

# **An information-theoretic account of human–computer interaction**



**Dari Trendafilov**

School of Computing Science  
College of Science and Engineering  
University of Glasgow

This dissertation is submitted for the degree of  
*Doctor of Philosophy*

September 2017



## **Declaration**

I hereby declare that except where explicit reference is made to the work of others, the contents of this dissertation are original and have not been submitted in whole or in part for consideration for any other degree or qualification in this, or any other institution. This dissertation is the result of my own work and contains no further contribution of others except as specified in Acknowledgements and Collaborations.

Dari Trendafilov  
September 2017



## Acknowledgements

This research has been conducted in multiple locations and teams over the years and I am very indebted to all my supervisors for their support, and for letting me be a part of such stimulating environments, which made the exploration of this field an exciting experience. Many people, directly or indirectly, contributed to this work and in the following list of acknowledgements I would like to express my gratitude to all of them, however knowing that such a list would always be incomplete.

First of all, I would like to thank my PhD supervisor Roderick Murray-Smith for empowering me with inspiration, for his guidance and advice. In the final stages of this work Daniel Polani has been an invariable source of invaluable insights, which helped reduce the entropy in my mind. I am grateful to my examiners for their valuable feedback, which improved the quality of the thesis, and in particular to Andrew Howes for bringing a fresh cognitive perspective to my work.

I would like to thank my collaborators at Nokia Research Center – Ákos, Saija, Kaj, Markus Virta and the Demola lab in Tampere – for their help and support during the initial user studies. John, Andy, Andrew and Steve Hughes from the Inference Dynamics and Interaction group at the University of Glasgow provided invaluable help in dealing with the SHAKES. I would also like to extend my thanks to the members of the GIST group in Glasgow and in particular to my co-authors, Yolanda and Eve, and my second supervisor Steven Brewster, for their critical insights. My colleagues from the Adaptive Systems Group at the University of Hertfordshire – Nicola, Christoph, Neil, Martin and Andrés – have been a source of countless fruitful discussions. Many thanks go to Jari, Vuokko, Viljakaisa and Jyri for their trust and support during my studies.

This work has been funded initially by Nokia Research Center Helsinki and subsequently by CORBYS FP7 ICT-270219 and H2020-641321 socSMCs FET Proactive projects.



## **Collaborations**

In chapter 6, Saija Lemmelä and I jointly developed the initial concept and the experimental protocol for Experiment II. I contributed the non-verbal negotiated interactive model and the implementation of the prototype. Saija contributed the game design and the user experience questionnaires. We jointly conducted the user study, which was facilitated by Demola/Suuntaamo technology incubator in Tampere, Finland. Subsequent analyses were performed independently.

In chapter 6, Yolanda Vazquez-Alvarez and I jointly developed the concept, the prototype and the experimental protocol for Experiment III. I contributed the visual and tactile user interface and the development of the prototype. Yolanda provided the auditory feedback and the user experience questionnaires. The holder for the integrated device was designed and 3D-printed by Markus Virta. I conducted the user study, which was facilitated by Demola/Suuntaamo technology incubator in Tampere, Finland. Subsequent analyses were performed independently.



## List of Contributing Publications

- [1] Lemmelä, S., Vetek, A., Mäkelä, K., and Trendafilov, D. (2008). Designing and evaluating multimodal interaction for mobile contexts. *In Proceedings of the 10th International Conference on Multimodal Interfaces (Crete, Greece), ICMI'08*, pages 265–272.
- [2] Trendafilov, D., Lemmelä, S., and Murray-Smith, R. (2014). Negotiation models for mobile tactile interaction. In Murray-Smith, R., editor, *Mobile Social Signal Processing, LNCS Volume 8045*, pages 64–73. Springer-Verlag Berlin Heidelberg.
- [3] Trendafilov, D., Maye, A., Polani, D., Murray-Smith, R., and Engel, A. (2015a). Mutual information as a measure of coordination in collaborative interaction. *British HCI 2015 Workshop on Ubiquitous and Collaborative Computing (iUBICOM)*.
- [4] Trendafilov, D. and Murray-Smith, R. (2013). Information-theoretic characterization of uncertainty in manual control. *IEEE International Conference on Systems, Man, and Cybernetics*, pages 4913–4918.
- [5] Trendafilov, D., Murray-Smith, R., and Polani, D. (2015b). Empowerment as a metric for optimization in hci. *CHI 2015 Workshop on Principles, Techniques and Perspectives on Optimization and HCI*.
- [6] Trendafilov, D., Polani, D., and Murray-Smith, R. (2015c). Model of coordination flow in remote collaborative interaction. *In Proceedings of the 17th IEEE UKSim-AMSS International Conference on Modelling and Simulation*, pages 361–366.
- [7] Trendafilov, D., Vazquez-Alvarez, Y., Lemmelä, S., and Murray-Smith, R. (2011). Can we work this out?: an evaluation of remote collaborative interaction in a mobile shared environment. *In Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services*, pages 499–502.



## **Abstract**

This thesis presents a theoretical framework for the study of interactive systems, using methods from information theory, machine learning and control theory. The framework builds on the information-theoretic capacities of empowerment, relevant information and mutual information, which I adapt and apply to the domain of human–computer interaction. Three user studies exploring dynamic interactive scenarios – one car-tracking and two collaborative target-acquisition experiments – provide empirical data for the development of probabilistic models, used in the characterisation of specific aspects of human performance, such as the level of control, the quality of decision-making, and the level of engagement in interpersonal coordination. Human control models are extended to accommodate for the inherent lags, characteristic for human–computer and human–human interaction, in a principled way. Optimal controllers, describing particular patterns of human behaviour, are built on these theoretical models, providing evidence for specific limits of human performance through simulations. The thesis describes the potential of empowerment, as a generic task-independent measure of control, to characterise the uncertainty in human–machine interfaces. This work builds an important bridge between theory and experiments, and suggests that the proposed information-theoretic concepts could provide analytical tools for supporting the design and evaluation of interactive systems, by elucidating novel aspects of human performance complementing standard measures. The thesis provides proof of concept examples for the application of such information-theoretic measures, and demonstrates how they can be treated naturally side-by-side along traditional metrics used in HCI research. It emphasises the acquisition cost of accurate theoretical models, necessary to ensure the reliability of such measures.



# Table of contents

<b>I Preliminaries</b>	<b>1</b>
<b>1 Introduction</b>	<b>3</b>
1.1 Usability Engineering . . . . .	4
1.1.1 Uncertainty . . . . .	5
1.1.2 The ‘Weak’ Science of HCI . . . . .	6
1.2 Towards Theoretical Foundations of HCI . . . . .	7
1.2.1 Measure of Control . . . . .	8
1.2.2 Measure of Relevant Information . . . . .	9
1.2.3 Measure of Coordination . . . . .	9
1.3 Methodology . . . . .	10
1.4 Contributions to the Field of HCI . . . . .	10
1.5 Thesis Organisation . . . . .	10
<b>2 Background</b>	<b>13</b>
2.1 Introduction . . . . .	13
2.2 Behaviour as Control of Perception . . . . .	13
2.3 The Probabilistic Mind . . . . .	14
2.4 Sensorimotor Control . . . . .	15
2.4.1 Limits of Human Performance . . . . .	15
2.4.2 Fitts’ Law . . . . .	16
2.5 Information and Control Theory . . . . .	16
2.6 Information as a Utility . . . . .	17
2.7 Summary . . . . .	19
<b>3 Concepts and Algorithms</b>	<b>21</b>
3.1 Introduction . . . . .	21
3.2 Information-theoretic Concepts . . . . .	21
3.3 Causal Bayesian Networks . . . . .	24

3.3.1	Perception–Action Loop . . . . .	24
3.4	The Principle of Empowerment . . . . .	25
3.4.1	Definition . . . . .	26
3.4.2	Interpretation . . . . .	27
3.4.3	Applications . . . . .	28
3.4.4	An HCI Perspective . . . . .	28
3.4.5	Numerical Example . . . . .	29
3.4.6	Empowerment and Fitts’ Law . . . . .	29
3.5	Markov Decision Processes . . . . .	30
3.5.1	Value Iteration . . . . .	31
3.6	Relevant Information . . . . .	31
3.6.1	Consistent State Distributions . . . . .	34
3.6.2	Look-Ahead Relevant Information . . . . .	34
3.7	Summary . . . . .	35
 <b>II Human–machine Interaction</b>		<b>37</b>
 <b>4 Empowerment as a Measure of Uncertainty</b>		<b>39</b>
4.1	Introduction . . . . .	39
4.2	Background . . . . .	40
4.3	Model . . . . .	40
4.3.1	Information-theoretic . . . . .	40
4.3.2	Control-theoretic . . . . .	41
4.3.3	Simulations . . . . .	42
4.4	Experiment I . . . . .	43
4.4.1	Experimental System . . . . .	43
4.4.2	Tasks . . . . .	45
4.4.3	Methodology . . . . .	45
4.4.4	Performance Measures . . . . .	45
4.4.5	Performance Analysis . . . . .	46
4.4.6	Behaviour Analysis . . . . .	47
4.4.7	Tracking Patterns . . . . .	48
4.5	Discussion . . . . .	51
4.6	Conclusion . . . . .	52

<b>5</b>	<b>Information Parsimony in Decision Making</b>	<b>53</b>
5.1	Introduction . . . . .	53
5.2	Background . . . . .	54
5.3	Stochastic Model . . . . .	54
5.3.1	Information vs. Utility Trade-off . . . . .	55
5.3.2	LA-RI vs. C-RI . . . . .	56
5.3.3	Empirical Trade-Off . . . . .	58
5.3.4	Soft vs. Sharp Policies . . . . .	60
5.4	Human Performance . . . . .	61
5.4.1	Lag Inference . . . . .	61
5.4.2	Empirical Trade-off . . . . .	62
5.5	Human vs. an Info-Parsimonious Optimal controller . . . . .	64
5.6	Model Fitting . . . . .	65
5.6.1	Look-Ahead Relevant Information . . . . .	66
5.6.2	Gaussian Process Regression . . . . .	68
5.6.3	Performance Characteristics . . . . .	70
5.7	Discussion . . . . .	71
5.8	Conclusion . . . . .	72
<b>III</b>	<b>Social Interaction</b>	<b>73</b>
<b>6</b>	<b>Nonverbal Negotiated Interaction</b>	<b>75</b>
6.1	Introduction . . . . .	75
6.2	Background . . . . .	76
6.3	Interactive Concept . . . . .	77
6.3.1	Distal Collaborative Scenario . . . . .	78
6.3.2	Membrane Metaphor . . . . .	79
6.4	Experiment II . . . . .	79
6.4.1	Experimental Design . . . . .	79
6.4.2	Implementation . . . . .	81
6.4.3	Tasks . . . . .	81
6.4.4	Methodology . . . . .	82
6.4.5	Performance Analysis . . . . .	83
6.4.6	User Experience . . . . .	84
6.4.7	Negotiation Strategies . . . . .	85
6.4.8	Discussion . . . . .	87

6.5	Experiment III . . . . .	88
6.5.1	Experimental Design . . . . .	88
6.5.2	Tasks . . . . .	90
6.5.3	Methodology . . . . .	91
6.5.4	Performance and User Experience . . . . .	91
6.5.5	Emerging Strategies . . . . .	93
6.5.6	Turn-taking Patterns . . . . .	95
6.5.7	Discussion . . . . .	97
6.6	Conclusion . . . . .	98
<b>7</b>	<b>Empowerment in Eyes-free Interaction</b>	<b>99</b>
7.1	Introduction . . . . .	99
7.2	Two Agents in a Line World . . . . .	100
7.2.1	Model . . . . .	100
7.2.2	Dyad Empowerment . . . . .	102
7.3	Experiment . . . . .	104
7.3.1	Methodology . . . . .	104
7.3.2	Results . . . . .	105
7.4	Empowerment as a Measure of Control . . . . .	109
7.5	Discussion . . . . .	111
7.6	Conclusion . . . . .	112
<b>8</b>	<b>Interpersonal Coordination in Collaborative Environments</b>	<b>113</b>
8.1	Introduction . . . . .	113
8.2	Background . . . . .	114
8.3	Model of Coordination . . . . .	115
8.4	Measure of Coordination . . . . .	117
8.4.1	Time-shift Inference . . . . .	117
8.4.2	Probability Distributions . . . . .	120
8.4.3	Delay Sensitivity . . . . .	121
8.4.4	Noise Sensitivity . . . . .	122
8.4.5	Simulations . . . . .	123
8.5	Evaluation . . . . .	125
8.5.1	Analysis of Variance . . . . .	127
8.5.2	Learning Effect . . . . .	127
8.6	Discussion . . . . .	129
8.7	Conclusion . . . . .	130

---

<b>IV</b>	<b>Conclusions</b>	<b>131</b>
<b>9</b>	<b>Conclusions and Discussion</b>	<b>133</b>
9.1	Contributions of the Thesis . . . . .	133
9.1.1	Empowerment in HCI . . . . .	134
9.1.2	Information Parsimony . . . . .	134
9.1.3	Coordination of Behaviour . . . . .	135
9.2	Scope and Limitations . . . . .	135
9.3	Future Work . . . . .	135
9.3.1	Continuous Models . . . . .	135
9.3.2	Discrete Interaction . . . . .	136
9.3.3	Non-parametric Methods . . . . .	136
9.4	Final Remarks . . . . .	136
	<b>References</b>	<b>137</b>
	<b>Appendix</b>	<b>145</b>
A	Experiment II – User Experience . . . . .	147
B	Experiment II – Negotiation Strategies . . . . .	149



# **Part I**

## **Preliminaries**



# Chapter 1

## Introduction

*‘Experience without theory is blind, but theory without experience is mere intellectual play.’*

---

Immanuel Kant

A growing number of increasingly powerful computational resources have become available in everyday life. Personal and pervasive devices are getting new sensing capabilities from physical and virtual sources, which provide richer access to streams of information and open the opportunity for users to engage in novel embodied interactions. Whether these technologies will become a constant source of distraction or a natural extension of our own abilities depends largely on the maturity of interactive systems to understand human behaviour, and infer our intentions and goals [112]. To design better interactive systems we need to draw on sound principles and formal models, provided by theoretical frameworks handling sensing, modelling and inference, as vital aspects of modern HCI. However, this has often been lacking in the HCI community [100], in part due to the perceived gap between the complexity of human behaviour and our ability to capture and model it. This thesis aims to address this gap by exploring the potential of novel information-theoretic utility measures, which tap into the perception–action loop of dynamic interactive systems, to characterise human performance.

The adoption of such utility measures could foster the foundation of a more solid theoretical framework for the study of human–computer interaction, as well as having a number of practical benefits. Rigorous measures could provide analytical tools, revealing the trends and the gradients in interactive models, and could give direct insight into the underlying properties and provide confidence regions for the system’s parameters. They could help provide a firm foundation for designers to treat and evaluate interactive systems in a general fashion. Novel

objective performance measures applied to specific trials of systems with human users could augment our current usability studies. Automatic optimisation of interaction measures with models of user behaviour could be used in the development of adaptive, learning interfaces.

## 1.1 Usability Engineering

Usability specifications and metrics ensure consistency, compatibility and exchangeability in the creation of efficient, effective and satisfactory interfaces, the primary goal of user interface design. Usability engineering provides tools and methods to identify and guarantee that user goals and needs are taken into account. However, as the market becomes increasingly discriminating and terms such as ‘usability’ and ‘ease of use’ are often too vague, it becomes more important to back claims with rigorous evidence. The challenges in usability are summarised by Norman [68] as the gulfs of execution and evaluation – corresponding to the two sides of the perception–action loop – the former reflecting errors in interaction mechanisms, while the latter addressing feedback interpretation issues.

Various evaluation techniques are used in the development of interactive systems. Some use scenarios of interaction to drive both the design and evaluation, however they tend to constrain designers’ attention on a small selection of tasks. Cooperative evaluation provides primarily a rough qualitative feedback in the form of opinions. Observational ethno-methodology gathers helpful feedback during an initial requirements analysis, however it requires a considerable amount of skill and hypotheses are not formally specified. Usability inspection methods, for example such as heuristic evaluation, are tools that can be applied at any stage of the development, since they do not require running systems nor representative test users, however they lack a model. Usability testing is widely recognised as the most reliable way to improve usability, however it is costly and time consuming, and on its own it is not effective in generating constructive recommendations, being rather a means to identify problems than provide solutions. It generally helps designers to get the design right rather than the right design, whereas testing multiple alternatives, along with comparative user ratings, could help designers in selecting the right design, before proceeding with getting the design right [103]. The more rigorous formal experimental techniques, which rely on empirical methods of scientific disciplines, provide measurable observations using an appropriate experimental method, however the subsequent statistical analysis of objective empirical data sometimes test limiting hypotheses, which might neglect important design issues.

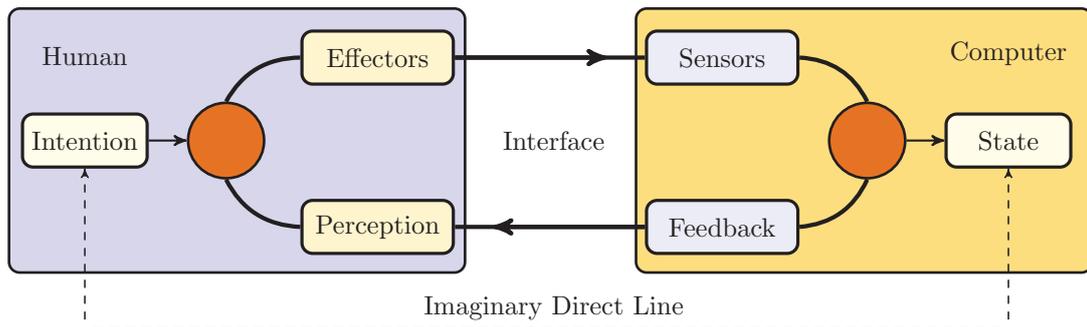


Fig. 1.1 Human–computer interaction as a closed-loop control process. Communication takes place through the interface. The dashed line represents the direct path between user’s intention and system’s state. The circles denote comparator units. (Courtesy of John Williamson).

### 1.1.1 Uncertainty

The closed loop between a user and a computer can be represented as a dynamic system (see Figure 1.1), in which the user can adapt to an extent and designers can alter the feedback mechanisms in order to create an appropriate closed-loop behaviour. Feedback in the perception–action loop is subject to disturbances, as transmission delays and measurement noise, which create uncertainty about the state of the human–machine system and adversely affect usability. Uncertainty has an impact on the design process, as it affects human perception of the environment in the process of receiving information from the world, followed by cognitive processing and action. Human–machine systems must operate in various complex environments of a largely unknown and uncertain world. Therefore, interfaces must reflect the uncertainty in system beliefs and not just filter it out [67], since quality of control depends on appropriate feedback. Increasingly growing computing power enables us to deal with uncertainty in more principled ways, beyond the current ad-hoc approaches. Appropriate use of uncertain feedback could regularise user behaviour and lead to smoother interaction [55]. Lag is known to degrade human performance in motor-sensory tasks of interactive systems, reflected by an increase in both completion time and error rate [62]. Several sources of lag are inevitable and can be attributed to human sensorimotor constraints, sampling rates of input and update rates of output devices (see Figure 1.2). Lag is increased further due to ‘software overhead’ – a loose expression for a variety of system-related factors as communication protocols, network configurations, number crunching, and application software. To compensate for such time-varying delays designers need

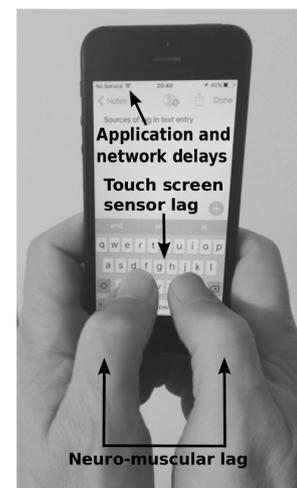


Fig. 1.2 Sources of lags.

to optimise systems for speed, reliability and overall user experience. There is a trade-off between these factors, since if we would only optimise for speed, the system would behave erratically as delay varies, or it would become sluggish if we artificially increased the inherent delay for the sake of stability. The detailed sensitivity analysis of this trade-off requires rigorous tools and measures, however, currently used metrics of human performance require extensive studies, which are costly and still pose certain risks regardless of the point density used for evaluation.

### 1.1.2 The ‘Weak’ Science of HCI <sup>1</sup>

Various theories and schools of thought in the field of HCI are concerned with studying and improving the factors that influence the effectiveness and efficiency of computer use by combining techniques from psychology, sociology, physiology, engineering, computer science, and linguistics. For example, mathematical theories predict human performance when interacting with computers (e.g. Fitts’ Law), theories of cognitive processes create abstract mental models, diversity in the human population is studied by social scientists, and there is also a major stream of art and design work. The collaborative and interdisciplinary structure of HCI research enables the integration of the true knowledge (behavioural models and theories) with the how knowledge (technical solutions) [119], however a unified integrated approach is still lacking [89]. Several authors emphasise the need for an adequate science of HCI [33, 34], in order to address the challenges due to risky hypotheses, difficulty of substantiating experiments through replication, and over-generalisation of experimental results. Fundamental theories could help us to avoid much of the trial and error currently required in HCI engineering.

Interaction designers create artefacts, which inform peers about gaps and opportunities, and thus inspiring new research. These design exemplars currently serve as a medium for the transfer of HCI research to the community of practitioners, instead of rigorous theoretical models. It requires work to be documented in a way convenient for peers to reproduce the results, nevertheless it is unlikely that two designers given the same problem could produce identical or even similar artefacts. HCI researchers often report facts based on hypotheses and collected data, which do not have much context outside of a single experiment, thus most results do not contribute to an underlying fundamental theory, which could serve as a basis for generalisation of results [33]. The multitude of differences between interfaces, devices, target groups, and tasks renders true replications rarely possible, and due to the

---

<sup>1</sup>The title is adopted from Greenberg and Thimbleby [34].

lack of underlying theory there is no means to predict the effect of even minor variations or transfer knowledge from one experimental result to another set of conditions.

While data collection and hypothesis testing helps inform short-term human factors decisions on system design, there is an urgent need for theory-based HCI research. Such research would sidestep replication, for the theory should predict the results under different circumstances. A human–computer system based on solid science should exhibit more predictable behaviour than one founded on intuition or engineering principles alone.

## 1.2 Towards Theoretical Foundations of HCI

One of the first fundamental laws of HCI, Fitts' Law [27], has been shown to apply under a variety of conditions. Recently, further more rigorous formal approaches to HCI research have started to emerge building on e.g. social science [23], bounded rationality [42, 59] and reinforcement learning [20]. MacKenzie's information-theoretic interpretation [63] links Fitts' index of difficulty to the transmission of Shannon information. However, information theory as originally formulated by Shannon is not an experimental science and is concerned primarily with providing theoretic limits on specific quantities assuming a particular world model – the more accurate the model the more informative the limits. A real physical system, however, can be too complex to expect a complete description of its behaviour, e.g. a sensory neuron. Nevertheless, neuroscientists investigating the links between different regions of the brain would like to quantify the information transmission in such complex physical systems [81]. For that purpose recently introduced non-parametric information-theoretic measures such as transfer entropy [86] and directed information [64] are used.

In this thesis I explore a family of novel model-based information-theoretic parametric methods in the context of one specific aspect of human behaviour – i.e. stochastic manual control – and investigate their potential to provide a calculus for identifying and formulating principles guiding biological perception–action loops [50]. Models provide the freedom to control system's parameters and means for deriving quantities in analytic closed-form at the cost of making certain simplifying assumptions and approximations. In the thesis I propose both (1) theories about the user's objective function and (2) measures of the utility of interaction design. The work consists of adaptations of recently introduced generic information-theoretic utilities and their applications for the characterisation of human performance and the evaluation of user interfaces. Building on the principle of empowerment [49], which reflects an agent's potential capacity to control and influence its environment, I develop a universal measure quantifying the uncertainty characteristic for a user interface, and explore the relation of that measure to standard performance metrics. Applying the

concept of relevant information [76], which provides lower limits on the information required on average for decision making in order to achieve a certain level of utility, I derive a model characterising a particular human manual control behaviour. Building on the notion of mutual information I develop a framework for quantifying interpersonal coordination and evaluate it in collaborative studies.

### 1.2.1 Measure of Control

Our view on the fundamentals of interaction is that users' behaviour is about them controlling their perceptions [78]. The more control they have over their perceptions, the more empowered they are by the user interface to achieve their goals. Conversely, they are less empowered the less they perceive the effect of their actions. Perception is tangled up with specific possibilities of action, called 'affordances' by Gibson [31]. Such affordances are the possibility for use, interpretation and action offered by the environment to a specific type of embodied agent. An ability to quantify affordances in a principled way could greatly benefit HCI design.

Klyubin et al [49] introduced the information-theoretic notion of empowerment as a universal task-independent utility, defined by the channel capacity between agent's actions and subsequent sensory observations. Conceptually empowerment is a quantity, which measures the uncertainty in the agent's perceptions related to its actions from the agent's point of view, which makes it a suitable optimality criterion for user interface design optimisation. It captures uncertainty attributed to various factors of human-machine systems (i.e. noise, delays, errors, etc.) in a generic theoretical measure, which reflects the level of control or influence a user has over the environment in the course of interaction. Models of empowerment could predict user performance and users' perception of their own performance.

In this thesis I adapt the original empowerment formalism to the domain of manual control, and suggest its potential in making predictions and providing theoretical bounds on standard performance metrics, based solely on properties of the environment. This approach enables the theoretical evaluation of system's usability in various environmental conditions, as it only requires the probabilistic model of the system's feedback loop. Empowerment could provide an analytical tool for performance tuning by revealing critical salient points in the system's design, before resorting to costly usability testing. Analysing the trends and the gradients of empowerment could give direct insight into the underlying properties, and provide confidence regions for system's parameters trade-off in the design optimisation phase. These insights will help designers to make a better choice for systems to evaluate in user tests. Using the empowerment measure as a first step in the system's analysis will improve quality of design, and at the same time reduce risk and evaluation costs.

### 1.2.2 Measure of Relevant Information

A recent extension of the classical dynamic programming approach to the framework of Markov Decision Processes used in Reinforcement Learning, introduced by Polani et al [76], builds on the hypothesis that the ability to trade off value and informational cost lies at the core of natural behaviour. This information-theoretic view models the agent–environment coupling as an MDP, in which the task is encoded as a cumulative reward maximised by the agent, and transforms the problem into an information–reward trade-off. This informational treatment of the cost required by the agent to attain a given level of performance adopts the view that biological systems implement an information parsimony principle [56, 75], i.e. achieve a given level of performance at the lowest informational cost possible (or perform as well as possible under a given informational bandwidth), which correspond to operating close to the optimal reward/information (decision complexity) trade-off curve, provided a suitable reward function. This hypothesis implies that humans attempt to realise valuable behaviours while minimising the cognitive cost, and may resort to suboptimal solutions if they save a significant amount of cognitive resources.

In order to explore the informational cost in the perspective of a performance-oriented study, and get insight into the structure of decision-making and information processing, I apply the relevant information formalism [76] to the domain of manual control and demonstrate its ability to characterise and model human behaviour.

### 1.2.3 Measure of Coordination

To identify the underlying psychological processes supporting human collaboration and understand how humans perceive, intend, learn, control, and coordinate complex behaviours requires a general framework connecting brain, mind and behaviour, and extending the physical concepts of self-organisation [48]. An example of such a generalisation is unifying dynamic systems theory with neuroscience in the study of cognition and action [99].

In this thesis I take a model-based information-theoretic approach to characterise interpersonal coordination between subjects in collaborative environments. A measure of coordination could reveal the level of mutual entrainment and synchrony between subjects, and thus provide insights about the quality of a particular user interface in real time. The approach is evaluated in collaborative studies, suggesting the potential of the proposed measure to predict learning effects, and revealing its relation and sensitivity to the smoothness of interaction. A detailed sensitivity analysis demonstrates the accuracy, the robustness and the coherence of the measure, emphasising its benefits and drawbacks.

### **1.3 Methodology**

The research methodology followed in this thesis consists of iterative design, simulation and validation of theoretic models capturing specific aspects of human behaviour – manual control, interpersonal coordination and decision making – using tools from information theory, reinforcement learning and control theory. To address uncertainty associated with a user interface in a uniform and principled way I employ entropy-based information-theoretic principles and derive probabilistic models, describing more appropriately the variability in human behaviour. In order to explore the properties of the constructed models and fine-tune the models' parameters I perform series of iterative simulations. To validate the theoretical models I conduct empirical studies and feed back the collected experimental data into the next iteration of model refinement.

