliberative processes under investigation.

### **Approach**

The Joint Control Framework (Lundberg & Johansson, 2021) was considered a natural fit to model temporal dynamics of intentional processes. Specifically, in Paper I and Paper II, the six levels of cognitive control served as a cornerstone for the operationalisation and subsequent modelling of situated intent. This provided the theoretical and methodological basis for moving from studying static ‘snapshots’ of situated intent (the why, what, and how) to representing and analysing the temporal dynamics of practical reasoning (the when) through which those states arise. However, as originally described, the Joint Control Framework primarily focuses on interaction between agents in shared control processes, rather than on intentional processes themselves. To address this misfit, two adaptations were made. First, the analytical

focus was shifted from modelling interactions to representing the practical reasoning underlying intent. Second, choices were explicitly represented alongside the framework's perception, decision, and action points, creating four core activities across the levels of cognitive control. To represent these processes temporally, score episode diagrams were adopted as the representational form, mapping activity points across levels and time. Together, these form a temporally organised structure through which the deliberative processes underlying agency in complex and dynamic situations can be represented and analysed.

Knowledge elicitation was conducted through scenario-based interviews with ten experienced fighter pilots. The scenario, further developed from Paper II, outlined a narrative around a reconnaissance mission involving Manned–Unmanned Teaming, focusing on data link disruptions and aircraft encounters. During the interviews, an interview protocol and low-fidelity simulations were used as complementary instruments. The protocol, designed around the six levels of cognitive control, structured the probes in relation to the applied framework. The simulations, in the form of animated presentations, provided situational cues and temporal reference points as the scenario unfolded. Data were recorded and subsequently analysed using a qualitative approach guided by a conceptual model of situated intent (see Papers I-II). This analysis included organising interview data into intent elements across the levels of cognitive control and over time. From this analysis, knowledge was represented in score episode diagrams by sequentially mapping perceptions, choices, decisions, and actions across the levels and time. This made the temporal structures of intentional processes analytically more accessible and communicable.

## Findings

In practice, the adapted methodology demonstrated its value for representing and analysing the temporal dynamics of intentional processes. In total, four empirically grounded models were represented and analysed, each describing identified temporal patterns of pilot reasoning in the studied situations. The first model represents an episode involving the onset of concurrent data link disruption and aircraft encounter events. The analysis indicated that pilots reorganised

these into a prioritised sequence rather than managing them in parallel, given situational constraints. The second model represents proactive intent formation to uphold data link status. Here, immediate adjustments to maintain system performance preceded a transition to deliberation over plans, criteria, and goals when these adjustments were no longer sufficient. This shift reflects a movement from individual to team-level coordination. The third model illustrates two episodes in response to link loss: an initial phase of rebuilding awareness, followed by reconnecting and reassessing. Together, these episodes reflect a shift from reactive to proactive intent formation. The last model depicts intent interpretation in an encounter with an unknown aircraft, unfolding through two episodes: aircraft identification and subsequent intent interpretation. Together, these illustrate a satisficing strategy in which classification and perspective-taking constrain the space of plausible inferences across levels of cognitive control.

Besides the models of practical reasoning in the studied situations, the findings reaffirm the relationship between prior conceptualisations and representations of situated intent. Whereas work domain models of situated intent may be seen as representing ‘snapshots,’ the resulting models illustrate how activities such as perceiving, choosing, deciding, and acting are organised over time. Consequently, intent is not merely a structure of commitments for action, but one that also includes commitments to future perception and deliberation, enabling coherence as situations evolve. This organisation further indicates a synchronising function, distributing activities across different time horizons to organise not only what to do but when.

Established intent structures were found to support both intent formation and interpretation. In intent formation, full or partial segments of established structures enabled enactment under bounded rationality. Two patterns were identified: a reactive control mode, characterised by rapid responses with limited deliberation, and a proactive control mode, where plans, criteria, and goals were reconsidered and recommitted more gradually across levels. At times, such enactment may occur without explicit awareness, reflecting habitual or routinised use. At the same time, these structures provide a basis for further deliberation, enabling incremental adaptation within

and across levels as situations evolve. In intent interpretation, established structures provide a basis for attributing commitments. Each assumption constrains subsequent assumptions, effectively reducing the space of plausible inferences across levels of cognitive control.

Such mechanisms are essential, as the analysis indicates multiple shifting and interacting layers of constraints that shape intent formation and modulate one another: contextual, situational, causal, intentional, and relational. These create dependencies within tasks and interdependencies across team processes, through which new requirements and opportunities for further activity may arise. In effect, this reflects the notion of an intent space as dynamically changing, structuring what can and should be done in terms of perceiving, choosing, deciding, and acting. The temporal organisation of these activities is essential for managing demands. For example, concurrent events within the mission situation may be reorganised into sub-situations to cope with time pressure, task difficulty, and risk exposure. This observation further substantiates the concept of situations-of-situations, helping to explain how pilots may structure complex problems into manageable units. Lastly, situated intent demonstrates its inherent coupling with intentionality and agency, as formative and interpretative aspects cannot be meaningfully separated. Intent emerges and evolves through ongoing activity in which perception, choice, decision, and action co-develop and continuously reshape one another.

## Contributions

The paper contributes theoretically, methodologically, and empirically toward modelling the temporal dynamics of intentional processes.

**Method.** The Joint Control Framework was reconceptualised in terms of intentional processes and applied within the Manned–Unmanned Teaming domain, demonstrating practical value for representing and analysing the temporal dynamics underlying pilot intent. Theoretically, this extends the prior conceptualisation of intent, from committed situated choices to a cross-temporal organisation of intentional activities, with the intent space correspondingly understood as dynamically changing. Operationally, intentional processes comprise four core activities: perceptions, choices, decisions, and actions, mapped

across levels of cognitive control and time in score episode diagrams. Such temporal structuring is fundamental for understanding how intent emerges and evolves in situations with changing demands.

**Models.** Through application of the adapted approach, four empirically grounded models of intentional processes in a Manned–Unmanned Teaming scenario were developed and analysed: three representing intent formation and one intent interpretation. Together, they demonstrate the value of the approach for identifying temporal patterns of practical reasoning.

## Paper IV

‘Fast but Fragile Intent Interpretations: Effects of Transparency on Awareness and Performance in MUM-T’ moves toward design and evaluation, instantiating the conceptual solutions and introducing team intentionality as the broader capacity the framework aims to support.

### Problem and Aim

The paper identifies the challenge of sustaining team intentionality when real-time communication is unavailable. Pilots must then interpret, understand, and predict the behaviour of their wingman through observation and expectation alone. This shifts transparency from a problem of information transfer toward a problem of interpretation. For design, this is recognised as the need to make agent state and intent accessible beyond explicit communication in hybrid teams. To this end, transparency is proposed as a mechanism to support the pilot’s awareness of the wingman’s state and intent. The paper aims to investigate the effects of two levels of intent transparency on a lead pilot’s performance and awareness in a link loss situation. Through a human-in-the-loop study, the paper contributes empirically grounded insights into how intent transparency may shape interpretation, along with design considerations for fostering team intentionality.

### Approach

The concept of team intentionality was introduced and conceptualised as a collective capacity to sustain coordinated action through shared

awareness, compatible goals, and meshed plans in hybrid teams. The conceptual model of situated intent presented in Paper I was then used to compare and align two established transparency models, revealing a substantial overlap across their elements. Notably, whereas such models emphasise conveying the agent's reasoning and projections, the alignment also spans the lower levels through which the agent becomes observable in the first place. With this as a basis, intent transparency was operationalised through the levels of cognitive control and common intent elements, which functioned as a design lens for identifying information needs. Paper II provided the specific information presented in the HMI to foster awareness of another agent's intent in link loss situations. Paper III informed the scenario and task structure, grounding them in the intentional processes: aircraft encounters were scripted as a priority task that also functioned as an attentional distractor, contributing to awareness degradation of the wingman during link loss episodes, reflecting patterns identified in Paper III. Paper I, finally, informed the development of the questionnaire, so that the design and the measures of awareness and performance derived from the same operationalisation of intent.

With this grounding, an exploratory human-in-the-loop experiment was conducted using a higher-fidelity simulation. A within-subject design was applied, in which participants were exposed to low and high levels of intent transparency. The conditions differed in the extent to which information about the wingman's state and intent was provided: the low level offered baseline mission and state information, whereas the high level added projective cues about the wingman's likely state and intent during link loss. Eight participants each completed two experimental trials following the same procedure. For each trial, video recordings and questionnaires were collected. Video recordings provided objective indicators of behavioural performance, focusing on task success and completion time. Questionnaires captured subjective assessments of performance and awareness of the wingman's state and intent.

## Findings

Building on the contributions of Papers I–III, the study presents an instantiation in which the conceptual solutions are developed, implemented, and evaluated through a human-in-the-loop simulation study. The results were mixed across conditions. In line with findings from Papers II and III, awareness of the wingman’s state and intent appeared to rest on correctly identifying aircraft and on matching expected against observed trajectories through established plans. Entity identification proved particularly important: correct identifications were followed by correct plan and goal selections, whereas incorrect identifications cascaded into subsequent errors across both conditions. However, the observed behavioural patterns point to a divergence from the sequential reasoning strategy identified in Paper III. Rather than constraining assumptions step by step, participants appeared to maintain multiple hypotheses in parallel, seeking a satisficing match between identity, plan, and goal. The conditions likewise diverged from expectations, indicating a trade-off. Low support led to higher success rates and faster entity detection, grounding early state awareness. High support instead accelerated entity identification and plan and goal selection, at the cost of more errors. Self-judged awareness and performance were generally similar across conditions. The exception was the low-support condition, where ratings dropped noticeably when task components were unsuccessful. In contrast, the high-support condition showed comparable ratings regardless of task outcome, indicating that participants maintained confidence even when their interpretations were incorrect. Together, these patterns suggest that increasing intent transparency may promote fast but fragile interpretations: formed quickly and confidently, yet difficult to revise as new information becomes available. Such interpretative lock-ins, accompanied by misaligned confidence, correspond to findings on complacency and automation bias. Given the exploratory design and small sample, the findings are read as indicative patterns, corroborating rather than extending broader research in human factors and transparency.

From these findings, four design considerations were proposed. Transparency should be adaptive, providing what is displayed, when, and how, according to situational needs, with particular care for the

lower levels through which another agent becomes observable. It should support comparison between parallel hypotheses. It should calibrate confidence against available data. Finally, it should preserve revisability, keeping interpretations open to challenge to prevent premature lock-ins. Together, these considerations aim to support the transparency needed for the interpretive processes required for continuous alignment, fostering team intentionality as situations change.

## Contributions

The paper provides empirical and methodological contributions toward designing for team intentionality, carrying the conceptual solutions of Papers I–III into instantiation and evaluation.

**Instantiation.** The conceptual solutions were instantiated in a working implementation, comprising the human–machine interface, the simulation environment, and the human-in-the-loop study. The instantiation demonstrates the constructs, models, and methods of the preceding papers in context, yielding the finding that intent transparency may promote fast but fragile interpretations.

**Method.** The paper demonstrates the use and conversion of models across design stages as a means of maintaining the design process, from problem definition and interface content to scenario and measures. Deriving the design and the measures from the same operationalisation maintained correspondence between what was designed and what was evaluated.

**Models.** Intent transparency was operationalised through the levels of cognitive control and common intent elements, forming an intent transparency model as a design lens for identifying information needs.

**Design considerations.** Four considerations for transparency mechanisms were identified: adaptive transparency, hypothesis comparison, confidence calibration, and revisability. While individually consistent with prior transparency research, they are here derived through the operationalisation of intent, grounding them in what must be interpreted to sustain team intentionality.



# Team Intentionality by Design

Designing for team intentionality confronts a complex landscape of interconnected factors, spanning the human, the technology, and their relation. To support the design of future joint human–technology activity, this chapter seeks to develop a practical framework of team intentionality, one that establishes central aspects and offers guidance for their use, grounded in the preceding theory and the studies across the papers. The framework is developed in two parts, addressing the two research questions in turn. It begins by framing team intentionality, setting out the concepts and relations that may be central to analysing it (RQ1). It then turns to realising it, setting out how those concepts and relations can be used in the design of hybrid teaming (RQ2). It closes with the case that motivated it, Manned–Unmanned Teaming, and what the framework offers for where its design stands and where it may go. Together, they seek to make team intentionality more tractable and traceable in the design of future joint human–technology activity.

## Framing Team Intentionality

Framing team intentionality concerns setting out its core concepts and their general relations, addressing what may be central to it (RQ1). It does so through eight aspects, each foregrounding a facet of the phenomenon, shaping the problem and solution spaces, and posing its own questions for analysis (see Appendix). It begins with the hybrid team, the unit whose intentionality is in view. It then turns to how intentional states are held and arise, through the strata and the situation as framed, and to how a unit regulates its activity and bounds its choices, through agency, the agentive unit, and autonomy. It closes with intent and aligned intentionality. A concluding synthesis then reads the aspects together — as one interdependent account rather than a checklist — and recasts each as a place where awareness or intent can fail, so that misfit becomes locatable.

# The Hybrid Team

As hybrid teaming emerges as a distinct field, recognising what constitutes such a unit becomes essential. As presented earlier (see The Evolution of Hybrid Teaming), its various conceptualisations converge toward traditional team definitions, with the distinction that they envision human and technology as peers in joint activity. To this end, rather than proposing new definitions that may add confusion to established ones, making implicit assumptions more explicit becomes more important. In this regard, the structural facets from team science (see Defining Team Structures), alongside the subsequent integration with intentionality research (see Structural Foundation), can provide direction. Specifically, they provide an analytical lens for examining human-technology configurations in terms of perceivable membership, defined boundaries, a shared purpose, established accountability, and activity interdependence. Figure 11 provides an overview of these facets and the analytical questions they pose for examining human-technology configurations. Each facet is elaborated after.




Figure 11. Five structural facets of a hybrid team with example questions.

**Membership** establishes who is recognised as a participant, fixing whose intentionality counts in joint activity. Operationally, this requires reciprocal recognition: not only humans perceiving technology as a participant rather than an instrument, but technology able to identify and orient toward its human partners. A question may be who is recognised as a member, and whether that recognition is reciprocal. In Manned–Unmanned Teaming, for example, this surfaces whether the

unmanned aircraft is treated as a wingman or an expendable tool, and whether it can orient toward its human partners in turn.

**Boundaries** position the team within its organisational context, defining the delegated authority and ethos from which purpose and constraints derive. The ethos grounds the motivational conditions of satisfaction for collective action, and may include beliefs, goals, values, norms, standards, traditions, and practices. Operationally, what is conferred from outside cannot be assumed to regulate conduct from within, as the two may diverge without explicit alignment. A question may be what ethos can be derived, and whether it is aligned across members. In Manned–Unmanned Teaming, for example, the Commander’s Intent is expressed against the ethos, surfacing whether it is aligned across human and technology.

**Purpose** defines why the team exists and what it is directed toward, anchoring both individual and collective agency. Where boundaries set conditions from outside, purpose requires members’ internal commitment in accordance with the accepted ethos. Such internal commitment cannot be assumed, and alignment is not fixed but must be actively maintained as situations evolve. A question may be what intent members hold, and whether it remains aligned. In Manned–Unmanned Teaming, for example, this concerns whether the pilot and wingman hold the same mission intent as conditions change.

**Accountability** structures how roles and responsibilities are distributed, with outcomes collectively owned. Responsibilities span role-specific tasks, overlapping contributions where roles meet, and collective outcomes for the team as a whole. In practice, this asymmetry operates across levels: what is distinctly held, what is shared, and what is collectively owned must each be recognised as such. A question may be who is responsible for what, and whether accountability is held individually, jointly, or collectively. Across papers, the pilot answers for the air picture and the wingman for surface reconnaissance, while both share responsibility for safe separation, and the team as a whole answers for the mission.

**Interdependence** defines the dynamic relations that bind members in teaming activity, ranging from pooled and sequential workflows to the reciprocal and multidirectional forms characteristic of teams. Managing these dependencies means organising mental and physical activities and resources over time, addressing both the demands they impose and the opportunities they afford. This relies not only on synchronising actions but on meshing members' intentional states, which is often left implicit yet essential to sustain team agency. A question may be what requirements and opportunities may arise, and how they are organised over time. In Manned–Unmanned Teaming, for example, the pilot and wingman are often expected to continuously share information to align perceptions that shape subsequent decisions and actions.

Broadly aligned with these facets, teamness (see Teaming as an Activity) concerns how strongly that structure is expressed in a given situation. Interdependence in activities, reflected in the interaction modes of coordination, cooperation, and collaboration, may provide a complementary analysis: each mode may be rated against the demands of a specific activity, and the gap between what an activity demands and what the configuration provides is what the analysis seeks to locate. Figure 12 presents such rating as an illustrative case for the pilot–wingman dyad in shared reconnaissance under normal conditions. The collaborative demand is rated high (7), as the shared mission goal requires reciprocal participation across air and surface tasks. The cooperative demand is rated low (3), as the relationship is supervisory rather than peer-like. The coordinative demand is rated moderate (4), as their tasks run largely in parallel, with little coupling to manage between them. Other activities would produce different profiles.




Figure 12. Interaction mode profile example for shared reconnaissance under normal conditions.

Together, the five facets, the teamness continuum, and the interaction modes provide a structured basis for examining human-technology configurations. Rather than asking whether something is a team, the framework asks what structural properties are present, to what degree they are expressed, and what the activity demands of its members. This shifts the question from classification to degree: not what to call a configuration, but what it requires and whether those requirements are met when needed.

### The Strata of Intentionality

Drawing on previous sections (see Individual and Collective Intentionality), intentionality refers to the epistemic aboutness and conative directedness that emerge and evolve within agency. The framework conceptualises intentionality across three strata: a repertoire in the Background, occurrent representations in the Network, and salient states in the Foreground. Further, a split is made between the nonintentional preconditions in the Background and the intentional states in the Network and Foreground. Figure 13 illustrates this conceptualisation of intentionality within agency, and it is further elaborated below.




Figure 13. The three strata of intentionality within agency.

**Background** broadly comprises the repertoire of performative and epistemic preconditions that enable intentional states. The performative (how to do things) concerns capacities, abilities, skills, and competences, among others, that ground agency, while the epistemic (how things are taken to be) comprises the underlying structures, such as mental models, frames, values, conventions, traditions, and practices. Together, they set the foundation and boundary conditions for team intentionality, posing questions such as: what can members do individually, jointly, and collectively given these? Further, how do they configure in a given situation, and what intentional states may then emerge? For example, Paper II identified that the synthetic wingman requires preconditions akin to those of a human teammate: performative, such as the ability to reason about others' intent, and epistemic, such as the underlying structures that support such reasoning.

**Network** comprises a holistic web of intentional states that give rise to salient aboutness and directedness, and defines the conditions under which those states are satisfied. Within this web, all states are available but not equally accessible, as some shape intentionality implicitly while others are open to reflection. More specifically, these states can be understood as nodes forming a gradient from the salient centre outward. Those closely connected are more accessible, while those further out shape reasoning without being readily apparent. Examining

these nodes then reveals why salient states take the form they do. This poses questions such as: what underlying states shape the salient ones, and do members draw on the same? For example, Paper II revealed that pilots held different underlying assumptions and priorities regarding the same situation, which may be explained by different nodes in their respective Networks (e.g., beliefs, assumptions) and different Background foundations (e.g., expertise, stances).

**Foreground** holds the situationally salient intentional states: those actively about and directed toward objects and states of affairs, whether real or envisioned. Each salient state draws on the Network's central nodes, pairing psychological modes with representative content: believing that something is the case, desiring a particular outcome, or intending a course of action. These pairings carry conditions of satisfaction that function as motivational criteria within agency, surfacing selective aspects of both aboutness and directedness. This poses questions such as: what aspects of the situation and intent are surfaced, what conditions must hold for the situation to be recognised and the intent to be achieved, and do members surface the same? For example, Paper III demonstrated how evolving situations shifted what pilots surfaced, leading to changes in both how they framed the situation and what they intended to do about it.

While described as separate, the strata work in conjunction, actively developed and maintained through mechanisms (see Representational Foundations) such as communication (explicit exchange), socialisation (tacit sharing), externalisation (implicit made explicit), and internalisation (explicit made implicit). Externalisation and internalisation mediate movement within the strata, converting between implicit and explicit, while communication and socialisation extend this interplay between members. This poses questions such as: what can be surfaced and communicated between members, what can only be developed through socialisation over time, and what remains implicit? For example, Paper II identified procedures as having a particular quality: once internalised through training, they are expected to function as shared implicit intent, available without communication.

Together, the strata provide a structured account of intentionality within agency: what is surfaced, what it draws on, and what enables it. For team intentionality, this allows salient states to be traced to their underlying conditions, and divergences between members to be located within the strata.

### The Situation as Framed

In the framework, a situation refers to the ascribed state of affairs an entity takes to be the case. As such, it is not a feature of the states alone, but an interplay between these and what human or technology take them to be. Here, it is proposed that frames and framing provide the mechanism through which the entity interprets and updates such intentional states through sensemaking (see Epistemic Aboutness). These frames are conceptualised as webs comprised of slots and relations that, within conditions of satisfaction, give rise to the situated epistemic aboutness and conative directedness. As a synthesis of schema, scripts, and mental models, their content may concern not only states of affairs but courses of action and other intentional entities. This situated intentionality is organised within the strata (see The Strata of Intentionality), where a situation may be understood as a particular instantiation. The Background repertoire provides the preconditions for frames and framing. The occurrent Network comprises the contextually relevant states, and the Foreground gives rise to the salient nodes the entity takes to be the case. This is illustrated in Figure 14, where the frame fits epistemic slots (circles) and conative slots (squares) from its sources. The filled frame then projects an ascribed situation, which may depart from the actual one. From the figure, intentionality may be examined along three aspects: orienting, situating, and nesting.




Figure 14. The frame projecting an ascribed situation against the actual.

**Orienting** can be understood as the activity through which conative directedness shapes situational framing. Figure 14 illustrates this, where what the framing is about, and what it leaves aside, follows from the motivations directing it. These motivations are pro-attitudes of varying strength, for example ranging from wishes (W) through desires (D) and intent (I). Thus, framing is never neutral: what is attended to, why it matters, how it is interpreted, and when it becomes relevant are always intentionally directed. Further, this directedness extends beyond the present and the epistemic: a situation may be framed across time, encompassing what is anticipated or remembered, and across psychological modes, encompassing what is imagined or desired, among others. A purpose may also orient framing toward different objects, whether a state of affairs of oneself or another, so that framing the situation and framing a teammate's intent are both shaped by what is directed. It follows that the purpose orienting a frame is not explicit in its content but determines which intentional states are occurrent in the Network and which arise in the Foreground. A question for analysis may therefore be what purposes make what frames available, which of these are adopted, and how this shapes the framing of both the situation and a teammate within it. For example, Paper III indicated that a mission frame oriented pilots toward both their own air reconnaissance and the wingman's surface reconnaissance, framing the wingman as a teammate

whose state and intent had to be understood for the team to realise desired mission outcomes.

**Situating** concerns the activity through which an entity adopts, elaborates, questions, sustains, or abandons a frame (see Epistemic Aboutness). Through these activities, content from different sources is fitted within the frame structure, establishing a web of intentional states in the Network, concerning, for example, objects and their attributes, courses of action, or the intentionality of team members. The conditions of satisfaction then determine what is salient in the situation and what stays in the surrounding Network, effectively functioning as a filter to determine what is surfaced and salient or submerged and occurrent. The boundary a frame draws reflects what the entity takes to be situational, which may differ from what is situational in fact. Figure 14 illustrates this, showing the ascribed boundary against the actual one, with slots fitted from the environment (E), assumptions (A), or prior knowledge (K), and empty where content is currently unsatisfied. The content fitting these slots may then be examined for its reliability and accuracy, and where slots are empty, whether the missing content can be obtained. This is apparent in Papers II-IV, where pilots represented the wingman against their own situation, fitting knowledge of its plans, assumptions about its behaviour, and observations of its position to interpret its state and intent. Here the assumptions need confirmation, for instance by checking identity and correspondence between behaviour and intent, while under link loss the position slots emptied and the wingman's state had to be inferred from what remained. From such cases, questions follow as to what sources a human or technology fits to each slot, how far these can be trusted, and how closely they match the actual state when sources fail, or content cannot be obtained.

**Nesting** concerns how a situation relates to its context, and how it may itself become the context for a situation through a sub-frame (see Epistemic Aboutness). Such nesting may be necessary, as for example, to decompose a problem into sub-problems. Papers I and III conceptualised this as an organisation of interdependent situations of varying granularity and temporality: their grain, from a broad context to finer sub-situations, and their extension across time. This conceptualisation suggests that intentionality of human and technology

may draw on situations nested within one another, where a context and the situations it contains feed both ways through context-situation relations. In this view, the granularity and temporality determine which situations can be, should be, or are currently salient within an interdependent ontology. Figure 15 illustrates such an ontology as a set of situations organised by granularity and temporality, showing a nesting in which situations are decomposed and stretched across time.




Figure 15, Situations nested by granularity and temporality within a context.


As indicated, nesting may point to a single situation ( $S_1$ ) within its broader situational context ( $S_o$ ), and to concurrent or overlapping situations ( $S_2$ ,  $S_3$ ) that may be interdependent (connector), where one shapes the other.  $S_3$  may itself be decomposed into sub-situations ( $s_{3.1}$ ,  $s_{3.2}$ ). Across this nesting, a human or technology may flow between situations by granularity and over time (dotted arrows). From these notions, a question for analysis may be which situations are salient within a context, at what granularity, how they shape one another, and how a human or technology moves between them over time. For example, Paper I suggested a mission context may be decomposed into salient situations, with a transfer of control further divided into phases. Paper III further suggested that concurrent situations may be reorganised to cope with demand, where attending to one defers another and their trade-offs shape how a pilot moves between them.

Taken together, the situation as framed provides a structured account of how a human or technology takes a state of affairs to be the

case: a frame directed by purpose, its content fitted within a bounded structure, and its situations nested within one another over time. The account is not tied to any one situation, and it can be applied to a given context to ask where a framing diverges across members, rests on what cannot be verified, or is held at the wrong grain, and where each threatens coordination.

### Agency as Regulating

Within the framework, intentionality is understood as existing within agency. Here, agency is a property of identifiable entities that arises through autonomous organisation of intentional processes in their coupling with the environment (Barandiaran et al., 2009). Paper III identified four central activities through which the entity modulates this coupling: perceiving, choosing, deciding, and acting. Figure 16 illustrates this coupling as a recursive loop, with the four activities modulating the entity's relation to the environment.



*Figure 16. Agency as a recursive loop of intentional processes coupled with the environment.*

Within this coupling, the activities form the interface between the entity's internal orientation (e.g., intent) and its external realisation (e.g., outcome). In use, they connect the two. As Paper III indicated, the loop is recursive rather than sequential, with shortcuts that bypass parts of

the cycle. The activities are enabled by the Background, operate in the Network, and may become salient in the Foreground as the situation demands (see The Strata of Intentionality). A possible question for analysis is how the loop operates in a given situation, and which activities surface, recur, or remain latent. Each activity is examined in turn.

**Perceive** can broadly be understood as the activity through which the human or technology takes in data, including the effects of its own action, and integrates it into awareness. It is conceptualised to operate through framing, where data is given meaning by the frames the team member has adopted (see The Strata of Intentionality, The Situation as Framed). Thus, what is perceived is selective, which bounds the loop that follows: what does not reach awareness cannot be chosen or acted on. One question for analysis is therefore how perceiving builds awareness, and how this bounds the choices a human or technology can consider. For example, Paper III indicates that pilots build awareness of the wingman by tracking it, while Paper IV indicates that what was perceived could bound the options considered, sustaining a confident but mistaken interpretation.

**Choose** refers to the evaluative activity through which the human or technology considers alternatives. The alternatives available are constrained by the current awareness, and the criteria for weighing them are drawn from the Network of intentional states (see The Strata of Intentionality, The Situation as Framed). In this sense, choice brings beliefs and desires together: an alternative is weighed by what the entity takes to be the case and what it seeks to bring about. A possible question for analysis is what conditions an alternative must satisfy to be considered, and what further conditions it must satisfy to be decided. For example, Paper III found that when an option no longer satisfied current criteria, such as maintaining data links within set bounds, pilots widened the space of possibilities, reconsidering the criteria by which alternatives were judged.

**Decide** can be understood as the activity through which the entity settles a choice, fixing what it takes to hold, whether epistemic or conative. The resulting decision may feed back into the Network of intentional states, where it can hold beyond the moment of deciding and remain open to reconsideration when new relevant beliefs or desires

arise. One question for analysis is what conditions prompt a decision, and what conditions prompt its reconsideration. For example, Paper III found that pilots decided on courses as situations developed, treating these decisions as stable references that could be reconsidered and revised as conditions changed. The same analysis suggests that interpretations could function as provisional decisions, held as the situation was understood and revised when they no longer fit.

**Act** concerns the activity through which the entity enacts a decision in the world. It enacts a practical decision in accordance with its conditions of satisfaction, toward the desired end. Of the four activities, action alone engages the environment: it is where the entity's internal orientation reaches outward conduct, and its conduct and outcome may feed back into perception. One question for analysis is whether the performance can be carried out and an outcome produced, and whether they meet the conditions of satisfaction set for it. For example, Paper III found that pilots acted to uphold degrading data links through navigation adjustments and system management, such as reducing distance or optimising power output, while monitoring the data link trend to judge whether these actions held.

Across these activities, the loop is driven by the distance between what is and what is desired, and works to reduce it (see *The Strata of Intentionality*, *The Situation as Framed*). This distance need not arise from perception alone, but may surface at any activity: perceiving may register a state far from what is desired; choosing may find no alternative that reduces it, or alternatives that conflict; deciding may settle on none, or on one that does not close it; and acting may fail to move the world closer, or widen the distance rather than close it. Any of these may start or restart the loop. A potential question for analysis is where the distance surfaces, and how the loop responds.

Taken together, agency provides a structured account of how an entity operates intentionally: a recursive loop of perceiving, choosing, deciding, and acting that works to reduce the distance between what is and what is desired. The account is not tied to any one kind of entity, and it can be applied to a given situation to ask where intentional processes hold and where they break.

# The Agentive Unit

In the framework, the agentive unit identifies the entity in view within a hybrid team configuration, scoping the intentionality and agency to be examined. Drawing on the generic agency definition (see Agency as Regulating), an agentive unit is any identifiable entity that autonomously regulates its activity dependencies to satisfy its intentionality. It draws on the earlier notion of team intentionality (see Structural Foundation), reflecting a possibility for perspectival shifts, from the individual to the collective. From this perspective, a selected agentive unit brings into view a configuration of agentive units, tied together across grains within a social ontology, from components to collectives. Figure 17 illustrates one such ontological configuration within an operational environment, where a collective ( $\Gamma$ ) is composed of constituents ( $\alpha$ ,  $\beta$ ), and each constituent has components ( $C$ ), tied to activities ( $A$ ) connected by uni- and bidirectional interdependencies ( $\rightarrow$ ,  $\leftrightarrow$ ).




Figure 17. An agentive unit configuration within an operational environment.

Examined from a command structure perspective, the agentive units may be Command and Control as an encompassing collective, with teams as constituents and their members as components. These are configured through interdependent activities, where the output of one feeds another. It is in such activity that intentionality and agency meet, so without it no binding holds, whether between members or teams. Activity is therefore the basis from which relations emerge across the ontology, examined here along three aspects: agentive participation, activity coupling, and configurational dynamics.

**Agentive participation** establishes how members take part in the activities, distributing the intentionality and agency they hold. It surfaces where their epistemic and conative states converge or diverge, as the demand for alignment rises from independent through interdependent to shared activities. For example, examining Figure 17, the three forms differ in the alignment they demand. An independent activity is held alone (disjoint,  $C_1-A_i$ ). Interdependent activities complement one another across constituents (relative complement,  $C_2-A_{ii}$ ,  $C_3-A_{iii}$ ). A shared activity overlaps in a common region (joint,  $C_n-A_n$ ). Drawing on these distinct relations, a question may be whether intentionality aligns to the degree each configuration demands. For example, Paper II indicate a mix of participation types: relative complement in their mission tasks, and joint in supervisory control over the wingman task.

**Activity coupling** establishes how the output of one activity becomes input to another, carrying the outcome its agency has realised. From an intentionality perspective, this coupling may take non-directional, unidirectional, bidirectional, or multidirectional forms, determined by the agentive unit under analysis. Such forms are illustrated in Figure 17, which indicates the directionality of outcomes across activities. A disjoint activity is non-directional, its outcome contributing to the team independently rather than to another activity ( $A_i$ ). A bidirectional interdependence is reciprocal, where two activities contribute outcomes to each other within constituent  $\alpha$  ( $A_{ii} \leftrightarrow A_{iii}$ ). A unidirectional interdependence is one way, where an activity in constituent  $\beta$  contributes its outcome to one in  $\alpha$ . It poses questions such as which intentional processes feed one another, what outcomes are

expected between them, and toward what shared end they converge. For example, Paper III shows the pilot holding two responsibilities, air reconnaissance and supervisory control, as separate processes, with activity couplings across them being about and directed at a state of affairs.

**Configurational dynamics** establishes how the agentive unit maintains its viability over time, in accordance with its intentionality. Operationally, this requires sufficient stability and adaptability of a configuration through, for example, resistance, redundancy, or reconfiguration in the operational environment. In Figure 17, for example, the disjoint coupling ( $C_1-A_i$ ) requires adequate robustness to changing demands, or, should that fail, the adaptability of constituent  $\alpha$  to reconfigure. Likewise, interdependent activities ( $C_2-A_{ii}$ ,  $C_3-A_{iii}$ ) require each constituent to hold its part, or the adaptability to redistribute it across them. In contrast, the shared activity (intersection,  $C_n-A_n$ ) suggests existing redundancy, where the overlap allows either constituent to sustain it. A question may be whether a configuration allows sufficient adaptability to uphold its intentionality as the situation changes. For example, Paper II shows such cases of team reconfiguration, where the pilot and wingman are coupled into one agentive unit by a transfer of control, or decoupled from it by link loss.

Together, the agentive unit sets whose intentionality and agency are in view, not as a property of a standing entity but of a configuration brought into focus for analytical purposes. Through participation, coupling, and dynamics, its agency establishes and adapts the configuration it exists within. Examining team intentionality thus starts not with what counts as a team, but with the unit whose intentionality and agency are in question.

### Autonomy as Relations

Autonomy here refers to the relational condition of an agentive unit's self-sufficiency and self-directedness over its own agency in its coupling with the environment and other agentive units (See Foundations in Hybrid Teaming, Agency as Regulating, The Agentive Unit). Thereby, the locus of the agentive unit under analysis determines its dependencies: independent, dependent, or interdependent with the environment and

interfacing units. From this, the framework suggests autonomy may be analysed as a degree of freedom: the set of choices connecting the unit’s perceptions to actions in an examined situation. Such sets may be represented as autonomy boundaries in matrices, as Figure 18 illustrates for two members,  $\alpha$  and  $\beta$ , whether human or technology, of a unit  $\Gamma$ .<sup>2</sup> For each member ( $\alpha$  and  $\beta$ ), the can and may boundaries delimit which choices remain open for deciding, connecting the perceptions (P) and actions (A) through which it is independent, dependent, or interdependent. Synthesised, the boundaries summarise the requirements and opportunities at some point open for  $\Gamma$ . The representation affords analysis at two grains: locally, a single choice under analysis, and globally, the full matrix as the complete span of the situation examined. From the representation, three aspects become analysable: open choices, boundary dynamics, and autonomy fit.




Figure 18. Autonomy boundaries for members  $\alpha$  and  $\beta$ , synthesised into unit  $\Gamma$ .

<sup>2</sup> See Modelling Intent Variation (Hammarbäck, 2022). The figure extends the intent matrix diagram to illustrate members’ can and may separately and their synthesis.