### **1.4 Contributions to the Field of HCI**

The main contributions of this thesis to the field of HCI are as follows:

- an extension of the empowerment formalism to accommodate for delays inherent in the perception–action loop of interactive systems, which enables its application to problems in the domain of HCI
- an application of the principle of empowerment to model uncertainty in a dynamic interaction task
- a concept for minimalist non-verbal multimodal remote collaboration
- a model of empowerment for minimalist collaborative interaction
- a model of human tracking behaviour and an information-theoretic measure of interpersonal coordination in collaborative interaction
- an application of the concept of relevant information providing theoretical bounds on human performance and a model of human behaviour in a dynamic interaction task

### **1.5 Thesis Organisation**

The thesis is split in four parts as depicted in Figure 1.3. Part I sets out the scope and the motivation for this work (Chapter 1), introduces the perception–action loop in human sensorimotor control (Chapter 2) and provides a brief overview of the technical background

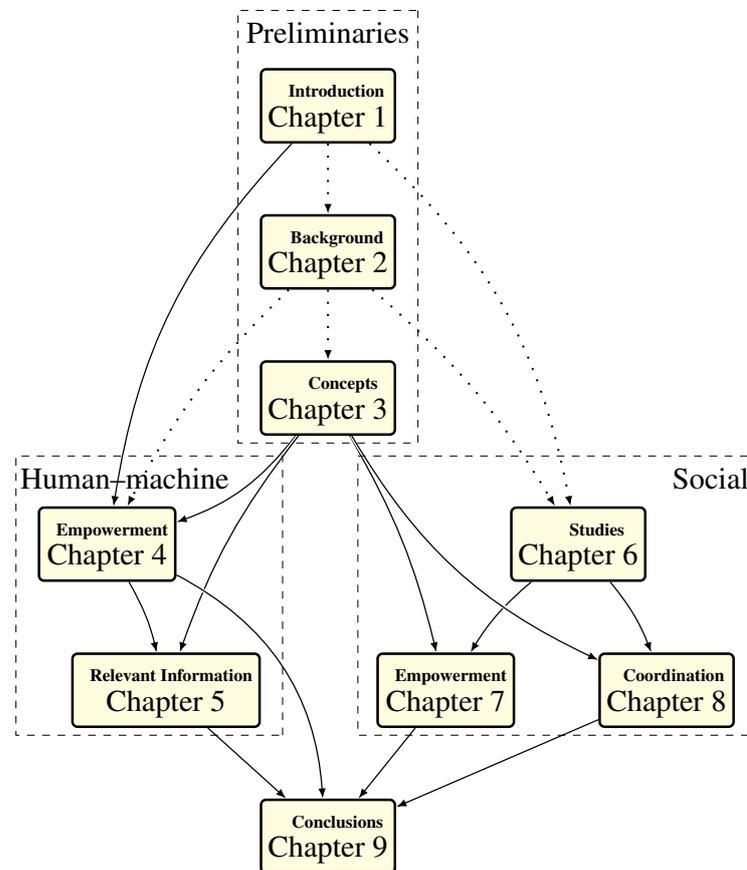


Fig. 1.3 Reader's Guide describing the chapter dependences. Continuous arrows depict a pre-requisite, and dashed arrows a co-requisite relationship between chapters.

needed to comprehend the thesis, including terminology, standard notation and recently introduced information-theoretic concepts (Chapter 3).

Part II presents two information-theoretic treatments of human-machine interaction in the context of an object tracking scenario based on empowerment (Chapter 4) and relevant information (Chapter 5). Chapter 4 introduces an extension of the empowerment formalism accommodating for delays inherent in HCI and its application to a particular interactive task revealing the relation of empowerment to standard performance metrics used in HCI research. Chapter 5 presents an application of the relevant information approach for modelling an object tracking manual control task, along with simulation results and model fitting to specific patterns of human behaviour.

Part III investigates the potential of empowerment and mutual information for characterising specific aspects of human behaviour in the context of dyadic social interaction. Chapter 6 introduces a novel concept for minimalist collaborative interaction and presents its initial feasibility exploration in two user studies – the first one investigating an abstract form of user interface, while the second one focusing on a particular calendar UI concept. Chapter 7 develops an empowerment model of tracking behaviour building on the minimalist collaborative interaction paradigm and explores its properties in simulations. The data collected in both studies is analysed using this model and the resulting average empowerment levels are related to subjective performance metrics. Chapter 8 develops models of tracking behaviour based on the experimental data and proposes an information-theoretic measure of interpersonal coordination building on mutual information along with its detailed sensitivity analysis and a complete coordination analysis of both user studies.

Part IV (Chapter 9) presents the conclusions, including the main contributions, the scope and the limitations of the thesis, and potential directions for future work.

# Chapter 2

## Background

*‘There is no truth. There is only perception.’*

---

Gustave Flaubert

### 2.1 Introduction

This chapter introduces the control system view of human behaviour adopted in this thesis, along with the Bayesian approach to human rationality. It presents specific limits of human sensorimotor performance characteristic for various parts of the perception–action loop, which are fundamental for this work. Furthermore, it introduces the concept of Shannon information as a complementary utility function to control-theoretic approaches, which will be explored throughout this thesis in modelling human performance.

### 2.2 Behaviour as Control of Perception

What an organism senses affects what it does and what it does affects what it senses. This circle of cause and effect is closed through the feedback mechanism described by Hebb [38] as follows: *“Any behavioural response to a single stimulation thus produces a sensory feedback which can act as the initiator of a second response, whose feedback initiates a third response, and so on.”* A living organism behaves in a smooth continuous manner with both responses and stimuli continually changing and interacting with each other. Action feedback occurs as soon a response begins and can affect the response while it is happening, whereas learning feedback is generally delayed enough to prevent the response from being directly affected by it [1].

Powers' perceptual control theory [78] views human behaviour as a feedback process organised around the control of perception and continuously reorganised in order to maintain itself close to a genetically-specified optimal condition. The perception–action loop is closed through the environment, therefore models of human performance should give consideration to contextual and ecological constraints.

Tomasello's control system approach views goal-oriented behaviour as an iterative feedback process, in which goal determines action, which changes perception (feedback), which (when compared to the goal) again determines action [104]. Utilising negative feedback as a capacity to perceive actions in the context of intentions is an important skill humans have for coping with the environment. This views human as a closed-loop controller responding to the relation between the reference input and the output.

## 2.3 The Probabilistic Mind

Almost every aspect of cognition including learning and motor control involves uncertainty and its resolution plays a key role in both perception and action. One form of uncertainty is the inherent ambiguity arising from situations with multiple interpretations. Another form is signal imprecision. In addition, noise at various processing stages limits the precision of perception and noise in the motor outputs limits the precision of control. Another important characteristic of physical control systems is the presence of time delays, which impose further constraints on efficiency [43]. The human brain, however, has developed remarkably effective methods for dealing with uncertainty in a long evolutionary process.

Rational human behaviour involves achieving goals in complex environments, in which the lack of crucial information and computationally intractable optimisation suggest that it cannot be explained with systems of logic [78]. Growing evidence reveal that in the face of a complex and uncertain world humans are guided by Bayesian rationality, supporting the probabilistic view of cognition as calculus of uncertain reasoning [70]. The human brain can combine multiple sources of information to form maximum likelihood estimates and can incorporate prior beliefs in order to generate maximum a posteriori estimates [117, 118]. The maximum likelihood estimates aim to increase the probability of seeing the data given the parameters, while initial knowledge about the parameters of the world (i.e. a prior) allows to estimate the posterior and find the most probable parameter settings using the Bayes rule (see Equation 2.1). Bayesian inference provides a probabilistic approach for describing the world in the face of uncertainty, where probabilities represent the degrees of belief in different propositions.

$$Posterior = \frac{Prior \times Likelihood}{Evidence} \quad (2.1)$$

## 2.4 Sensorimotor Control

Manual control theory typically studies the human as an operator of dynamic systems performing continuous tracking tasks, in which the goal is to minimise the error between a controlled object and a target, i.e. nullify a disturbance. One basic control model, known as ‘observer’, estimates the current state using both sensory feedback and forward models. The human brain runs simultaneously multiple forward models that predict the sensory consequences of particular actions and if the prediction of a forward model closely matches the actual sensory feedback, then its paired controller is selected to determine subsequent motor commands [118]. This basic model can compensate for sensorimotor delays by predicting ahead and can reduce uncertainty arising from noise inherent in both sensory and motor signals. Various means for penalising errors, which in different fields have different names – loss, cost, utility or reward, are used in the action-selection decision-making process, with the same goal of minimising the expected loss/cost or maximising the expected reward/utility given prior beliefs.

Motor control problems generally require the consideration of dynamic constraints of the process being controlled, physical constraints and value constraints or performance criteria [43]. While classical control theory is usually concerned with stability and tracking error, certain cases require different performance criteria. For example, the smoothness of response could be characterised by combining mean squared tracking error with the rate of change of control movement.

### 2.4.1 Limits of Human Performance

The limits of human performance, related to features of the neural circuitry and musculoskeletal system, significantly affect our ability to produce accurate and fast movements. The human motor control system, which is high-dimensional and operates on over 600 muscles [116], is adapted to a natural environment, in which disturbances occur with limited rate of change. Our error correction mechanisms operate fast enough to prevent natural disturbances from having significant effects on what we perceive and control. Every control action, however small the disturbance, is limited in speed of error correction. An absolute limit is set by transmission lags inside and outside the control system.

There are time delays in both the transduction and transport of sensory signals to the Central Nervous System (CNS). Visual feedback, for example, can take up to 100 ms to be processed. Cumulative sensory and efferent delays associated with movement result in appreciable levels of lag. Intrinsic neural noise in sensory inputs and motor commands further limits our ability to perform rapid and accurate movements.

In the face of uncertainty humans seem to allow variability in redundant (task-irrelevant) dimensions and achieve goals reliably and repeatedly with movements rarely reproducible in their detail [102]. Such a control model does not imply a single optimal trajectory, but instead uses feedback dynamically to correct only deviations that interfere with task goals. Whenever the task allows redundant solutions, movement duration exceeds the shortest sensorimotor delay, and optimal performance is achieved by a feedback control that resolves redundancy moment-by-moment. By postponing decisions regarding movement details until the last possible moment, this control model takes advantage of the opportunities that are constantly created by unpredictable fluctuations away from the average trajectory. This ‘minimal intervention’ principle suggests why humans are sometimes ‘lazy’, ‘sloppy’, and perform below their peak abilities. Such behaviour can only be optimal if it saves valuable resources part of the cost function. What is often interpreted as a sign of sloppy control by the brain may reflect an optimal strategy accounting for the information processing cost required for decision making, since every motor task needs a specific control law and the CNS must select the appropriate one in each case.

### 2.4.2 Fitts’ Law

The fact that motor commands are corrupted proportionally by noise implies that faster movements are less accurate. This proportion, known as Fitts’ law [27], has been extensively studied in the field of HCI [35]. Fitts’ law characterises performance in terms of endpoint error and explains the inverse relationship between speed and accuracy. In its original form it quantifies the (index of) difficulty  $ID$  of a target selection task, measured in bits (see Equation 2.2), where  $D$  is the distance to target and  $W$  is the width of the target.

$$ID = \log \frac{2D}{W} \quad (2.2)$$

Performance indices based on trajectory details rather than outcome alone have also been proposed, since results related to the smoothness of arm trajectories seemed impossible to explain with purely outcome-based indices.

## 2.5 Information and Control Theory

Both information and control-theoretic approaches can be applied for modelling specific aspects of human performance, however certain features of such complex control systems can only be apprehended by taking an integrated perspective on the phenomenon [43]. For example, information theory can be used to characterise the communication channel

bandwidth in human perception, while control theory can model the action dynamics. Such an integrated view could address effectively the key question of how constraints on both sides of the feedback loop interact to bound performance.

The interplay between sensors and actuators is commonly described as a transfer of information involving three steps: estimation, decision and actuation. In the first step sensors are used to gather information from the controlled system regarding its state. This information is then processed according to a control strategy in order to determine which control dynamics is to be applied, to be finally transferred to the actuators which modify system's dynamics, typically aiming at decreasing uncertainty. In closed-loop or feedback control techniques, actuators rely explicitly on the information provided by sensors to apply the actuation dynamics, and this relation is shown to be a zero-sum game, i.e. each bit of information gathered from a dynamical system by a control device can decrease the entropy of that system by at most one bit additional to the reduction of entropy attainable without such information [107, 108].

Information and uncertainty represent complementary aspects of control. Closed-loop methods obtain information about system variables, and use that information to decrease the uncertainty about the values of those variables. Therefore, in a control process information must constantly be acquired, processed and used in order to maintain the system trajectory. Entropy is a suitable candidate for characterising uncertainty as it offers a precise measure of disorderliness or missing information by quantifying the minimum amount of resources (bits) required to encode unambiguously the ensemble describing the system. Decreasing entropy stabilises dynamical systems from disturbances associated with environmental noise, motion instabilities, and incomplete specification of control conditions. The goal in system control is typically to reach a low-entropy final state, starting from a high-entropy (random) initial state. Performance criteria, such as distance to target or energy consumption, are used to determine the optimal control, however they do not consider the information processing cost required for decision making.

## 2.6 Information as a Utility

Perception, information processing and actuation are usually treated as individual processes, however considering them as an ensemble introduces the perception–action feedback cycle, which is described by Fuster [28, 29] as ‘the circular flow of information between an organism and its environment in the course of a sensory guided sequence of actions towards a goal’.

There are various ways of modelling the perception–action cycle, which poses a challenge for a unified quantitative treatment. One universal approach is the information-theoretic view

inspired by Ashby [3] and developed further by many others [14, 50, 52, 107, 108]. It is general, conceptually transparent and can be post-hoc imbued with the specific constraints of particular models. On the informational level scenarios with different computational models can be directly compared with each other. The informational treatment also allows to impose constraints on the information processing capacity, and enables the consideration of the informational cost.

The concept of information, as introduced by Shannon, derives its power from the strict rejection of semantic elements in its formalism, and initially it has been doubtful whether semantics could be treated within such framework at all [7, 31]. However, emerging evidence suggests that exploring the intrinsic dynamics of information can provide various interesting utility concepts. Important information about a system's structure can be obtained by measuring to what extent individual components contribute to information production and at what rate they exchange information between each other. Several authors emphasise the utility of having a measure for a flow of information [5, 49, 50, 113].

Various methods for studying the dynamics of information shared between processes have been proposed, which aim to detect the directionality of coupling and quantify the degree of asymmetry [83, 93]. Their goal is to assess the interaction between two subsystems by analysing the interrelation between the two signals at their outputs. In order to determine the direction of causality between two variables Granger introduced measures of causal lag and causal strength in an explicit and testable fashion [32]. Mutual information has been used widely to measure the overlap of information content between two (sub-) systems, as a natural way to quantify for deviation from independence of two processes, however it contains neither dynamical nor directional information. Shannon entropy and mutual information are properties of the static probability distributions, while the process dynamics are represented by the transition probabilities. Another utility, introduced by Schreiber [86] and called transfer entropy, aims to quantify the information transfer between two systems, by characterising the statistical coherence of the systems evolving in time.

Ay and Polani [6] proposed a measure for the strength of a causal effect, which captures essential properties of a Shannon-type quantity, while realising a flow-like philosophy different from the correlative nature of mutual information. Their concept of information flow, based on causal Bayesian networks, can be seen as an information-theoretic counterpart of the probabilistic formalism of Pearl [74]. A measure of causal flow of information could enable the quantification of a number of phenomena in the areas of synchronisation, game dynamics and the perception–action loop. For example, cooperative behaviour of coupled complex systems is related to synchronisation phenomenon. When two players adapt their strategies over time, the game dynamics could move towards cooperative or antagonistic

behaviour, and information flow could reveal a given player's contribution for the emergence of a particular cooperative or antagonistic strategy.

Shannon information theory can quantify the amount of information involved in communication, which depends only on the probabilistic structure of the communication process, however it does not take into consideration the economic impact of uncertainty. The theory of the value of information addresses both the probabilistic and economic factors that affect decisions [41]. Placing a value on the reduction of uncertainty is the first step in experimental design, which provides a basis for allocating resources to reduce uncertainty.

## **2.7 Summary**

This chapter provided the rationale for exploring probabilistic information-theoretic models of human–computer interaction and a link to related control-theoretic aspects. In the next chapter I give a concise introduction to basic information-theoretic concepts used in the thesis and present a detailed description of two key principles, i.e. empowerment and relevant information, which I apply to human–computer interaction.



# Chapter 3

## Concepts and Algorithms

*‘Nobody knows what entropy really is, so in a debate you will always have the advantage.’*

---

John von Neumann to Claude Shannon

### 3.1 Introduction

The formalisms presented in this thesis rely extensively on the combination of different fields of science – statistics, Markov decision processes (MDP), information and control theory. MDPs describe the decision-making process of an agent interacting in a stochastic world. Information theory provides means for modelling, analysing and optimising informational properties of such processes. Control theory describes the human operator as a controller. Furthermore, tools from human–computer interaction theory and statistics are utilised in the experimental design and the evaluation of user studies. In this chapter I give a brief overview of relevant theoretic frameworks, introduce the respective terminology and elaborate on the recently introduced principles of empowerment and relevant information, which are applied in this work. More in-depth review of MDPs can be found in [12, 97], of information theory in [24, 61, 87], and of manual control in [43, 47, 77].

### 3.2 Information-theoretic Concepts

The most basic concept in probability theory is the notion of a random variable. It is a variable which can take on values from a set, called an alphabet, according to some probability distribution. Random variables are denoted with capital letters, such as  $X$ ,  $Y$ ,  $S$ ,

and  $A$ , and their alphabets with the corresponding calligraphic capitals,  $\mathcal{X}$ ,  $\mathcal{Y}$ ,  $\mathcal{S}$ , and  $\mathcal{A}$ . The size or the cardinality of an alphabet is denoted as  $|\mathcal{X}|$ . Events are specific instantiations of random variables, denoted with the corresponding lower case letters, i.e.  $x$ ,  $y$ ,  $s$ , and  $a$ .

The probability distribution of a random variable  $Pr(X = x)$  is a function over all events in an alphabet. To simplify notation, I will use  $p(x)$  for the probability  $Pr(X = x)$ , whenever the random variable in question is obvious. The uncertainty about the value of a random variable is quantified by the information-theoretic concept of entropy, which is defined as

$$H(X) = - \sum_x p(x) \log p(x).$$

Depending on the base of the logarithm entropy can vary by a multiplicative constant implying specific units. For example, if a binary logarithm is applied the entropy is measured in bits, whereas in the case of natural logarithm the units are called nats. The type of units does not affect the generality of the theoretical treatment and across this thesis bits and nats are explicitly specified and used consistently.

Entropy can be interpreted as the average amount of information gained when a variable's value is revealed. When multiple variables are correlated, then knowing the value of one reduces the uncertainty about the other. The average uncertainty about  $Y$  left after revealing the value of  $X$  is quantified by the conditional entropy

$$H(Y|X) = - \sum_{x,y} p(x,y) \log p(y|x).$$

The average reduction in uncertainty, interpreted as the amount of information that knowing  $X$  gives about  $Y$ , is defined as the mutual information between the two variables

$$I(X;Y) = H(Y) - H(Y|X) = \sum_x p(x) \sum_y p(y|x) \log \frac{p(y|x)}{p(y)}.$$

The mutual information is

- non-negative:  $I(X;Y) \geq 0$ ,
- zero, if and only if the variables are fully independent,
- symmetric:  $I(X;Y) = I(Y;X)$ ,
- upper bounded by the entropy of each variable:  $I(X;Y) \leq \min(H(X), H(Y))$ .

Furthermore, adding a third variable  $Z$ , the information  $X$  provides about  $Y$  given knowledge of the value of  $Z$  defines the conditional mutual information

$$I(X;Y|Z) = H(Y|Z) - H(Y|X,Z) = \sum_z p(z) \sum_x p(x|z) \sum_y p(y|x,z) \log \frac{p(y|x,z)}{p(y|z)}.$$

The conditional mutual information depends on how the variables interact. Three random variables  $X$ ,  $Y$ , and  $Z$  form a Markov chain  $X \rightarrow Y \rightarrow Z$ , if  $X$  and  $Z$  are independent given  $Y$ . This implies that the joint probability  $p(x,y,z) = p(x)p(y|x)p(z|y)$  and  $I(X;Z|Y) = 0$ . The relative entropy of two distributions, i.e. Kullback-Leibler divergence, is defined as

$$D_{KL}(p(x)||q(x)) = \sum_x p(x) \log \frac{p(x)}{q(x)}.$$

It reflects the distance between two distributions, being zero when they are identical, and growing when their probabilities diverge. However, it is not a proper metric, as it is not symmetric (generally,  $D_{KL}(p||q) \neq D_{KL}(q||p)$ ). Furthermore, if  $q(x) = 0$  for some  $x$  for which  $p(x) > 0$  a division by 0 occurs. A symmetric derivative, bounded between zero and one, is defined as Jensen-Shannon divergence (where  $m = (p + q)/2$ )

$$D_{JS}(p||q) = (D_{KL}(p||m) + D_{KL}(q||m))/2.$$

Communication channels are used to describe the transmission of information. A channel is modelled with two random variables  $X$  and  $Y$  and their corresponding alphabets  $\mathcal{X}$  and  $\mathcal{Y}$ , describing the possible inputs and outputs, and a transition probability function  $p(y|x)$ , where  $x$  is the emitted and  $y$  the received message. Channels come in several flavours, determined by the possible interactions between the input and the output. One basic example of a channel is when the output depends only on the immediate input – past inputs/outputs do not have a direct effect:  $p(y_t|x_t, y_{t-1}, x_{t-1}, \dots, y_{t-k}, x_{t-k}) = p(y_t|x_t)$  for any  $k > 0$ . Such a channel is called memoryless, and if this equality does not hold the channel has memory. When the channel's output affects the next input, i.e.  $p(x_{t+1}|y_t) \neq p(x_{t+1})$ , the channel has feedback.

One fundamental information-theoretic measure is the channel capacity, which is the maximum amount of information that can be transmitted through a channel, and is defined as the maximum mutual information between the input  $X$  and the output  $Y$ , computed with respect to all input distributions  $p(x)$

$$C = \sup_{p(x)} I(X;Y).$$

### 3.3 Causal Bayesian Networks

In order to formalise qualitative assumptions about cause-effect relationships and derive causal inferences from a combination of assumptions, experiments and data, I will use graphical models in the form of directed acyclic graphs (DAG) [73], which serve as a mathematical language for integrating statistical and subject-matter information. A Bayesian network is a type of directed acyclic graph, which requires that no vertex is a descendant or an ancestor of itself, hence its vertices can be ordered chronologically, from ancestors to descendants, defining a causal ordering and a timeline on the graph directed from left to right. More than one Bayesian network may fit observed data. The key benefit of DAG is the ability to describe causal, rather than statistical associations, since causal models provide information about the effects of actions. In essence, a joint distribution tells us how probable events are and how probabilities would change with subsequent observations, but a causal model also tells us how these probabilities would change as a result of external interventions. A causal Bayesian network allows to determine the effects of interventions, i.e. changes in the underlying mechanisms [74]. Examples of interventions include fixing a variable to a particular value or ‘injecting’ information into the system.

#### 3.3.1 Perception–Action Loop

Interactions between an agent and its environment form a closed perception–action loop. The world is in a state, part of which is captured by the agent’s sensors. Based on this, the agent selects an action, which is combined with the inherent development of the world resulting in the next world state, which closes the loop. Unrolling the perception–action loop in time and constructing a graph with the random variables as vertices and their interactions indicated with edges provide a causal Bayesian network description of the system. A general model of an agent with a controller and memory is presented in Figure 3.1, in which any node, given its parents, is conditionally independent from any other node that is not its parent or successor. Such a system can be partitioned into subsystems in different ways. Here the partitioning is done from the perspective of the agent’s controller, which has direct access

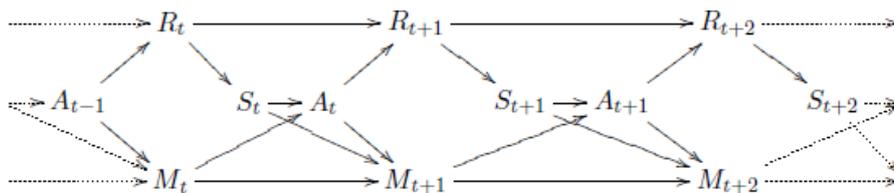


Fig. 3.1 Section of the perception–action loop as a causal Bayesian network.

only to its sensor, its actuator, and its memory. Everything else is denoted as the rest of the system. All constituents of the perception–action loop are modelled as random variables:

- S – sensor state,
- A – action performed by the actuator,
- M – controller’s memory state,
- R – state of the environment.

Sensors and actuators enable information flows between the agent and its environment, and introducing time into the model facilitates to address the temporal aspects of such flows. The states of the sensor, the actuator, the controller’s memory, and the environment at discrete time  $t$  are denoted by the random variables  $S_t$ ,  $A_t$ ,  $M_t$ , and  $R_t$ , respectively. Modelling the relations between the variables at different time steps unrolls the perception–action loop in time. Assuming that the pattern of relations between the variables is time-invariant, and thus holds for any  $t$ , implies that the graph in Figure 3.1 is one section of the network. The agent’s controller is modelled with minimal assumptions about its architecture, using a discrete state, discrete time, and obeying the laws of classical probabilities. In the most general case, the controller performs a probabilistic mapping  $(S_t, M_t) \rightarrow (A_t, M_{t+1})$ .

### 3.4 The Principle of Empowerment

The notion of empowerment is based on an information-theoretic model of the perception–action loop of an embodied agent and its environment. It is a task and representation independent utility function that quantifies the maximum potential information flow from the agent’s actuators to its sensors through the environment. Empowerment was introduced in [49], based on the causal Bayesian network representation of the perception–action loop of a memoryless agent (see Figure 3.2), as a measure of information that can be injected by the agent into its environment and then perceived by its sensors. It depends on the temporal horizon and reflects the environment and the agent’s morphology (embodiment).

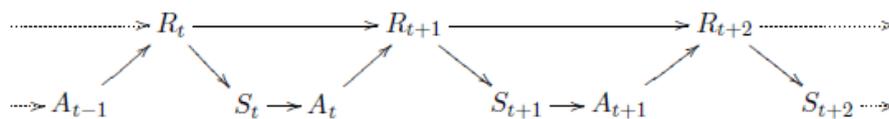


Fig. 3.2 Memoryless perception–action loop as a causal Bayesian network.

### 3.4.1 Definition

Formally, empowerment is defined for stochastic dynamic systems, where transitions arise as the result of making a decision, such as an agent interacting with an environment, as the channel capacity from the sequence of actions  $A_t, A_{t+1}, \dots, A_{t+n-1}$  to the perceptions  $S_{t+n}$  after an arbitrary number of time-steps

$$C(A_t, \dots, A_{t+n-1} \rightarrow S_{t+n}) = \sup_{p(\vec{a})} I(A_t, \dots, A_{t+n-1}; S_{t+n}),$$

where  $p(\vec{a})$  is the probability distribution over the action sequences  $\vec{a} = (a_t, \dots, a_{t+n-1})$ . In the description of the formalism and the algorithm computing empowerment, for simplicity, I will assume a discrete vector-valued state space  $\mathcal{X} \subset \mathbb{Z}^D$  and a discrete action space  $\mathcal{A} = \{1, \dots, N_A\}$ . Furthermore, assuming that the system and the sensor state collapse, I will use the variables  $X$  and  $S$  interchangeably, without loss of rigour. Let  $X'$  denote the random variable associated with  $x'$  given  $x$ . Assume that the choice of a particular action  $a$  is also random and modelled by the random variable  $A$ . The transition function, given in terms of a density  $p(x_{t+1}|x_t, a_t)$ , specifies the probability of going from state  $x_t$  to  $x_{t+1}$  after making a decision  $a_t$ . Assuming the system is fully defined in terms of 1-step transitions, consider the general  $n$ -step interactions, defined by the sequence  $\vec{a}_t = (a_t, \dots, a_{t+n-1})$  of  $n$  single-step actions and the corresponding probability density  $p(x_{t+n}|x_t, \vec{a}_t)$  of making the respective  $n$ -step transitions.

The  $n$ -step empowerment  $\mathfrak{E}(x)$  in state  $x$  is defined as the Shannon channel capacity between the choice of an action sequence  $A^n$  and the resulting successor state  $X'$  (using here the differential entropy form)

$$\mathfrak{E}(x) := \max_{p(\vec{a})} I(X', A^n | x) = \max_{p(\vec{a})} \{H(X' | x) - H(X' | A^n, x)\}. \quad (3.1)$$

The maximisation of mutual information is with respect to all possible distributions  $p(\vec{a})$  over  $\mathcal{A}^n$ . The conditional entropies are given by

$$H(X' | x) = - \sum_{x' \in \mathcal{X}} p(x' | x) \log p(x' | x)$$

and

$$H(X' | A^n, x) = - \sum_{\vec{a} \in \mathcal{A}^n} p(\vec{a}) H(X' | \vec{a}, x) = - \sum_{\vec{a} \in \mathcal{A}^n} p(\vec{a}) \sum_{x' \in \mathcal{X}} p(x' | x, \vec{a}) \log p(x' | x, \vec{a}).$$

Using  $p(x'|x) = \sum_{\vec{a} \in \mathcal{A}^n} p(x'|x, \vec{a})p(\vec{a})$ , Equation 3.1 can be written as

$$\mathfrak{E}(x) := \max_{p(\vec{a})} \sum_{\vec{a} \in \mathcal{A}^n} p(\vec{a}) \sum_{x' \in \mathcal{X}} p(x'|x, \vec{a}) \log \left\{ \frac{p(x'|x, \vec{a})}{\sum_{\vec{a}' \in \mathcal{A}^n} p(x'|x, \vec{a}')p(\vec{a}')} \right\}. \quad (3.2)$$

Hence, given the probability for making the  $n$ -step transitions  $p(x'|x, \vec{a})$ , empowerment  $\mathfrak{E} : \mathcal{X} \rightarrow \mathbb{R}^{\geq 0}$  is a function that maps an arbitrary state  $x$  to a non-negative value  $\mathfrak{E}(x)$ . Numerically, the channel capacity can be computed with the Blahut-Arimoto Algorithm 1.