**Open choices** concern the activities afforded to the agentic unit's deliberation, those choices it both can and may decide in joint human–technology activity (See Agency as Regulating). In Figure 18, each open (available) choice traces locally to who holds it: to one member alone ( $\alpha/\beta$ ), marking its independence or relative complement; to both ( $\alpha\beta$ ), marking a shared opportunity; or only jointly ( $\alpha\rightarrow\beta$ ), marking an dependence. Closed (unavailable) choices trace likewise to whether the can, the may, or both are lacking, and for whom; a row closed across the examined situation (P4) marks a situation without global autonomy for the unit ( $\Gamma$ ). Following such tracing, questions arise as to whether and to what degree autonomy exists, and at what agentic unit in hybrid teaming. For example, Paper III showed how pilots, adapting to Commander's Intent, imposed plan boundaries on the wingman, constraining it at the plan level while leaving the choices below open to its own deciding.

**Boundary dynamics** concern when choices open and close, changing who can and may perform what in the examined situation. While Figure 18 presents the autonomy boundaries globally by summarising local choices, these choices are also bound by time. The framework posits that choices may stand open or closed throughout, or open in windows while their can and may overlap, affording requirements and opportunities for the constituent members ( $\alpha$ ,  $\beta$ ) or the unit ( $\Gamma$ ) as a whole. Windows may further follow decisions, as one member's settling may open or close choices for another ( $\alpha\rightarrow\beta$ ). One question for analysis is thus when and how long each choice stands open, in what order, and whether the openings mark a temporal independence, dependence, or interdependence. For example, Paper III showed how pilots held set criteria for data links, first adjusting activities as links degraded and, only when adjustments no longer sufficed, transitioning to deliberation over plans, criteria, and goals, exhausting choices bottom-up across the levels of cognitive control.

**Autonomy fit** concerns whether the boundaries under analysis hold, locally and globally, toward the world and across agentic units (see The Situation as Framed, The Agentic Unit). Toward the world, fit is the epistemic condition of the unit's intentionality: the boundaries it ascribes must be satisfied by the real state of affairs, as open choices may

stand only in accurate perceptions (P). Such ascriptions may be inaccurate from the first, or become so as boundaries move. Between units, the demand for fit follows the traced forms: choices held by one member alone ( $\alpha/\beta$ ) demand that those depending on them fit the holder; shared choices ( $\alpha\beta$ ) demand that members fit one another on who takes them; and jointly held choices ( $\alpha\rightarrow\beta$ ) that one member's may fit the other's can. Closed choices matter likewise, as a choice presumed unavailable may stand open, withholding the unit's autonomy. The framework thereby posits that incongruences may create false alignments, where members act on boundaries that do not hold in hybrid teaming. One question for analysis is thus how well the boundaries fit, toward the world and between units, and how the fit holds as situations unfold. For example, both Paper IV and Paper II showed pilots presuming divergent boundaries: in Paper IV, from confidently held yet inaccurate interpretations of the synthetic wingman's state, and in Paper II, from framing the same link loss differently.

Across these aspects, autonomy is analysable through the choices open to an agentive unit: which stand open and to whom, when they open and close, and whether the boundaries presumed hold toward the world and between units. Each shifts with the locus under analysis, as what a member depends on another for, the encompassing unit may keep open from within. Thus, from a team intentionality perspective, a team's autonomy is analytically distinct from its members'.

### Intent as Commitment

Drawing on previous sections, intent may be understood as the conative directedness of intentionality that coordinates and guides agency to reduce the distance between what a state is believed to be and what it is desired to be (see Conative Directedness, Agency as Regulating). Where agency settles decisions in the moment, intent is the structure those decisions form and persist within, from which agency continues to act. To provide an ontological account of what intent is and how it can be analysed, the framework adopts Paper I's definition and operationalisation, expressed as:

**Definition:** Intent is a structure comprised of choices characterised by commitment, with the functional roles of coordinating and guiding cognitive and physical activities personally and socially.

**Operationalisation:** Intent is deliberated and connoted choices connected by context–situation and means–end reasoning within an intent space with six levels of cognitive control.

For team intentionality, the definition and operationalisation point to two dimensions: what intent is constituted of and how it functions in joint human–technology activity. The constitutive dimension concerns what intent comprises; the functional dimension concerns how it operates. The following sections elaborate each in turn.

### The Constitutive Dimension

The constitutive dimension concerns the building blocks of intent: the choices it may comprise, the commitments that settle them, and the temporal structure that arranges them. These are illustrated in Figure 19 and elaborated in the following.




Figure 19. An illustrative cross-temporal intent structure.

**Choice content** is what an agentive unit considers within situated intentionality and agency, presupposing the epistemic beliefs and conative desires from which intent is formed (see Conative Directedness). In the framework, it is not predefined but bounded by the situation, where the intentionality strata constrain what becomes salient and accessible for awareness (see The Strata of Intentionality, The Situation as Framed). It may concern any aspect through which intent feasibly becomes actionable. For example, Paper I and Paper II suggested content across the six levels of cognitive control within an intent space (see Figure 19). Across these levels, Paper III further indicated that content may also concern agentic activities such as what and when to perceive, choose, decide, and act (*italics in the figure*). As such, choice content is specified across the levels and time, from the ascribed situation through the ends and the means that transition toward a desired state of affairs (see Conative Directedness). A question for analysis may be what aspects of intent are or should be under consideration, and whether the content of these is shared or complementary between members in joint human–technology activity. For example, Paper III found that new perceptions reshaped what was under consideration as the situation evolved, including whether to act at once or defer to later.

**Conative commitment** concerns those choices, among those considered, that have been decided to establish the motivation and orientation of intentionality. In this view, the psychological mode intend fixes the choices considered, establishing the commitment characteristics that make it distinct from other conative attitudes, such as wants and desires. Further, a commitment may be deliberated in the Foreground, or connoted by structural association within the Network (see The Strata of Intentionality). In the framework, both deliberated and connoted commitments are intended by default, though at a resolution that further reasoning may refine. Paper I further suggests that commitment carries a ‘strength,’ ranging from tentative to firm, and conditions under which it may be revised or revoked (see Conative Directedness). A question for analysis may be which choices are committed, whether deliberated or connoted, how firmly, and under what conditions they may change in joint human–technology activity.

For example, Paper III suggests that demands such as time pressure, task difficulty, and risk exposure may require prior commitments to be revised or revoked to cope with ongoing demands.

**Temporal structure** concerns the organisation of commitments that coordinate and guide mental and physical activities. Paper I and Paper III suggest such structures may be full or partial, yet need to be internally coherent and sufficiently complete to be actionable. Paper I further suggests commitments are connected through context–situation and means–end relations, making the reasoning between decisions traceable (see Figure 19). Paper III extends these relations across past, present, and projected episodes, where the temporal ordering of commitments becomes a further dimension for analysis (time envelope in Figure 19). A question for analysis may be how commitments are organised across time and levels of cognitive control, whether this organisation is coherent within and between members, and where partial structures are reused or missing in joint human–technology activity. For example, Paper III indicated that pilots drew on full or partial segments of established intent structures to enact intent as the situation evolved, rather than deliberating each anew in the present. Figure 19 illustrates this, as the activities at lower levels of cognitive control specify established intent after a trigger cue.

Together, these constitute intent as a structure: content under consideration, settled by commitment, and arranged across levels and time. What matters for analysis is less any single choice than whether the three cohere, whether what is considered, what is committed, and how it is organised hold together as an actionable whole, within a member and between members in joint human–technology activity.

### *The Functional Dimension*

Beyond its constitution, intent functions in two ways. Coordinating concerns how commitments relate to one another, keeping the structure coherent; guiding concerns how a commitment controls its enactment, directing conduct. The following elaborates each.

**Coordinating commitments** here refers to the management of intent interdependencies within and between participants in joint human–technology activity. In the framework, the purpose of

coordinating commitments is to keep intent coherent and actionable within the agentic unit. As Paper I and Paper III suggest, a settled commitment both anchors and constrains further reasoning (see Conative Directedness). It anchors as stable input for continued deliberation, whether as a full or partial intent structure. It constrains as a filter, focusing the intent space to available choices while reducing the need to reconsider what is already committed. The same decision anchoring and choice admissibility holds between intent structures, personally or socially: decisions anchored in one structure may filter inadmissible choices in another. Figure 20 exemplifies decision (D) and choice (C) pairings, D–D and D–C, by admissibility.




Figure 20. Paired commitments by status and admissibility.

Each pairing either admits or demands resolution. An admissible D–D pair holds together, and partial intent may be assembled incrementally into a fuller structure. An admissible D–C pair lets the prior commitment frame the choices admitted, an anchor for further consideration and decision. Where it demands resolution, this follows the pairing: an inadmissible D–D pair requires one to be revoked to restore consistency or compatibility, while a D–C pair is asymmetric, the choice filtered or the decision reconsidered. A question for analysis may be where a human or technology holds commitments that are inadmissible together, and whether their status leaves the conflict

resolvable in joint activity. For example, Paper III found that pilots coordinated commitments with the wingman in anticipation of degrading data links, revising prior commitments to keep intent admissible across the team as the situation changed.

**Guiding conduct** here refers to how intent becomes realised, as deliberated and connoted commitments direct joint human–technology activity (see Conative Directedness). Drawing on Paper I and Paper III, the framework conceptualises intent not only as a temporal structure of coordinated commitments, but also as a control structure that directs mental and physical activities. Thus, whether a commitment is to perceive, choose, decide, or act, each guide conduct toward its end. Further, such guidance is suggested to be enabled by anticipation, preparation, and recognition of a situation. From this, a commitment is specified from coarse to finer-grained across the levels as the situation nears, where recognition may trigger enactment in the world, as Figure 21 illustrates. It is here suggested that this may largely function by nesting (see The Situation as Framed), whereby a commitment opens a sub-situation carrying its own levels, within which the same context-situation and means-end reasoning runs at finer grain.




Figure 21. Specification and enactment of a commitment over time.

Drawing on the figure, guidance extends from prior- and future-directed intent to present-directed intent-in-action, allowing commitments to be distributed over time, from coarse anticipated

situations to the specific object configurations. Further, recognising that the anchored situation may trigger a transition to a performative state, effectively shifting the conduct from a deliberation mode to a monitoring mode, where the established commitments are enacted and adjusted.

Together, this distribution over time and the shift to monitoring let intent be performed under constraints: its load falls across the horizon rather than at the moment of action, which can relieve situational demands in hybrid teaming. A question for analysis may be how the tempo of shifting between deliberation and monitoring aligns between a human and technology, whether they hold commitments at different grains, and whether the shift runs one way or moves back and forth as the situation unfolds. For example, Paper III found that pilots moved between modes of reactive and proactive control: a recognised cue, such as a data link disturbance, triggered established commitments immediately into navigation and system adjustments, while longer-horizon deliberation over plans and goals unfolded when these responses no longer sufficed.

Across the two functions, intent is analysable as commitment held in two senses at once: coordinating keeps it coherent as settled content, guiding carries it into conduct as the situation arrives. Because the two are orthogonal, joint activity may fail in either way: a coherent structure left unenacted, or conduct running on commitments that never cohered. Thus, from a team intentionality perspective, a breakdown is locatable as one or the other, and the two call for different responses in hybrid teaming: coherence restored among commitments, or timing restored between deliberation and the moment that demands it.

### Aligned Intentionality

In the framework, aligned intentionality concerns the sharedness among team members and fit with the world. It is specified along the epistemic and conative sides of intentionality, as awareness and intent, whereby discrepancy in either affords a need for realignment (see *The Situation as Framed*, *Intent as Commitment*, *Agency as Regulating*). Such realignments, it is suggested, rest on perspectival shifts, as members move between individual and team modes to uphold their ethos (see *The Hybrid Team*, *Structural Foundation*). This shift is here conceived as

members taking a bird’s-eye view, holding not only their own frame but that of each teammate and of the team as a whole: the me, the you, and the we (see Mindreading). Drawing on these notions, Figure 22 conceptualises this as an alignment plane with two orthogonal dimensions: the sharedness among members (vertical) and the fit with the world (horizontal). Placing a member’s awareness and intent within it, where they stand and where they are heading, grounds the satisficing alignment and positional trajectory that follow.




Figure 22. The alignment plane: awareness and intent positioned by sharedness and world-fit.

**Satisficing alignment** concerns the awareness and intent required of members, in the sharedness and world-fit a given end calls for. Drawing on the two dimensions, sharedness sets requirements and offers opportunities through independent, dependent, or interdependent activities. In Figure 22, these largely correspond to the modes, from the individual to the team, addressing the me, the you, and the we. For example, comparing awareness (circle) and intent (square) indicates where each is (empty) and where it should be (filled), signalling

whether a perspectival shift is warranted. Similarly, world-fit may place a different demand on awareness and intent, where satisficing may serve better than reaching for the optimal. For example, using the figure, members need to recognise the unsatisfied, whether awareness or intent, to motivate realignment. If awareness (epistemic), it must be brought to fit the state of affairs; if intent (conative), the state of affairs must be brought to fit it (see Epistemic Aboutness, Conative Directedness). One question for analysis is whether a member holds the mode the activity warrants, and whether its world-fit is satisficed and any divergence recognised. For example, Paper III indicates pilots may, working independently, use a satisficing strategy holding stances of a me and a you during link loss.

**Positional trajectory** places an agentive unit's awareness and intent within the plane and traces where each is heading, so the distance to where they should stand, and whether the unit is closing it, become analysable. From Figure 22, a member's awareness and intent (square) are marked at their present position, empty where unsatisfied and filled at the satisfied target. A solid arrow projects where each is actually heading; a dashed arrow marks the optimal path it should take and where that path could reach. The figure here affords two analyses. Where the solid departs from the dashed, the unit is not on its optimal path, heading from where it should for alignment. Where the solid follows the dashed yet falls short, the unit cannot presently arrive on its own (see *Autonomy as Relations*). It poses questions such as whether a unit's projected course meets where it should hold, whether its optimal path can reach at all, and where it cannot, what it takes to be the case. For example, Paper IV found participants confidently holding wrong interpretations of the wingman, suggesting their awareness could have reached its proper position had support been sufficient.

Across these aspects, aligned intentionality is analysable through awareness and intent placed within the plane: how far each stand from where the situation asks, and whether the unit is on course to close it. Where they are, where they should be, and where they can reach may all be moving at once, the target shifting as the situation does. Thus, from a team intentionality perspective, alignment is a standing achievement, upheld as situations shift rather than secured once.

## The Analytical Account

The framing developed team intentionality through eight aspects, from the hybrid team to aligned intentionality, each posing its own questions for analysis. Together they are the framework's answer to RQ1: the core concepts and relations through which team intentionality becomes analysable and communicable. The aspects form an account rather than a checklist, their parts conditioning one another. The unit must be fixed before its intentionality can be examined, its agency before what it may do, its intent before how that intent aligns. None reduces to another, and removing any one leaves a distinct question of analysis unasked. The same running situation may be traced through all eight aspects: the pilot and wingman on a reconnaissance mission, through degrading data links and link loss. Each examines it from its own viewpoint, and the questions for analysis converge on the one situation rather than dispersing. Even so, other cases would serve as well; this single situation only makes the aspects easier to see together.

Each aspect may also be understood as a place where a unit's awareness or intent can fail to hold: toward the world, where what it takes to be the case departs from what is, or toward one another, where what one holds departs from another's. From this, such misfits become locatable at a chosen unit: which awareness rests on what cannot be verified, which choice is presumed open that is not, or which alignment is assumed that does not hold. Thus, from a team intentionality perspective, the account makes these misfits analysable before they manifest in action, since a misfit may be held as confidently as a fit. What the analysis locates, design may then work to reduce: the same misfits traceable from here to the artefact released.

## Realising Team Intentionality

Realising team intentionality concerns the practical means by which the framing is brought into design, addressing how the framework's concepts and relations can be used in the design of hybrid teaming (RQ2). Where the previous section set out the concepts and relations central to analysing the phenomenon, this one turns to using that

account in design: first built into the models a designer works with, then applied to examine what must be aligned between members. A conceptual account aids design only once it can be operationalised in this way, so that the misfit located in analysis remains traceable to the artefact released. This section sets out these means in two moves: building a toolkit, the models that carry the framing, and then putting it to work, applying those models to what should be disclosed and clarified between members. Across both, questions are posed using design thinking (see Appendix), and a concluding synthesis draws the means together.

### Building a Toolkit

Building a toolkit shapes the models a designer works with, carrying the framing's aspects into the Analyse stage to locate what a design need to support. It does so in three moves. It begins with establishing reference points, the shared objects that hold the design process to a common account. It then selects what to represent, reducing the phenomenon to what a purpose requires, and builds and manipulates models, extending representations to foreground aspects. A concluding note then draws these together.

#### *Establishing Reference Points*

A reference point here functions similar to a boundary object that supports shared understanding across participants in the design process, taking the form of artefacts such as definitions and operationalisations, methods, models, and instantiations (see Modelling Intentionality). The framework posits that established reference points give the design process a stable object to work against, so that what is recognised, analysed, conceptualised, actualised, and finalised can be traced to the same object. Figure 23 illustrates this through Paper IV: the definition and operationalisation of intent seed the situated intent and intentional process models, which converge on the intent transparency operationalisation and model, informing both the HMI design and its evaluation.




Figure 23. Reference points traced across the papers, converging on the intent transparency instantiation in Paper IV.

A similar tracing can be made through the approaches explored across Papers I–III. Though differing in purpose and focus, they share many aspects: the conceptualisation of intent, the levels of cognitive control, and core elements such as choices and decisions. This shared basis enables a coherent mapping between methods and models, establishing reference points that keep translations and transformations between them aligned. These reference points should therefore be established as coherent and shared among participants when designing hybrid teaming. One question for analysis is therefore which reference points to establish, whether they are coherent, and how they are kept aligned among participants as realisations translate and transform across the design stages. For example, Paper IV drew on a single operationalisation of intent transparency to inform both the HMI design and the evaluation instrument, so that awareness was measured against the same reference the design was built to support.