---

**Algorithm 1: BLAHUT-ARIMOTO**


---

**Data:**  $\mathcal{X}, \mathcal{A}, P_{xx'}^a, \varepsilon$   
**Result:**  $C_k(x) = \sum_{\vec{a}} p_k(\vec{a}) \cdot d_k(\vec{a})$

- 1  $p_0(\vec{a}) \leftarrow \frac{1}{|\mathcal{A}^n|}$
- 2  $d_0(\vec{a}) \leftarrow 0$
- 3  $C_0(x) \leftarrow 0$
- 4  $k = 0$
- 5 **repeat**
- 6      $z_k \leftarrow \sum_{\vec{a}} p_k(\vec{a}) \exp[d_k(\vec{a})]$
- 7      $p_{k+1}(\vec{a}) \leftarrow z_k^{-1} p_k(\vec{a}) \exp[d_k(\vec{a})]$
- 8      $d_{k+1}(\vec{a}) \leftarrow \sum_{x'} p(x'|x, \vec{a}) \log \left[ \frac{p(x'|x, \vec{a})}{\sum_{\vec{a}'} p(x'|x, \vec{a}') p_k(\vec{a}')} \right]$
- 9      $C_{k+1}(x) \leftarrow \sum_{\vec{a}} p_{k+1}(\vec{a}) \cdot d_{k+1}(\vec{a})$
- 10     $k = k + 1$
- 11 **until**  $|C_k(x) - C_{k-1}(x)| < \varepsilon$

---

### 3.4.2 Interpretation

The empowerment model represents the world as a communication channel – when the agent performs an action, it injects Shannon information into the environment, which may or may not be modified, and subsequently reacquires part of this information from the environment via its sensors. By definition empowerment quantifies the maximal potential information an agent can transfer into its own sensors through the environment – the higher the empowerment the more information it can potentially inject. Unlike other task-independent utility functions, it characterises the potential rather than actual behaviour of the agent. Empowerment can be interpreted as the capacity of an agent to control or influence its environment as perceived by its sensors. It can be intuitively thought of as a measure of how many observable adjustments an embodied agent can make to its environment, either immediately, or in the case of  $n$ -step empowerment, over a given period of time. Empowerment maximisation guides the agent to

places in the world where it can benefit most from its sensors and actuators, i.e. the places optimising the options for influencing its relationship with the environment.

### 3.4.3 Applications

Previous work focused on exploring empowerment as a natural utility function for survival-type scenarios in abstract domains, where states with high empowerment coincide with the natural and intuitive choice of a goal state. The smoothness of the system informs the local empowerment gradients around the agent's state of where the most 'alive' states are, property that many dynamical systems have. In goal-directed scenarios when empowerment maximisation is based on the agent's position in space, it mimics the graph-theoretic notion of centrality, i.e. it minimises the expected distance to a random goal state [2].

Empowerment maximisation has provided optimal solutions for continuous problems well-known in the domain of control theory, such as balancing an inverted pendulum, a bicycle and an acrobat [45]. Furthermore, it has shown its potential as a fitness function for learning unknown system dynamics by exploration [45].

In the context of multi-agent systems, empowerment provides a general utility measure facilitating the generation of complex collective behaviour [19]. Different agent embodiments, i.e. sensoric apparatus and motoric abilities, reveal the existence of a trade-off between the freedom of the agents and the constraints they impose on each other, i.e. agents must have enough freedom to act, but they also must be constrained by what the other agents are doing.

### 3.4.4 An HCI Perspective

In order to gain an intuitive insight into the implications of empowerment for HCI consider the tilt-controlled marble maze game on a mobile device. In such a low-latency perception-action loop feedback reliability is crucial. When the player is in a static environment without exposure to major disturbances, the chances for good performance are higher, since there are no significant external factors creating uncertainty in the feedback loop and thus decreasing the empowerment. However, playing the game while walking on the street adds an additional complexity to the interaction, due to changes in gait and posture, which require compensation. In such an environment, the player's empowerment will drop initially and then steadily rise back as the player learns to compensate effectively for the disturbances. This is the process of learning the noisy dynamics, or filtering the noise from the dynamics. Playing the game on the subway might turn out to be an impossible task, since the player loses control over the interaction, i.e. empowerment drops to zero. If the device is able to compensate for the subway jerk, it could make the interaction easier. However, it is not enough for only the

computer to adapt, the player needs to adapt as well, in order to compensate for the erratic hand movements induced by the external disturbance. When the player and the computer take into account the context, they will optimise the contextual empowerment, i.e. empowerment with side information [53]. Adaptive interfaces need to respond to changes in the context and the user model, which requires dealing with a range of uncertainties that affect both the environmental and the user state. These factors are inherently stochastic and can only be dealt effectively with probabilistic tools, which provides grounds for using empowerment.

### 3.4.5 Numerical Example

For a concrete example<sup>1</sup> in the context of HCI consider a pop-up window with ‘OK’ and ‘Cancel’ action buttons and a uniform distribution of actions  $p(a)$  over  $\mathcal{A}$ . In the deterministic case when both actions have distinct outcomes, e.g.  $p(x_1|x, a_1) = p(x_2|x, a_2) = 1$ , using  $p(x'|x) = \sum_a p(x'|x, a)p(a)$  and a natural logarithm gives  $H(X'|x) = \log 2$  and  $H(X'|A', x) = 0$  and hence empowerment  $\mathfrak{E}(x) = I(X', A'|x) = \log 2 \approx 0.69$  nats. In the case of stochastic state transitions, e.g.  $p(x_1|x, a_2) = p(x_2|x, a_1) = 0.1$  and  $p(x_1|x, a_1) = p(x_2|x, a_2) = 0.9$ , leads to  $H(X'|x) \approx 0.69$ ,  $H(X'|A', x) \approx 0.16$  and a decrease of empowerment to 0.53 nats. Further increasing the uncertainty to the maximum level  $p(x_1|x, a_1) = p(x_2|x, a_2) = p(x_1|x, a_2) = p(x_2|x, a_1) = \frac{1}{2}$  gives  $H(X'|x) = H(X'|A', x) = \log \frac{1}{2}$  or  $I(X', A'|x) = 0$ . Zero empowerment means that regardless of the action taken no outcome difference can be perceived, since both choices lead to the same states as the transition probabilities of the two actions completely overlap. Intuitively, empowerment measures the number of available actions on a logarithmic scale, the outcome of which can be perceived. It is zero if, regardless of the action, the outcome will be the same, and is maximal if every action has a distinct outcome.

### 3.4.6 Empowerment and Fitts’ Law

One intuitive interpretation of Fitts’ Law by the principle of empowerment rests on the bounded rationality of the human agent and in particular on the information processing capacity of the human brain (CNS). Faster movements imply higher uncertainty about the effect of every action due to information bandwidth constraints, resulting in lower empowerment and lower acquisition accuracy. Vice versa, slower movements allow for a more accurate information processing, which decreases the entropy in the action feedback and increases empowerment leading to higher precision. This interpretation considers the movement as a closed-loop feedback process consisting of multiple short (in the range of milliseconds) actions executed in a sequence.

<sup>1</sup>A more elaborate numerical example can be found in [45].

### 3.5 Markov Decision Processes

The Markov decision processes (MDP) framework formulates agent–environment systems and their interactions. An MDP model describes how the state of the system is transformed as a result of decisions made by the acting agent. Furthermore, it defines the costs associated with performing the actions, which leads to the question of finding an optimal decision making policy that minimises future costs.

Assuming that the agent–environment system state space  $\mathcal{S}$  is finite, and both time and states are discrete, let the agent in a state  $s_t \in \mathcal{S}$  at time  $t$  choose an action  $a_t$  from its set of available actions  $\mathcal{A}$ . The model dynamics, which describe the evolution of the system’s state, is determined by the state transition probabilities,

$$p(s'|s, a) = p(s_{t+1} = s' | s_t = s, a_t = a),$$

i.e. the probability of ending in a successor state  $s'$  after applying action  $a$  in state  $s$ . The agent receives immediate feedback after performing every action, in the form of a quantifiable reward  $r(s, a, s')$ , or a cost in case of a negative reward.

The goal of a decision-making agent is to act in a way that maximises the expected reward, or minimises the expected cost incurred. The decision-making process is modelled by a policy  $\pi$ , which gives the probability of performing an action given the state

$$\pi(a|s) = P(a_t = a | s_t = s).$$

Assume, for simplicity, that the policy is time-invariant – i.e. the probability of choosing an action is constant over time. The dynamic model and the policy completely specify the sequence of collected rewards, which are accumulated in the value function  $V^\pi(s)$ , defining the expected reward when starting in state  $s$  and following policy  $\pi$  indefinitely

$$V^\pi(s) = \sum_a \pi(a|s) \sum_{s'} p(s'|s, a) [r(s, a, s') + \gamma V^\pi(s')]. \quad (3.3)$$

This recursive definition is called Bellman equation, after one of the first authors to describe MDPs Richard Bellman [10]. The discount factor  $\gamma \in [0, 1]$  reflects the importance of future rewards vs. short-term rewards – if  $\gamma$  is large, future rewards are weighed higher in the expected sum, and if it is small, immediate rewards have higher relative weight. In general,  $\gamma$  serves as a safeguard against divergence.

### 3.5.1 Value Iteration

Finding a policy that maximises the expected future reward is equivalent to solving the following problem

$$\pi^* = \operatorname{argmax}_{\pi} V^{\pi}(s),$$

where the optimal policy  $\pi^*$  is not necessarily unique. Furthermore, any convex combination  $\alpha\pi_1^* + (1 - \alpha)\pi_2^*$  ( $\alpha \in (0, 1)$ ) of two optimal policies  $\pi_1^*$  and  $\pi_2^*$  results in a new optimal policy. Having multiple optimal policies poses the question of whether one is preferred over the others. One natural criterion is the cognitive processing associated with executing a specific policy, i.e. the cost of making a decision. The standard value iteration algorithm finds an optimal policy by computing first the optimal value function  $V^*$

$$V^*(s) = \max_a p(s'|s, a)[r(s, a, s') + \gamma V^*(s')],$$

starting from an arbitrary  $V$  and iterating Equation 3.3 until convergence. From dynamic programming this iteration is known to converge to a global optimal solution, providing a numerical method for its computation. Given  $V^{\pi}$  the state–action value function is defined as

$$U^{\pi}(s, a) = \sum_{s'} p(s'|s, a)[r(s, a, s') + \gamma V^{\pi}(s')], \quad (3.4)$$

referred to as the utility function.

## 3.6 Relevant Information

A recently introduced approach to apply information theory to behaviour, termed relevant information [76], is based on the mutual information between the world state<sup>2</sup> and the agent action variables  $I(S;A)$ . To compute it, it is necessary to specify the joint distribution  $p(s, a)$ . Given an a priori probability distribution  $p(s)$  over the state space, the joint probability  $p(s, a) = \pi(a|s)p(s)$  is fully determined by the strategy  $\pi$ , i.e. the conditional probability of actions given the state. In order to formalise the agent's preferences about  $\pi$  the utility  $U(s, a)$  of taking an action  $a$  in a state  $s$  is used, according which the agent prefers high to low utility values. If the same action is optimal in every state, then the relevant information vanishes, since that optimal action can be selected in every state, without the need to elicit any information about the state. On the opposite, if every state leads to a different optimal action, then all states have to be distinguished in order to select the correct action, and hence

<sup>2</sup>Assuming a fully observable MDP model, in which the sensor state and the world state collapse.

the relevant information is equal to the entropy of the state space  $H(S)$ , i.e. the uncertainty about the current state.

The problem of informational parsimony can be formulated as search for a value-optimal strategy  $\pi(a|s)$ , which at the same time minimises the required relevant information  $I(S;A)$ , i.e. is also information-optimal

$$I^*(S;A) = \min_{\pi^*(a|s)} I(S;A).$$

The optimal strategies  $\pi^*(a|s)$  are exactly those for which the expected utility  $E_{\pi}[U(S,A)] = \sum_{s,a} \pi(a|s)p(s)U(s,a)$  attains its maximal value. This double optimisation is transformed into an unconstrained minimisation problem via Lagrange multipliers

$$\min_{\pi(a|s)} (I(S;A) - \beta E[U(S,A)]). \quad (3.5)$$

The  $\beta$  multiplier<sup>3</sup> implicitly enforces the maximisation of  $U(S;A)$ . For  $\beta \rightarrow \infty$  the optimisation restricts the policy to a value-optimal one. A reduction of  $\beta$ , however, makes the minimisation less sensitive to a drop in utility. This minimax optimisation problem corresponds to trading in utility for a reduction of relevant information.

The relevant information (RI) is given by the solution to this problem, i.e. it is the minimal level of information required for achieving a certain level of performance as measured by the expected utility [76]. Constructing the Lagrangian equation

$$\Lambda(\pi, \beta) = I(S,A) - \beta E[U(S,A)],$$

partially differentiating it

$$\frac{\delta}{\delta \pi} \Lambda(\pi, \beta) = p(s) \log \frac{\pi(a|s)}{p(a)} - \beta p(s)U(s,a),$$

and equating it to zero provides the self-consistent solution

$$\pi(a|s) = \frac{1}{\zeta} p(a) \exp[\beta U(s,a)], \quad (3.6)$$

where  $\zeta$  is a normalising partitioning function.

The standard value iteration method, introduced in the previous section, is extended to accommodate for the double optimisation, by adding an interleaved operation of computing

---

<sup>3</sup>Known from statistical mechanics as thermodynamic beta or ‘inverse temperature’.

the strategy  $\pi'$ , which minimises the relevant information for the current utility level  $U^\pi$ . This leads to the following iteration

$$\pi \xrightarrow{(3.3)} V^\pi \xrightarrow{(3.4)} U^\pi \xrightarrow{(3.5)} \pi', \quad (3.7)$$

which generates a sequence of consistent policy–utility ( $\pi$ – $U$ ) pairs for specific  $\beta$ , by interlacing the updates corresponding to Equations 3.3 and 3.5, and performing these iterations until convergence of both policy and utility [76] (see Algorithm 2).

This approach allows the computation of optimal strategies, while trading off utility and relevant information. However, it does not address the cost involved with obtaining information, but focuses only on how much (relevant) information needs to be acquired (and processed) in order to achieve a certain level of utility, ignoring the possible acquisition cost.

The relevant information reflects the amount of information an agent takes in and processes on average per time-step in the action selection process. This amount depends on the policy – different policies require different informational bandwidth, i.e. processing capacity. It is hypothesised that the required capacity is correlated to the metabolic cost of information processing and constitutes a quantitative measure of cognitive burden [75]. The parsimony pressure tries to minimise this quantity, while the efficiency pressure drives towards higher performance.

---

**Algorithm 2: RELEVANT INFORMATION**


---

**Data:**  $\mathcal{S}, \mathcal{A}, p(s'|s, a), r(s, a, s'), \beta, \gamma, \varepsilon_\pi, \varepsilon_U$   
**Result:**  $\pi_k = \operatorname{argmin}_\pi [I(\mathcal{S}; \mathcal{A}) - \beta E[U(\mathcal{S}, \mathcal{A})]]$

- 1  $p(s) \leftarrow \frac{1}{|\mathcal{S}|}$
- 2  $\pi_0(a|s) \leftarrow \frac{1}{|\mathcal{A}|}$
- 3  $U_0(s, a) \leftarrow 0$
- 4  $k = 0$
- 5 **repeat**
- 6      $p_k(a) \leftarrow \sum_s p(s) \pi_k(a|s)$
- 7      $\pi_{k+1}(a|s) \leftarrow \frac{1}{\zeta} p_k(a) \exp[\beta U_k(s, a)]$
- 8      $U_{k+1}(s, a) \leftarrow \sum_{s'} p(s'|s, a) [r(s, a, s') + \gamma \sum_{a'} \pi_k(a'|s') U_k(s', a')]$
- 9      $k = k + 1$
- 10 **until**  $D_{JS}(\pi_k || \pi_{k-1}) < \varepsilon_\pi$  and  $\|U_k - U_{k-1}\|_\infty < \varepsilon_U$

---

### 3.6.1 Consistent State Distributions

Given a fixed policy  $\pi(a|s)$  and the transition model  $p(s'|s, a)$ , the development of the state in time is described by a first-order Markov chain with transition probabilities

$$p(s'|s) = \sum_a \pi(a|s)p(s'|s, a).$$

When all states in an MDP are positive recurrent, i.e. the average return time for each state is finite, the process has a stationary distribution, which satisfies

$$p(s) = \sum_{s'} p(s'|s)p(s). \quad (3.8)$$

This defines a state distribution consistent with the policy, which results in a variant of Algorithm 2 by replacing the static default uniform distribution in Step 1 with the dynamically updated in every iteration Equation 3.8, and is called the consistent relevant information method (C-RI) [110].

### 3.6.2 Look-Ahead Relevant Information

Introducing a time variable into the model provides the ability to take into account the future. Let  $X_t$  be the state of variable  $X$  at time  $t$ , and  $X_T$  be the state at a random point in time. Making time explicit provides the following informational quantity

$$I(S_T; A_T) = \sum_{s_T} p(s_T)I(S_T = s_T; A_T) = \sum_{t=0}^{\infty} p(t) \sum_{s_t} p(s_t)I(S_T = s_t; A_T).$$

The look-ahead relevant information (LA-RI) [109] in state  $s_t$  under policy  $\pi$ , denoted as  $I^\pi(s_t)$ , is defined by the total expected future relevant information intake when starting from state  $s_t$  and following  $\pi$  indefinitely, where  $p(s_{t+1}|s_t) = \sum_{a_t} \pi(a_t|s_t)p(s_{t+1}|s_t, a_t)$ ,

$$I^\pi(s_t) = I(S_T = s_t; A_T) + \gamma \sum_{s_{t+1}} p(s_{t+1}|s_t)I^\pi(s_{t+1}).$$

Constructing the Lagrangian equation

$$\Lambda_{LA}(\pi, \beta) = p(t=0) \sum_{s_0} p(s_0)I^\pi(s_0) - \beta E[U(S_T, A_T)],$$

partially differentiating it and equating it to zero provides for the optimal policy

$$\pi(a_t|s_t) = \frac{1}{\zeta} p(a_t) \exp[-\gamma \sum_{s_{t+1}} p(s_{t+1}|s_t, a_t) I^\pi(s_{t+1}) + \beta U^\pi(s_t, a_t)]$$

and the following modified Algorithm 3 [109].

---

**Algorithm 3: LOOK-AHEAD RELEVANT INFORMATION**


---

**Data:**  $\mathcal{S}, \mathcal{A}, p(s'|s, a), r(s, a, s'), \beta, \gamma, \epsilon_\pi, \epsilon_U, \epsilon_I^\pi$   
**Result:**  $\pi_k = \operatorname{argmin}_\pi [I^\pi(S) - \beta E[U(S, A)]]$

- 1  $p_0(s) \leftarrow \frac{1}{|\mathcal{S}|}$
- 2  $\pi_0(a|s) \leftarrow \frac{1}{|\mathcal{A}|}$
- 3  $U_0(s, a) \leftarrow 0$
- 4  $I_0^\pi(s) \leftarrow 0$
- 5  $k = 0$
- 6 **repeat**
- 7      $p_k(s) \leftarrow \text{StaticDistribution}(\pi_k, p(s'|s, a))$
- 8      $p_k(a) \leftarrow \sum_s p_k(s) \pi_k(a|s)$
- 9      $\pi_{k+1}(a|s) \leftarrow \frac{1}{\zeta} p_k(a) \exp[-\gamma \sum_{s_{t+1}} p(s_{t+1}|s_t, a_t) I^\pi(s_{t+1}) + \beta U_k(s, a)]$
- 10      $U_{k+1}(s, a) \leftarrow \sum_{s'} p(s'|s, a) [r(s, a, s') + \gamma \sum_{a'} \pi_k(a'|s') U_k(s', a')]$
- 11      $I_{k+1}^\pi(s) \leftarrow D_{KL}(\pi_k(a|s) || p_k(a)) + \gamma \sum_{a, s'} \pi_k(a|s) p(s'|s, a) I_k^\pi(s')$
- 12      $k = k + 1$
- 13 **until**  $D_{JS}(\pi_k || \pi_{k-1}) < \epsilon_\pi$  and  $\|U_k - U_{k-1}\|_\infty < \epsilon_U$  and  $\|I_k^\pi - I_{k-1}^\pi\|_\infty < \epsilon_I^\pi$

---

## 3.7 Summary

This chapter provided the information-theoretic background, which is a pre-requisite for the rest of the thesis. In Part II I adapt and apply the principles of empowerment and relevant information to a human–machine interaction scenario. The theoretic models are built for a specific interactive task and the respective information-theoretic quantities are computed using the presented algorithms.



## **Part II**

# **Human-machine Interaction**



# Chapter 4

## Empowerment as a Measure of Uncertainty

*‘Information is the resolution of  
uncertainty.’*

---

Claude Shannon

### 4.1 Introduction

One of the main goals of this thesis is to demonstrate an application of the principle of empowerment in the field of human–computer interaction and to propose a novel framework for supporting user interface design. This chapter introduces empowerment both as (1) a theory about the user’s objective function and (2) a measure of the utility of interaction design. I adapt the original empowerment formalism and demonstrate its ability to capture uncertainty attributed to various factors in a manual control task (e.g. noise, delays) in a generic theoretical measure. Empowerment was introduced by Klyubin et al. [49] as an information-theoretic universal utility based on the maximisation of control an agent has on its environment. It is fully specified by the dynamics of the agent–environment coupling, i.e. the state transition probabilities, and can serve as an implicit intrinsic reward facilitating the identification of salient states. Initial research on empowerment focused on discrete gridworlds [49] and later expanded to continuous control tasks in the area of reinforcement learning [45]. Previous work [45, 49, 51] aimed to provide a natural utility function for driving behaviour and found that empowerment maximisation directs the agent to ‘interesting’ states in various survival-type scenarios.

## 4.2 Background

The Heidegger-Gibson view of human behaviour as ecological with respect to the environment postulates that perception is the acquisition of information supporting action with regard to constraints on action, where actions and the environment affect each other in complete reciprocity [31]. Humans respond to changing uncertainty in their environment through adaptation influenced by the constraints they observe and the needs that are apparent or intrinsic to them [88]. An example of such adaptation is the McRuer's model, which shows that the human naturally and unconsciously changes to become the inverse of the controlled system dynamics achieving compromise between stability and fast transient response [66]. In mediated interaction the coupling between the human and the interface can affect the coupling between the interface and the controlled process and vice versa, thus influencing interactivity. Key factors contributing to interactivity are the rate at which input can be assimilated, the number of possible actions at any given time, and the way human actions are mapped to changes in the mediated environment [94]. This chapter will present a generic information-theoretic treatment of such factors and will demonstrate how the availability and predictability of human actions are characterised by the principle of empowerment. This will relate system's interactivity to a universal utility measure such as empowerment.

## 4.3 Model

### 4.3.1 Information-theoretic

Originally, empowerment was defined as a function of the system's state, which can often be unknown or uncertain. Therefore, this model considers the prediction of the exact state – the more disturbances affecting the system the more uncertain the prediction. For simplicity, the stochastic transition model is built on discrete action and state spaces.

Consider the following Bayesian network with introduced time delay  $d$  between the actuation and the sensing channel (see Figure 4.1), reflecting that the effect of the action  $a_t$  will only be perceivable at time  $t + d$ . Assuming the system is fully defined in terms of the 1-

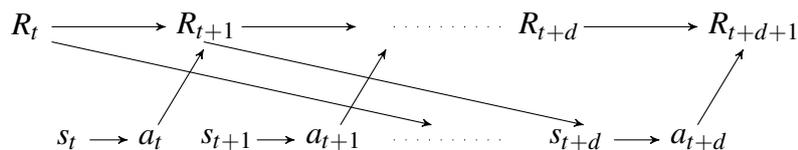


Fig. 4.1 Section of the perception–action loop represented as a Bayesian network with fixed time delay ( $d$ ) between actuation ( $a$ ) and sensing ( $s$ ), closed through the environment ( $R$ ).

step transitions, we are interested in more general  $d$ -step interactions, where  $d \geq 1$  is a random variable from a given distribution. Considering the sequence  $\vec{a}_{t-1} = (a_{t-d}, a_{t-d+1}, \dots, a_{t-1})$  of  $d$  single-step actions, the corresponding probability density  $p(x|x_t, \vec{a}_{t-1})$  of making the respective  $d$ -step transition and assuming that the complete action history is known let  $p(\vec{a}_{t-1})$  denote the probability distribution of the delay  $d$ . Given the latest observed state  $x_t$  and the 1-step  $p(x_{t+1}|x_t, a_t)$  and  $d$ -step  $p(x|x_t, \vec{a}_{t-1})$  transitions define the 1-step transitions for every  $x_{t+1} \in \mathcal{X}$  and  $a_t \in \mathcal{A}$  from the uncertain state  $\tilde{x}_t$  as follows:

$$p(x_{t+1}|\tilde{x}_t, a_t) = \sum_{x \in \mathcal{X}} \sum_{\vec{a}_{t-1}} p(\vec{a}_{t-1}) p(x|x_t, \vec{a}_{t-1}) p(x_{t+1}|x, a_t). \quad (4.1)$$

Equation 4.1 sums up the probability of going from state  $x_t$  to  $x$  under the action sequence  $\vec{a}_{t-1}$  and then from  $x$  to  $x_{t+1}$  under the action  $a_t$ . Substituting the transition density of Equation 4.1 in the original definition of empowerment (Equation 3.2) let

$$\mathfrak{E}(\tilde{x}_t) := \max_{p(a_t)} \sum_{a_t} p(a_t) \sum_{x_{t+1} \in \mathcal{X}} p(x_{t+1}|\tilde{x}_t, a_t) \cdot \log \left\{ \frac{p(x_{t+1}|\tilde{x}_t, a_t)}{\sum_{a'_t} p(x_{t+1}|\tilde{x}_t, a'_t) p(a'_t)} \right\} \quad (4.2)$$

define the 1-step empowerment in the uncertain state  $\tilde{x}_t$ . This formalism extends empowerment to cases in which the exact state is unknown/unobservable and the perception–action loop is affected by time-varying delays (constant delays being a special case).

### 4.3.2 Control-theoretic

Building on the above model I introduced two types of disturbances into the feedback loop of the control system, i.e. communication delays and state noise (see Figure 4.2). The actions are delayed in time, the system state is corrupted by an additive Gaussian noise, and the observation is based on a predictive display.

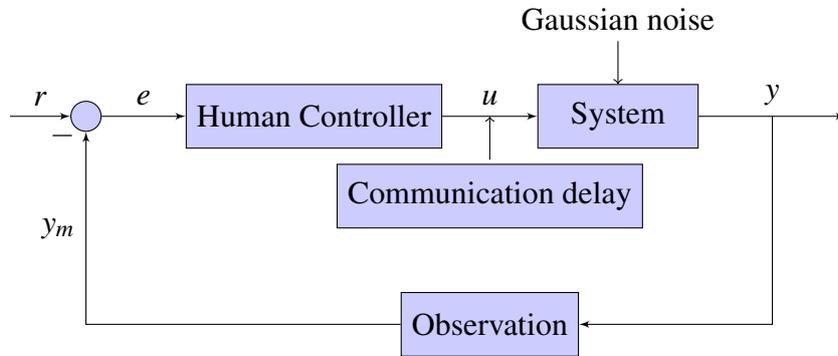


Fig. 4.2 Human as an operator – a control system view of the HCI feedback loop.

the difference between the observation and the reference point, in a given range. Following the notation in Figure 4.2 the human–machine system dynamics is defined as follows

$$x_t = f(x_{t-1}, h(u_t)), \quad (4.3)$$

$$y_t = x_t,$$

where

$$f(x, U) = x + U + \mathcal{N}(0, \sigma),$$

$$h(u_t) = U_t,$$

$$y_m = g(y_{t-d}).$$

The system’s transfer function  $f$  is a pure displacement perturbed with a Gaussian noise. The control transfer function  $h$  is a mapping of the continuous control to a set of discrete actions. Furthermore, the predictive feedback display  $g$  reflects a prediction based on the last available observation  $y_{t-d}$  ( $d$  being the communication delay) and the full control history.

### 4.3.3 Simulations

I performed series of simulations exploring the properties of this model, in which the transition probability functions were generated with various levels of constant delays and Gaussian noise, and the corresponding empowerment levels were computed<sup>1</sup> using the control model of Equation 4.3 with a set of three discrete actions  $\mathcal{A} = \{left, idle, right\}$ . The results of one simulation (see Figure 4.3) present empowerment as a function of disturbances (delay and noise). Empowerment increases as uncertainty decreases and as uncertainty vanishes approaches its upper-bound of  $\log|\mathcal{A}| = \log 3 = 1.1$  nats. The figure reveals that pure delay does not affect empowerment as much as pure noise, yet delay emphasises noise as the two jointly deteriorate empowerment faster than each factor alone. This reverberates with the fact that humans can adapt to deterministic systems featuring fixed levels of delay, since they do not impose such a high cognitive burden as systems with stochastic noise [54]. An initial exploration of the interplay between constant delays and state noise at a low granularity level allowing for a small scale study, in which the uncertainty is captured by the empowerment measure, is presented in the next section. Unlike frequent expectable delays, to which humans can adapt successfully over time, irregular delays undermine system’s reliability and increase uncertainty [54]. For simplicity, however, this study focuses on constant levels of delay as time-varying delays can be viewed as another stochastic type of noise.

<sup>1</sup>Blahut-Arimoto algorithm [15] is used for computing the channel capacity in Equation 4.2.

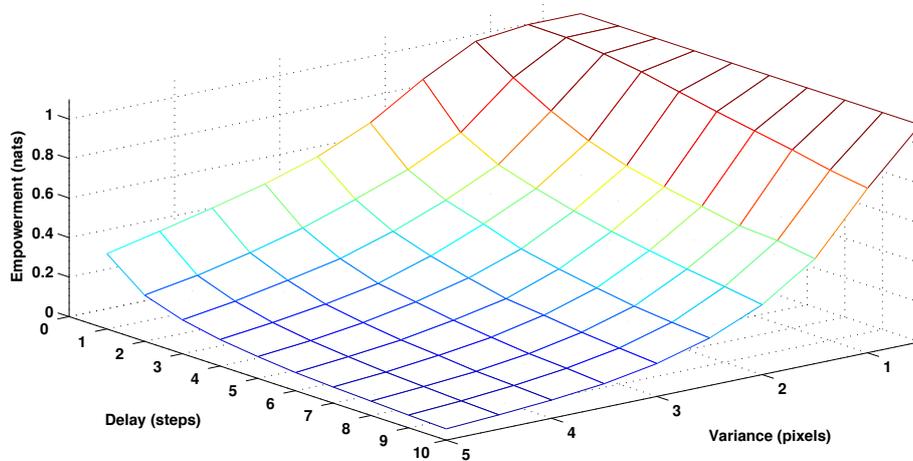


Fig. 4.3 Empowerment as a function of disturbances in the feedback loop – constant delays and Gaussian noise  $\mathcal{N}(0, \text{Variance})$ . Delay is measured by the number of discrete steps in the feedback loop and the zero-mean Gaussian noise is represented by its variance. Empowerment is non-negative and upper-bounded by  $\log 3 = 1.1$  nats. The effect of delay is not as significant as the effect of noise, however, the combination of the two deteriorates empowerment to a higher degree than each factor alone.

## 4.4 Experiment I

The interactive system used in the user study explores the effect uncertain feedback has on human behaviour and performance in an object tracking task with the aim to investigate the relation of empowerment to standard metrics of human performance and to explore its potential as a theoretical measure for supporting UI design.

### 4.4.1 Experimental System

The prototype represents a simple car-simulator game implemented on Nokia N9 mobile phone. The one-dimensional touch input enables the subject to steer a car horizontally with a 3-way joystick on a randomly moving road with added artificial disturbances to the feedback channel and to the control input (see Figure 4.4). The touch-sensitive control area is divided into three segments corresponding to the three discrete actions – moving horizontally the car to the left, to the right, or keeping it static. The goal is to maintain the car in the middle of the road and leaving the road is considered a failure. Whenever the car is off the road a beeping audio cue is played. The control loop is executed at 33Hz, which allows for real-time interaction.

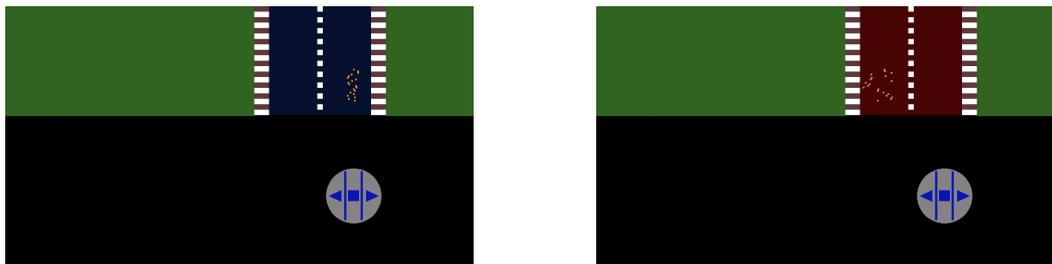


Fig. 4.4 Visual display in the High/Medium (left) and in the Low (right) empowerment condition. The car is represented with a single random point in a narrow (left) and a wider point cloud (right). The grey sphere represents the touch-sensitive control area, which maps the finger position to the three discrete actions – left, idle and right.

Various levels of artificial noise and delay introduced into the system to control the uncertainty of the feedback loop and the corresponding theoretic levels of empowerment are presented in Table 4.1, for each experimental condition. The uncertain visual display represents the car with a single random point hidden in a point cloud, which has different properties across conditions. The motion of the point cloud is affected by artificial Gaussian noise and its size reflects the communication delay. In the case of non-zero delay a predictive display propagates the data points in respect to control inputs, which increases the cloud's size as a function of the accumulated noise.

The visual displays of the three conditions corresponding to High, Medium and Low levels of empowerment (see Table 4.1) are shown in Figure 4.4. The difference between the visual feedback in the High/Medium empowerment conditions (Figure 4.4/left, zero delay) and the Low empowerment condition (Figure 4.4/right, non-zero delay) is in the spread of the cloud, which is wider in the latter case. The motion of the point cloud is smoother and more predictable in the High empowerment condition (zero noise) and more erratic and unpredictable in the Medium/Low empowerment conditions (non-zero noise). I limited the number of conditions to three in order to entrain participants through in reasonable time. Future studies could explore the controlled variables in higher resolution, thus providing more data points on the regression curves.

Table 4.1 Controlled variables and empowerment levels

Condition (empowerment)	Delay (steps)	Noise ( $\sigma$ )	Empowerment (nats)
High	0	0	1.1
Medium	0	5	0.34
Low	10	5	0.04

### 4.4.2 Tasks

For each trial per condition a subject performed a car steering task consisting of following a randomly moving road. In order to achieve this subjects had to use the 1-DOF ‘joystick’ and shift the car horizontally along the road. In the events of failure, i.e. when the car was off the road, an audio cue was played. Each session lasted two minutes and the complete interaction, including the positions of the finger, the target, the car, the cloud particles, and the action taken, was logged.

### 4.4.3 Methodology

Twelve participants (9 male, aged 18 to 48) took part in the study; all but one self-reported right dominant hand. A counterbalanced within-subjects design was used. Trials were always presented in the same order. Conditions were tested in a sitting lab environment. The experiment consisted of a short introduction followed by a training session and the three conditions. The training session, which was exclusively devoted to familiarise subjects with the controls, lasted ten minutes and was divided into four parts: in the first part participants experienced the joystick controls and then practised all three versions of the game. In all three conditions subjects had to perform the same tasks of tracking identical random sequences of a horizontally moving road for two minutes. At the end of each condition they were asked to complete a NASA-TLX [37] and a free-form questionnaire measuring perceived workload and level of engagement respectively. Once all conditions were completed a user experience questionnaire was also filled in. The experiment took about half an hour in total and participants were allowed to rest between conditions.

### 4.4.4 Performance Measures

The following performance measures were used for the analysis of the experimental data:

- Average duration on-target –  $\overline{\sum_{|e_t| \leq \varepsilon} \Delta t}$ ,
- Total time off-target –  $\sum_{|e_t| > \varepsilon} \Delta t$ ,
- Expended finger movement effort –  $\Delta u_t = \frac{|u_t - u_{t-1}|}{\Delta t}$ ,
- Resulting change in control input –  $\Delta U_t = \frac{|U_t - U_{t-1}|}{\Delta t}$ ,
- Total number of errors –  $\sum_{|e_t| > \varepsilon, |e_{t-1}| \leq \varepsilon} 1$ .

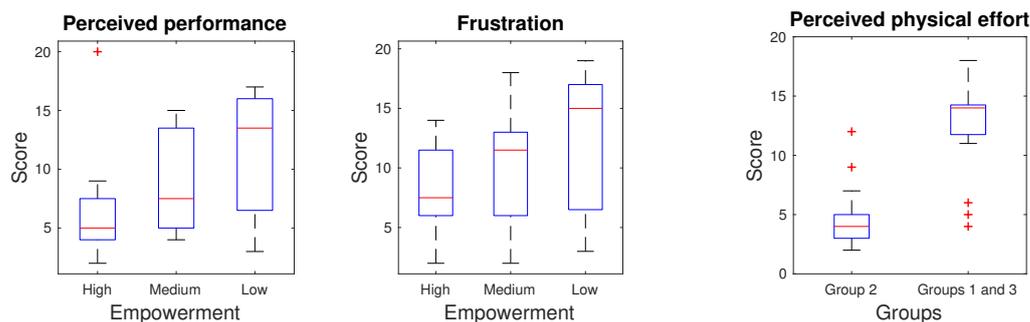


Fig. 4.5 Perceived workload – lower scores correspond to better ratings. Performance and frustration (left) are significantly correlated with empowerment ( $\rho = -0.4, p = .015$  and  $\rho = -0.33, p = .048$ ). Between-groups physical effort (right) – Group 2 is significantly lower than Groups 1 and 3 combined, all conditions combined (Wilcoxon rank sum  $p < .000$ ).

#### 4.4.5 Performance Analysis

A Pearson correlation test revealed significant negative correlation between empowerment levels and NASA-TLX (see Figure 4.5/left) perceived performance ( $\rho = -0.4, p = .015$ ) and frustration ( $\rho = -0.33, p = .048$ ), and perceived uncertainty ( $\rho = -0.55, p = .000$ ) (see Figure 4.6/right). Among objective measures listed above only total time off-target (see Figure 4.6/left) was significantly correlated with empowerment ( $\rho = -0.47, p = .004$ ).

Pair-wise Wilcoxon signed ranks tests show that subjects performed significantly better in High than Medium and Low empowerment conditions – i.e. they had lower total number of errors ( $p < .001$ ), higher average duration on-target ( $p < .001$ ), lower total time off-target (Medium  $p = .024$  with Bonferroni correction and Low  $p < .005$ ). Perceived uncertainty was significantly higher in Low than High and Medium empowerment conditions ( $p < .005$ ) and perceived performance was significantly better in High than Low empowerment condition ( $p = .014$  with Bonferroni correction), which is consistent with measured performance.

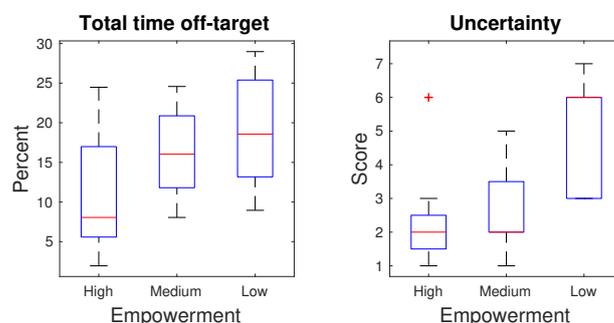


Fig. 4.6 Total time spent off the road (left) and perceived uncertainty about the exact location of the car measured on a 7-point Likert scale (right) are significantly correlated with empowerment levels (Pearson  $\rho = -0.47, p = .004$  and  $\rho = -0.55, p = .000$  respectively).

### 4.4.6 Behaviour Analysis

Subjects exhibit various ways of coping with noise reflected by changes in their activity levels influenced by varying levels of disturbances, which are visible in the rate of change in finger movement effort and the resulting change in control input (see Figure 4.7). Classifying subjects into Group 1 (1,2,3,4), Group 2 (5,6,7,8,9) and Group 3 (10,11,12), based on these trends and using the ranks from Figure 4.7, leads to the following observation. Group 1 react to increased disturbances by decreasing the expended physical effort and actual control by nearly 20%, i.e. by adopting a more relaxed behaviour, while Group 3 react by increasing them by 20% and more, i.e. by putting more effort into dealing with disturbances. This suggests that Group 1 expend more effort the more empowered they are to control the environment. Group 2 expend lower levels of effort on average and have nearly constant levels in all conditions, however the changes in finger movement effort and in resulting control input are not reciprocal. Interestingly, higher levels of finger movement effort result in lower levels of control input. This incongruity can be attributed to loss of empowerment due to the higher level of disturbances, i.e. subjects did not perceive accurately the effect

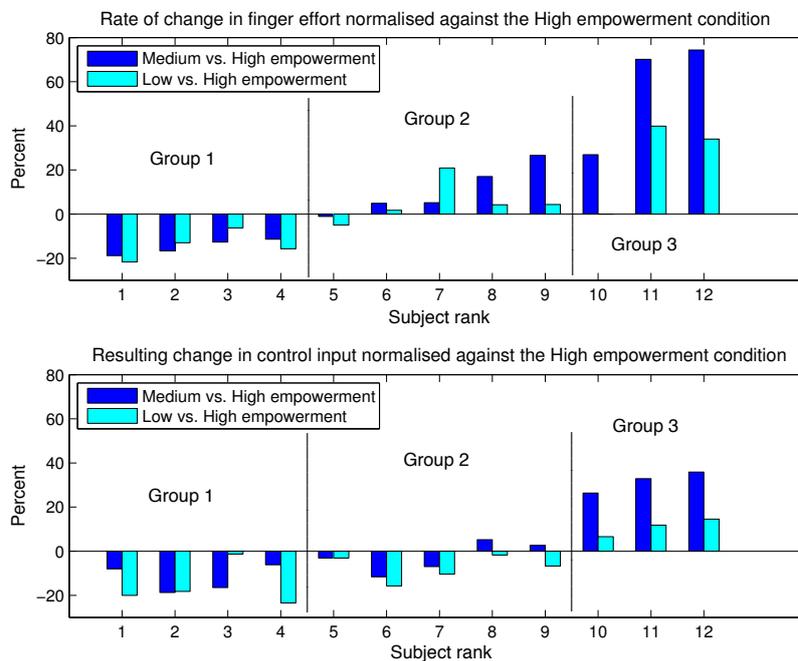


Fig. 4.7 Change in finger movement effort (top) and the resulting change in control input (bottom) – subjects are ranked according to performance in the Medium vs. High empowerment condition in the top and the same ranking is used in the bottom. The Medium and Low empowerment conditions are normalised against the High.

Table 4.2 User preferences

Preference	High	Medium	Low
Most preferred	9	3	0
Least preferred	1	4	7

of their actions in the Medium and Low empowerment conditions. Furthermore, the user preferences in Table 4.2 show that only subjects from Group 2 (all but one) disliked mostly the Medium empowerment condition, which exhibits more erratic visual feedback. Group 2 is also significantly lower in perceived physical effort (see Figure 4.5/right) than Groups 1 and 3 combined, all conditions combined (pair-wise Wilcoxon rank sum test  $p < .000$ ).

#### 4.4.7 Tracking Patterns

Most participants tried to anticipate the switches in the target's direction and put an effort into minimising their reaction time. The task required only basic skills, still the response time varied significantly across subjects – some were faster and others slower to react. Eager anticipation of more active subjects led to cases of oscillation at repeated turning points, while more relaxed ones showed smoother tracking patterns. A more detailed qualitative analysis of the time series of two characteristic subjects reveals patterns of lead and lag behaviour, the first one – subject A – being more active (see Figure 4.8), while the second

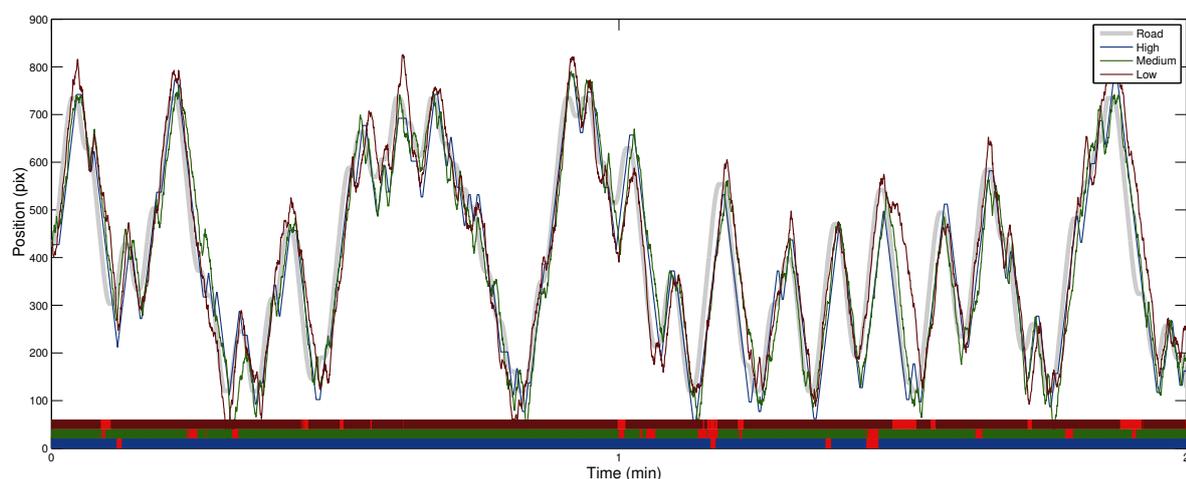


Fig. 4.8 The time series of subject A in the High (blue), Medium (green) and Low (brown) empowerment conditions – an example of high expended effort. Lead patterns of target-anticipating behaviour are visible at the turning points. The colour strips underneath represent the events on/off-target. The red blocks show few periods off-target, more frequent in the Medium and the Low empowerment conditions as well as the second half of the sessions.

one – subject B – more relaxed (see Figure 4.9). The time series of subject A, presented in Figure 4.8, show an example of high expended physical effort in all conditions, leading to close tracking performance. The off-target events, highlighted in red at the bottom, show higher frequency in the Medium and the Low empowerment conditions as well as the second half of the sessions, the latter suggesting a sign of boredom. The time series of subject B, presented in Figure 4.9, show an example of low expended physical effort resulting in a loose tracking performance reflected by the significant number of longer periods off-target highlighted in red at the bottom.

Further analysis reveals specific aspects of the measured and perceived performance. Table 4.3/A show that subjects A and B spent considerably different amounts of time off-target, while for A this indicator is lower than 5-10%, for B it is higher than 20%. This is consistent with the NASA-TLX results (see Table 4.3/B) where lower values reflect better ratings. Table 4.3/C and 4.3/D present further aspects of the error as measured by the total number of leaving-the-road events and by the distribution of off-target periods duration. They reveal details of the interaction dynamics indicating how often subjects went off the road and how quickly they were eager and able to come back on the track. The figures show that subject B left the track significantly more often than subject A and spent considerably longer periods of time off the road. The activity levels of both subjects (see Table 4.3/E) show that subject A applied significantly higher physical effort than subject B measured by the direct control effort derivative of finger motion as detected by the touch-sensitive screen,

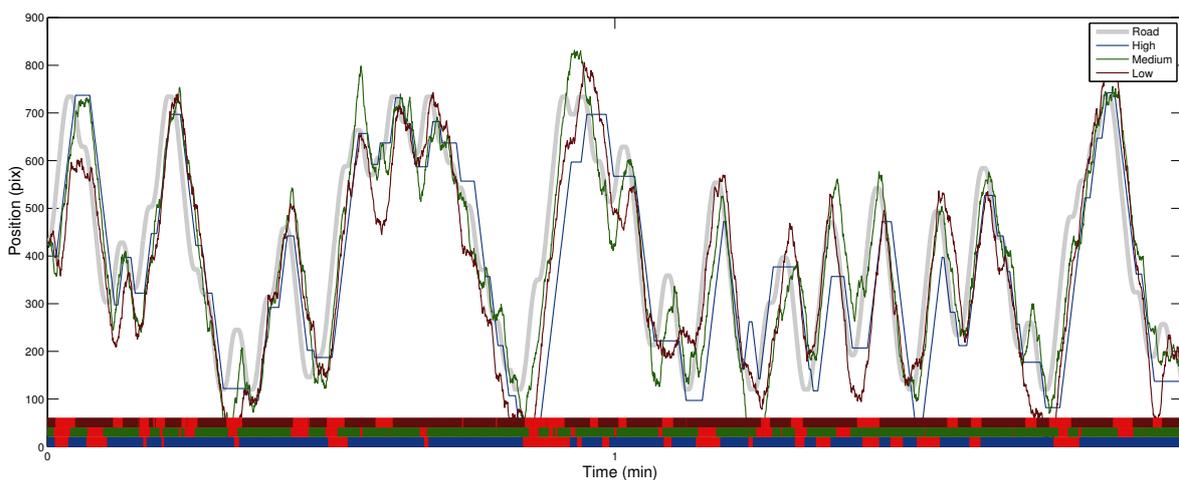
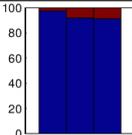
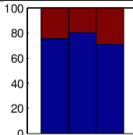
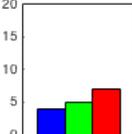
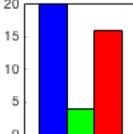
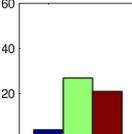
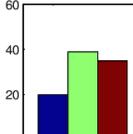
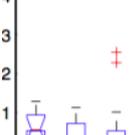
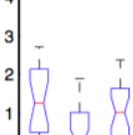
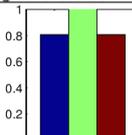
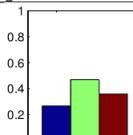
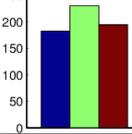
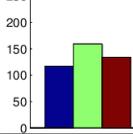
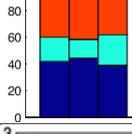
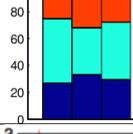
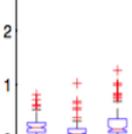
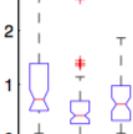


Fig. 4.9 Time series of subject B, an example of low expended effort, in the High (blue), Medium (green) and Low (brown) empowerment conditions. Lag patterns of relaxed target-following are visible at the peaks. The colour strips underneath depict events on/off-target. The red blocks reveal a significant number of periods off-target of considerable length.

which correspond to higher (A) and lower (B) number of action alternations (see Table 4.3/F). The distributions of actions applied by both subjects (see Table 4.3/G) reveal that subject B applied semi-uniformly all actions and most frequently the idle one, while subject A alternated between the left and the right action barely using the idle action. The distributions of idle periods duration (see Table 4.3/H) reflect the more relaxed behaviour of subject B,

Table 4.3 Various performance indicators for the subjects A and B in the High, Medium and Low empowerment conditions (presented in this order).

Code	Measure	Subject A	Subject B
A	Relative time on/off-target (%)		
B	Perceived performance (the lower the better)		
C	Total number of leaving-the-road events		
D	Distribution of off-target periods duration (sec.)		
E	Normalised expended finger movement effort		
F	Total number of action alternations		
G	Action distribution – left/idle/right (%)		
H	Distribution of idle periods duration (sec.)		

who kept idle for significantly longer periods compared to the extremely brief idle patterns of subject A (~200 ms), revealing the more erratic and sharp transitions between left and right actions and approaching the limits of human sensorimotor abilities. Both subjects expended the most effort in the Medium empowerment condition (see Table 4.3/E,F,G), which however resulted in the worst performance in terms of the indicator presented in Table 4.3/C. This highlights specific differences between the Medium and the Low empowerment conditions and suggests that the predictive display of the latter helps keeping the car in the safe zone at a lower cost compared to the direct, but more erratic feedback of the former.

## 4.5 Discussion

The results show that performance dropped significantly in Low and Medium compared to High empowerment condition, as expected, however this was not linked to a consistent trend in control effort. Some subjects increased while others decreased their effort, which reveal alternative forms of human adaptation in the face of uncertainty. The analysis identified a particular group of subjects, who kept the control effort steady across conditions and disliked most the Medium empowerment condition, and furthermore had significantly lower perceived physical effort than the rest. Low and Medium empowerment conditions were significantly different in perceived uncertainty, which is a reflection of the system design, however this did not translate proportionally into performance difference as they were close both in measured and perceived performance. As expected, uncertainty increased frustration, nevertheless participants were able to adapt appropriately to the uncertain environment (i.e. increasingly rely on prior knowledge as suggested in [55]) and to deliver a good performance.

Empowerment correlates significantly with total time off-target, perceived uncertainty, perceived performance and frustration, which demonstrate its overlap with standard measures. However, the large data variability suggests that further work is required to validate statistically the proposed measure and establish the exact regression curves compared to baselines. The loss of empowerment implicates interesting trends in the activity levels indicating another area for future research.

Technology and design-related usability issues affected the prototype system. Most subjects lacked the tactile feedback from the control area covered by their finger, which increased the difficulty in controlling the discrete transitions at this specific resolution level. All subjects found it easy to learn how to use the system, however few used predominantly the two extremes, left and right, and rarely the middle, perhaps due to the narrow range.

This work built on certain simplifying assumptions in the representation of the continuous interactive task with a relatively low-granularity discrete model, which inevitably affects the

model accuracy. There is a trade-off in the design process, which implicates that the more accurate the models the more reliable the measure they imply, however at a higher modelling cost. The purpose of the discretisation applied here was primarily to keep the complexity of the presentation low and thus reduce the threshold for readers not adept with the employed information-theoretic concepts.

## 4.6 Conclusion

This chapter introduced a measure for quantifying uncertainty in dynamic interactive systems built on the principle of empowerment and suggested its potential for making predictions and providing theoretical bounds on standard performance metrics based solely on properties of the environment. The formal extension of the original empowerment framework to incorporate lags in the perception–action loop in a principled way captures uncertainty attributed to various factors common in interactive systems (e.g. noise, delays, errors, etc.) in a generic theoretical measure and expands the capability of empowerment to model human–machine systems enabling its application in the domain of HCI.

As a universal, generic and embodied utility measure empowerment could provide an analytical tool facilitating the theoretical evaluation of system’s usability in various environmental conditions as it only requires the probabilistic model of the system’s feedback loop and could serve as a criterion in user interface design optimisation. This, along with the theoretical coherence of empowerment, suggests its potential as part of a future toolset for understanding interactive systems.

The adoption of entropy-based information-theoretic principles could address more appropriately the variability in human–machine systems and could foster the foundation of a more solid theoretical framework for the study of HCI as well as provide a number of practical benefits. Such principles could help designers treat and evaluate interactive systems in a general fashion and could augment current usability studies improving quality of design, while at the same time reducing risk and evaluation costs. Applying empowerment, however, involves a certain amount of modelling costs, which may be too high for particular systems.

The application of empowerment to other domains of HCI, e.g. to the traditional discrete interaction paradigm, is an exciting topic for future work. Another interesting direction is the exploration of time-varying irregular delays, which are supported by the proposed empowerment model, however their investigation was beyond the scope of this thesis. Furthermore, the application of the recently developed continuous approximation of empowerment could enable the treatment of more realistic continuous models of dynamic interactive systems.

# Chapter 5

## Information Parsimony in Decision Making

*‘As far as the laws of mathematics refer to reality, they are not certain; and as far as they are certain, they do not refer to reality.’*

---

Albert Einstein

### 5.1 Introduction

This chapter presents a characterisation of human information processing with respect to decision making in the domain of dynamic manual control building on the concept of Relevant Information, which is arguably linked directly to cognitive workload [76, 101]. The approach makes certain simplifying assumptions about how we process information and models that based on the external manifestations of the perception–action loop, i.e. the feedback sent to the human sensors and the actions applied in response to that feedback, without diving into the underlying neurophysiological processes running in the human brain. The chapter explores the relationship between the three variants of the Relevant Information concept – RI, C-RI and LA-RI (Section 3.6) – in an attempt to find a model that best fits human data recorded in Experiment I (Section 4.4). Numerical simulations and an in-depth analysis reveal the relation between a human operator and an information-parsimonious optimal controller. A further benchmark, elucidating the capacity of the Relevant Information approach to explain specific types of human behaviour, is provided by Gaussian Process machine learning technique regressing on the same set of empirical data.

## 5.2 Background

Perception, information processing and actuation are often considered as individual processes, however combining them introduces the perception–action feedback cycle, which is often defined as ‘the circular flow of information between an organism and its environment in the course of a sensory guided sequence of actions towards a goal’ [28, 29]. The Markovian Decision Processes (MDPs) provide a framework, which enables finding a value-optimal policy for a given task measured by a reward function, however, does not consider the decision-making cost associated with executing that policy. Thus, an optimal policy for an MDP may be found, which maximises the reward achieved by the agent, but does not heed possible computational costs or limits imposed on the agent. The combination of information theory and MDPs offers the possibility to control the information processing cost required by an agent to achieve a certain level of performance in terms of rewards [76, 101].