### *Selecting What to Represent*

Selecting what to represent here concerns making explicit which aspects of team intentionality a model carries, rather than leaving them

implicitly assumed (see Modelling Intentionality). In the design process, the complexity of joint human–technology activity may need reducing to make a model useful, focusing on what matters from the questions the framework may pose (see Appendix). Given this plurality, a design may draw on several complementary but related representations. Across the papers, the Joint Control Framework has been the cornerstone from which adaptations were made to represent intentional states (Papers I–II) and processes (Paper III), sharing a basis while providing complementary views. For example, both share the levels of cognitive control and the choices considered and decisions made; however, each affords a different analysis. Figure 24 illustrates this: the situated structure holds the paths taken (decisions) and not taken (choices), breaking where a branch cannot hold; the intentional process takes the decided path and decomposes it into activity points over time.




Figure 24. The situated intent model and intentional process model as complementary representations.

Thus, the framework does not rest on one representation but may combine complementary model types, each exposing a different aspect of intentionality. Different purposes may therefore warrant different representations, while sharing reference points to stay coherent. One

question for analysis is therefore which aspects to represent, under which analysis, and how the representations may stay connected through shared reference points. For example, Paper III represented intent as a deliberative process, addressing the static view of the situated intent model in Papers I–II. In doing so, it also recognised that its own representation, by the method selected, foregrounds certain aspects of intentional states while omitting others, pointing to complementary representations for what it cannot hold.

### *Building and Manipulating Models*

As established, modelling team intentionality is not the pursuit of a single complete model but a plural activity, where a model is iteratively built and manipulated, each manipulation affording a different means of learning from it (see *Modelling Intentionality*). Building and manipulating models is, in this sense, how the implicit is made explicit, and an object of study is foregrounded. As Table 6 indicates, designers may need to foreground several aspects.

*Table 6. Aspects of team intentionality a model may represent.*

| <b>Aspect</b>      | <b>Representational content</b>                                                                                                                                  |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Ethos              | Collective beliefs, goals, values, norms, standards, traditions, practices.                                                                                      |
| Team               | Structures and configuration; degree of teamness.                                                                                                                |
| Members            | Skills, abilities, preferences, dispositions; roles and responsibilities.                                                                                        |
| Environment        | Conditions and constraints; systems and equipment.                                                                                                               |
| Task               | Processes, strategies, procedures, contingencies; interdependencies; interaction modes; the perceiving, choosing, deciding, and acting through which it unfolds. |
| Autonomy           | The agentive unit and its relations; what it can and may do; the choices these leave open or closed.                                                             |
| Intentional states | What is believed, desired, and intended; awareness and intent; alignment against conditions of satisfaction.                                                     |

Drawing on Human Factors methods, as noted, there are several avenues for representing such aspects, such as causal cognitive maps of beliefs, task diagrams of goals and plans, or hierarchical graphs of goals and subgoals, with or without adaptation (see Representation). Such manipulation is not tied to any one approach. One question for analysis is therefore which method can represent the selected aspects, and whether they serve as they are or require adaptation to foreground what a design needs. As the framework is grounded in the methods applied across Papers I–IV, these serve here to demonstrate how a representation can be manipulated to foreground other aspects: first, decomposing situated intent models to foreground belief and desire; second, relating intentional process models to foreground dependencies between activities.

**Decomposing situated intent** models separates an intent into the belief and desire it rests on. It draws on the conception of intent as a belief paired with a desire, held with a strength understood as the willingness to exert effort at the threshold (see Conative Directedness). Figure 25 illustrates this, with strength shown by border weight. For example, what is believed may draw on epistemic states such as knowledge or assumptions; what is desired on likes and needs. From this, the framework suggests that rooting an intent in what it rests on may locate why it holds or fails: a belief not matched to the situation, a desire without a believed means, or a commitment held by necessity rather than desirability. In this manner, what a model of intent leaves implicit, its formation and strength, may be made available to design.




Figure 25. A situated intent model decomposed into belief and desire.

**Relating intentional process models** connects the couplings between members’ score episode diagrams to expose their autonomy. It draws on activity coupling, where a process’s outcome may feed another, and on autonomy as the complement of dependence. Figure 26 illustrates this, showing couplings within a process, unidirectional between processes, and bidirectional between them. For example, a coupling may be read as independent where processes proceed alone, dependent where one feeds another, or interdependent where each feeds the other. From this, the framework suggests that relating scores episode diagrams may locate where a process’s autonomy is bounded by another: a reliance on an outcome not yet produced, or a reciprocal coupling that neither can settle alone. In this manner, what a single process leaves implicit, how one process’s autonomy depends on another over time, may be made available to design.




Figure 26. Relations within and between intentional process models.

Beyond these, the framework's own illustrations offer further representational examples: the interaction models illustrating teamness requirements (Figure 12); the framed situation as believed and real

(Figure 14); the intent matrix as autonomy boundaries (Figure 18); and the alignment plane as a means to trace epistemic and conative states (Figure 22). Each offers an example representational form that may support analysis and communication, even where no paper applied it as such.

Together, such tools help realise team intentionality. Reference points hold the design process to a shared object, so realisations translate and transform without losing it. Selecting what to represent reduces the phenomenon to what a purpose requires, drawing on several representations that stay connected through that shared basis. Building and manipulating extends those representations to foreground aspects they were not built for, from the belief and desire beneath an intent to the autonomy between processes. From this, the framework suggests that representing team intentionality is not the pursuit of one complete model, but the coherent use and extension of several. What differs across them is therefore not the phenomenon but what each is made to foreground, and what a design must connect between them.

### Identifying What to Convey

Identifying what to convey puts the toolkit to work, examining what intentional content is, can, or should be aligned between members. Elements may be disclosed to make activity interpretable, and clarified through their relations to make it understandable (see Transparency and Explainability). The two are ordered: what is clarified must first be disclosed, though what is disclosed need not be clarified. Drawing on the operationalisation of intent (Paper I), its extension to intent transparency (Paper IV), and the activity points (Paper III), the framework generalises to analyse transparency and explainability in three moves, each cutting differently at one question, what to disclose and clarify between members. Sharing intent structures uses score episode diagrams; aligning situations and intent uses situated intent models; bounding what is open uses intent matrices. This chiefly informs the Conceptualise stage, where a design's content is shaped by what members must convey to align their intentionality.

Sharing Intent Structures

Intentional process models here utilise score episode diagrams to position a member’s elements across levels and time. Figure 27 illustrates such a structure, one a member may hold, and another may need to interpret and understand. From it, three operations follow: tracing questions to answers, positioning elements in time, and making structures available.




Figure 27. Tracing explanatory questions within an intent structure (illustrative).

**Tracing questions to answers** locates elements and follows them to other elements by their relations. It draws on Paper I, suggesting that relations provide explanatory power, and Paper III, suggesting that integrating activities across levels and time is necessary for convergence at a decision point within an intent structure. From this, the framework suggests that such structures may support both positions: disclosure, where a question is answered by content disclosed at the point itself, and clarification, where it is answered by content disclosed at a further point traced from it. For example, in Figure 27, one member may hold the structure that another needs to interpret and understand. To this end, Table 7 exemplifies eleven questions the interpreting member may pose, of which the first five are answered by disclosing content and the remainder by clarifying content.

Table 7. Questions for tracing disclosure and clarification within an intent structure.

| Question     | Meaning                                           |
|--------------|---------------------------------------------------|
| 1. What?     | The intentional content committed                 |
| 2. Who?      | The member holding the commitment                 |
| 3. Who for?  | The member needing to interpret and understand it |
| 4. When?     | The time it applies, and the duration it holds    |
| 5. How firm? | The strength with which the commitment is held    |
| 6. Why?      | The motivation, as reasons or causes              |
| 7. How come? | What the commitment followed from                 |
| 8. What for? | The end it is directed toward                     |
| 9. How to?   | The means through which the end is met            |
| 10. Why not? | The alternatives held open when it was settled    |
| 11. What if? | What would follow were conditions otherwise       |

Drawing on such questions, participants across Papers I to IV faced hypothetical situations where they arise. Encountering an aircraft during a link loss episode, a pilot may ask who it is, what it is committed to, why it is committed to that, and how to resolve the event. Others may be left unposed. One matter for analysis is therefore which elements are, can, or should be disclosed and clarified, whether as a single element or a joined segment, as the selection may help align team intentionality through transparency and explainability.

**Positioning elements in time** places them as past, present, or projected in joint human–technology activity. It builds on transparency and explainability as means of proactive preparation and reactive repair, through which members disclose elements about themselves as situations change (see Transparency and Explainability). Paper III illustrates this after link loss, where reconnection synchronised information, from which pilots filled gaps in the frame with what had occurred and elaborated it with what needed doing next. From this, the framework suggests that where an element sits in time shapes what its

disclosure may support. A past element may support repair, a present one adaptation, and a projected one preparation. In this manner, the centre of investigation may be moved to any of these positions, and since each carries its own elements and relations, the same questions return different answers from each. One concern for analysis is therefore where the centre is placed, since it determines what repair and preparation may draw on to align team intentionality.

**Making structures available** presents a member's intent structure to another, so that it may be interpreted and understood when disclosure and clarification are impractical or impossible. It draws on Paper IV, suggesting that transparency shifts from a problem of information transfer to one of interpretation under link loss, and Paper III, suggesting that communicated commitments and established intent structures provide the basis for interpreting another member. From this, the framework suggests that representations and simulations of a member's intent structure may stand available across the team. In that manner, below the divide between mind and world, a member is observable in its own right; above it, only what has been made available. Interpretations so formed are provisional, and may require correction once conveying resumes. For example, Paper IV represented the wingman's plans, goals, and activities, and projected its likely position, against which pilots interpreted what they observed. The added cues sped interpretation yet raised errors and left confidence intact. One question for analysis is therefore which elements should stand available in advance, and how they are reconciled with observation, which may mislead.

Together, the same intent structure in score episode diagrams affords three analyses. Traced through its relations, it reveals answers to questions about a commitment, some at the point itself and others at further points reached from it. Analysed by position, it reveals what a disclosure may support, whether repair of what has passed or preparation for what is ahead. Standing available, it reveals what a member holds without being asked. What differs is therefore not the structure but what is selected from it. Paper IV proposed such selection as adaptive transparency, providing what is conveyed, when, and how according to situational needs rather than exposing the structure in full.

*Aligning Situation and Intent*

Situated intent models here nest members’ intent structures under a shared team frame. Figure 28 illustrates this for a unit  $\Gamma$  with constitutive members  $\alpha$  and  $\beta$ . From it, two analyses follow: relating members’ situations under the shared frame, and comparing their structures where commitments meet.




Figure 28. Two members’ intent structures under a shared team frame.

**Relating members’ situations** here nests members’ frames under a shared team frame to locate what must be disclosed and clarified for another to interpret its situation. It draws on Papers I to III, suggesting situations-of-situations, and posits that individual and team frames may be used for perspectival shifts between the me, the you, and the we (see *The Situation as Framed, Aligned Intentionality*). Figure 28 illustrates this: a team frame (parent) relating to members’ individual frames (children). The team frame marks what members hold in common. The connection between their individual frames marks where awareness is shared. From this, the framework suggests that locating the parent and child frames by their relations sets the grain at which they align, so members may disclose and clarify at the level the situation requires. In this manner, independent, dependent, and interdependent activities may call for individual, shared, or team awareness, each requiring a different perspectival shift. For example, Paper IV indicated

participants operating under the mission frame, assuming the wingman did so too, using their own framing to understand the wingman's situation and intent. One question for analysis is therefore at what grain each member places the other's situation under the shared team frame, whether a perspectival shift is warranted to reach that grain, and what must be disclosed and clarified to achieve it.

**Comparing intent structures** allows tracing members' reasoning in joint human–technology activity. It draws on Paper I, suggesting that analysis of situated intent models can identify decisions inconsistent in fact and conflicting in effect. Choices, likewise, may be unused in fact and underutilised in effect. From this, the framework suggests that comparing structures and tracing their decisions and choices pairs them as admissible or inadmissible, within and between members (see The Functional Dimension). Figure 28 illustrates this. Between members, inadmissible pairings from the generic level down mark the origin of divergence, what must be disclosed, and which member holds the choices open to resolve it ( $\beta$ ), clarified as how to realign. Within a member, an inadmissible pairing ( $\alpha$  at physical) calls for disclosure only where it traces up to the team level. For example, Paper III found a pilot replanning at the team level, accounting for both members, once local adjustments no longer met what the situation required. One question for analysis is therefore what must be disclosed and clarified to realign: what the origin is and who holds it, why it diverged, and how to resolve it through the choices left open.

Together, situated intent models here afford two analyses. The grain at which members situate one another sets what a member must convey to be interpreted. Where their commitments fail to admit one another marks where realignment must reach, and who holds the means. From this, the framework suggests that a member's structure cannot be read from the situation alone, but must be conveyed and its frame located against another's. Paper IV proposed such conveyance as ensuring revisability, keeping interpretations open as situations and framings shift.

# Bounding What Is Open

Autonomy boundary models here utilise intent matrices to connect choices and decisions to perceptions and actions (see Autonomy as Relations). Figure 29 illustrates this for members  $\alpha$  and  $\beta$ , and their synthesis as a unit  $\Gamma$ . From it, three operations follow: contrasting decisions with open choices, matching boundaries, and locating what a decision requires of another.




Figure 29. Autonomy boundary models for members  $\alpha$  and  $\beta$ , synthesised into unit  $\Gamma$ .

**Contrasting decisions with open choices** places a settled choice against those that stood open around it. It draws on the framework's conceptualisation of autonomy, where what a member can and may do bounds its deliberation, and on agency, where deciding settles among alternatives found (see Autonomy as Relations, Agency as Regulating). Figure 29 illustrates this for unit  $\Gamma$ , where member  $\alpha$  and  $\beta$  have each settled one choice among those open to them. Locally, a settled decision stands against the cells open around it, answering why not those instead. Globally, the cells open under other perceptions show what would stand open had the situation been ascribed otherwise. Both may be clarified by the members' boundaries, showing whether the can, the

may, or both were lacking where a cell stood closed. For example, Paper IV found participants holding several hypotheses about the wingman in parallel, seeking a match between entity, activity, plan, and goal. From this, the framework suggests that disclosing the decision alone leaves hidden the choices it was settled against — the open cells that answer why not those. One question for analysis is therefore which alternatives must accompany a disclosed decision, and when a closure requires clarifying by the can or by the may.

**Matching boundaries** asks whether the can and may a member ascribes correspond to the actual state of affairs. It draws on situations as framed and autonomy (see *The Situation as Framed, Autonomy as Relations*). In Figure 29, members may draw on boundaries that do not align with the world or between themselves, leaving them to act on false premises. For example,  $\alpha$  and  $\beta$  may ascribe a cause or reason differently, affording different sets of open choices in unit  $\Gamma$  (e.g.,  $\alpha$  at  $P1A1-A2$ ,  $\beta$  at  $P3A3$  alone). Inconsistent or conflicting decisions and actions may then follow. In this manner, each ascription stands as a condition under which one set of choices is open, and another closed. From this, the framework suggests that choices and decisions may be varied as what-if, locally at a cell and globally across a perception or an action, to test what would stand open were a condition otherwise. For example, Paper IV found participants evaluating competing hypotheses against the unfolding situation, testing which ascription of entity, plan, and goal fit what was the case. One question for analysis is therefore which boundaries must be matched between members and against the world, and which must be left to what if where matching is not possible.

**Locating what a decision requires of another** places a settled choice against the type of cell it occupies. It draws on the agentic unit, where dependence holds in three values by direction (see *The Agentic Unit*). Figure 29 illustrates this, where each type follows from how the can and may compose across members. Where a choice stands open to one member alone ( $\alpha/\beta$ ), the holdings are disjoint, suggesting independence. Where it stands open to both ( $\alpha\beta$ ), they intersect, suggesting a shared opportunity. Where it stands open only jointly ( $\alpha\rightarrow\beta$ ), one member's may meet the other's can; the holding is one-way, suggesting a dependence. From this, the framework suggests that the

situation under analysis may require one member's decision, or a combination of decisions across unit  $\Gamma$ . Following a decision, what members can and may do might change, thereby opening and closing choices for the unit. For example, Paper III found pilots first adjusting navigation and systems alone as data links degraded, and only when these no longer sufficed, replanning across the team. Such replanning leaves the wingman's earlier boundaries no longer current. One question for analysis is therefore which decisions change what another can and may do, and which of these merit disclosure and clarification.

Together, intent matrices here afford three analyses. Contrasted with the choices around it, a decision reveals the alternatives declined when it was settled. Matched against the state of affairs, it reveals where the boundaries drawn on misfit what is the case, leaving members to act on false premises. Located by the cell it occupies, it reveals how the choice is held between members, and what settling it changes for them. From this, the framework suggests that no member holds the matrix alone, since the synthesis stands only once each has disclosed its can and may. What differs is therefore not what a member is capable of doing, but the unit's ability to verify the boundaries under which it acts. Paper IV proposed such verification as the comparison of parallel hypotheses, evaluating competing interpretations rather than committing to one.

## The Design Account

Realising set out how the framing is put to use, in two moves. A toolkit builds the models that carry team intentionality into the Analyse stage, foregrounding what a design must support; identifying what to convey then reads those models for the Conceptualise stage, locating what members must disclose and clarify to align. Together they answer RQ2: the framework's concepts and relations become not only describable but workable, a means to represent team intentionality in a design situation and trace what it requires. Where the framing made the phenomenon analysable, the realising seek to makes it operable in the design process.

Beneath these models sit shared reference points: whichever aspect a model foregrounds, it is built from the same operationalisation of intent and the same levels of cognitive control, so several representations trace to one account rather than many. In this sense,

what differs across them is not the phenomenon they carry but what each is made to foreground, and what is selected to convey from it. Even so, other representations would serve, provided they share those reference points; the coherence lies in the basis, not in any one model.

The second move turns on this: a member's awareness and intent cannot be read from the situation alone, so what is believed, committed, or held open must be disclosed to be interpretable, and clarified through its relations to be understandable. Thus, from a team intentionality perspective, what the analysis locates, design may then convey and carry into the artefact, so that the misfit recognised at the outset is reduced in what is released.