The information-theoretic treatment of decision and action sequences introduced by [101] reveals a direct analogy to codes in communication and their complexity – the minimal number of decisions required for reaching a goal bounded by informational measures. In addition, this information-theoretic view is universal, general and can be post-hoc imbued with the specific constraints of particular models. On the informational level scenarios with different computational models can be directly compared with each other. At the same time the informational treatment enables the incorporation of limits on the information processing capacity, which are fundamental properties of agent–environment systems. This allows the consideration of the informational cost for performing a task or whether the task is feasible if constraints in the form of limited informational bandwidth are imposed.

## 5.3 Stochastic Model

Here the car tracking system (see Chapter 4) is modelled through the standard Reinforcement Learning approach where the target is a recurring state in a non-episodic task and the agent receives at each step a ‘reward’ (or a ‘penalty’). The following model of a fully observable MDP tries to preserve the system dynamics while keeping the complexity and the dimensionality low. It is defined over the discrete state space

$$\mathcal{S} = \mathcal{P} \times \mathcal{V},$$

corresponding to two random variables – the car’s position relative to the middle of the track ( $P \in \mathcal{P}$ ) and the car’s velocity relative to the track velocity ( $V \in \mathcal{V}$ ) – varying in the range

of  $[-60, 60]$  and  $[-6, 6]$  respectively. The action space consists of three discrete actions

$$\mathcal{A} = \{left, idle, right\},$$

corresponding to the selected direction of motion. The state transition function is approximated empirically using simulations based on the deterministic action dynamics and the stochastic road trajectory. Since the goal is to keep the car in the middle of the road and leaving the road is considered an error the reward function is designed so as to reach its maximum of one when the car is in the middle of the road, decrease linearly to zero towards both edges of the road, and keep the zero level beyond. The discount factor used in the computation of the three relevant information methods was  $\gamma = 0.95$ , while the ‘inverse temperature’  $\beta$  varied in the range of  $[10^{-1}, 10^6]$ , the higher bound serving as an approximation of infinity as larger  $\beta$  values lead to numerical instability due to exponent divergence.

### 5.3.1 Information vs. Utility Trade-off

Assuming the above model the level of mutual information between states and actions  $I(S, A)$  is upper-bounded by the entropies of both spaces, which in turn are bounded by the logarithm of the cardinality of the respective space, i.e.  $\log |\mathcal{X}|$  and  $\log |\mathcal{A}| = \log 3 = 1.58$  bits. The accumulated utility level given the maximum reward of one unit and  $\gamma = 0.95$  discount factor is upper-bounded by 20 units, which is an upper limit for the achievable utility given the specific model. These bounds provide theoretical constraints on the trade-off curves and characterise the distance between a particular solution and the theoretical optimum.

Figure 5.1 presents the information vs. utility trade-off curves computed with the three variants of the relevant information method using Algorithms 2 and 3. It shows that an increase in relevant information leads to an increase in expected utility, i.e. investing more in information bandwidth allows for higher performance. Note that when  $\beta$  vanishes this attempts to save on relevant information while being indifferent to the achievement of a high utility. As opposed to that letting  $\beta \rightarrow \infty$  aims for a policy that is optimal in its utility, however unlike the conventional Reinforcement Learning method of value iteration the relevant information algorithm will select a soft-max policy among the optimal policies, more precisely the one that minimises the relevant information.

Figure 5.1 also reveals that the standard RI method yields lower levels of utility due to the assumption of a static uniform state distribution, which averages uniformly the utility accumulated in every state. This corresponds to the assumption that the agent has no prior knowledge about the state it is most likely to observe. In contrast, the C-RI and LA-RI methods constantly update the underlying state distribution, which is consistent with the

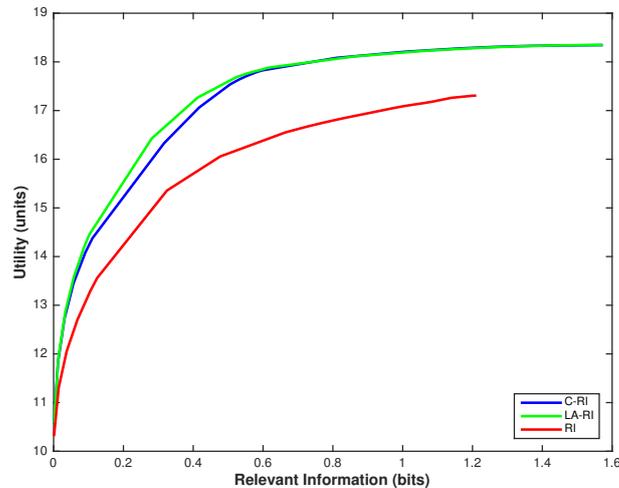


Fig. 5.1 The theoretical information vs. utility trade-off curves generated by the look-ahead (LA-RI), the consistent (C-RI) and the standard (RI) relevant information methods in green, blue and red respectively. RI has lower utility due to the uniform state distribution. LA-RI achieves higher utility than C-RI for a range of finite  $\beta$  levels.

actual policy at all times and therefore achieve a higher overall utility level. In this case the C-RI and LA-RI methods generate the same trade-off curves asymptotically as  $\beta \rightarrow 0$  and  $\beta \rightarrow \infty$  with more pronounced deviations particularly in the range of  $\beta \in [0.5, 100]$  corresponding to relevant information  $I(S,A) \in [0.1, 0.5]$  bits, which are explored in more detail below.

This view adopts the hypothesis that humans implement an information-parsimony principle [56, 75] implying that they achieve a given level of performance at the lowest informational cost possible. This corresponds to operating close to the optimal reward/information trade-off curve provided that a suitable reward function is available [13, 98]. This treatment uses Shannon information with the specific interpretation as an information processing cost – the minimal informational cost of the decision process that the agent has to undergo in order to achieve a particular reward.

### 5.3.2 LA-RI vs. C-RI

The average distances between the optimal policies provided by the LA-RI and C-RI algorithms measured in Jensen-Shannon divergence are presented in Figure 5.2. The figure reveals that for very small ( $\beta < 0.1$ ) and large ( $\beta > 1000$ ) values of  $\beta$  both methods provide identical optimal solutions. For values in-between a policy difference appears, which peaks for  $\beta = 2$ . Diving into more detail Figure 5.3 presents the Jensen-Shannon divergence

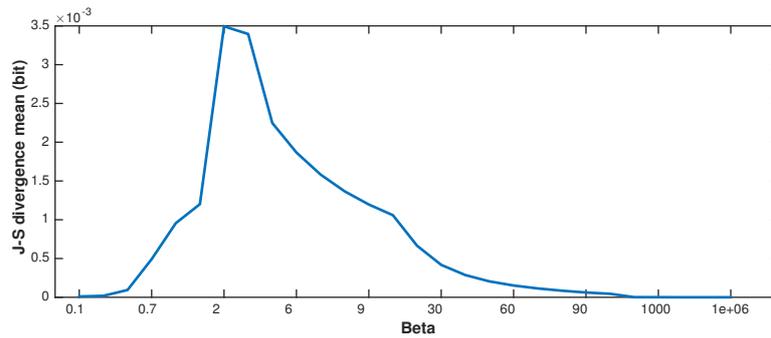


Fig. 5.2 The Jensen-Shannon divergence mean between the optimal LA-RI and C-RI policies computed for various levels of  $\beta$ . The policies are most distinct for  $\beta = 2$  and coincide at the two extremes.

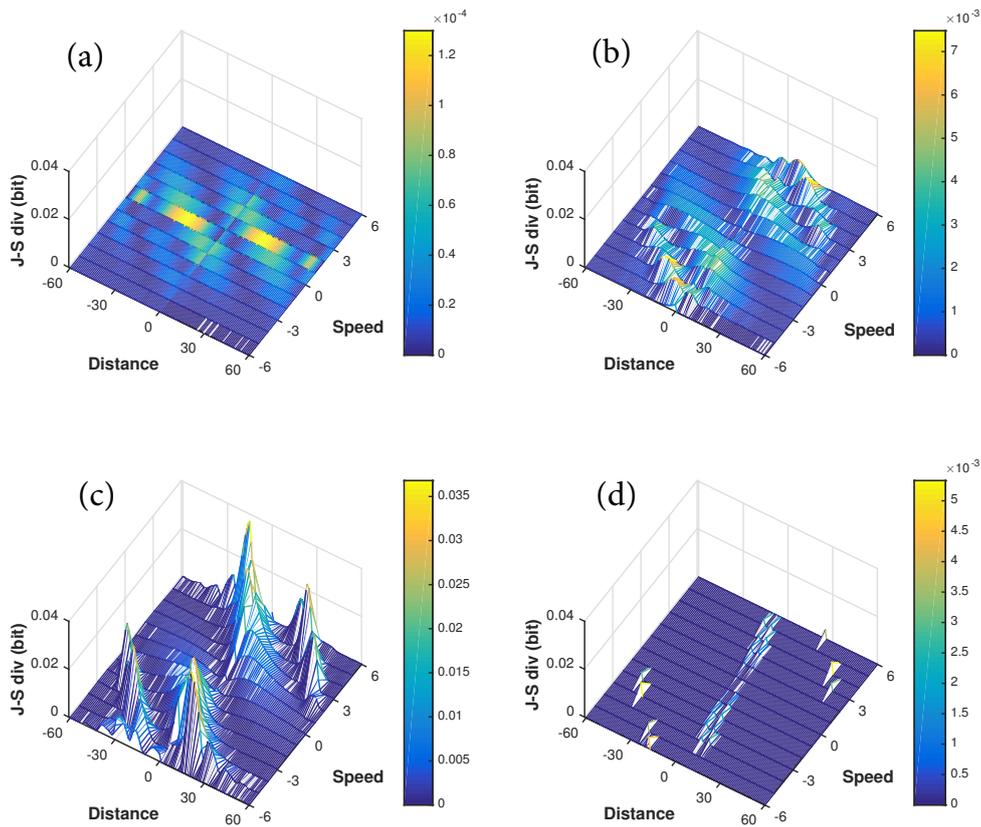


Fig. 5.3 The Jensen-Shannon divergence between the optimal LA-RI and C-RI policies of a moving action derived for four specific levels of  $\beta$ : .3, 1, 3 and 100, revealing more pronounced variabilities for higher relative velocity or higher error.

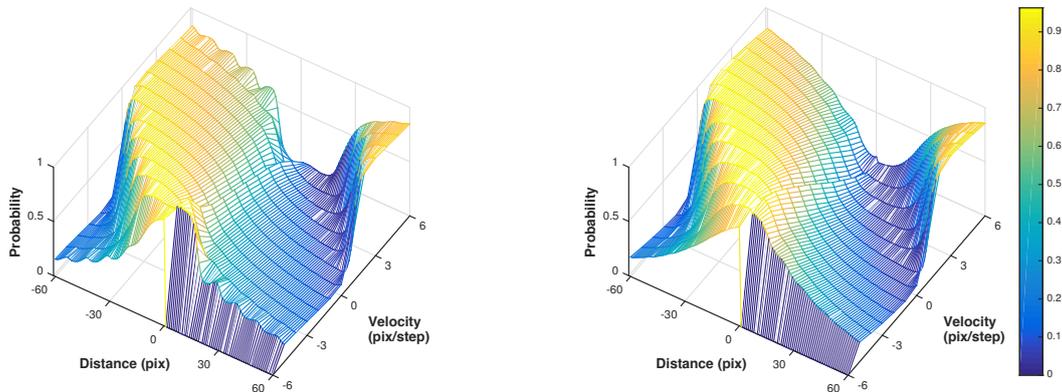


Fig. 5.4 The optimal C-RI (left) and LA-RI (right) policies of a moving action derived for  $\beta = 2$ . The probabilistic profile of the opposite action is symmetrical with an identical shape. For higher velocities LA-RI provides smooth policy transitions unlike C-RI.

profiles between the optimal C-RI and LA-RI policies for a moving action and four specific levels of  $\beta$ . It reveals that the differences are more pronounced in areas with higher relative velocity or higher error. The complete policy profiles corresponding to Figure 5.3/c ( $\beta = 2$ ) are presented in Figure 5.4, in which the landscape reflects the probability map of taking the action in a specific state. The figure shows that for higher relative velocities the larger predictive horizon of the LA-RI method ensures smooth monotonic policy transitions unlike the single-step C-RI approach. These translate into better solutions found by the LA-RI method, which achieve higher utility levels than the C-RI (see Figure 5.1).

### 5.3.3 Empirical Trade-Off

In order to explore the behaviour of an artificial agent performing the car tracking task in settings close to the user study I conducted a series of simulations using the optimal policies of the three RI methods derived for a various levels of  $\beta$ . All simulations were initialised with the fixed  $(0,0)$  starting state as in the user study, which renders certain states in the space unreachable in the course of interaction due to the discrete action set  $\mathcal{A} = \{-5, 0, 5\}$ .

From the simulated data I computed the empirical utility levels and estimated the mutual information between states and actions using a discrete calculator. The resulting empirical trade-off curves, presented in Figure 5.5, reveal certain deviations from the theoretically computed curves. For example, the maximal utility level does not exceed 15 units, which is lower than the theoretical limit presented in Figure 5.1. The reason for these deviations is that the computation of the theoretical limits considers every possible initial state while all simulations started from the same state.

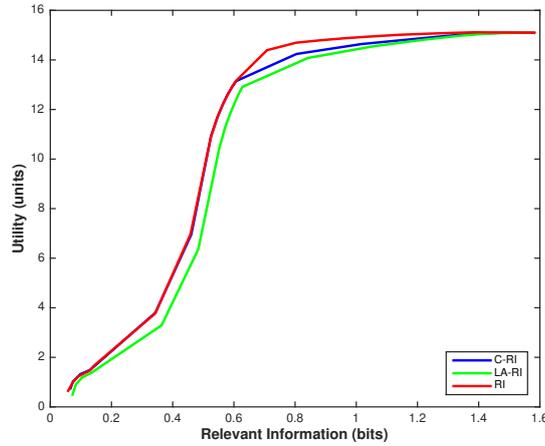


Fig. 5.5 The empirical trade-off curves computed on simulated data obtained from executing the optimal RI, C-RI and LA-RI policies corresponding to various levels of  $\beta$ . The deviations from the theoretical curves are due to the fixed initial starting state used in all simulations, which restricts the number of visited states.

The absolute error histogram of an optimal  $P_{LA-RI}^*$  policy simulation (see Figure 5.6) reveals that due to the fixed initial starting point and the discrete action step used in the simulations the smallest possible errors account for 5%, 10% and 15% of the interaction duration respectively. For most of the simulation run (bar 5%) the optimal agent is not able to nullify the error because of getting trapped by the policy one or two cells away from the target, since jumping +5/-5 cells would further distance it from the centre of the target.

This explains the deviation from the theoretically achievable utility levels where the different starting points increase the probability of getting closer to the target. The empirical mean error – in red in this figure – provides a characterisation for the task difficulty from an optimal controller point of view.

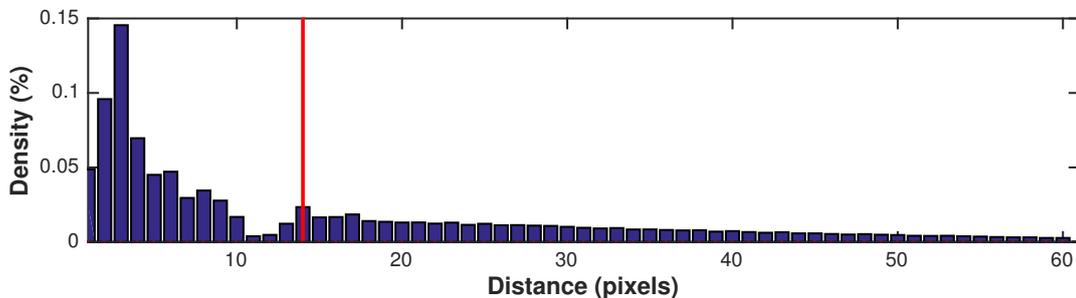


Fig. 5.6 The absolute distance from target histogram provided by a simulation using the optimal  $P_{LA-RI}^*$  policy derived for  $\beta = 10^6$ . The red line depicts the mean distance. The agent gets frequently trapped one or two cells away from the target due to the specific initial state.

Another obvious difference between the theoretical and the empirical trade-off curves is the amount of information used by the standard RI method, which in the simulation is computed over the empirical (hence consistent) state distribution by default and approaches the limit of 1.58 bits unlike its theoretical limit of 1.2 bits computed over a uniform state distribution. For the same reason of consistent state distributions the three methods achieve empirically identical asymptotic levels of utility unlike theoretically (see Figure 5.2).

### 5.3.4 Soft vs. Sharp Policies

The simulated trajectories of 500 000 trials using the optimal  $P_{LA-RI}^*$  policies corresponding to six levels of  $\beta$  are aggregated in Figure 5.7, in which  $\beta$  increases from left to right and from the top down. For small  $\beta$  the optimal policy is close to a random walk at little informational cost and low utility levels since it does not track the target and the trajectories look like a diffusion process. As  $\beta$  and therefore the available information capacity increases sharper, more refined and accurate policies emerge, which enable the gradual convergence to the target trajectory. The curves for intermediate levels of  $\beta$  split in certain locations where the agent loses the target with a relatively high probability. For higher levels of  $\beta$  the optimal controller closely tracks the target.

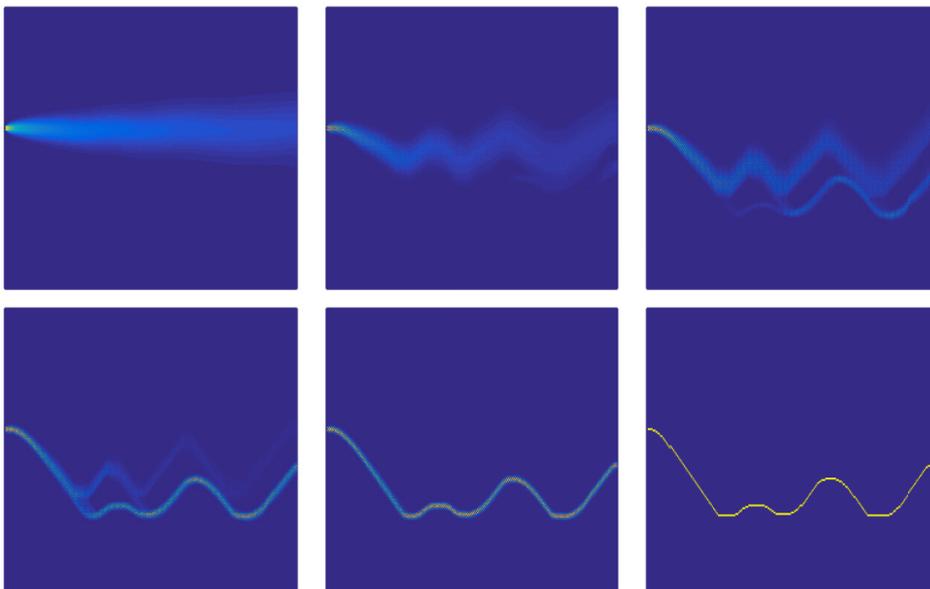


Fig. 5.7 The tracking trajectories aggregated from 500 000 simulations using the optimal  $P_{LA-RI}^*$  policies for six levels of  $\beta$ : .1, .9, 2, 3, 5 and  $10^6$ . As  $\beta$  increases from left to right and from the top down the initial random walk transforms into loose tracking occasionally losing the target in certain locations and converging to close tracking for the highest  $\beta$ .

## 5.4 Human Performance

This section presents further aspects of human performance inferred from data recorded in the ‘High’ empowerment condition of Experiment I, which complements the analysis presented in Chapter 4 with new insights. In order to relate human performance to the theoretical limits presented earlier in this chapter I computed the empirical utility vs. relevant information trade-off per subject and characterised the results with the corresponding levels achieved by an optimal  $P_{LA-RI}^*$  artificial agent. Prior to determining the best estimate of mutual information from the time series I investigated the time-shift characteristics between perception and action using two different methods.

### 5.4.1 Lag Inference

A standard cross correlation analysis performed on the visual feedback and the actions in the recorded experimental data (see Figure 5.8/top) reveals high levels of correlation for all subjects with a relatively low variance. The corresponding lag distribution (see Figure 5.8/middle) reflects a delay of 300-500 ms characteristic for the human sensorimotor loop with few exceptions.

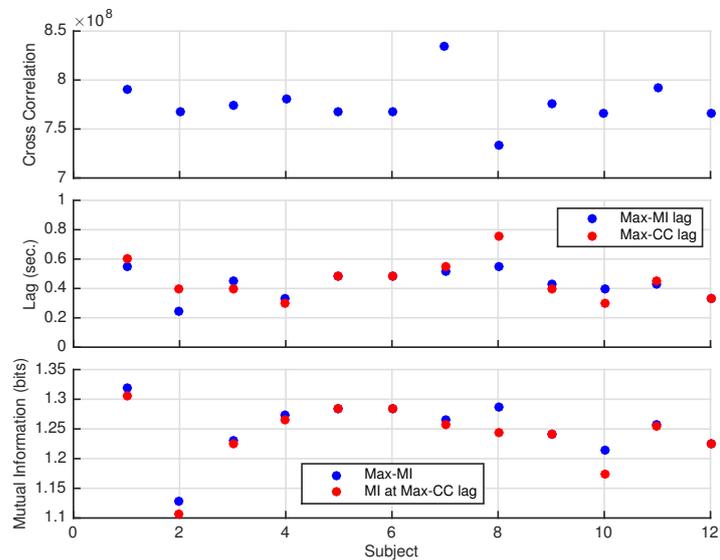


Fig. 5.8 Cross Correlation (top) and Mutual Information (bottom) of perception and action in the condition ‘High’ of Experiment I. High levels of cross correlation with low variability (top). The lags between perception and action maximising cross correlation (red) and mutual information (blue) reflect the human sensorimotor delay of 300-500 ms (middle). The levels of mutual information between perception and action maximised over lags (blue) and computed for maximal cross correlation lags (red) are upper-bounded by 1.58 bits (bottom).

Further insight into the correlation between perception and action is provided by their mutual information computed with a discrete estimator. Maximising the mutual information by time-shifting the data series over reasonable levels of delay resulted in similar lag profiles as shown in Figure 5.8/middle with few exceptions for subjects  $S_2$ ,  $S_8$  and  $S_{10}$ .

This provides evidence for considering the mutual information maximised over the lag as an accurate estimate to be used in the utility-information trade-off chart. The maximised levels of mutual information and their counterparts computed with time-shifting the data by the maximal cross correlation lags (see Figure 5.8/bottom) reveal that the corresponding levels of mutual information are very close to each other, as expected, with the above exceptions for subjects  $S_2$ ,  $S_8$  and  $S_{10}$ .

### 5.4.2 Empirical Trade-off

Applying the earlier specified reward function and  $\gamma = 0.95$  discount factor used in the theoretical computation I calculated the utility levels per subject from the experimental data. The corresponding trade-offs between these utility levels and the two levels of mutual information – the maximal and the one computed with the maximal cross correlation lags – per subject show interesting trends (see Figure 5.9).

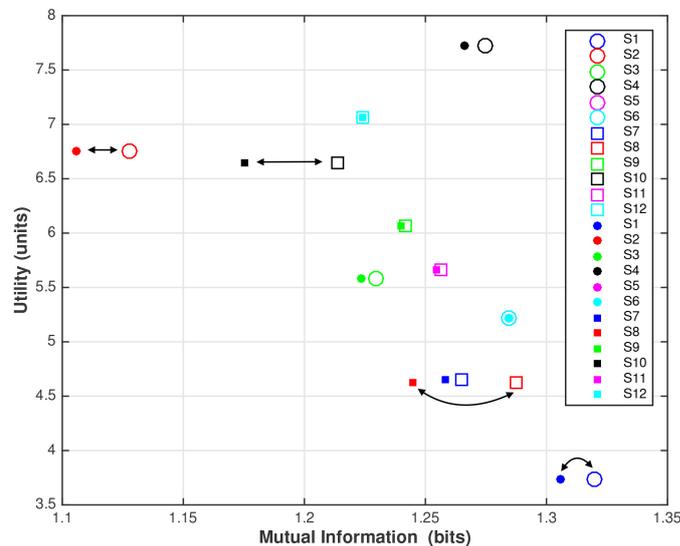


Fig. 5.9 Mutual information vs. utility levels per subject derived from empirical data recorded in the ‘High’ empowerment condition of Experiment I using: (empty) maximal mutual information, and (filled) mutual information computed with maximal cross correlation lags. Large variability in utility levels is achieved with similar levels of mutual information, suggesting the structural differences between the applied action strategies.

Subject  $S_2$  achieved a relatively high utility level by exploiting the two moving actions for ~80% of the time (see Figure 5.10), which pushed the mutual information level down to almost 1.1 bits, i.e. very close to the theoretical limit corresponding to using two actions only  $\log 2 = 1$  bit. Subject  $S_{10}$  achieved the same level of utility by applying a similar distribution of actions (see Figure 5.10), however at the slightly higher informational cost of ~1.2 bits due to a more consistent action selection.

This example demonstrates the potential of the applied information-theoretic indicator to reveal the degree of structure in human behaviour and discriminate two otherwise identical – in terms of the applied action distribution and the resulting utility level – patterns of human performance. The highest utility level was achieved by subject  $S_4$ , who demonstrated a more relaxed behaviour than subjects  $S_2$  and  $S_{10}$  and applied the three actions more uniformly, however also more consistently, which incurred a slightly higher informational cost of 1.27 bits. On the other extreme of the chart are subjects  $S_1$  and  $S_8$ , who stayed idle more than the others and achieved the lowest utility levels, however at the highest informational cost of ~1.3 bits – the level of information at which the empirical trade-off curve saturates at the highest utility level (see Figure 5.5). This suggests that subjects  $S_1$  and  $S_8$  underperformed at this specific informational level, i.e. they wasted cognitive resources in achieving the respective (low) level of utility.

This novel information-theoretic aspect characterises the underlying structure of the experimental data and links that structure (or the lack of it) to purely performance-oriented indicators. It could provide an index for the consistency of a particular strategy facilitating the more insightful description of human behaviour and complementing standard performance metrics.

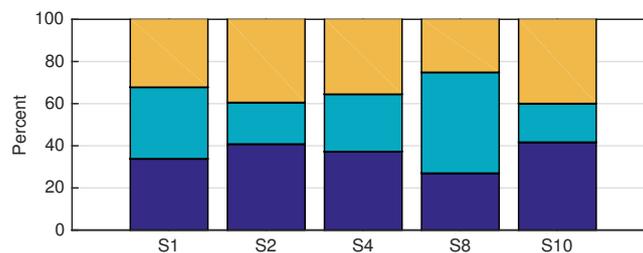


Fig. 5.10 Action distributions in blue/green/red corresponding to the left/idle/right actions respectively derived from data recorded in the ‘High’ empowerment condition of Experiment I for five characteristic subjects, which reveal that  $S_1$  and  $S_8$  stayed idle twice longer overall than  $S_2$  and  $S_{10}$ .

## 5.5 Human vs. an Info-Parsimonious Optimal controller

In order to find a fit for the model of the human operator I explored the behaviour of an optimal controller in series of simulations using settings close to the user study. In all simulations an agent executed the optimal  $P_{LA-RI}^*$  policy ( $\beta = 10^6$ ) starting from the  $(0,0)$  initial state and tracking the same target trajectory as in the user study. The settings corresponded to the ‘High’ empowerment condition, with no disturbances in the feedback loop in the form of noise or delays. However, in order to emulate the inherent human sensorimotor lag various constant levels of artificial delay were added to the actuation channel.

From the simulation data I computed the empirical trade-off curves for several levels of actuation delay in the range of  $[0,750]$  ms, presented in Figure 5.11. The figure relates the empirical utility vs. mutual information levels provided by the user study with the simulated trade-off curves revealing that the human levels are located around the extensions of the curves corresponding to delay levels of 300-450 ms, known as characteristic for human reaction time. Subjects  $S_1$  and  $S_8$ , whose response time is in the range of 600 ms (see Figure 5.8/middle) achieve lower utility, slightly exceeding the level of an optimal controller performing at 600 ms constant lag, however at a higher informational cost of 1.3 bits as opposed to 1.1 bits. Subject  $S_2$ , whose response time is at the other extreme ( $\sim 200$  ms) achieves a higher utility than the curve corresponding to 300 ms delay. This is the

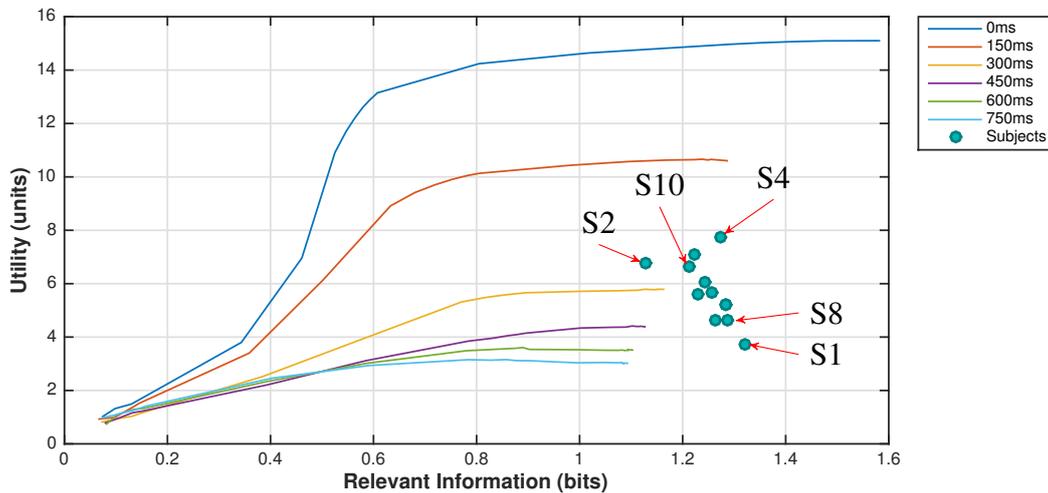


Fig. 5.11 The empirical trade-off curves generated from simulation data of  $P_{LA-RI}^*$  optimal controller for various levels of  $\beta$  and actuation delays in the range of  $[0,750]$  ms. The utility vs. mutual information levels of all participants in the ‘High’ empowerment condition of Experiment I are depicted with green dots. The human levels are in the region of an optimal agent operating with 300-450 ms response time with a certain excess of informational cost.

only subject performing at an informational level closely matching an optimal controller. Subject  $S_{10}$ , who has a moderate response time of 300-400 ms, exceeds the utility of an agent performing at 300 ms delay, however at the price of a higher informational cost of 1.2 bits. It needs to be emphasised that the agent policy is optimised with respect to no delays in the perception–action loop and thus can be suboptimal and even inferior to human performance when executed with delays. Subject  $S_4$  provides an example of human performance characteristic of 300 ms lag achieving a higher utility, however at a higher informational cost compared to an agent operating at the same lag level.

The empirical trade-off curves generated with constant levels of actuation delays are not necessarily optimal for the proposed MDP model and are presented only for providing an insight in the search for a good fit for the model of the human operator. More accurate models of optimal controllers operating with delays in the perception–action loop require a different framework, which however exceeds the scope of the presented information-theoretic treatment of MDPs.

## 5.6 Model Fitting

This section presents the application of two approaches for fitting models to specific types of tracking performance observed in Experiment I. The first approach is based on the LA-RI method, while the second one is using Gaussian Process regression machine learning technique. The key characteristics of both approaches are demonstrated through an empirical evaluation of the fitted models. The following analysis was performed on the experimental data of two particular subjects, each representing a distinct type of human behaviour.

The first one (subject  $S_8$ ) is an example of very relaxed tracking performance with large varying response time in the range of 600 ms (see Figure 5.12). The smooth edges around the

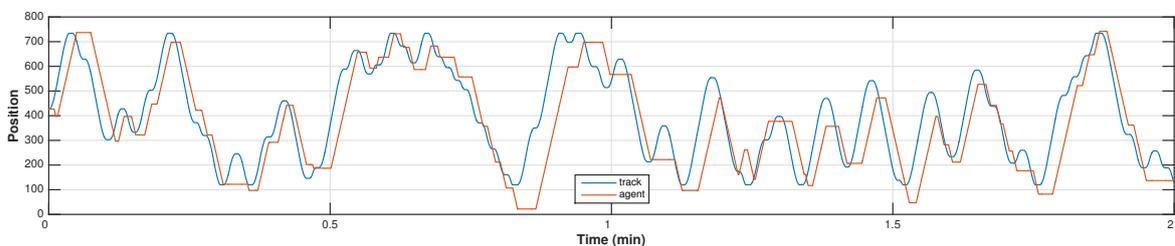


Fig. 5.12 Time series of subject  $S_8$ , an example of very relaxed tracking performance characterised by a large varying response time (~600 ms) and extended idle periods yielding a low utility level.

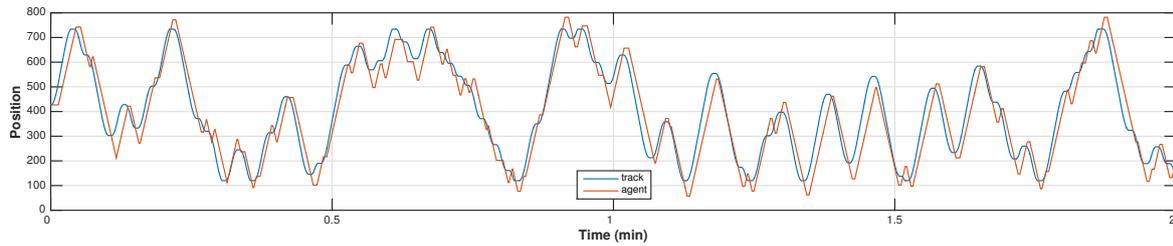


Fig. 5.13 Time series of subject  $S_{10}$ , an example of close tracking performance characterised by a high level of activity around the turning points, an average response time (300-400 ms) and short idle periods yielding a high utility level.

turning points with rare overshoot suggest a substantial degree of anticipation characteristic for higher order models, which differ from the properties of the purely reactive controller introduced in this chapter. As expected, no sufficiently good fit was found for this pattern of human behaviour among the controllers described in the previous section.

The second one (subject  $S_{10}$ ) is an example of close tracking performance, characterised by short response time (300-400 ms), high level of activity around the turning points and systematic overshoot (see Figure 5.13). Qualitatively, this tracking pattern appears similar to the reactive behaviour generated by an optimal LA-RI controller with the obvious lack of spring dynamics, which is not included in the model for the sake of simplicity.

### 5.6.1 Look-Ahead Relevant Information

The simulated trajectories of two types of parameterisation of an optimal LA-RI controller, the first one generated for five specific levels of  $\beta$  and a constant actuation delay of 300 ms, while the second one for the maximal  $\beta = 10^6$  and various constant delay levels, are presented in Figures 5.14 and 5.15 respectively.

Figure 5.14/top reveals how the agent struggles to track the target for low levels of  $\beta$  at this level of actuation delay, whereas a higher level of  $\beta$  provides a more accurate policy enabling the compensation for the delay (middle plot), and demonstrates a close tracking performance for even higher  $\beta$  (bottom plot).

The trajectory presented in Figure 5.15/top reveals a very close tracking performance corresponding to no actuation delay in the feedback loop, while increasing the delay introduces an obvious lag, resulting in overshoot and oscillations around the turning points. The tracking pattern corresponding to the maximal  $\beta = 10^6$  and an actuation delay of 300 ms (see Figures 5.14/bottom and 5.15/middle) appears qualitatively as the closest match for the experimental trajectory of subject  $S_{10}$ .

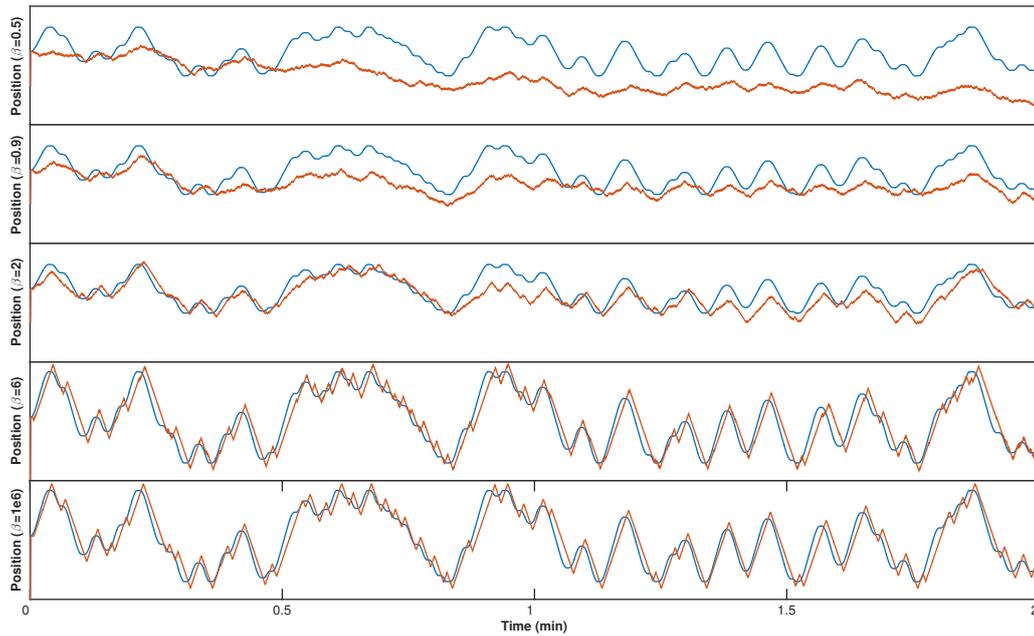


Fig. 5.14 Tracking trajectories of a LA-RI optimal controller for five levels of  $\beta$ : .5, .9, 2, 6 and  $10^6$  (from the top down) and actuation delay of 300 ms. Higher levels of  $\beta$  result in sharper policies enabling closer tracking.

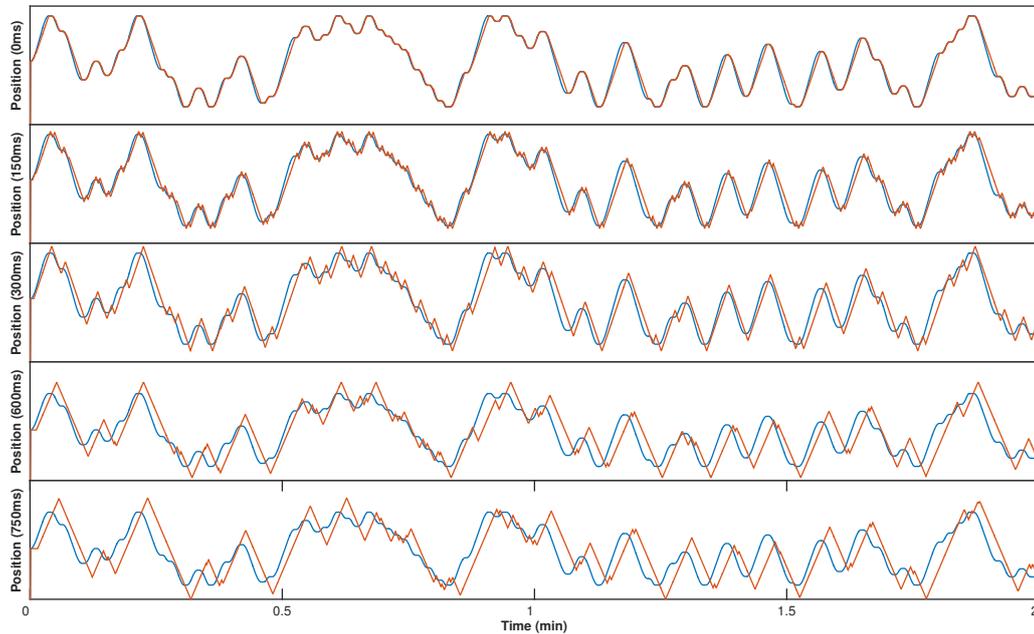


Fig. 5.15 Tracking trajectories of a LA-RI optimal controller ( $\beta = 10^6$ ) with actuation delay of various magnitudes: 0, 150, 300, 600 and 750 ms (from the top down). Increasing delays introduce lags, overshoot and oscillations around turning points.

## 5.6.2 Gaussian Process Regression

To further elucidate the properties of the information-theoretic approach I contrasted the above model to models derived with Gaussian Process (GP) regression machine learning technique [61, 79]. The GP models regressed the next car position based on the last ten car and track positions. The recorded experimental data was split into training and validation data sets per subject. Since the exact implementations of Gaussian Processes scale with  $\mathcal{O}(n^2)$  and  $\mathcal{O}(n^3)$  in computational and memory complexity respectively, this constrains their practical application. To address these issues a number of computationally efficient approximations of GPs have been proposed, which reduce memory requirements to  $\mathcal{O}(nm)$  and the computational complexity to  $\mathcal{O}(nm^2)$  ( $m \ll n$ ). The results presented in this section are based on the Sparse Spectrum Gaussian Process (SSGP) approximation method, which was proposed in [57].

The precision of the SSGP model fit measured by the normalised mean square error and the mean negative log-probability is presented in Figure 5.16 for a wide range of hyperparameters (basis functions) for subjects  $S_8$  and  $S_{10}$ . To minimise the probability of getting stuck in a local minima, which is typical for such non-convex optimisation methods, I initialised each trial 100 times with a different random seed. These indicators reveal a better fit for subject  $S_8$ , for which SSGP minimises the error for as few as six basis functions, while for subject  $S_{10}$  it requires 20 basis functions. The regression means (see Figure 5.17) provided by SSGP models derived with 100 basis functions follow closely the validation curves of both subjects, with a minor overshoot around the turning points.

Figures 5.18 and 5.19 demonstrate the predictive power of the SSGP models of subjects  $S_{10}$  and  $S_8$  respectively, derived with 100 basis functions, in simulations with previously

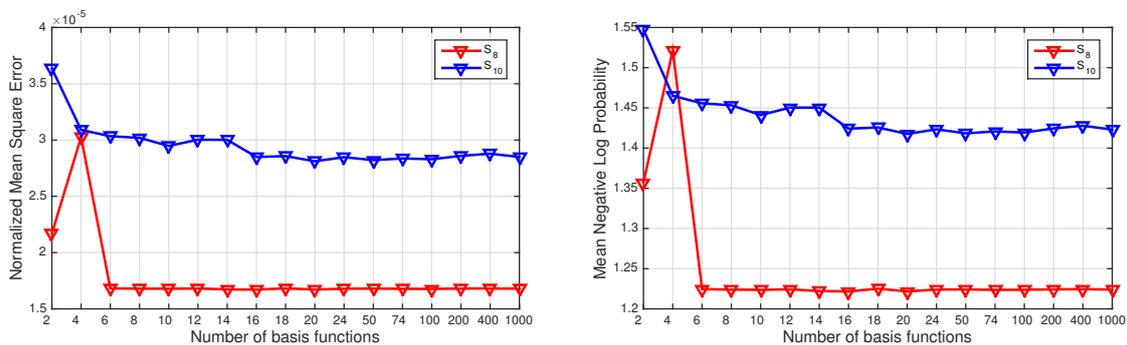


Fig. 5.16 Normalised mean square error (left) and mean negative log-probability (right) of SSGP models trained on experimental data of subjects  $S_8$  (red) and  $S_{10}$  (blue) for various numbers of basis functions. The model accuracy is significantly higher for  $S_8$  than for  $S_{10}$ , and furthermore saturates the indicators for six vs. 20 basis functions respectively.

unseen randomly generated track trajectory. The model of subject  $S_{10}$  performs with a short lag, smoother transitions and less overshoot around the turning points than observed in the corresponding experimental data (see Figure 5.13), while the model of subject  $S_8$  exhibits a larger lag with smoother transitions than characteristic for the trajectory in Figure 5.12.

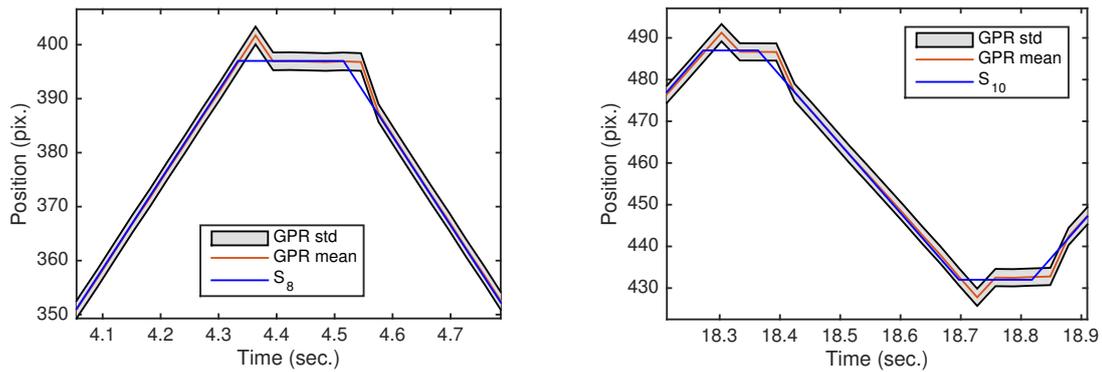


Fig. 5.17 Regression mean (red) and standard deviation (grey) of SSGP models for subjects  $S_8$  (left) and  $S_{10}$  (right) derived with 100 basis functions, revealing a close fit with a minor overshoot around the turning points.

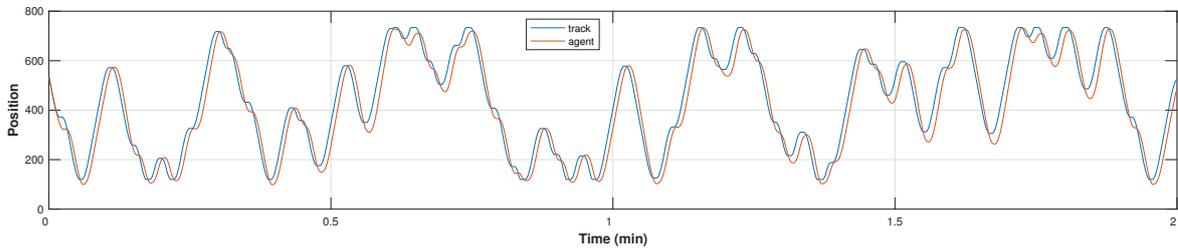


Fig. 5.18 Tracking trajectory of the SSGP model of subject  $S_{10}$ , derived with 100 basis functions, generated in simulation with previously unseen track in predictive mode, revealing short lags, smooth transitions and minor overshoot around the turning points.

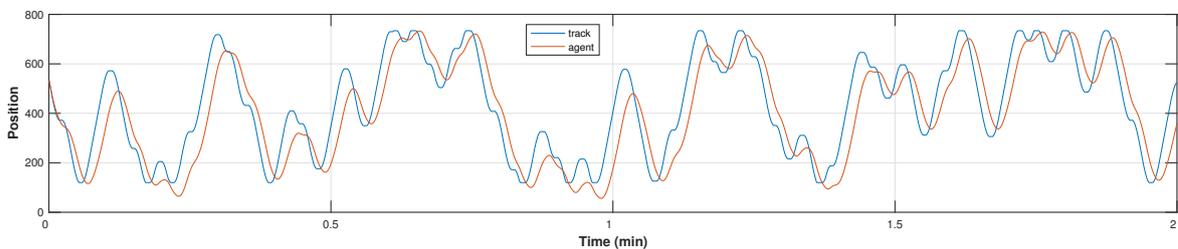


Fig. 5.19 Tracking trajectory of the SSGP model of subject  $S_8$ , derived with 100 basis functions, generated in simulation with previously unseen track in predictive mode, characterised by larger lags and smooth transitions.

### 5.6.3 Performance Characteristics

The following performance analysis of experimental and simulated data provides important insights into the properties of the two modelling methods. The GP models were trained on experimental data of subjects  $S_{10}$  and  $S_8$  with 100 basis functions, whereas the generative LA-RI model is subject-independent. All performance indicators provided by GP and LA-RI simulations use an identical previously unseen track (see Figures 5.18 and 5.19), which is comparable in characteristics, however is of different shape than the experimental track.

Figure 5.20 reveals that subject  $S_{10}$  and a LA-RI controller operating with 300 ms actuation delay have similar performance characteristics measured by the frequency of action alternations, movement periods duration, number of leaving-the-road events, relative time on/off target and on/off-target periods duration. The performance deviations in action distribution and idle periods duration are due to the unconstrained speed in alternating the two moving actions, which benefits the LA-RI controller and results in a nearly binary action policy and extremely short idle periods. This speed is however limited in the physical prototype by the specific design of the control device and the sensorimotor constraints of the human operator. The results demonstrate that this particular type of human performance is approximated well by the LA-RI controller. In contrast, the  $GP_{10}$  model delivers a higher – continuously on-target – performance with a lower effort, characterised by less frequent action alternations and longer idle periods, providing a human-like action distribution.

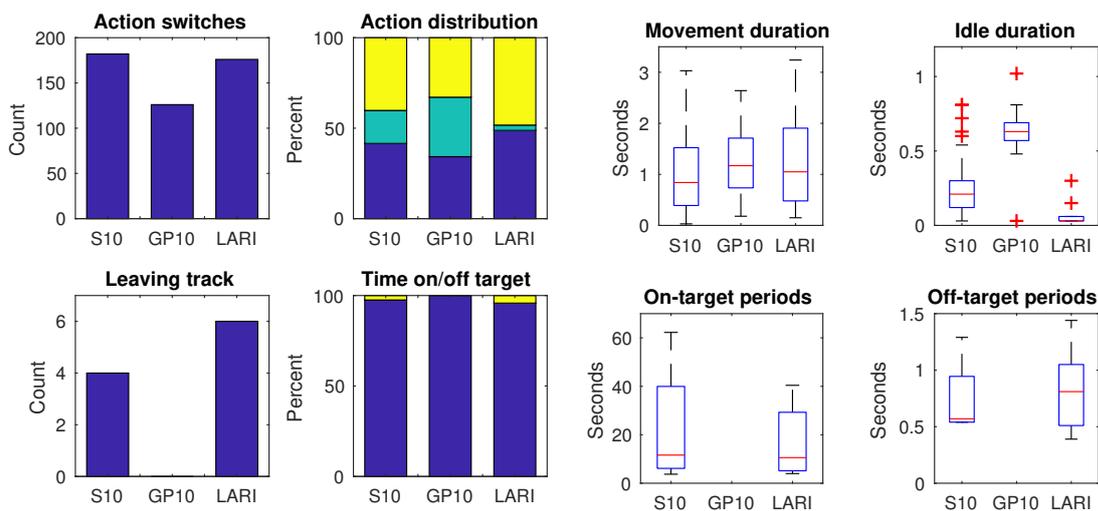


Fig. 5.20 Performance indicators for subject  $S_{10}$ , the corresponding Gaussian Process regression model  $GP_{10}$  and a LA-RI information-parsimonious model operating with 300 ms actuation delay.  $S_{10}$  and LA-RI show similar levels for all indicators, but action distribution and idle periods duration, while  $GP_{10}$  achieves the best performance with the least effort.

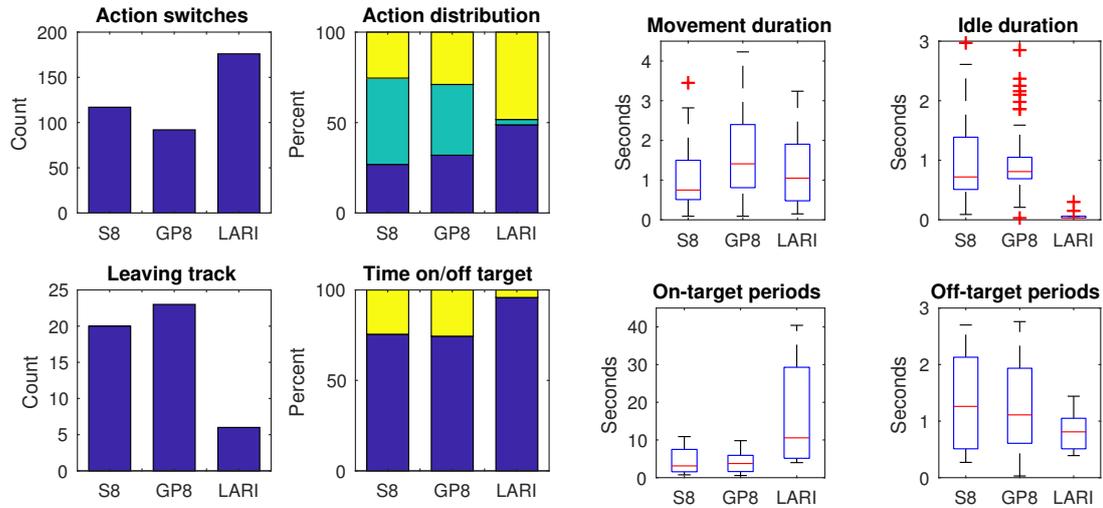


Fig. 5.21 Performance indicators for subject  $S_8$ , the corresponding Gaussian Process regression model  $GP_8$  and a LA-RI information-parsimonious model operating with 300 ms actuation delay.  $S_8$  and  $GP_8$  achieve comparable levels for all indicators, which are significantly different than the LA-RI levels, except for movement periods duration.

Figure 5.21 suggests that  $GP_8$  is a good fit for subject  $S_8$  across all indicators, while LA-RI deviates substantially in all but one measure (movement duration). The results reveal fundamental differences between the two approaches, emphasising the relation of the LA-RI model to purely reactive behaviour and the good GP approximation of smoother performance.

## 5.7 Discussion

This analysis highlighted a novel aspect of human behaviour reflecting the structure of decision making, which is not explicitly captured by standard performance metrics. It provides a new perspective for characterising performance and could augment current user evaluations. Albeit confined by certain simplifying assumptions the information-parsimonious optimal controller provided a particular fit to human performance when executed with 300-450 ms actuation delay, the range typically observed in the user study. The lack of physical constraints in the LA-RI model (e.g. spring dynamics) allowed the controller to alternate instantaneously between left and right, and ignore the idle action when it was not necessary. The particular design of the control device and the limits of human performance, however, hindered such a binary interaction style in the user study as subjects had to move inevitably their fingers across the idle zone and activate the idle action when switching direction.

The information vs. utility trade-off curves generated by an optimal controller operating with 300-450 ms actuation delay reach maximal utility levels at a cost close to 1 bit of relevant

information. Intuitively, this key implication can be interpreted as the binary knowledge about the position of the car relative to the target (left/right), which is sufficient to make an optimal decision and requires  $\log 2 = 1$  bit on average. One subject operating at the limits of human performance with a response time of 200 ms approached the informational level of an optimal controller by achieving 1.1 bits on average, i.e. an excess of just 0.1 bit of informational cost on decision making. Action policies of other subjects reveal further structures and higher informational levels of up to 1.3 bits, which is typically due to more frequent use of the idle action. The upper bound of this informational property given by the cardinality of the action space  $\log 3 = 1.58$  bits reflects the exact knowledge of the uniform partitions of the state space corresponding to the three available actions. An optimal controller operating with no delay achieves maximal utility level at a cost close to the theoretical limit of 1.58 bits, i.e. by trading off more information for a significantly higher performance.

Despite its limitations the information-parsimonious model of tracking behaviour provided a sufficiently close fit (in informational cost and several performance indicators) for one characteristic pattern observed in the user study. A direct comparison with GP models revealed key performance differences primarily due to the purely reactive nature of the information-theoretic model, and further elucidated its properties.

## 5.8 Conclusion

This chapter presented an application of a novel information-theoretic framework to a reward-driven decision process in the perception–action cycle of an agent. It proposed a particular MDP representation of a specific manual control problem encoded with a hypothetical reward function, which is maximised by the agent. The information-theoretic view reframed this problem into a trade-off between the reward achieved and the informational cost of performing a task by incorporating limits on the information processing capacity, which are fundamental properties of agent–environment systems. The results related standard performance metrics to the capacity of relevant information and demonstrated its potential for characterising and modelling human behaviour in a case study. This opens up the possibility for evaluating the complexity of human–machine systems and the consistency of human performance from a new perspective, and could provide a new analytic tool expanding the existing set of standard measures applied in HCI research. The information-theoretic treatment is universal, general and could enable the direct comparison of scenarios with different computational models. However, further work is required to provide more accurate models of human behaviour and to validate the approach in the domain of HCI.

## **Part III**

# **Social Interaction**



# Chapter 6

## Nonverbal Negotiated Interaction

*'Diplomacy is the art of letting someone else have your way.'*

---

American Proverb

### 6.1 Introduction

This chapter presents a novel method for negotiated interaction, inspired by the minimalist interaction paradigm described in [4] and explores its feasibility in collaborative studies investigating how people interact in dyads using restricted means for communication in various environmental conditions. The experimental systems enable participants to communicate and collaborate at a distance with the emission and reception of tactile signals. Two collaborative experiments provide empirical data for (1) exploring the effect limited modality feedback has on human behaviour, performance and experience in eyes-free environments as captured by the empowerment model developed in Chapter 7 and (2) investigating the collaborative strategies and roles emerging over time and examining in detail how people coordinate their actions as characterised by the information-theoretic measure of interpersonal coordination developed and validated in Chapter 8. Research on human communication typically uses pre-established languages, however, it would greatly benefit from elucidating how such systems emerge and develop in the context of joint human activities in the absence of pre-established communication systems [30]. Therefore, in the first study participants were allowed to agree on a strategy prior to the experiment and in the second they were not, which forced them to negotiate with the limited nonverbal means for communication available during the experiment, elucidating the emergence of collaborative strategies.