## **The Case of Manned–Unmanned Teaming**

With the purpose of supporting the design of future joint human–technology activity, Manned–Unmanned Teaming is the case in focus, a locus characterised as both an envisioned world problem and a wicked one. Manned–Unmanned Teaming has been studied for almost three decades (see Manned–Unmanned Teaming), and remains largely envisioned. Its realisation places demands on workload, awareness, and trust, among others. Understanding and sharing intent was proposed to help meet these, placing aligned intentionality at the centre of the teaming. With the framework now at hand, two questions follow: where the design of Manned–Unmanned Teaming stands, and where it goes from here.

Setting the agentic unit for the joint human–technology activity in view poses the question of how much teamness the configuration requires (see The Agentic Unit, The Hybrid Team). This makes explicit whether the unmanned aircraft is designed as an tool or a teammate, rather than leaving it an assumption. Presently, the unmanned aircraft is envisioned as something between: a collaborator under supervision executing its delegated tasks autonomously within a bounded solution space in the command chain (see The Operational Domain). Whilst hybrid teaming research often promotes perceiving the technology as a teammate, there is merit in asking whether it warrants that status here, where the command chain holds and the pilot retains oversight and

command. Full teammate status is thus questionable: the aircraft acts within delegated authority, not as an equal or true replacement.

Even where the unmanned aircraft is not a full teammate, understanding and sharing intent remains relevant: in the framework's terms, the alignment of awareness and intent the team depends on. Achieving that alignment, as the framework presents it, is where several challenges remain across the human, the technology, and their relation. Among these challenges is developing the repertoire and underlying models on which team intentionality rests (see *The Strata of Intentionality*). For the pilot, it requires making the unmanned aircraft interpretable and understandable, through mechanisms such as communication and socialisation (see *Representational Foundations*). Transparency and explainability, noted earlier to share much with team intentionality (see *Transparency and Explainability*), offer another avenue: the framework sets out means to support the interpretability and understandability of the unmanned aircraft (see *Identifying What to Convey*). For the technology, relying on these mechanisms is harder. Mindreading is a naturally evolved human capacity (see *Mindreading*), whereas its artificial counterpart remains in its infancy. It must instead be implemented as models, where the difficulty lies in making explicit what for humans often remains implicit. As a case in point, the ethos (the goals, values, norms, and standards) is likely among the hardest to fully capture and implement. Nevertheless, whether the design supports the pilot through the interface or equips the unmanned aircraft to comprehend pilot intent, the modelling remains the key point, and the toolkit offers a starting point (see *Building a Toolkit*).

These two threads meet at one point. What the unmanned aircraft is, and what aligning intentionality demands of it, both turn on the same thing: how far the envisioned teamness can be built. The framework does not settle that, but it makes the gap examinable, locating what a given configuration would require against what can presently be modelled and realised. Closing that gap need not be all or nothing. It may proceed in increments, each configuration supporting a little more of the teamness sought, from control toward the collaboration envisioned.

.

# Limitations and Future Work

Three limitations follow from the delimitations set at the outset (see Scope and Delimitations), with a fourth in how far the framework has been taken. The concepts and relations are a bounded selection from a broad and interdisciplinary field. In line with its aim, these are those central to analysing and designing for team intentionality, not an exhaustive account, and a different selection could yield a different framework. Future work may extend the selection with concepts and relations from adjacent fields, or refine those included where design use shows them insufficient. The studied case is drawn from Manned–Unmanned Teaming alone, and whether the account holds beyond this setting has not been shown across domains. It was developed within the applied setting of this research, a salient and representative instance of hybrid teaming. Future work may apply the framework through cases in other forms of hybrid teaming, testing whether its concepts and relations hold as the domain shifts. The intent modelled in the papers is envisioned rather than observed in action, since the teaming does not yet exist operationally and it was accessible only as espoused. The accounts were elicited from experienced domain experts, though what is espoused need not match what is enacted. Future work may study the intentionality in more realistic settings, such as simulation or human-in-the-loop trials, to verify whether the account holds closer to enactment.

The framework was also developed here rather than applied in practice. Its parts are argued and illustrated through the papers, in constructed settings, rather than demonstrated in the field (see Framework Development). Future work may apply the framework in real design projects, where its value can be shown in use rather than argued.



# Conclusion

‘Do you understand me?’ The thesis began with an old question, and an observation: for millennia it passed only between people, intent flowing in one direction, from the maker into the made. As technology is envisioned not as an tool but as a teammate, the question comes to pass both ways. Its purpose was to support the design of future joint human–technology activity in operational settings, a long-term end no single thesis settles. Toward it, the aim was to develop a framework of team intentionality, one that establishes central aspects and provides guidance for their use across the design process. That aim is met: the framework has been developed and grounded in the studies reported across the four papers, framing what team intentionality consists of and how it may be designed for. What it offers the purpose of this thesis is a means, not a settlement; its guidance is argued and illustrated here, with its value in use awaiting confirmation in the field.

The framework answers two questions. To the first, it offers the analytical account: eight aspects that make team intentionality analysable and communicable. Their force is cumulative rather than separate. Each foregrounds a different thing an analyst weighs — whose intentionality is in view, what bounds its choices, how these hold across time, and whether they align with the world and with one another, so that together they bring a single configuration into focus. Because the unit in view is chosen rather than fixed, the account is not tied to Manned–Unmanned Teaming but extends to other human–technology configurations. In it, team intentionality is introduced as a construct in its own right, synthesised from team science and intentionality research, offering a shared vocabulary where the field held fragments.

To the second question, how these concepts and relations may be used in the design of hybrid teaming, it offers the design account: two moves, establishing a toolkit and identifying what to convey. Together they put the framing to work across the design process, as an analytical and evaluative lens on design situations. The account also contributes methodologically: the human-centric methods by which the models are built and analysed, and the operationalisation of intent they carry. Beneath the models sit shared reference points, that operationalisation

and the levels of cognitive control, so that many representations trace to one account rather than many. Put to use, this gives design teams a way to locate misfits at a chosen unit before they manifest in action, and evaluators a way to trace each representation back to that shared account. Thus the framework does not add another model to the field so much as a basis on which its models, and those who design and evaluate with them, may cohere.

Across both accounts runs a single thread: misfit, located, then reduced. The analysis makes misfits locatable at a chosen unit before they surface in action; design then works to close them, traceable to the artefact released. Read this way, team intentionality is never secured once and kept. It is held in place as situations shift, and the framework's work is to make it examinable before it is built for.

The framework also turns back on the question that motivated it. The field takes as its premise that technology is no longer a tool but a teammate. That is not presumed here. Faced with a misfit, the framework lets one ask what the situation actually needs: a teammate, which may still fail to understand; a tool, which may lack the reach the work demands; or neither, a change in how the work itself is organised. Take Manned–Unmanned Teaming as the case: is the unmanned aircraft a wingman, an asset under control, or something between? From this position, team intentionality may be read through the degree of teamness a configuration calls for, at times peer-like, at times closer to control. Both have their place, and the analytical question is when one gives way to the other. However, whichever it is, peer-like or control, one demand does not change: what a member holds must still be made available to be understood. That need for conveyance is what endures across every configuration, and it is what the framework is for.

# References

- Ajzen, I. (1991). The theory of planned behavior. *Organizational Behavior and Human Decision Processes*, 50(2), 179–211. [https://doi.org/10.1016/0749-5978\(91\)90020-T](https://doi.org/10.1016/0749-5978(91)90020-T)
- Alberts, D. S., & Hayes, R. E. (2006). *Understanding command and control*. CCRP Publications.
- Albrecht, S. V., & Stone, P. (2018). Autonomous agents modelling other agents: A comprehensive survey and open problems. *Artificial Intelligence*, 258, 66–95. <https://doi.org/10.1016/j.artint.2018.01.002>
- Alexander, C. (1964). *Notes on the synthesis of form*. Harvard University Press.
- Alonso, V., & De La Puente, P. (2018). System transparency in shared autonomy: A mini review. *Frontiers in Neurobotics*, 12, 1–11. <https://doi.org/10.3389/fnbot.2018.00083>
- Alshenqeeti, H. (2014). Interviewing as a data collection method: A critical review. *English Linguistics Research*, 3(1), 39–45. <https://doi.org/10.5430/elr.v3n1p39>
- Andrews, J. M., Rusnock, C. F., Miller, M. E., & Meador, D. P. (2020). Reshaping airpower: Development of an imprint model to analyze the effects of manned-unmanned teaming on operator mental workload. *2020 Winter Simulation Conference (WSC)*, 170–181. <https://doi.org/10.1109/WSC48552.2020.9383859>
- Andrews, R. W., Lilly, J. M., Srivastava, D., & Feigh, K. M. (2023). The role of shared mental models in human-AI teams: A theoretical review. *Theoretical Issues in Ergonomics Science*, 24(2), 129–175. <https://doi.org/10.1080/1463922X.2022.2061080>
- Anjomshoe, S., Calvaresi, D., Najjar, A., & Främling, K. (2019). Explainable agents and robots: Results from a systematic literature review. *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS '19)*, 1078–1088. <https://dl.acm.org/doi/abs/10.5555/3306127.3331806>
- Anscombe, G. E. (2000). *Intention*. Harvard University Press.
- Argyris, C., & Schön, D. A. (1974). *Theory in practice: Increasing professional effectiveness*. Jossey-Bass.
- Athanasiadou, M., Franzini, G., & Metge, A. (2026). Beyond automation levels: A framework for Human–Autonomy and Manned–Unmanned Teaming. *Automation*, 7(4), 105. <https://doi.org/10.3390/automation7040105>
- Azungah, T. (2018). Qualitative research: Deductive and inductive approaches to data analysis. *Qualitative Research Journal*, 18(4),

- 383–400. <https://doi.org/10.1108/QRJ-D-18-00035>
- Balcom, S., Doucet, S., & Dubé, A. (2021). Observation and institutional ethnography: Helping us to see better. *Qualitative Health Research*, 31(8), 1534–1541. <https://doi.org/10.1177/10497323211015966>
- Baltrusch, S. J., Krause, F., de Vries, A. W., van Dijk, W., & de Looze, M. P. (2022). What about the human in human robot collaboration? *Ergonomics*, 65(5), 719–740. <https://doi.org/10.1080/00140139.2021.1984585>
- Banas, J., Mehling, T., Paul, T., & Cords, A. (2020). Manned-unmanned teaming challenges in the maritime environment. *Vertical Flight Society's 76th Annual Forum and Technology Display*. <https://doi.org/10.4050/F-0076-2020-16314>
- Barandiaran, X. E., Di Paolo, E., & Rohde, M. (2009). Defining agency: Individuality, normativity, asymmetry, and spatio-temporality in action. *Adaptive Behavior*, 17(5), 367–386. <https://doi.org/10.1177/1059712309343819>
- Baratta, A., Cimino, A., Gnoni, M. G., & Longo, F. (2023). Human Robot Collaboration in Industry 4.0: A literature review. *Procedia Computer Science*, 217, 1887–1895. <https://doi.org/10.1016/j.procs.2022.12.389>
- Baskerville, R., Baiyere, A., Gergor, S., Hevner, A., & Rossi, M. (2018). Design Science Research contributions: Finding a balance between artifact and theory. *Journal of the Association for Information Systems*, 19(5), 358–376. <https://doi.org/10.17705/1jais.00495>
- Baumeister, R. F., Vohs, K. D., & Funder, D. C. (2007). Psychology as the science of self-reports and finger movements: Whatever happened to actual behavior? *Perspectives on Psychological Science*, 2(4), 396–403. <https://doi.org/10.1111/j.1745-6916.2007.00051.x>
- Bedwell, W. L., Wildman, J. L., DiazGranados, D., Salazar, M., Kramer, W. S., & Salas, E. (2012). Collaboration at work: An integrative multilevel conceptualization. *Human Resource Management Review*, 22(2), 128–145. <https://doi.org/10.1016/j.hrmr.2011.11.007>
- Bendell, R., Williams, J., Fiore, S. M., & Jentsch, F. (2024). Individual and team profiling to support theory of mind in artificial social intelligence. *Scientific Reports*, 14(1), 12635. <https://doi.org/10.1038/s41598-024-63122-8>
- Benishek, L. E., & Lazzara, E. H. (2019). Teams in a new era: Some considerations and implications. *Frontiers in Psychology*, 10(MAY), 1–15. <https://doi.org/10.3389/fpsyg.2019.01006>
- Berberian, B. (2019). Man-machine teaming: A problem of agency. *IFAC-PapersOnLine*, 51(34), 118–123.

- <https://doi.org/10.1016/j.ifacol.2019.01.049>
- Berretta, S., Tausch, A., Ontrup, G., Gilles, B., Peifer, C., & Kluge, A. (2023). Defining human-AI teaming the human-centered way: A scoping review and network analysis. *Frontiers in Artificial Intelligence*, 6. <https://doi.org/10.3389/frai.2023.1250725>
- Bhaskara, A., Skinner, M., & Loft, S. (2020). Agent transparency: A review of current theory and evidence. *IEEE Transactions on Human-Machine Systems*, 50(3), 215–224. <https://doi.org/10.1109/THMS.2020.2965529>
- Bishop, P., Hines, A., & Collins, T. (2007). The current state of scenario development: An overview of techniques. *Foresight*, 9(1), 5–25. <https://doi.org/10.1108/14636680710727516>
- Bobbe, T., Krzywinski, J., & Woelfel, C. (2016). A comparison of design process models from academic theory and professional practice. In D. Marjanović, M. Štorga, N. Pavković, N. Bojčetić, & S. Škec (Eds.), *DS 84: Proceedings of the DESIGN 2016 14th International Design Conference* (pp. 1205–1214). The Design Society.
- Börjeson, L., Höjer, M., Dreborg, K.-H., Ekvall, T., & Finnveden, G. (2006). Scenario types and techniques: Towards a user's guide. *Futures*, 38(7), 723–739. <https://doi.org/10.1016/j.futures.2005.12.002>
- Box, G. E. P. (1976). Science and statistics. *Journal of the American Statistical Association*, 71(356), 791–799. <https://doi.org/10.1080/01621459.1976.10480949>
- Bradshaw, J. M., Feltovich, P. J., Jung, H., Kulkarni, S., Taysom, W., & Uszok, A. (2004). Dimensions of adjustable autonomy and mixed-initiative interaction. In M. Nickles, M. Rovatsos, & G. Weiss (Eds.), *Agents and computational autonomy* (pp. 17–39). Springer. [https://doi.org/10.1007/978-3-540-25928-2\\_3](https://doi.org/10.1007/978-3-540-25928-2_3)
- Bratman, M. E. (1987). *Intention, plans, and practical reason*. Harvard University Press.
- Bratman, M. E. (1992). Shared cooperative activity. *The Philosophical Review*, 101(2), 327–341. <https://doi.org/10.2307/2185537>
- Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101. <https://doi.org/10.1191/1478088706qp0630a>
- Brentano, F. (1995). *Psychology from an empirical standpoint* (2nd ed.). Routledge.
- Burnette, C. (2002). Intentionality and design. In D. Durling & J. Shackleton (Eds.), *Common Ground - DRS International Conference 2002*. Staffordshire University Press. <https://doi.org/10.21606/drs.2002.4>
- Cabour, G., Morales-Forero, A., Ledoux, É., & Bassetto, S. (2023). An

- explanation space to align user studies with the technical development of Explainable AI. *AI and Society*, 38(2), 869–887. <https://doi.org/10.1007/s00146-022-01536-6>
- Calhoun, G. L., Ruff, H. A., Behymer, K. J., & Frost, E. M. (2018). Human-autonomy teaming interface design considerations for multi-unmanned vehicle control. *Theoretical Issues in Ergonomics Science*, 19(3), 321–352. <https://doi.org/10.1080/1463922X.2017.1315751>
- Cannon-Bowers, J. A., Salas, E., & Converse, S. A. (1993). Shared mental models in expert team decision making. In N. J. Castellan Jr. (Ed.), *Individual and group decision making: Current issues* (pp. 221–246). Lawrence Erlbaum Associates.
- Carlile, P. R. (2002). A pragmatic view of knowledge and boundaries: Boundary objects in new product development. *Organization Science*, 13(4), 442–455. <https://doi.org/10.1287/orsc.13.4.442.2953>
- Carroll, J. M. (1999). Five reasons for scenario-based design. *Proceedings of the 32nd Annual Hawaii International Conference on System Sciences*. <https://doi.org/10.1109/HICSS.1999.772890>
- Castañer, X., & Oliveira, N. (2020). Collaboration, coordination, and cooperation among organizations: Establishing the distinctive meanings of these terms through a systematic literature review. *Journal of Management*, 46(6), 965–1001. <https://doi.org/10.1177/0149206320901565>
- Castelfranchi, C. (2000). Founding agents’ “autonomy” on dependence theory. *The 14th European Conference on Artificial Intelligence (ECAI’00)*, 353–357.
- Chen, J. Y. C., & Barnes, M. J. (2014). Human-Agent Teaming for multirobot control: A review of human factors issues. *IEEE Transactions on Human-Machine Systems*, 44(1), 13–29. <https://doi.org/10.1109/THMS.2013.2293535>
- Chen, J. Y. C., Lakhmani, S. G., Stowers, K., Selkowitz, A. R., Wright, J. L., & Barnes, M. (2018). Situation awareness-based agent transparency and human-autonomy teaming effectiveness. *Theoretical Issues in Ergonomics Science*, 19(3), 259–282. <https://doi.org/10.1080/1463922X.2017.1315750>
- Chen, J. Y. C., Procci, K., Boyce, M., Wright, J., Garcia, A., & Barnes, M. J. (2014). Situation awareness-based agent transparency (ARL-TR-6905). U.S. Army Research Laboratory.
- Chung, H., Holder, T., Shah, J. A., & Yang, X. J. (2026). A systematic review and taxonomy of Human-Agent Teaming testbeds. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 68(2), 197–238.