## 6.2 Background

Social interaction is an act of sharing meaning [25] in a complex two-way communication. In a control system view, where the controlled object is another person, time delays can be of the order of seconds to minutes or even days. Furthermore, there is a complex, noisy and nonlinear relationship between our actions and the consequences, which makes the prediction of another person's response hard. There is noise in the perception of actions at both ends of the loop and a stochasticity in the response due to fluctuations in internal states. The human factor further increases the degrees of freedom as different people may react differently to the same input [116]. To deal with such complexity humans use their capacity to infer the mental states of others, understand their beliefs and read their intentions and goals, an ability Tomasello described as putting oneself in someone else's 'cognitive shoes' [105].

In collaborative interaction agents engage in a coordinated activity, in which they change and/or influence the behaviour of one another by affecting each other's psychological states. Initiating a coordinated activity and starting a conversation requires sharing an experience and establishing joint attention, which has the following four prerequisites: attention detection, attention manipulation, social coordination and intentional understanding. In joint attentional scenes humans jointly attend to the same aspect of the environment and to one another's attention to it [46]. In order to share experiences with one another humans are able to acquire symbolic conventions allowing to conceptualise things in ways that would not be possible otherwise: 'thinking for speaking'. Substitutes for linguistic symbols, such as manual sign languages, are equally effective in directing attention and cognition when based on shared conventional symbols. Cooperative behaviour requiring specific cognitive structures facilitates the development of higher social skills like language [16, 39, 91, 92, 105], which in turn is a key element in the emergence of cooperation [22, 106]. However, the question addressing how shared interaction routines necessary for coordinating behaviour develop and what dynamics lead to the formation of turns during the interaction is largely unexplored. The feeling of sharing a common space with another intentional being can emerge by switching between two kinds of perception: perceiving the other as part of the environment versus perceiving the activity of the other perceiving me [4]. Joint activities necessitate the coordination of both content and process mediated by common ground, i.e. a 'state of mutual understanding among conversational partners about the topic at hand' [21]. They further require people to have shared awareness [22], providing a context for one's own activity [26].

A sense of shared space and physical presence in a virtual world, i.e. the subjective experience of being in one place or environment even when physically situated in another [115], is facilitated by low latency and high degree of interactivity [84] provided by Collaborative Virtual Environments (CVE) [11]. The appropriate synchronisation at fine-grained time scales

is critical in order to establish a natural flow and rhythm to the computer-mediated social interaction that is characteristic of real-world interaction and resembles a fluid dance [16]. Research on cooperative physical tasks in CVE suggests that haptics may facilitate collaboration and improve the sense of presence, performance and engagement [69, 82, 85, 95, 96], and furthermore may provide interpersonal link, i.e. sense of togetherness, between participants [9, 44]. However, existing CVE typically offer audio and video means for remote communication, which constrains the richness of the communication experience and limits their mobile usage. Most systems are highly visual and non-visual interaction methods are rare, although they could bring benefits when the visual attention is compromised, e.g. in a mobile context [18, 58, 72, 80, 111]. Various devices have been invented to compensate deficits in one sensory modality with another, however such modality substitutions are challenging tasks, since the underlying sensorimotor contingencies in different sensory domains are subject to different properties, i.e. what differentiates vision from audition or touch is the structure of the rules governing the sensory changes produced by various motor actions [71].

Existing social networking applications enable users to communicate via regularly updated context information enriching the dynamics within a social group. However, they are typically restricted to asynchronous static status updates emphasising the indirectness and remote nature of the interaction, whereas a fluid synchronous communication style would create the sense of immediateness and engagement and bridge the physical distance gap. One issue that routinely emerges from this kind of applications is privacy. How do users choose to display or hide their current status or availability in a flexible way that does not hinder this more embodied style of interaction? Dynamic models can act as a mediating mechanism, which enables users to maintain their privacy if desired and can create the possibility for a continuous transition in one's status from completely transparent to completely opaque.

### **6.3 Interactive Concept**

This section presents a concept of a novel bidirectional negotiated interaction style, in which users probe objects and friends in a continuous dynamic manner. It provides the benefits of haptic and (non-verbal) audio communication modes in creating a higher sense of presence and engagement in realistic collaborative settings. The concept builds on a metaphor for remote human communication, which enables a form of touch at a distance and allows for a more embodied interaction style digitally bridging the gap with others in social networks. A mechanical system metaphor driven by signals from capacitive sensors could provide new means for users to interact with and probe others in their social networks. The aim was to explore the feasibility of the new interaction modes and the limits of human performance.

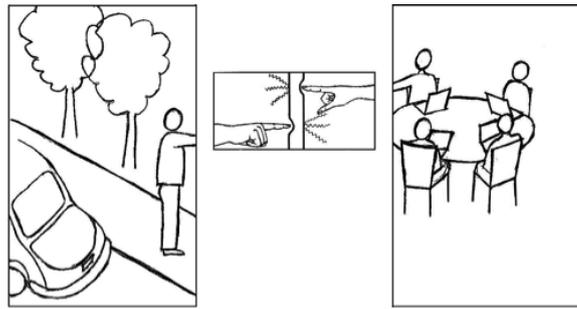


Fig. 6.1 An abstract concept of a membrane as a medium to convey the sense of touch between remote parties. The environmental conditions prohibit the use of other interaction methods than more subtle ones, e.g. touch and tactile. (Courtesy of Saija Lemmelä).

### 6.3.1 Distal Collaborative Scenario

Let us consider the following scenario. *"Andy is in a meeting room with other colleagues while his friend John is walking along a busy street (see Figure 6.1). They certainly cannot have a phone conversation at the moment, but would like to agree on a specific time for a call. Since they are not aware of each other's schedule they would have to negotiate. Ringing each other up intrusively is not an option; texting extensively is not too convenient either. Instead they could work this out in a more dynamic and fluid manner by probing each other on their shared membrane and agreeing on the time. In this situation one important concern is privacy. They would not like the other one to have full visibility of their own schedule, which makes the task more difficult than if they would have shared calendars (see Figure 6.2). Therefore, they will have to negotiate a time slot suitable for both without revealing too much private information."*

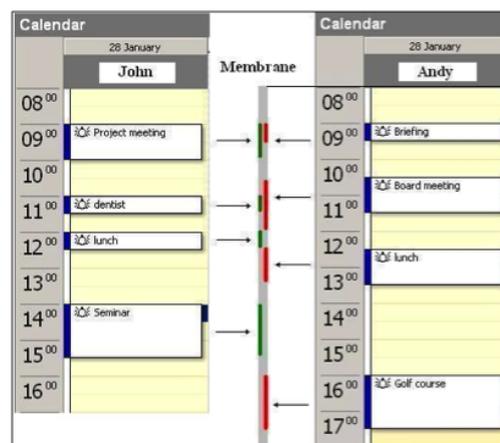


Fig. 6.2 The concept for diary alignment represented as a shared membrane. The diary entries are mapped on both sides of the virtual membrane, allowing for privacy considerations.

### 6.3.2 Membrane Metaphor

One possible interaction mechanism takes the simulation of a membrane dynamic system driven by capacitive touch to facilitate and enrich interaction in scenarios as described above. The dynamics of such systems are similar to that of winding a clock, twisting a door knob or turning a key, all everyday tangible physical metaphors, for which people possess a natural intuition for the effects their actions have on the system [100]. The membrane metaphor enables people to touch each other remotely and engage in a continuous dynamic interaction by sliding their fingers and pushing gently on both sides. The visual, audio and tactile feedback provides rendering of the internal states of the simulated dynamic system. Perceiving the changing physical characteristics of this system could convey much richer information about the current state of a person via continuous interaction and rich feedback than a static event-based technique would. Such a haptic ‘membrane’ interface between users opens up the opportunity for a less intrusive touch-based rich negotiation. The membrane metaphor provides a proxy for the status of the person with whom we are attempting to create a dialog. Building on this metaphor I developed two interactive prototypes and conducted user studies with participants interacting in pairs in different modality feedback conditions.

## 6.4 Experiment II

### 6.4.1 Experimental Design

The collaborative system prototype builds on the membrane concept for remote communication and provides a form of human touch at a distance (see Figure 6.3) utilising capacitive sensing and vibro-tactile devices in three feedback conditions: 1) Visual, 2) Tactile and



Fig. 6.3 Conveying a remote sense of touch using haptic devices instrumented with capacitive sensing and vibro-tactile feedback (left) and a SHAKE SK7 device (right).

3) Combined (visual&tactile). The membrane is designed with a certain number of holes on both sides. The interaction concept consists of two subjects exploring simultaneously the membrane from their respective side, while trying to find a hole through which they could touch each other. The system illustrates an example of how shared environments can be created with low-latency capacitive sensing and multimodal feedback. The interaction design aimed to present the same information consistently in both modalities, reflecting cross-modal interaction principles [40]. This results in redundant presentation of visual and tactile feedback, however facilitates modality comparison.

The feedback displays were designed so as to allow subjects to sense each other whenever their fingers meet on the shared membrane, and to sense the holes in their respective side of the membrane whenever they locate one. The membrane is displayed in a section in Figure 6.4 as a vertical grey strip measuring  $400 \times 80$  pixels. The visual cue of a finger is a bell-shaped marker and the tactile one is a fast and sharp vibration pattern delivered by a linear actuator. The visual cue of a hole is a black square marker and the tactile one is a slow pulsing vibration delivered by a pager motor. The size of the markers defines the granularity of the multimodal feedback and is fine-tuned to enhance user experience. Too small markers would be difficult to locate and hold on to, while large ones would hinder the smooth interaction. The right balance of 50 pixels for bell-shapes and 40 pixels for squares (vertical size) was found after several trials.

The system consists of capacitive sensing device for finger touch input, a dual vibrotactile engine and a separate visual display. Using the input device subjects can probe the membrane up and down vertically, and search for holes and for the remote partner, which can be sensed only when probing in their vicinity (see Figure 6.4), beyond which they are

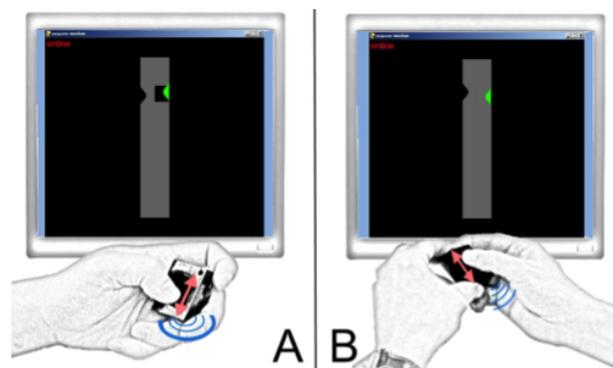


Fig. 6.4 Subject A (in green) has found a hole represented by a black square and a slow pulsing vibration. He can see and feel his partner (in black) on the other side of the membrane (left). Subject B (in green) feels and sees only his partner (in black), since there is no hole on his side of the membrane in this location (right). (Courtesy of Saija Lemmelä).

masked. Only one's own green marker is visible at all times. Subjects obtain information by observing impact events (via visual and/or haptic feedback), i.e. when their fingers collide with objects in the shared environment. The task requires active exploration of the membrane and locating a hole which is shared by both sides. A target hole is acquired when both partners locate it simultaneously and hold on still for 0.5 sec, which minimises incidental acquisitions. Both sides have three holes each, of which one is shared. The holes can be perceived only from the respective side of the membrane.

Assuming a uniform distribution and considering the impact range provides 0.35 as an estimate for the probability of randomly placing a finger on the right hole, and 0.375 for the respective probability of randomly placing the two fingers in touch. Assuming stochastic independence between the two events yields 0.13 for the probability of accidentally placing both fingers on the right hole, which can serve as a representative static descriptor for the complexity of this dynamic task. More accurate estimate could be provided from simulations using artificial agents built on models of human behaviour, which could give a better insight into the underlying stochastic process.

### 6.4.2 Implementation

The prototype system was constructed using two IBM ThinkPad X31 laptops, in conjunction with two SHAKE SK7 (see Figure 6.3/right) sensor packs (SAMH Engineering Services). The Bluetooth-enabled sensor pack provides capacitive sensing from 12 small square pads in a  $4 \times 3$  configuration, measuring  $26 \times 20$  mm and providing 8-bit resolution at 100Hz. It provides a dual vibro-tactile feedback display – one pager-style vibration motor and one pulsed resonant actuator (ALPS ForceReactor™ S-type). The two vibration devices provide a wider range of fidelity than is provided by either device alone – the rotary motor provides good low frequency actuation and the ForceReactor provides excellent high frequency (click/tick) actuation. The former is used for sensing the holes and the latter for sensing the fingers. The ForceReactor is run at the lower frequency of 22Hz in order to avoid the sound artefacts it generates in the higher frequency range around 150Hz. The system is implemented in Python and links the laptops over WiFi, which in turn are connected to the respective SK7 devices over Bluetooth. Data is sampled and transmitted at a frequency of 100Hz, which enables real-time interaction despite the lags in the wireless networks.

### 6.4.3 Tasks

For each trial per condition, a pair of participants performed a collaborative target acquisition task, consisting of finding a shared hole in the membrane. In order to achieve this, subjects

had to individually explore their side of the membrane and suggest a hole to their partner by maintaining the cursor at its location and waiting until their partner found the place. At this point a ‘success’ vibro-tactile cue was played if the hole was shared. For each trial both subjects had three holes on their side of the membrane, of which only one was shared. The positions of the holes were randomised and each session lasted five minutes. The number of trials per session was not limited, and a new one started automatically five seconds after the previous was completed successfully. The complete interaction, consisting of finger position, contact events, target position, localisation and acquisition, was logged.

#### 6.4.4 Methodology

Twenty-six people (18 males and eight females, average age 32 years) were allocated into 13 pairs; 21 self-reported right dominant hand. Four pairs knew each other well (friends, couple), two pairs were colleagues, and seven pairs did not know each other well previously (see Table 6.1). All participants had a prior experience and eleven reported a daily use of touch screen devices. A counterbalanced within-subjects design was used with an extra pair. The trials were always presented in the same order. The conditions were tested in a sitting lab environment with the pairs separated into two different rooms. The participants had no other way to communicate except via the shared virtual environment. The performance depended on the cooperation between the partners, which in turn required some sort of communication. Therefore, in order to succeed in their task, the participants were enforced to converge onto working collaborative strategies using the available means for interaction.

Table 6.1 Pair allocation

Pair	Acquaintance	Gender	Exposure to touch screen
P1	friends	male	daily
P2	none	female/male	occasional/daily
P3	friends	female/male	occasional
P4	none	male	daily users
P5	colleagues	male	occasional/daily
P6	couple	female/male	daily
P7	none	male	occasional/daily
P8	none	male	occasional
P9	colleagues	female/male	occasional/daily
P10	none	male	occasional
P11	friends	female	occasional/daily
P12	none	female	daily
P13	none	male	daily

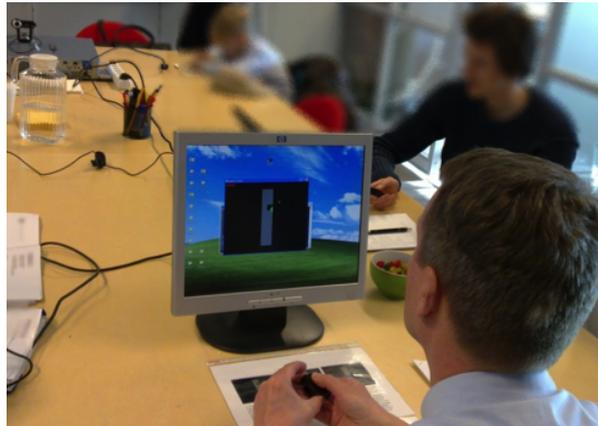


Fig. 6.5 Study participants familiarising with the system while sitting in the same room. During the training session verbal communication was allowed and discussing potential collaborative non-verbal communication strategies was encouraged.

The experiment consisted of a short introduction, followed by a training session and the three conditions. In the training session, which was devoted to familiarise subjects with the interaction method, both participants were sitting in the same room and interacting via the Combined version of the system, which supported both visual and tactile feedback (see Figure 6.5). The training session lasted ten minutes and participants were encouraged to discuss potential collaborative interaction strategies. They were informed about the three versions of the system and that after the practice, they would be separated in different rooms, not being able to communicate verbally. In all three conditions the pairs performed the same tasks for five minutes. At the end of each condition they were asked to complete an extended NASA Task Load Index (NASA TLX) questionnaire [17, 37] measuring perceived workload on a 20-point scale, and answer a specific set of questions. They were also asked to associate each interaction method with three to five positive and negative words and in the analysis these words were clustered according to their meaning. Once all conditions were completed the subjects filled in a user experience and an open-ended questionnaire surveying the user preferences and the applied collaborative strategies. The experiment took about one hour in total and the participants were allowed to rest between different conditions.

### 6.4.5 Performance Analysis

Figure 6.6 presents the numbers of targets acquired by each pair per condition (left) and reveals the learning effect between the sessions utilising visual feedback ranked by their execution order (right). A non-parametric Friedman test shows a significant effect ( $p < 4.8e - 5$ ) of the type of feedback on the number of targets acquired. Pair-wise Wilcoxon signed ranks

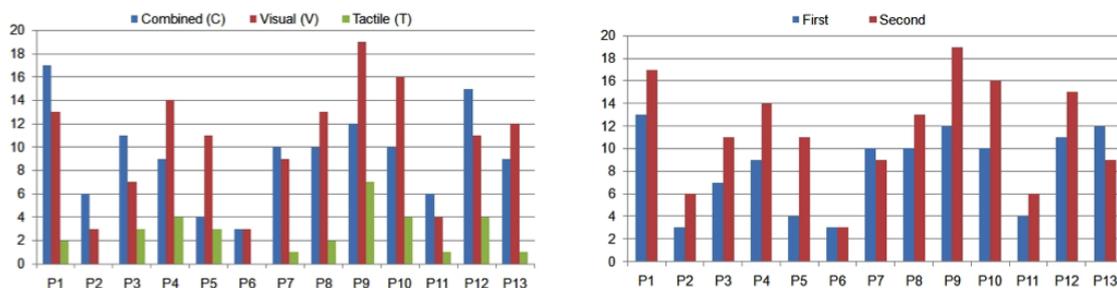


Fig. 6.6 The total numbers of targets acquired by each pair per condition, revealing a large gap between the Tactile and the other two conditions (left). The total numbers of targets acquired in Combined and Visual condition ranked by their respective execution order, revealing a significant (Wilcoxon signed ranks test  $p < 0.008$ ) learning effect (right).

tests show a significant difference between Tactile and Combined/Visual conditions (both  $p < 0.0048$  with Bonferroni correction). The grand totals per condition (see Table 6.2) show a large gap between visual and eyes-free conditions. The trial order of Visual and Combined conditions reveals a significant learning effect (Wilcoxon signed ranks test  $p < 0.008$ ) – i.e. most pairs scored much less in the first than in the second condition. In this respect the Tactile condition does not show a clear interference with the rest, nevertheless, few pairs performed relatively well when eyes-free, even though the scores were much lower than in the other two conditions. Top scoring pairs – e.g. P1, P8, P9 and P10 – found working strategies (analysed in more detail later on), which were executed well in Visual and Combined, but less successfully in Tactile condition. When strategies broke down, particularly in Tactile condition, the participants tried but only few pairs managed to converge to a new one.

#### 6.4.6 User Experience

The analysis of variance shows that Tactile is significantly different than Combined and Visual on all factors of NASA-TLX (see Appendix A for details). Combined is significantly higher than Visual in Time (Wilcoxon signed ranks  $p < 0.013$  with Bonferroni correction), arguably due to the intrusive nature of vibro-tactile feedback. This resonates well with free-form questionnaire feedback, in which Visual was considered unhurried and slow and Combined was found more responsive and active. Tactile was seen as emphasising collaboration and contact between the partners, and the sense of togetherness as suggested in [4].

Table 6.2 The total numbers of targets acquired by all pairs per condition.

Condition	Combined	Visual	Tactile
Total score	122	135	32

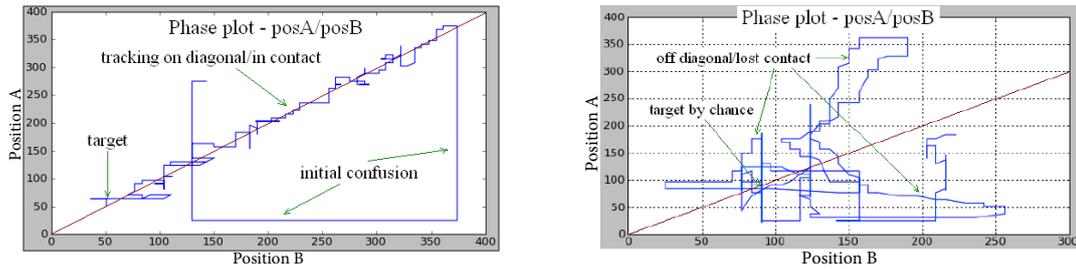


Fig. 6.7 Positions of A vs. B revealing tight tracking (left) and random behaviour (right).

### 6.4.7 Negotiation Strategies

Various types of negotiation strategies can be observed during the collaborative interaction. Although subjects were allowed to verbally agree on a specific solution prior to the experiment, not all pairs had a working strategy in place. One pair described their working strategy as moving ‘together from the top down’. Figure 6.9 presents their time series in the Visual condition, which is an example of continuous tracking – i.e. exploring the virtual space hand-in-hand together. The yellow spheres highlight the target acquisition events. More details from this figure can be found in Appendix B. This pair did not have clearly defined leader and follower roles, but instead implemented a turn-taking leadership strategy resembling more to a smooth dance than a command-and-control behaviour. The phase plot of their close tracking performance (see Figure 6.7/left) shows that after ‘getting in touch’ the pair continuously stayed in contact until reaching the target. Further successful strategies, different from continuous tracking, are shown in Figure 6.8 and in Appendix B. A pair lacking a joint strategy found it difficult to get the other to move to the same direction and achieved only a fraction of the top performers’ results. Given the limited means for exercising a command-and-control style behaviour this pair was unable to interact successfully as can be seen in Figure 6.7/right. The minimalist interaction paradigm, which was the purpose of this study, did not allow them to literally drag or control, but only to perceive each other.

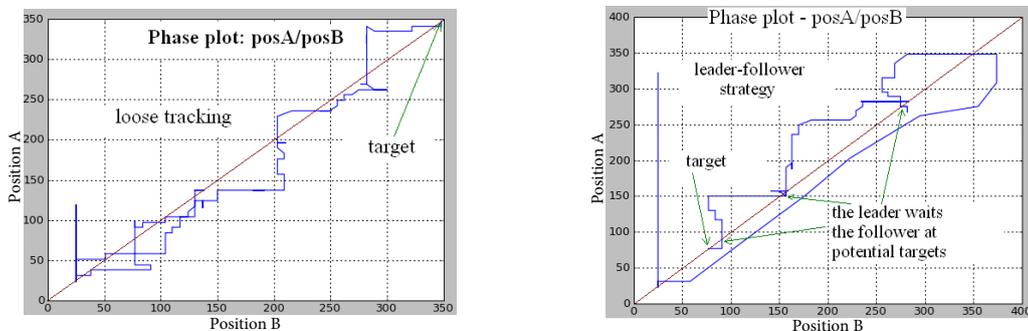


Fig. 6.8 Positions of A vs. B revealing loose tracking (left) and leader–follower roles (right).

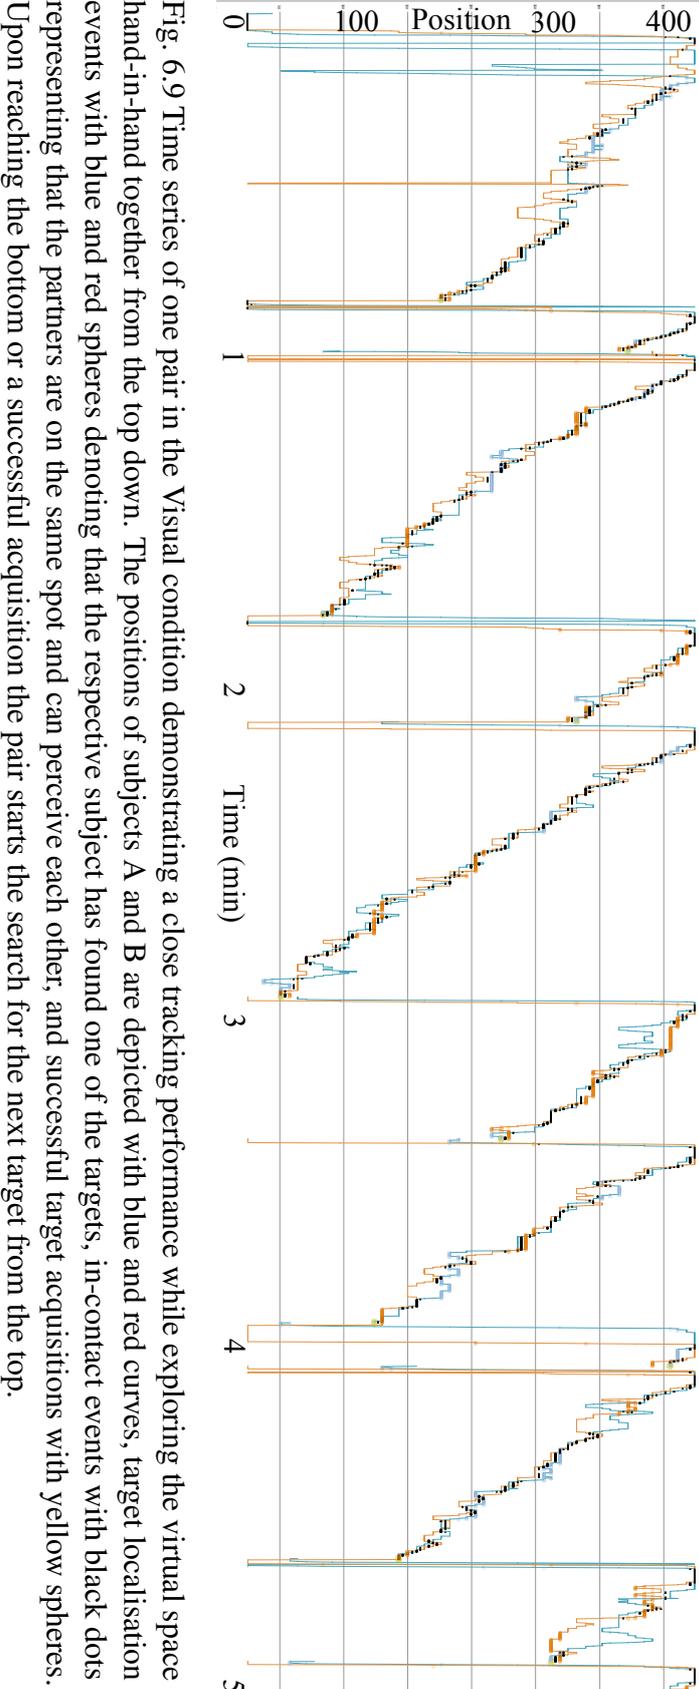


Fig. 6.9 Time series of one pair in the Visual condition demonstrating a close tracking performance while exploring the virtual space hand-in-hand together from the top down. The positions of subjects A and B are depicted with blue and red curves, target localisation events with blue and red spheres denoting that the respective subject has found one of the targets, in-contact events with black dots representing that the partners are on the same spot and can perceive each other, and successful target acquisitions with yellow spheres. Upon reaching the bottom or a successful acquisition the pair starts the search for the next target from the top.

### 6.4.8 Discussion

This exploratory study suggests that such minimalist embodied interaction methods could provide dynamic features enhancing current UI technology. The results demonstrate that participants with little experience or practice with such a system were able to collaborate with each other and reveal interesting behaviours in the social coordination task. The second modality in the Combined condition was completely redundant, still, the majority of the participants preferred that condition over the other two conditions. However, the total score in the Visual condition was higher (see Table 6.2), perhaps because when subjects did not have to divide their attention between the tactile and the visual representation of the same feedback they were faster and more accurate as suggested in [90]. The data reveals that subjects considered the Visual slower and more relaxed than the Combined condition, which however is not correlated with the interaction speed, as participants generally operated at similar speeds in both conditions. In the Tactile condition the perceived workload was significantly higher and the lack of strategy created a feeling of isolation. The performance figures show a clear improvement over time even in this short experiment. The learning effect is particularly noticeable in the Visual and the Combined conditions, in which most pairs increased their scores significantly between the two sessions providing visual feedback.

The top scoring pairs managed to successfully establish a joint strategy prior to the experiment and executed it consistently during the trials. Even when certain strategies failed to materialise some pairs tried and few succeeded in creating new ones that eventually worked. There was no significant difference in performance between pairs who knew each other well and those who did not. The biggest challenge for the subjects was to get the partner to follow them and to find a common way to proceed in the task given the limited means for control. Due to the minimalist interaction style subjects could only perceive, but not drag or control each other directly.