- <https://doi.org/10.1177/00187208251376898>
- Cohen, S. (1997). What makes teams work: Group effectiveness research from the shop floor to the executive suite. *Journal of Management*, 23(3), 239–290. [https://doi.org/10.1016/S0149-2063\(97\)90034-9](https://doi.org/10.1016/S0149-2063(97)90034-9)
- Cooke, N. J., Cohen, M. C., Fazio, W. C., Inderberg, L. H., Johnson, C. J., Lematta, G. J., Peel, M., & Teo, A. (2024). From teams to teamness: Future directions in the science of team cognition. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 66(6), 1669–1680. <https://doi.org/10.1177/00187208231162449>
- Costa, D. G., Macul, V. C., Costa, J. M. H., Exner, K., Pfortner, A., Stark, R., & Rozenfeld, H. (2015). Towards the next generation of design process models: A gap analysis of existing models. In C. Weber, S. Husung, M. Cantamessa, G. Cascini, D. Marjanović, & S. Venkataraman (Eds.), *Proceedings of the 20th International Conference on Engineering Design (ICED15)* (Vol. 2, pp. 441–450). The Design Society.
- Crandall, B., Klein, G., & Hoffman, R. R. (2006). *Working minds: A practitioner's guide to cognitive task analysis*. MIT Press.
- Cuevas, H. M., Fiore, S. M., Caldwell, B. S., & Strater, L. (2007). Augmenting team cognition in human-automation teams performing in complex operational environments. *Aviation, Space, and Environmental Medicine*, 78(5), B63–70.
- Davidson, D. (1963). Actions, reasons, and causes. *The Journal of Philosophy*, 60(23), 685–700. <https://doi.org/10.2307/2023177>
- de Visser, E. J., Peeters, M. M. M., Jung, M. F., Kohn, S., Shaw, T. H., Pak, R., & Neerincx, M. A. (2020). Towards a theory of longitudinal trust calibration in Human–Robot Teams. *International Journal of Social Robotics*, 12(2), 459–478. <https://doi.org/10.1007/s12369-019-00596-x>
- Dennett, D. C. (1987). *The intentional stance*. MIT Press.
- Department of Defense. (2013). *Unmanned systems integration roadmap*. FY2013-2038. <https://irp.fas.org/program/collect/usroadmap2013.pdf>
- Department of Defense. (2018). *Unmanned systems integrated roadmap*. 2017-2042. <https://apps.dtic.mil/sti/citations/AD1059546>
- Desmedt, M., Ulenaers, D., Grosemans, J., Hellings, J., & Bergs, J. (2021). Clinical handover and handoff in healthcare: a systematic review of systematic reviews. *International Journal for Quality in Health Care: Journal of the International Society for Quality in Health Care*, 33(1), 1–24.

- <https://doi.org/10.1093/intqhc/mzaa170>
- DiCicco-Bloom, B., & Crabtree, B. F. (2006). The qualitative research interview. *Medical Education*, 40(4), 314–321. <https://doi.org/10.1111/j.1365-2929.2006.02418.x>
- Dorst, K., & Cross, N. (2001). Creativity in the design process: Co-evolution of problem–solution. *Design Studies*, 22(5), 425–437. [https://doi.org/10.1016/S0142-694X\(01\)00009-6](https://doi.org/10.1016/S0142-694X(01)00009-6)
- Edmondson, A. C. (2012). *Teaming: How organizations learn, innovate, and compete in the knowledge economy*. Jossey-Bass.
- Eiband, M., Schneider, H., Bilandzic, M., Fazekas-Con, J., Haug, M., & Hussmann, H. (2018). Bringing transparency design into practice. *23rd International Conference on Intelligent User Interfaces*, 211–223. <https://doi.org/10.1145/3172944.3172961>
- El Alami, H., Nwosu, M., & Rawat, D. B. (2023). Joint human and autonomy teaming for defense: Status, challenges, and perspectives. In L. Solomon & P. J. Schwartz (Eds.), *Artificial intelligence and machine learning for multi-domain operations applications V* (Vol. 12538, pp. 144–158). SPIE. <https://doi.org/10.1117/12.2665399>
- Elliott, G., Crawford, J., Watson, M., & Sanderson, P. (2000). Knowledge elicitation techniques for modelling intentional systems with cognitive work analysis. *Proceedings of the Fifth Australian Aviation Psychology Symposium*.
- Endsley, M. R. (1995). Toward a theory of Situation Awareness in dynamic systems. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 37(1), 32–64. <https://doi.org/10.1518/001872095779049543>
- Endsley, M. R. (2015). Situation awareness misconceptions and misunderstandings. *Journal of Cognitive Engineering and Decision Making*, 9(1), 4–32. <https://doi.org/10.1177/1555343415572631>
- Endsley, M. R. (2023). Supporting human-AI teams: Transparency, explainability, and situation awareness. *Computers in Human Behavior*, 140, 107574. <https://doi.org/10.1016/j.chb.2022.107574>
- Epstein, J. M. (2008). Why model? *Journal of Artificial Societies and Social Simulation*, 11(4), 1–11.
- Feinstein, A. H., & Cannon, H. M. (2002). Constructs of simulation evaluation. *Simulation & Gaming*, 33(4), 425–440. <https://doi.org/10.1177/1046878102238606>
- Flemisch, F., Abbink, D. A., Itoh, M., Pacaux-Lemoine, M.-P., & Weßel, G. (2019). Joining the blunt and the pointy end of the spear: Towards a common framework of joint action, Human–Machine

- Cooperation, cooperative guidance and control, shared, traded and supervisory control. *Cognition, Technology & Work*, 21(4), 555–568. <https://doi.org/10.1007/s10111-019-00576-1>
- Ford, K. M., Bradshaw, J. M., Adams-Webber, J. R., & Agnew, N. M. (1993). Knowledge acquisition as a constructive modeling activity. *International Journal of Intelligent Systems*, 8(1), 9–32. <https://doi.org/10.1002/int.4550080103>
- Frank, G. A., Helms, R., & Voor, D. (2000). Determining the right mix of live, virtual, and constructive training. *Proceedings of the 21st Interservice/Industry Training Systems and Education Conference*, 1–13.
- Freedman, R. G., & Zilberstein, S. (2019). A unifying perspective of plan, activity, and intent recognition. *Proceedings of the AAAI Workshops: Plan, Activity, Intent Recognition*, 1–8.
- Frith, C. D., & Frith, U. (2006). The neural basis of mentalizing. *Neuron*, 50(4), 531–534. <https://doi.org/10.1016/j.neuron.2006.05.001>
- Geddes, N. D. (1994). A model for intent interpretation for multiple agents with conflicts. *Proceedings of IEEE International Conference on Systems, Man and Cybernetics*, 3, 2080–2085. <https://doi.org/10.1109/ICSMC.1994.400170>
- Geddes, N. D. (1997). Large scale models of cooperative and hostile intentions. *Proceedings International Conference and Workshop on Engineering of Computer-Based Systems*, 142–147. <https://doi.org/10.1109/ECBS.1997.581841>
- Geddes, N. D., & Lizza, C. S. (2001). Practical applications of a real time, dynamic model of intentions. *AAAI Fall Symposium*.
- Gericke, K., & Blessing, L. (2012). An analysis of design process models across disciplines. In D. Marjanović, M. Štorga, N. Pavković, & N. Bojčetić (Eds.), *Proceedings of the 12th International Design Conference (DESIGN 2012)* (pp. 171–180). The Design Society.
- Gold, N., & Sugden, R. (2007). Collective intentions and team agency. *Journal of Philosophy*, 104(3), 109–137. <https://doi.org/10.5840/jphil2007104328>
- Gollwitzer, P. M. (1993). Goal achievement: The role of intentions. *European Review of Social Psychology*, 4(1), 141–185. <https://doi.org/10.1080/14792779343000059>
- Grindley, B., Parnell, K. J., Cherett, T., Scanlan, J., & Plant, K. L. (2025). Understanding the human factors challenge of handover between levels of automation for uncrewed air systems: A systematic literature review. *Transportation Planning and Technology*, 48(6), 1383–1408. <https://doi.org/10.1080/03081060.2024.2375645>
- Gustavsson, P. M., Hieb, M. R., Moore, P., Eriksson, P., & Niklasson, L.

- (2011). Operations intent and effects model. *The Journal of Defense Modeling and Simulation: Applications, Methodology, Technology*, 8(1), 37–59. <https://doi.org/10.1177/1548512910379477>
- Hackman, J. R. (2011). *Collaborative intelligence: Using teams to solve hard problems*. Berrett-Koehler Publishers.
- Hajdukiewicz, J. R., Burns, C. M., Vicente, K. J., & Eggleston, R. G. (1999). Work Domain Analysis for intentional systems. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 43(3), 333–337. <https://doi.org/10.1177/154193129904300343>
- Hammarbäck, J. (2022). *Modelling intent for manned-unmanned teaming: Exploring human-centric approaches for future combat aircraft systems* (Vol. 1946). Linköping University Electronic Press. <https://doi.org/10.3384/9789179295486>
- Han, T. A., Pereira, L. M., & Santos, F. C. (2011). Intention recognition promotes the emergence of cooperation. *Adaptive Behavior*, 19(4), 264–279. <https://doi.org/10.1177/1059712311410896>
- Han, T. A., Santos, F. C., Lenaerts, T., & Pereira, L. M. (2015). Synergy between intention recognition and commitments in cooperation dilemmas. *Scientific Reports*, 5(1), 9312. <https://doi.org/10.1038/srep09312>
- Harris-Watson, A. M., & Larson, L. E. (2026). When tech makes the team: Mind perception as a unifying framework for human and human-agent teams. *Computers in Human Behavior*, 181, 109000. <https://doi.org/10.1016/j.chb.2026.109000>
- Hassall, M. E., & Sanderson, P. M. (2014). A formative approach to the strategies analysis phase of cognitive work analysis. *Theoretical Issues in Ergonomics Science*, 15(3), 215–261. <https://doi.org/10.1080/1463922X.2012.725781>
- Heider, F. (1958). *The psychology of interpersonal relations*. John Wiley & Sons.
- Heider, F., & Simmel, M. (1944). An experimental study of apparent behavior. *The American Journal of Psychology*, 57(2), 243. <https://doi.org/10.2307/1416950>
- Heilemann, F., Lindner, S., & Schulte, A. (2021). Experimental evaluation of tasking and teaming design patterns for human delegation of unmanned vehicles. *Human-Intelligent Systems Integration*, 3(3), 223–240. <https://doi.org/10.1007/s42454-021-00036-0>
- Heinze, C. (2004). *Modelling intention recognition for intelligent agent systems*. Doctoral thesis. University of Melbourne.
- Heskett, J. (2005). *Design: A very short introduction*. Oxford University

- Press.
- Hevner, A. R. (2007). A three cycle view of Design Science Research. *Scandinavian Journal of Information Systems*, 19(2), 87–92. <https://doi.org/http://aisel.aisnet.org/sjis/vol19/iss2/4>
- Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design Science in Information Systems Research. *MIS Quarterly*, 28(1), 75–106. <https://doi.org/10.2307/25148625>
- Hobbs, A., & Lyall, B. (2016). Human factors guidelines for unmanned aircraft systems. *Ergonomics in Design: The Quarterly of Human Factors Applications*, 24(3), 23–28. <https://doi.org/10.1177/1064804616640632>
- Hochstein, E. (2017). Why one model is never enough: a defense of explanatory holism. *Biology & Philosophy*, 32(6), 1105–1125. <https://doi.org/10.1007/s10539-017-9595-x>
- Hodson, D. D., & Baldwin, R. O. (2009). Characterizing, measuring, and validating the temporal consistency of live–virtual–constructive environments. *Simulation*, 85(10), 671–682. <https://doi.org/10.1177/0037549709340732>
- Hoffman, R. R., & Lintern, G. (2006). Eliciting and representing the knowledge of experts. In K. A. Ericsson, N. Charness, P. J. Feltovich, & R. R. Hoffman (Eds.), *The Cambridge handbook of expertise and expert performance* (pp. 203–222). Cambridge University Press.
- Hollnagel, E., & Woods, D. D. (2005). *Joint Cognitive Systems: Foundations of cognitive systems engineering*. CRC Press.
- Holstein, J. A., & Gubrium, J. F. (1995). *The active interview*. Sage.
- Howard, N., & Cambria, E. (2013). Intention awareness: improving upon situation awareness in human-centric environments. *Human-Centric Computing and Information Sciences*, 3(1), 9. <https://doi.org/10.1186/2192-1962-3-9>
- Howard, T. J., Culley, S. J., & Dekoninck, E. (2008). Describing the creative design process by the integration of engineering design and cognitive psychology literature. *Design Studies*, 29(2), 160–180. <https://doi.org/10.1016/j.destud.2008.01.001>
- Huffman, J. L., McNeil, G., Bismilla, Z., & Lai, A. (2016). Essentials of scenario building for simulation-based education. In V. Grant & A. Cheng (Eds.), *Comprehensive healthcare simulation: Pediatrics* (pp. 19–29). Springer.
- Iftikhar, R., Chiu, Y.-T., Khan, M. S., & Caudwell, C. (2024). Human-Agent Team dynamics: A review and future research opportunities. *IEEE Transactions on Engineering Management*, 71, 10139–10154. <https://doi.org/10.1109/TEM.2023.3331369>
- Jameson, S., Franke, J., Szczerba, R., & Stockdale, S. (2005).

- Collaborative autonomy for manned/unmanned teams. *Proceedings of the American Helicopter Society 61st Annual Forum*.
- Jenkins, N., Bloor, M., Fischer, J., Berney, L., & Neale, J. (2010). Putting it in context: the use of vignettes in qualitative interviewing. *Qualitative Research*, 10(2), 175–198. <https://doi.org/10.1177/1468794109356737>
- Johnson-Laird, P. N. (1983). *Mental models: Towards a cognitive science of language, inference, and consciousness*. Harvard University Press.
- Johnson, M., Bradshaw, J. M., Feltovich, P. J., Jonker, C. M., van Riemsdijk, B., & Sierhuis, M. (2011). The fundamental principle of Coactive Design: Interdependence must shape autonomy. In M. De Vos, N. Fornara, J. V. Pitt, & G. Vouros (Eds.), *Coordination, organizations, institutions, and norms in agent systems VI* (pp. 172–191). Springer. [https://doi.org/10.1007/978-3-642-21268-0\\_10](https://doi.org/10.1007/978-3-642-21268-0_10)
- Johnson, M., Bradshaw, J. M., Feltovich, P. J., Jonker, C. M., Van Riemsdijk, M. B., & Sierhuis, M. (2014). Coactive Design: Designing support for interdependence in joint activity. *Journal of Human-Robot Interaction*, 3(1), 43. <https://doi.org/10.5898/JHRI.3.1.Johnson>
- Jones, E. E., & Davis, K. E. (1965). From acts to dispositions: The attribution process in person perception. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 2, pp. 219–266). Academic Press.
- Jordan, J. (2021). The future of unmanned combat aerial vehicles: An analysis using the Three Horizons framework. *Futures*, 134, 102848. <https://doi.org/10.1016/j.futures.2021.102848>
- Kaber, D. B. (2018). A conceptual framework of autonomous and automated agents. *Theoretical Issues in Ergonomics Science*, 19(4), 406–430. <https://doi.org/10.1080/1463922X.2017.1363314>
- Kahn, H., & Wiener, A. J. (1967). *The year 2000: A framework for speculation on the next thirty-three years*. Macmillan.
- Kaliardos, W. N., Holbrook, J., & Hobbs, A. (2024). Identifying Human Factors research for unmanned aircraft systems and advanced air mobility. *2024 AIAA DATC/IEEE 43rd Digital Avionics Systems Conference (DASC)*, 1–9. <https://doi.org/10.1109/DASC62030.2024.10748711>
- Kawulich, B. B. (2005). Participant observation as a data collection method. *Forum Qualitative Sozialforschung Forum: Qualitative Social Research*, 6(2).

- <https://doi.org/https://doi.org/10.17169/fqs-6.2.466>
- Kelley, H. H., & Michela, J. L. (1980). Attribution theory and research. *Annual Review of Psychology*, 31(1), 457–501. <https://doi.org/10.1146/annurev.ps.31.020180.002325>
- Kim, B. W., & Choi, G. E. (2024). Level and program analytics of MUM-T system. *International Journal of Aeronautical and Space Sciences*, 25(2), 593–604. <https://doi.org/10.1007/s42405-023-00675-4>
- Klein, G. (1994). A script for the Commander's Intent. In A. H. Levis & I. S. Levis (Eds.), *Science of command and control: Part III: Coping with change* (pp. 75–85). AFCEA International Press.
- Klein, G. (1999). *Sources of power: How people make decisions*. The MIT Press.
- Klein, G., Feltovich, P. J., Bradshaw, J. M., & Woods, D. D. (2005). Common ground and coordination in joint activity. In W. B. Rouse & K. R. Boff (Eds.), *Organizational simulation* (pp. 139–184). Wiley. <https://doi.org/10.1002/0471739448.ch6>
- Klein, G., Phillips, J. K., Rall, E. L., & Peluso, D. A. (2007). A data-frame theory of sensemaking. In R. R. Hoffman (Ed.), *Expertise out of context: Proceedings of the sixth international conference on naturalistic decision making* (pp. 113–155). Lawrence Erlbaum Associates.
- Klein, G., Woods, D. D., Bradshaw, J. M., Hoffman, R. R., & Feltovich, P. J. (2004). Ten challenges for making automation a “team player” in joint human-agent activity. *IEEE Intelligent Systems*, 19(6), 91–95. <https://doi.org/10.1109/MIS.2004.74>
- Klein, Moon, & Hoffman. (2006). Making sense of sensemaking 2: A macrocognitive model. *IEEE Intelligent Systems*, 21(5), 88–92. <https://doi.org/10.1109/MIS.2006.100>
- Kozlowski, S. W. J., & Ilgen, D. R. (2006). Enhancing the effectiveness of work groups and teams. *Psychological Science in the Public Interest*, 7(3), 77–124. <https://doi.org/10.1111/j.1529-1006.2006.00030.x>
- Kridalukmana, R., Lu, H. Y., & Naderpour, M. (2020). A supportive situation awareness model for human-autonomy teaming in collaborative driving. *Theoretical Issues in Ergonomics Science*, 21(6), 658–683. <https://doi.org/10.1080/1463922X.2020.1729443>
- Krosnick, J. A. (1999). Survey research. *Annual Review of Psychology*, 50(1), 537–567. <https://doi.org/10.1146/annurev.psych.50.1.537>
- Lakhmani, S. G., Neubauer, C., Krausman, A., Fitzhugh, S. M., Berg, S. K., Wright, J. L., Rovira, E., Blackman, J. J., & Schaefer, K. E. (2022). Cohesion in Human–Autonomy Teams: an approach for

- future research. *Theoretical Issues in Ergonomics Science*, 23(6), 687–724. <https://doi.org/10.1080/1463922X.2022.2033876>
- Langfield-Smith, K. (1992). Exploring the need for a shared cognitive map. *Journal of Management Studies*, 29(3), 349–368. <https://doi.org/10.1111/j.1467-6486.1992.tb00669.x>
- Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 46(1), 50–80. [https://doi.org/10.1518/hfes.46.1.50\\_30392](https://doi.org/10.1518/hfes.46.1.50_30392)
- Leu, G., & Abbass, H. (2016). A multi-disciplinary review of knowledge acquisition methods: From human to autonomous eliciting agents. *Knowledge-Based Systems*, 105, 1–22. <https://doi.org/10.1016/j.knosys.2016.02.012>
- Leveson, N. G. (2000). Intent specifications: an approach to building human-centered specifications. *IEEE Transactions on Software Engineering*, 26(1), 15–35. <https://doi.org/10.1109/32.825764>
- Li, J.-M., Wu, T.-J., Wu, Y. J., & Goh, M. (2023). Systematic literature review of human–machine collaboration in organizations using bibliometric analysis. *Management Decision*, 61(10), 2920–2944. <https://doi.org/10.1108/MD-09-2022-1183>
- Lim, B. Y., & Dey, A. K. (2009). Assessing demand for intelligibility in context-aware applications. *Proceedings of the 11th International Conference on Ubiquitous Computing*, 195–204. <https://doi.org/10.1145/1620545.1620576>
- Lim, B. Y., Dey, A. K., & Avrahami, D. (2009). Why and why not explanations improve the intelligibility of context-aware intelligent systems. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2119–2128. <https://doi.org/10.1145/1518701.1519023>
- Lim, Y., Gardi, A., Sabatini, Roberto, Lim, Y., Gardi, A., Sabatini, R., Ramasamy, S., Kistan, T., Ezer, N., Vince, J., & Bolia, R. (2018). Avionics human-machine interfaces and interactions for manned and unmanned aircraft. *Progress in Aerospace Sciences*, 102, 1–46. <https://doi.org/10.1016/j.paerosci.2018.05.002>
- Lim, Y., Pongsakornsathien, N., Gardi, A., Sabatini, R., Kistan, T., Ezer, N., & Bursch, D. J. (2021). Adaptive human-robot interactions for multiple unmanned aerial vehicles. *Robotics*, 10(1), 12. <https://doi.org/10.3390/robotics10010012>
- Lim, Y., Ramasamy, S., Gardi, A., Kistan, T., & Sabatini, R. (2018). Cognitive human-machine interfaces and interactions for unmanned aircraft. *Journal of Intelligent & Robotic Systems*, 91(3–4), 755–774. <https://doi.org/10.1007/s10846-017-0648-9>
- Lui, F., & Watson, M. (2002). Mapping cognitive work analysis (CWA)

- to an intelligent agents software architecture: Command agents. *Proceedings of the Defence Human Factors Special Interest Group (DHFSIG)* 2002, 1–5.
- Lum, H. C., & Phillips, E. K. (2024). Understanding Human–Autonomy Teams through a Human–Animal Teaming model. *Topics in Cognitive Science*, 16(3), 554–567. <https://doi.org/10.1111/tops.12713>
- Lundberg, J. (2015). Situation awareness systems, states and processes: A holistic framework. *Theoretical Issues in Ergonomics Science*, 16(5), 447–473. <https://doi.org/10.1080/1463922X.2015.1008601>
- Lundberg, J., Ducheve, A., Spyrou, E. D., Kappatos, V., Villani, A., Reis, R., Westin, C., Kirwan, B., Arrigoni, V., Pozzi, S., & Bång, M. (2026). Levels of Autonomy in Cognitive Control (LACC-LOA): Bidimensional assessment method applied to six AI assistant concepts and scenarios in the aviation domain. *International Journal of Human–Computer Interaction*, 1–22. <https://doi.org/10.1080/10447318.2026.2662524>
- Lundberg, J., & Johansson, B. J. E. (2021). A framework for describing interaction between human operators and autonomous, automated, and manual control systems. *Cognition, Technology & Work*, 23(3), 381–401. <https://doi.org/10.1007/s10111-020-00637-w>
- Lyons, J. B. (2013). Being transparent about transparency: A model for human-robot interaction. *Trust and Autonomous Systems: Papers from the AAAI Spring Symposium*, SS-13-07, 48–53.
- Lyons, J. B., & Havig, P. R. (2014). Transparency in a human-machine context: Approaches for fostering shared awareness/intent. In R. Shumaker & S. Lackey (Eds.), *Virtual, augmented and mixed reality: Designing and developing virtual and augmented environments* (pp. 181–190). Springer. [https://doi.org/10.1007/978-3-319-07458-0\\_18](https://doi.org/10.1007/978-3-319-07458-0_18)
- Lyons, J. B., Sycara, K., Lewis, M., & Capiola, A. (2021). Human–Autonomy Teaming: Definitions, debates, and directions. *Frontiers in Psychology*, 12, 1–15. <https://doi.org/10.3389/fpsyg.2021.589585>
- Malle, B. F. (1999). How people explain behavior: A new theoretical framework. *Personality and Social Psychology Review*, 3(1), 23–48. [https://doi.org/10.1207/s15327957pspr0301\\_2](https://doi.org/10.1207/s15327957pspr0301_2)
- Malle, B. F. (2004). *How the mind explains behavior: Folk explanations, meaning, and social Interaction*. MIT Press.
- Malle, B. F., & Knobe, J. (1997). The folk concept of intentionality. *Journal of Experimental Social Psychology*, 33(2), 101–121.

- <https://doi.org/10.1006/jesp.1996.1314>
- Malone, T. W., & Crowston, K. (1994). The interdisciplinary study of coordination. *ACM Computing Surveys*, 26(1), 87–119. <https://doi.org/10.1145/174666.174668>
- March, S. T., & Smith, G. F. (1995). Design and natural science research on information technology. *Decision Support Systems*, 15(4), 251–266. [https://doi.org/10.1016/0167-9236\(94\)00041-2](https://doi.org/10.1016/0167-9236(94)00041-2)
- Matta, D. (2026). Artificial intelligence and theory of mind. *Journal of Psychology and AI*, 2(1). <https://doi.org/10.1080/29974100.2026.2628373>
- McMeekin, N., Wu, O., Germen, E., & Briggs, A. (2020). How methodological frameworks are being developed: evidence from a scoping review. *BMC Medical Research Methodology*, 20(1), 173. <https://doi.org/10.1186/s12874-020-01061-4>
- McNeese, N., Demir, M., Chiou, E., Cooke, N., & Yanikian, G. (2019). Understanding the role of trust in human-autonomy teaming. *Proceedings of the Annual Hawaii International Conference on System Sciences*, 254–263. <https://doi.org/10.24251/HICSS.2019.032>
- McNeese, N. J., Flathmann, C., O'Neill, T. A., & Salas, E. (2023). Stepping out of the shadow of human-human teaming: Crafting a unique identity for human-autonomy teams. *Computers in Human Behavior*, 148, 107874. <https://doi.org/10.1016/j.chb.2023.107874>
- McRuer, D., & Krendel, E. (1962). The man-machine system concept. *Proceedings of the IRE*, 50(5), 1117–1123. <https://doi.org/10.1109/JRPROC.1962.288016>
- Merry, M., Riddle, P., & Warren, J. (2021). A mental models approach for defining explainable artificial intelligence. *BMC Medical Informatics and Decision Making*, 21(1), 344. <https://doi.org/10.1186/s12911-021-01703-7>
- Meyer, C., & Schulte, A. (2020). Operator controlled, reactive UAV behaviors in manned-unmanned teaming scenarios with selective datalink availability. *2020 International Conference on Unmanned Aircraft Systems (ICUAS)*, 1673–1679. <https://doi.org/10.1109/ICUAS48674.2020.9214018>
- Militello, L. G., & Hutton, R. J. B. (1998). Applied cognitive task analysis (ACTA): A practitioner's toolkit for understanding cognitive task demands. *Ergonomics*, 41(11), 1618–1641. <https://doi.org/10.1080/001401398186108>
- Miller, T. (2019). Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence*, 267, 1–38. <https://doi.org/10.1016/j.artint.2018.07.007>

- Minsky, M. (1974). *A framework for representing knowledge* (AI Memo 306). MIT Artificial Intelligence Laboratory.
- Minsky, M. (1975). Minsky's frame system theory. In R. Schank & B. L. Nash-Webber (Eds.), *Proceedings of the 1975 Workshop on Theoretical Issues in Natural Language Processing (TINLAP '75)* (pp. 104–116). Association for Computational Linguistics. <https://doi.org/10.3115/980190.980222>
- Moens, L. L., Lydon, S., Cucurachi, S., O'Connor, P., Sauter, T. C., Töndury, G.-A., & Manser, T. (2026). Measuring situation awareness: A meta-review across domains. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 68(5), 632–672. <https://doi.org/10.1177/00187208251412110>
- Mohammed, S., Rico, R., & Alipour, K. K. (2021). Team cognition at a crossroad: Toward conceptual integration and network configurations. *Academy of Management Annals*, 15(2), 455–501. <https://doi.org/10.5465/annals.2018.0159>
- Mohseni, S., Zarei, N., & Ragan, E. D. (2021). A multidisciplinary survey and framework for design and evaluation of explainable AI systems. *ACM Transactions on Interactive Intelligent Systems*, 11(3–4), 1–45. <https://doi.org/10.1145/3387166>
- Morgan, M. S., & Morrison, M. (1999). *Models as mediators*. Cambridge University Press.
- Naikar, N. (2013). *Work Domain Analysis: Concept, guidelines, and cases*. CRC Press.
- Norling, E. (2012). *Modelling human behaviour with BDI agents*. Doctoral thesis. University of Melbourne.
- Nowak, M. A. (2006). Five rules for the evolution of cooperation. *Science*, 314(5805), 1560–1563. <https://doi.org/10.1126/science.1133755>
- O'Hare, D., Wiggins, M., Williams, A., & Wong, W. (1998). Cognitive task analyses for decision centred design and training. *Ergonomics*, 41(11), 1698–1718. <https://doi.org/10.1080/001401398186144>
- O'Neill, T. A., Flathmann, C., McNeese, N. J., & Salas, E. (2023). Human-autonomy Teaming: Need for a guiding team-based framework? *Computers in Human Behavior*, 146, 107762. <https://doi.org/10.1016/j.chb.2023.107762>
- O'Neill, T., McNeese, N., Barron, A., & Schelble, B. (2022). Human–Autonomy Teaming: A review and analysis of the empirical literature. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 64(5), 904–938. <https://doi.org/10.1177/0018720820960865>
- Okcu, H. (2016). Operational requirements of unmanned aircraft systems data link and communication systems. *Journal of*

- Advances in Computer Networks*, 4(1), 28–32.  
<https://doi.org/10.18178/JACN.2016.4.1.199>
- Oppl, S. (2017). Supporting the collaborative construction of a shared understanding about work with a guided conceptual modeling technique. *Group Decision and Negotiation*, 26(2), 247–283.  
<https://doi.org/10.1007/s10726-016-9485-7>
- Ososky, S., Schuster, D., Jentsch, F., Fiore, S., Shumaker, R., Lebiere, C., Kurup, U., Oh, J., & Stentz, A. (2012). The importance of shared mental models and shared situation awareness for transforming robots from tools to teammates. *Unmanned Systems Technology XIV*, 8387(May 2012), 838710. <https://doi.org/10.1117/12.923283>
- Pacherie, E. (2008). The phenomenology of action: A conceptual framework. *Cognition*, 107(1), 179–217.  
<https://doi.org/10.1016/j.cognition.2007.09.003>
- Pankok, C., & Bass, E. J. (2018). Information content requirements for remote pilot handover of control of unmanned aircraft systems in the national airspace. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 62(1), 81–85.  
<https://doi.org/10.1177/1541931218621018>
- Parker, W. S. (2020). Model evaluation: An adequacy-for-purpose view. *Philosophy of Science*, 87(3), 457–477.  
<https://doi.org/10.1086/708691>
- Partelow, S. (2023). What is a framework? Understanding their purpose, value, development and use. *Journal of Environmental Studies and Sciences*, 13(3), 510–519. <https://doi.org/10.1007/s13412-023-00833-w>
- Patel, J., Boardman, M., Files, B., Gregory, F., Lamb, S., Sarkadi, S., Tešić, M., & Yeung, N. (2025). Give us a hand, mate! A holistic review of research on human-machine teaming. *BMJ Military Health*, 171(5), 392–396. <https://doi.org/10.1136/military-2024-002737>
- Peffer, K., Rothenberger, M., Tuunanen, T., & Vaezi, R. (2012). Design Science Research Evaluation. In K. Peffer, M. Rothenberger, & B. Kuechler (Eds.), *Design science research in information systems: Advances in theory and practice* (pp. 398–410). Springer.  
[https://doi.org/10.1007/978-3-642-29863-9\\_29](https://doi.org/10.1007/978-3-642-29863-9_29)
- Peffer, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A Design Science Research Methodology for information systems research. *Journal of Management Information Systems*, 24(3), 45–77. <https://doi.org/10.2753/MIS0742-1222240302>
- Perez-Osorio, J., & Wykowska, A. (2020). Adopting the intentional stance toward natural and artificial agents. *Philosophical Psychology*, 33(3), 369–395.

- <https://doi.org/10.1080/09515089.2019.1688778>
- Pezzulo, G. (2011). Shared representations as coordination tools for interaction. *Review of Philosophy and Psychology*, 2(2), 303–333. <https://doi.org/10.1007/s13164-011-0060-5>
- Phillips, E. K., Schaefer, K., Billings, D. R., Jentsch, F., & Hancock, P. A. (2015). Human-animal teams as an analog for future human-robot teams: Influencing design and fostering trust. *Journal of Human-Robot Interaction*, 5(1), 100. <https://doi.org/10.5898/jhri.5.1.phillips>
- Pigeau, R., & McCann, C. (2000). Redefining command and control. In C. McCann & R. Pigeau (Eds.), *The human in command: Exploring the modern military experience* (pp. 163–184). Springer.
- Pigeau, R., & McCann, C. (2006). Establishing common intent: The key to co-ordinated military action. In A. English (Ed.), *The operational art: Canadian perspectives: Leadership and command* (pp. 85–108). Canadian Defence Academy Press.
- Premack, D., & Woodruff, G. (1978). Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences*, 4, 515–526.
- Quan, S. X., Lam, C., Schabram, K., & Yam, K. C. (2024). All creatures great and small: A review and typology of employee-animal interactions. *Journal of Management*, 50(1), 380–411. <https://doi.org/10.1177/01492063231191090>
- Rajabiyazdi, F., & Jamieson, G. A. (2020). A review of transparency (seeing-into) models. 2020 IEEE International Conference on Systems, Man, and Cybernetics (SMC), 302–308. <https://doi.org/10.1109/SMC42975.2020.9282970>
- Ramirez, R., Mukherjee, M., Vezzoli, S., & Kramer, A. M. (2015). Scenarios as a scholarly methodology to produce “interesting research.” *Futures*, 71, 70–87. <https://doi.org/10.1016/j.futures.2015.06.006>
- Ranganathan, P., & Caduff, C. (2023). Designing and validating a research questionnaire - Part 1. *Perspectives in Clinical Research*, 14(3), 152–155. [https://doi.org/10.4103/picr.picr\\_140\\_23](https://doi.org/10.4103/picr.picr_140_23)
- Rasmussen, J. (1986). *Information processing and human-machine interaction: An approach to cognitive engineering*. North-Holland.
- Renger, M., Kolfshoten, G. L., & Vreede, G. J. De. (2008). Challenges in collaborative modelling: a literature review and research agenda. *International Journal of Simulation and Process Modelling*, 4(3/4), 248. <https://doi.org/10.1504/IJSPM.2008.023686>
- Rittel, H. W. J., & Webber, M. M. (1973). Dilemmas in a general theory of planning. *Policy Sciences*, 4(2), 155–169. <https://doi.org/10.1007/BF01405730>

- Rix, J. (2022). From tools to teammates: Conceptualizing humans' perception of machines as teammates with a systematic literature review. *Proceedings of the 55th Hawaii International Conference on System Sciences*, 398–407. <https://doi.org/10.24251/HICSS.2022.048>
- Roth, E. M., Sushereba, C., Militello, L. G., DiIulio, J., & Ernst, K. (2019). Function allocation considerations in the era of Human Autonomy Teaming. *Journal of Cognitive Engineering and Decision Making*, 13(4), 199–220. <https://doi.org/10.1177/1555343419878038>
- Roth, G., Schulte, A., Schmitt, F., & Brand, Y. (2020). Transparency for a workload-adaptive cognitive agent in a Manned–Unmanned Teaming application. *IEEE Transactions on Human-Machine Systems*, 50(3), 225–233. <https://doi.org/10.1109/THMS.2019.2914667>
- Rouse, W. B., & Morris, N. M. (1986). On looking into the black box: Prospects and limits in the search for mental models. *Psychological Bulletin*, 100(3), 349–363. <https://doi.org/10.1037/0033-2909.100.3.349>
- Rybing, J. (2018). *Studying simulations with distributed cognition* [Doctoral dissertation, Linköping University]. Linköping University Electronic Press. <https://doi.org/10.3384/diss.diva-145307>
- Saavedra, R., Earley, P. C., & Van Dyne, L. (1993). Complex interdependence in task-performing groups. *Journal of Applied Psychology*, 78(1), 61–72. <https://doi.org/10.1037/0021-9010.78.1.61>
- Salas, E., Dickinson, T. L., Converse, S. A., & Tannenbaum, S. I. (1992). Toward an understanding of team performance and training. In R. W. Swezey & E. Salas (Eds.), *Teams: Their training and performance* (pp. 3–29). Ablex.
- Salas, E., Linhardt, R., & Fernández Castillo, G. (2024). The science (and practice) of teamwork: A commentary on forty years of progress. *Small Group Research*, 1–34. <https://doi.org/10.1177/10464964241274119>
- Sanneman, L., & Shah, J. A. (2022). The Situation Awareness Framework for Explainable AI (SAFE-AI) and Human Factors Considerations for XAI Systems. *International Journal of Human–Computer Interaction*, 38(18–20), 1772–1788. <https://doi.org/10.1080/10447318.2022.2081282>
- Schaefer, K. E., Straub, E. R., Chen, J. Y. C., Putney, J., & Evans, A. W. (2017). Communicating intent to develop shared situation awareness and engender trust in human-agent teams. *Cognitive Systems Research*, 46, 26–39.

- <https://doi.org/10.1016/j.cogsys.2017.02.002>
- Schelble, B. G., Flathmann, C., & McNeese, N. (2020). Towards meaningfully integrating human-autonomy teaming in applied settings. *Proceedings of the 8th International Conference on Human-Agent Interaction*, 149–156. <https://doi.org/10.1145/3406499.3415077>
- Schneider, M. F., & Miller, M. E. (2018). Operationalized intent for communication in human-agent teams. 2018 IEEE Conference on Cognitive and Computational Aspects of Situation Management (CogSIMA), 117–123. <https://doi.org/10.1109/COGSIMA.2018.8423992>
- Schneider, M. F., Miller, M. E., Ford, T. C., Peterson, G., & Jacques, D. (2022). Intent integration for human-agent teaming. *Systems Engineering*, 1–13. <https://doi.org/10.1002/sys.21616>
- Schneider, M. F., Miller, M. E., & McGuirl, J. (2023). Assessing quality goal rankings as a method for communicating operator intent. *Journal of Cognitive Engineering and Decision Making*, 17(1), 26–48. <https://doi.org/10.1177/15553434221131665>
- Searle, J. R. (1983). *Intentionality: An essay in the philosophy of mind*. Cambridge University Press.
- Sheridan, T. B. (2002). *Humans and automation: System design and research issues*. Wiley.
- Simon, H. A. (1996). *The sciences of the artificial (3rd ed.)* The MIT Press.
- Singh, S., Rantanen, E., & Heard, J. (2026). Human–Robot Teaming: A comprehensive survey on collaboration, communication, and cognition. *ACM Transactions on Human-Robot Interaction*, 15(2), 1–48. <https://doi.org/10.1145/3776548>
- Snyder, H. (2019). Literature review as a research methodology: An overview and guidelines. *Journal of Business Research*, 104, 333–339. <https://doi.org/10.1016/j.jbusres.2019.07.039>
- Söderlund, M. (2002). Customer satisfaction and its influence on different behavioural intention constructs. *Journal of Customer Behaviour*, 1(2), 145–166. <https://doi.org/10.1362/14753920260595192>
- Stansbury, R. S., Vyas, M. A., & Wilson, T. A. (2009). A survey of UAS technologies for Command, Control, and Communication (C3). *Journal of Intelligent and Robotic Systems*, 54(1–3), 61–78. <https://doi.org/10.1007/s10846-008-9261-2>
- Stanton, N. A., Salmon, P. M., Rafferty, L. A., Walker, G. H., Baber, C., & Jenkins, D. P. (2017). *Human factors methods*. CRC Press. <https://doi.org/10.1201/9781315587394>
- Star, S. L. (2010). This is not a boundary object: Reflections on the origin

- of a concept. *Science, Technology, & Human Values*, 35(5), 601–617. <https://doi.org/10.1177/0162243910377624>
- Stensrud, R., Mikkelsen, B., & Valaker, S. (2024). Exploring human-autonomy teaming methods in challenging environments: The case of fighter pilots and loyal wingmen. *Human-Intelligent Systems Integration*, 6(1), 1–14. <https://doi.org/10.1007/s42454-024-00050-y>
- Stowers, K., Brady, L. L., MacLellan, C., Wohleber, R., & Salas, E. (2021). Improving teamwork competencies in human-machine teams: Perspectives from team science. *Frontiers in Psychology*, 12(May), 1–6. <https://doi.org/10.3389/fpsyg.2021.590290>
- Strong, J., Rogers, H., Sun, E., Todsen, A. L., Ede, J., Lumley, C., Yeung, N., Higham, H., & Noble, J. A. (2026). Human-AI collaboration in healthcare: A scoping review. *npj Digital Medicine*. <https://doi.org/10.1038/s41746-026-02918-6>
- Suddaby, R. (2010). Editor's comments: Construct clarity in theories of management and organization. *Academy of Management Review*, 35(3), 346–357. <https://doi.org/10.5465/amr.35.3.zok346>
- Sun, J., & Qiu, L. (2025). Human-machine collaboration: Definitions, models, and socio-economic impacts. *International Journal of Crowd Science*, 9(3), 149–163. <https://doi.org/10.26599/IJCS.2025.9100004>
- Tabrez, A., Luebbers, M. B., & Hayes, B. (2020). A survey of mental modeling techniques in Human–Robot Teaming. *Current Robotics Reports*, 1(4), 259–267. <https://doi.org/10.1007/s43154-020-00019-0>
- Taylor, G., & Turpin, T. (2015). Army Aviation manned-unmanned teaming (MUM-T): Past, Present, and Future. *18th International Symposium on Aviation Psychology*, 560–565.
- Theissing, N., & Schulte, A. (2013). Intent-Based UAV Mission Management Using an Adaptive Mixed-Initiative Operator Assistant System. *ALAA Infotech@Aerospace (I@A) Conference*, 1–17. <https://doi.org/10.2514/6.2013-4802>
- Thellman, S., de Graaf, M., & Ziemke, T. (2022). Mental state attribution to robots: A systematic review of conceptions, methods, and findings. *ACM Transactions on Human-Robot Interaction*, 11(4), 1–51. <https://doi.org/10.1145/3526112>
- Thellman, S., Silvervarg, A., & Ziemke, T. (2017). Folk-psychological interpretation of human vs. humanoid robot behavior: Exploring the intentional stance toward robots. *Frontiers in Psychology*, 8(NOV), 1–14. <https://doi.org/10.3389/fpsyg.2017.01962>
- Tomasello, M. (2014). *A natural history of human thinking*. Harvard University Press. <https://doi.org/10.4159/9780674726369>

- Tomasello, M., Carpenter, M., Call, J., Behne, T., & Moll, H. (2005). Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and Brain Sciences*, 28(5), 675–691. <https://doi.org/10.1017/S0140525X05000129>
- Tsamados, A., Floridi, L., & Taddeo, M. (2025). Human control of AI systems: from supervision to teaming. *AI and Ethics*, 5(2), 1535–1548. <https://doi.org/10.1007/s43681-024-00489-4>
- Tubbs, M. E., & Ekeberg, S. E. (1991). The role of intentions in work motivation: Implications for goal-setting theory and research. *Academy of Management Review*, 16(1), 180–199. <https://doi.org/10.5465/amr.1991.4279004>
- Tuomela, R. (2007). *The philosophy of sociality: the shared point of view*. Oxford University Press.
- Vagia, M., Transeth, A. A., & Fjerdingen, S. A. (2016). A literature review on the levels of automation during the years. What are the different taxonomies that have been proposed? *Applied Ergonomics*, 53, 190–202. <https://doi.org/10.1016/j.apergo.2015.09.013>
- Van-Horenbeke, F. A., & Peer, A. (2021). Activity, plan, and goal recognition: A review. *Frontiers in Robotics and AI*, 8, 1–18. <https://doi.org/10.3389/frobt.2021.643010>
- Van Creveld, M. (1985). *Command in war*. Harvard University Press.
- van de Merwe, K., Mallam, S., & Nazir, S. (2024). Agent transparency, situation awareness, mental workload, and operator performance: A systematic literature review. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 66(1), 180–208. <https://doi.org/10.1177/00187208221077804>
- Venable, J., Pries-Heje, J., & Baskerville, R. (2016). FEDS: a Framework for Evaluation in Design Science Research. *European Journal of Information Systems*, 25(1), 77–89. <https://doi.org/10.1057/ejis.2014.36>
- Verhagen, R. S., Neerinx, M. A., & Tielman, M. L. (2021). A two-dimensional explanation framework to classify AI as incomprehensible, interpretable, or understandable. In D. Calvaresi, A. Najjar, M. Winikoff, & K. Främling (Eds.), *Explainable and transparent AI and multi-agent systems* (pp. 119–138). Springer. [https://doi.org/10.1007/978-3-030-82017-6\\_8](https://doi.org/10.1007/978-3-030-82017-6_8)
- Vicente, K. J. (1999). *Cognitive work analysis: Toward safe, productive, and healthy computer-based work*. CRC Press. <https://doi.org/10.1201/b12457>
- Vincenzi, D. A., Moloua, M., Hancock, P. A., Pharmer, J. A., & Ferraro, J. C. (2024). *Human factors in simulation and training: Theory and methods* (2nd ed.). CRC Press. <https://doi.org/10.1201/9781003401360>

- Whittemore, R., & Knafl, K. (2005). The integrative review: updated methodology. *Journal of Advanced Nursing*, 52(5), 546–553. <https://doi.org/10.1111/j.1365-2648.2005.03621.x>
- Wickens, C. D., Lee, J., Gordon-Becker, S., & Liu, Y. (2014). *An introduction to human factors engineering* (2nd ed.). Pearson.
- Williams, J., Fiore, S. M., & Jentsch, F. (2022). Supporting artificial social intelligence with theory of mind. *Frontiers in Artificial Intelligence*, 5. <https://doi.org/10.3389/frai.2022.750763>
- Wolf, F. D., & Stock-Homburg, R. (2020). Human-robot teams: A review. In A. R. Wagner, D. Feil-Seifer, K. S. Haring, S. Rossi, T. Williams, H. He, & S. Sam Ge (Eds.), *Social robotics* (pp. 246–258). Springer. [https://doi.org/10.1007/978-3-030-62056-1\\_21](https://doi.org/10.1007/978-3-030-62056-1_21)
- Wolf, F. D., & Stock-Homburg, R. M. (2023). How and when can robots be team members? Three decades of research on Human–Robot Teams. *Group & Organization Management*, 48(6), 1666–1744. <https://doi.org/10.1177/10596011221076636>
- Woods, D., & Dekker, S. (2000). Anticipating the effects of technological change: A new era of dynamics for human factors. *Theoretical Issues in Ergonomics Science*, 1(3), 272–282. <https://doi.org/10.1080/14639220110037452>
- Wynn, D. C., & Clarkson, P. J. (2018). Process models in design and development. *Research in Engineering Design*, 29(2), 161–202. <https://doi.org/10.1007/s00163-017-0262-7>
- Wynne, K. T., & Lyons, J. B. (2018). An integrative model of autonomous agent teammate-likeness. *Theoretical Issues in Ergonomics Science*, 19(3), 353–374. <https://doi.org/10.1080/1463922X.2016.1260181>
- Yoshida, Y. (2021). Multiple-models juxtaposition and trade-offs among modeling desiderata. *Philosophy of Science*, 88(1), 103–123. <https://doi.org/10.1086/710054>
- Zimmerman, J., Forlizzi, J., & Evenson, S. (2007). Research through design as a method for interaction design research in HCI. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 493–502. <https://doi.org/10.1145/1240624.1240704>

# Appendix

This appendix collects the questions posed across the framework. Those from the framing (Table 8) concern what team intentionality consists of; those from the realising (Table 9) concern how it may be supported across the design process.

Table 8. Questions for analysis from the framing of team intentionality.

|                 |  | Concept Aspect     | Question for analysis                                                                                                                                                          |
|-----------------|--|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| The hybrid team |  | Membership         | Who is recognised as a member, and whether that recognition is reciprocal.                                                                                                     |
|                 |  | Boundaries         | What ethos can be derived, and whether it is aligned across members.                                                                                                           |
|                 |  | Purpose            | What intent members hold, and whether it remains aligned.                                                                                                                      |
|                 |  | Accountability     | Who is responsible for what, and whether accountability is held individually, jointly, or collectively.                                                                        |
|                 |  | Interdependence    | What requirements and opportunities may arise, and how they are organised over time.                                                                                           |
|                 |  | Teamness and modes | What the activity demands of each, and whether the total mode configuration meets those demands.                                                                               |
| The strata      |  | Background         | What members can do individually, jointly, and collectively, how these configure in a situation, and what intentional states may then emerge.                                  |
|                 |  | Network            | Which nodes shape a salient state, and why it takes the form it does.                                                                                                          |
|                 |  | Foreground         | What aspects of situation and intent are surfaced, what conditions must hold for the situation to be recognised and the intent achieved, and whether members surface the same. |

|                         |                        |                                                                                                                                                                  |
|-------------------------|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| The situation as framed | Mechanisms             | What can be surfaced and communicated, what can only be developed through socialisation over time, and what remains implicit.                                    |
|                         | Orienting              | What purposes make what frames available, which are adopted, and how this shapes the framing of both the situation and a teammate within it.                     |
|                         | Situating              | What sources are fitted to each slot, how far these can be trusted, and how closely they match the actual state when sources fail or content cannot be obtained. |
|                         | Nesting                | Which situations are salient within a context, at what granularity, how they shape one another, and how a member moves between them over time.                   |
| Agency                  | The loop               | How the loop operates in a situation, and which activities surface, recur, or remain latent.                                                                     |
|                         | Perceive               | How perceiving builds awareness, and how this bounds the choices a member can consider.                                                                          |
|                         | Choose                 | What conditions an alternative must satisfy to be considered, and what further conditions to be decided.                                                         |
|                         | Decide                 | What conditions prompt a decision, and what conditions prompt its reconsideration.                                                                               |
| The agentive unit       | Act                    | Whether the performance can be carried out and an outcome produced, and whether they meet the conditions of satisfaction set for it.                             |
|                         | Distance               | Where the distance between what is and what is desired surfaces, and how the loop responds.                                                                      |
|                         | Agentive participation | Whether intentionality aligns to the degree each configuration demands.                                                                                          |
|                         | Activity coupling      | Which intentional processes feed one another, what outcomes are expected between them, and toward what shared end they converge.                                 |

---

|                            |                          |                                                                                                                                                                                                      |
|----------------------------|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Autonomy boundaries</b> | Configurational dynamics | Whether a configuration allows sufficient adaptability to uphold its intentionality as the situation changes.                                                                                        |
|                            | Open choices             | Whether and to what degree autonomy exists, and at what agentive unit.                                                                                                                               |
|                            | Boundary dynamics        | When and how long each choice stands open, in what order, and whether the openings mark a temporal independence, dependence, or interdependence.                                                     |
|                            | Autonomy fit             | How well the boundaries fit, toward the world and between units, and how the fit holds as situations unfold.                                                                                         |
| <b>Intent</b>              | Choice content           | What aspects of intent are or should be under consideration, and whether their content is shared or complementary between members.                                                                   |
|                            | Conative commitment      | Which choices are committed, whether deliberated or connoted, how firmly, and under what conditions they may change.                                                                                 |
|                            | Temporal structure       | How commitments are organised across time and levels of cognitive control, whether this is coherent within and between members, and where partial structures are reused or missing.                  |
|                            | Coordinating commitments | Where a member holds commitments that are inadmissible together, and whether their status leaves the conflict resolvable.                                                                            |
| <b>Aligned</b>             | Guiding conduct          | How the tempo of shifting between deliberation and monitoring aligns between members, whether they hold commitments at different grains, and whether the shift runs one way or moves back and forth. |
|                            | Satisficing alignment    | Whether a member holds the mode the activity warrants, and whether its world-fit is satisfied and any divergence recognised.                                                                         |

---

|                       |                                                                                                                                                             |
|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Positional trajectory | Whether a unit's projected course meets where it should hold, whether its optimal path can reach at all, and where it cannot, what it takes to be the case. |
|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|

Table 9. Questions for analysis from the realising of team intentionality.

| Concept Aspect             |                                         | Question for design                                                                                                                                                                                        |
|----------------------------|-----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Establishing a toolkit     | Establishing reference points           | Which reference points to establish, whether they are coherent, and how they are kept aligned among participants as realisations translate and transform across the design stages.                         |
|                            | Selecting what to represent             | Which aspects to represent, under which analysis, and how the representations may stay connected through shared reference points.                                                                          |
|                            | Building and manipulating models        | Which approaches can represent the selected aspects, and whether they serve as they are or require adaptation to foreground what a design needs.                                                           |
|                            | Sharing intent structures               | Which elements are, can, or should be disclosed and clarified, whether as a single element or a joined segment.                                                                                            |
| Identifying what to convey | Making structures available             | Which elements should stand available in advance, and how they are reconciled with observation, which may mislead.                                                                                         |
|                            | Relating members' situations            | At what grain each member places the other's situation under the shared team frame, whether a perspectival shift is warranted to reach that grain, and what must be disclosed and clarified to achieve it. |
|                            | Comparing intent structures             | What must be disclosed and clarified to realign: what the origin is and who holds it, why it diverged, and how to resolve it through the choices left open.                                                |
|                            | Contrasting decisions with open choices | Which alternatives must accompany a disclosed decision, and when a closure requires clarifying by can or by may.                                                                                           |

---

Matching  
boundaries

Which boundaries must be matched between members and against the world, and which must be left to what if where matching is not possible.

Locating what a  
decision requires of  
another

Which decisions change what another can and may do, and which of these merit disclosure and clarification.

---



# Papers

The papers associated with this thesis have been removed for copyright reasons. For more details about these see:

<https://doi.org/10.3384/9789181186598>

## **FACULTY OF SCIENCE AND ENGINEERING**

Linköping Studies in Science and Technology. Dissertation No. 2542, 2026  
Department of Science and Technology

Linköping University  
SE-581 83 Linköping, Sweden

[www.liu.se](http://www.liu.se)