## 6.5 Experiment III

In a follow-up study I explored further the feasibility of the negotiation interaction method introduced in the previous section, redesigned for hand-held devices and evaluated in a mobile environment. In this study participants interacted in pairs remotely and through touch while walking in three feedback conditions: 1) Visual, 2) Audio-tactile and 3) Spatial Audio-tactile. The goals were to:

- Explore human behaviour and performance in mobile virtual environments;
- Reveal the extent to which users felt connected and present in the shared environment;
- Provide guidelines to designers for implementing mobile CVE.

The study aimed to answer the following research questions:

- Can users achieve shared awareness in the eyes-free environment?
- Is spatial audio a superior audio technique in such an environment?
- How efficient and usable is such an interface?

### 6.5.1 Experimental Design

The prototype systems were implemented on Nokia N900 mobile phones linked via WiFi and connected to SK7 sensor packs over Bluetooth. The sensor pack provides one-dimensional scanning touch input for exploring the interaction space. The study consisted of one visual-only and two eyes-free conditions. For the visual-only prototype the mobile phone and the sensor pack were fitted in a custom-made case (see Figure 6.10/left). For the eyes-free prototypes the phone was placed on a lanyard around the user's neck with the sensor pack in hand while wearing a pair of Sennheiser M@B40 headphones (see Figure 6.10/right).



Fig. 6.10 The prototype system consisting of a Nokia N900 and a SHAKE SK7 device used in the Visual condition (left), and the corresponding eyes-free configuration (right).

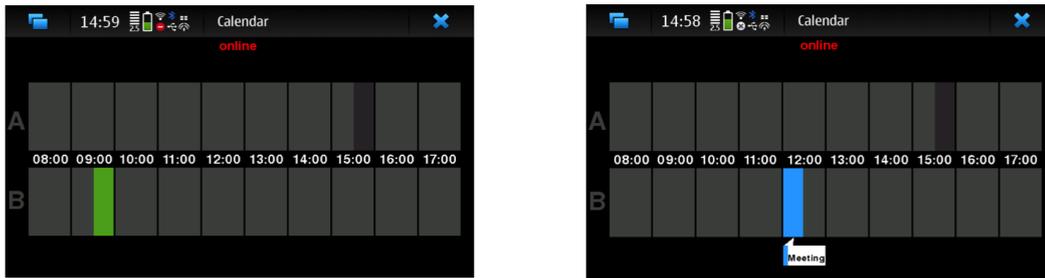


Fig. 6.11 The calendar-like visual display presents both subjects' diaries aligned on the timeline. Free calendar slots are displayed in green (left) and busy ones in blue (right). On the opposite side the partner's calendar is opaque and his cursor is shown in black.

Figure 6.11 presents the calendar-like display used in the Visual condition, which enables subjects to browse appointments on one specific day using the one-dimensional capacitive touch input, and agree on a time for an appointment with their partner. Both subjects' calendars are aligned on the timeline and the cursor movements are discretised at 30-minute-slot steps from 8am to 5pm (20 slots in total). The participants are able to visually follow the position of both cursors on their own mobile phone display. Free calendar slots are shown in green (see Figure 6.11/left) and busy ones in blue (see Figure 6.11/right). The calendar content is always opaque as the focus is on the collaborative interaction and not on how the content would affect the interaction.

In the Audio-tactile prototype a pitch-tone audio cue indicates both the location and the direction of movement of the partner's cursor. The SoundTouch audio processing library ([www.surina.net/soundtouch](http://www.surina.net/soundtouch)) was used to create 22 different sounds forming a 'chromatic scale' with 11 semitones in ascending and 11 semitones in descending order. This chromatic sound scale was mapped to the 20 slots available in the calendar (see Figure 6.12/top). The resulting sounds are 16-bit stereo, sampled at 44 kHz and normalised to 70% of the audio dynamic range, which equals to a normal conversation – typically 60-70dB. Using this scale

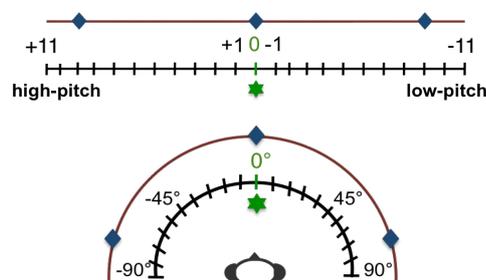


Fig. 6.12 The auditory feedback design in the Audio-tactile (top) and in the Spatial Audio-tactile (bottom) conditions. The hypothetical partner's cursor locations are represented with diamonds. The star at '0' denotes a proposed free slot.

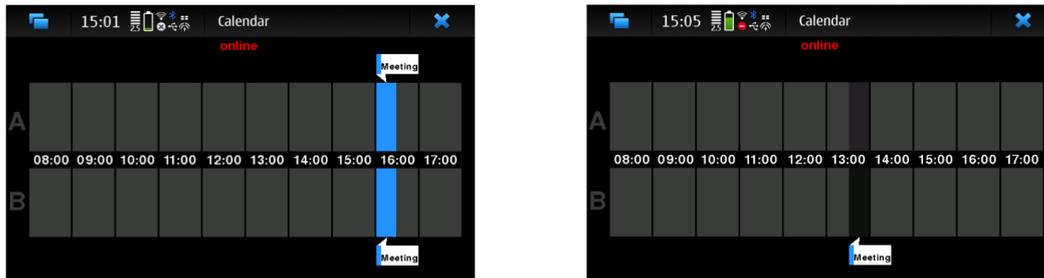


Fig. 6.13 The visual feedback representing a successful target acquisition (meeting arrangement) when both cursors are suggesting the shared free slot (left), and a false target exploration when both cursors are suggesting not shared free slot (right).

a high-pitch sound indicates the partner's cursor is to the left and a low-pitch sound indicates the partner's cursor is to the right. The higher the pitch sound the further to the left the cursor is and vice versa. To avoid overload with sound cues the audio is active only when one of the subjects stops browsing and suggests a free slot. At this point audio is played back at two predefined speeds – quickly, indicating that the partner is browsing, and slowly, indicating that the partner's cursor is stationary. Additional audio cues are played to both subjects in case of success and failure (see Figure 6.13). Auditory feedback is complemented with a dual vibro-tactile display, using slow pulsing pager motor vibration to identify available free slots, and a sharp and fast vibration from a pulse resonant actuator to indicate that both cursors are located in the same slot, i.e. users met in the shared environment.

The design of the Spatial Audio-tactile prototype is identical to the Audio-tactile except for the mapping of the location and direction of movement of the partner's cursor. Only one low-pitch sound source is used to indicate the location of the partner's cursor, instead of a chromatic scale, and the cursor movement direction is mapped to a location in a 3D space around the subject's head ( $-90^\circ$  to  $+90^\circ$  azimuth in 20 discrete steps) always 1m away in the frontal horizontal plane (see Figure 6.12/bottom). The sound source is 16-bit mono, sampled at 16 kHz and normalised to a conversation audio level. The library used to position the audio sources was a port of the JAVA JSR-234 Advanced Multimedia Supplements API to Maemo ([www.maemo.org](http://www.maemo.org)).

## 6.5.2 Tasks

For each trial per condition a pair of participants performed a collaborative target acquisition task consisting of finding a shared free slot on their calendars. In order to achieve this subjects had to individually explore their own free slots and suggest one to their partner by maintaining the cursor at the slot location and waiting until their partner found the slot.

At this point a success or a failure audio cue was played depending on whether that time slot was free for both partners, i.e. shared, or not. For each trial both subjects had three 30-minute free slots in their calendars, of which only one was shared. The positions of the free slots were randomised per trial. Each session lasted up to five minutes or until five trials were successfully completed. The complete interaction, consisting of finger position, contact events, target position, localisation and acquisition, was logged.

### 6.5.3 Methodology

Twenty-six participants (16 male, 10 female, aged 18 to 54) were allocated into 13 pairs, all reported normal hearing. A counterbalanced within-subjects design was used with an extra pair. The trials were always presented in the same order. The conditions were tested in a walking lab environment with the pairs separated into two different rooms. The participants had no other way to communicate except via the shared virtual environment. The experiment consisted of a short introduction followed by a training session and the three conditions. In the training session, which was exclusively devoted to familiarise the subjects with the interaction method, each participant was sitting in a separate room and interacting in pairs using the integrated version of the system providing a combination of visual, audio and vibro-tactile feedback. The training session lasted ten minutes and was divided into two parts: in the first half users could experience the pitch-tone audio and in the second half the spatial audio. In all three conditions the pairs had to perform the same tasks for up to five minutes. At the end of each condition they were asked to complete questionnaires for sense of togetherness [44], affective benefits (ABC-Q) [8], perceived social presence and perceived task performance [84]. Once all conditions were completed a user experience questionnaire was also filled in. The experiment took about one hour in total and the participants were allowed to rest between different conditions.

### 6.5.4 Performance and User Experience

A non-parametric Friedman test shows a significant effect of the type of feedback on the perceived sense of togetherness ( $\chi^2 = 19.5$ ,  $df = 2$ ,  $p < .001$ ,  $N = 26$ ), the affective benefits ( $\chi^2 = 27.2$ ,  $df = 2$ ,  $p < .001$ ,  $N = 26$ ), the perceived social presence ( $\chi^2 = 13.2$ ,  $df = 2$ ,  $p < .005$ ,  $N = 26$ ), and on a 5-point Likert scale on the perceived task performance ( $\chi^2 = 38.7$ ,  $df = 2$ ,  $p < .001$ ,  $N = 26$ ). Pair-wise Wilcoxon signed ranks tests show that the participants had a significantly higher sense of togetherness ( $p < .001$ ), social presence ( $p < .005$ ), affective benefits ( $p < .001$ ) and perceived task performance ( $p < .001$ ) in the visual than in the eyes-free conditions (see Table 6.3).

Table 6.3 The perceived sense of togetherness, social presence and affective benefits.

Condition	Togetherness Median (IQR)	Social Presence Median (IQR)	Affective Benefits Median (IQR)
Visual	4 (3 to 5)	4 (3 to 5)	5 (5 to 5)
Audio-Tactile	2 (1 to 3.25)	3 (2 to 4)	4 (3 to 5)
Spatial Audio-Tactile	2 (1 to 3)	3 (2 to 4)	4 (2.75 to 4.5)

Pair-wise Wilcoxon signed ranks test shows a significantly higher (see Figure 6.14/left) perceived task performance in the Spatial Audio-tactile condition than in the Audio-tactile condition ( $p = .036$  with Bonferroni correction). This is consistent with the measured performance presented in Figure 6.14/right and although that difference is not significant a more positive trend was observed for the Spatial Audio-tactile condition. Informal feedback from the free form questionnaire also show the same positive trend in user preferences towards Spatial Audio-tactile as opposed to Audio-tactile (11:5). These findings provide an affirmative answer to RQ2.

The results from the user experience questionnaire (see Questions 2, 3 and 5 in Figure 6.15) suggest the potential in learning this novel interaction method. However, they also reveal specific usability issues, particularly in the control of the input device (see Question 4, high rating is good) perhaps due to its small size. The results show that the visual baseline system provided higher shared awareness, efficiency and a strong learning effect. However, and although very challenging, the eyes-free systems still offered the ability to build shared awareness in the remote collaborative environments, particularly in the Spatial Audio-tactile one, addressing RQ1.

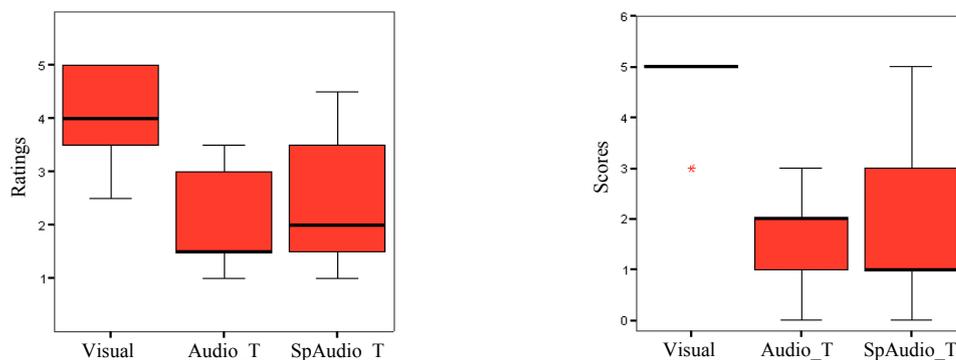


Fig. 6.14 The perceived (left) and the measured performance – number of targets acquired (right) in all three conditions. The perceived performance in the Spatial Audio-tactile is significantly higher than the Audio-tactile condition ( $p = .036$  with Bonferroni correction).

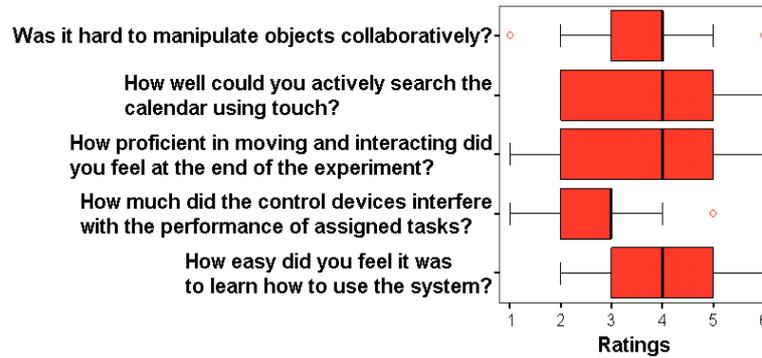


Fig. 6.15 User experience ratings on a 6-point Likert scale.

These results give an insight into and suggest the potential of different feedback mechanisms in the design of future mobile collaborative environments. The performance in the eyes-free systems dropped significantly compared to the visual one, addressing RQ3, which was expected, however the social presence and the affective benefits dropped less dramatically (see Table 6.3).

### 6.5.5 Emerging Strategies

Since the focus was placed on the spontaneous emergence of negotiation strategies and roles the participants were not allowed to discuss and agree verbally on a strategy prior to the experiment. The results show that the pairs were able to perform the task through negotiation. In the beginning the participants tend to explore the space on their own until they converge onto a common strategy and execute that for the rest of the session. Some subjects exhibit higher level of activity and others lower, leading to potential difficulties in establishing a negotiation strategy. Figure 6.16/top shows the emergence of a tracking strategy in the visual condition, where the pair moves together without losing contact, while Figure 6.16/bottom reveals the emergence of a dragging-like behaviour in the eyes-free mode, where after having found a target A gets in contact with B and then tries repeatedly to drag B back to the location of that target. This dragging strategy is presumably inflicted by the challenges for precise tracking without visual feedback.

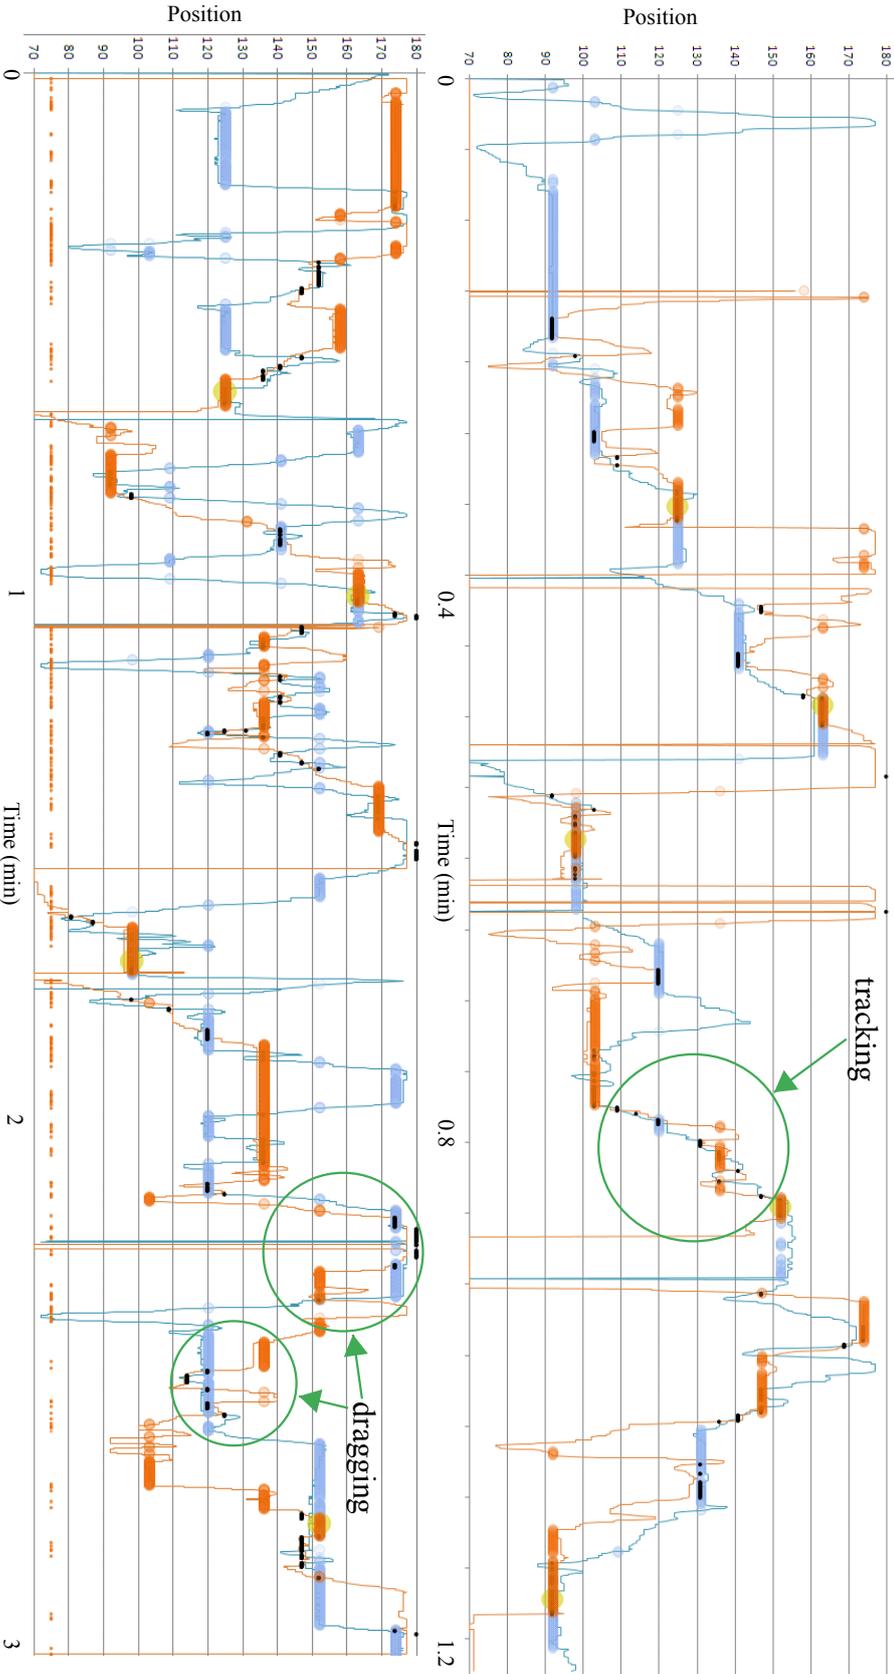


Fig. 6.16 Time series of a pair in the Visual (top) and Spatial Audio-Tactile (bottom) condition. Positions of subjects A and B are depicted with blue and red curves, target localisation events with blue and red spheres denoting that the respective subject has found a target, in-contact events with black dots depicting that A and B are on the same spot and can perceive each other, and successful target acquisitions with yellow spheres. The green circle (top) reveals the emergence of a tracking strategy, where the pair moves together without loosing contact depicted by consecutive black dots. The green circles (bottom) show the emergence of a dragging strategy, where having found a (red) target A gets in contact with B (black dots) and tries repeatedly to drag B back to the location of that target.

### 6.5.6 Turn-taking Patterns

Further analysis reveals specific turn-taking patterns by splitting the interaction phases into four categories:

- both subjects move (browse),
- both subjects are static (suggest),
- A moves, B is static,
- A is static, B moves.

The behavioural dynamics of well performing pairs show more frequent alternations between different phases of interaction, whereas less successful pairs exhibit less variability and longer phases.

The durations of consecutive phases taken by a pair performing relatively well when eyes-free (see Figure 6.17) reveal a dense pattern of exchanged turn-taking leadership in all three conditions. The phase durations show a decreasing trend over time, with a sign of convergence into the [2,4] seconds range, achieved in the second half of the Visual and the Spatial Audio-tactile sessions, reflecting the learning effect for this pair, who negotiated and mastered their collaborative strategy rather quickly even when eyes-free. However, this trend is less pronounced in the Audio-tactile condition, which resonates well with the corresponding measured performance, revealing lower levels for that condition than for the other two conditions. The figure also reveals the significant increase in phase alternations in the eyes-free compared to the visual condition, which can be attributed to the increased uncertainty leading to higher interactivity. This short experiment, however, did not provide a strong evidence of a clear learning effect across pairs due to the high data variability.

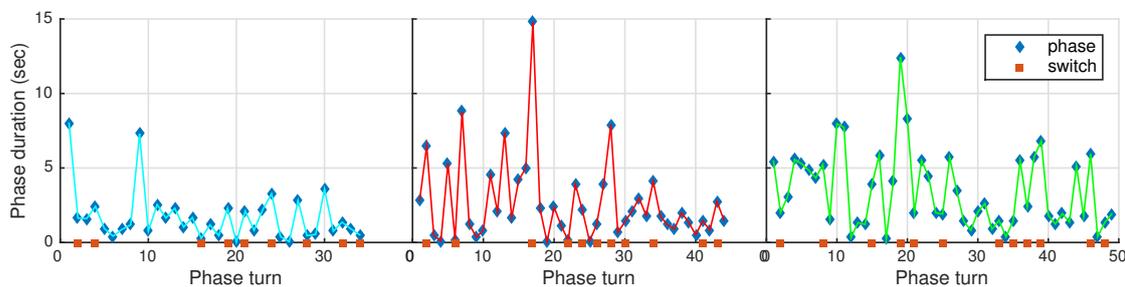


Fig. 6.17 The duration of consecutive phases (blue diamonds) and changes in turn-taking leadership (red squares) in the Visual (left), the Spatial Audio-tactile (middle) and the Audio-tactile (right) condition of a pair performing relatively well in the eyes-free environments. Performance is characterised by intensified interaction and reducing phases.

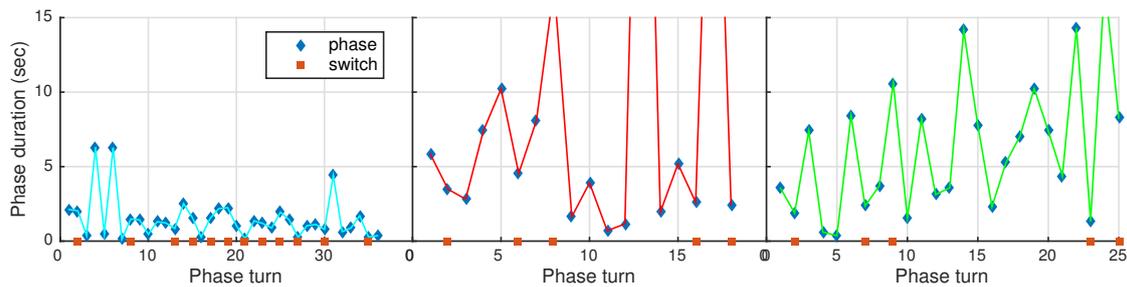


Fig. 6.18 The duration of consecutive phases (blue diamonds) and changes in turn-taking leadership (red squares) in the Visual (left), the Spatial Audio-tactile (middle) and the Audio-tactile (right) condition of a pair struggling in the eyes-free environments. The patterns reveal frequent disengaging from interaction and expanding phases.

The corresponding phase durations of a pair performing well in the Visual condition, however struggling in the eyes-free modes, are presented in Figure 6.18. For this particular pair the increased uncertainty led to continuously expanding phases and disengaging from interaction, reflecting the breakdown of collaboration. The significant behavioural change between the Visual and the eyes-free conditions emphasised by the numbers of phase alternations and turn-taking leadership exchanges in this figure suggests that this pair could not adapt successfully in order to establish a working strategy in these highly uncertain environments.

Figures 6.19 and 6.20 present the phase alternation dynamics of all pairs by stacking the phases on the timeline by their duration in all three conditions. They reveal that the interaction typically starts with a phase of pure exploration (in red), the duration of which, however, varies significantly between pairs. Later on, and particularly in the Visual condition, the phases of pure exploration become less frequent and shorter, suggesting the convergence

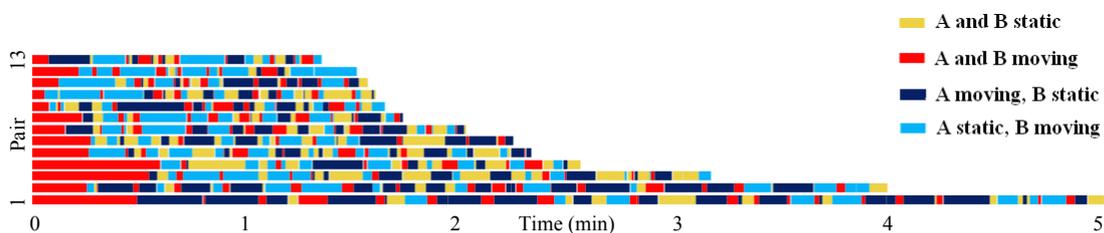


Fig. 6.19 The phase alternations of all 13 pairs (ranked in decreasing order of their session duration) in the Visual condition stacked by the phase duration on the timeline corresponding to the four categories – both move (red), both are static (yellow), one moves the other is static (dark and light blue). The interaction typically starts with a relatively long phase of pure exploration (in red), which later becomes shorter and less frequent.

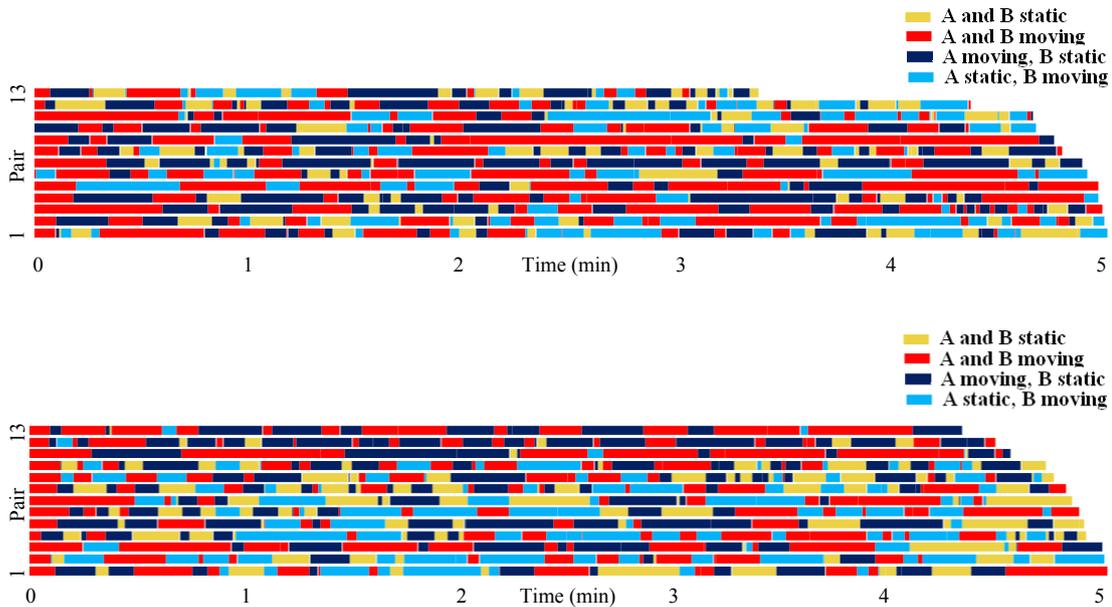


Fig. 6.20 The phase alternations of all 13 pairs (ranked in decreasing order of their session duration) in the Spatial Audio-tactile (top) and the Audio-tactile (bottom) condition stacked by the phase duration on the timeline corresponding to the four categories – both move (red), both are static (yellow), one moves the other is static (dark and light blue). The phases of pure exploration (in red) are relatively long and recurrent throughout the sessions.

to and the execution of specific negotiated strategies. The duration and the periodicity of the other three phases reflect the level of engagement between the partners and the smoothness of the interaction. The figures also reflect the variability in the task completion times between different pairs and across conditions, revealing that most participants could not acquire five targets within the five-minute long eyes-free sessions, whereas in the Visual condition all pairs, but one, succeeded in completing the task.

### 6.5.7 Discussion

This study demonstrated the feasibility of the minimalist interaction paradigm using a realistic collaborative task on current mobile devices. It revealed that users with little experience with such a system were able to establish means for collaboration with each other and provided an insight into the spontaneous emergence of coordination strategies and behavioural roles. The best performing pairs managed to develop a collaborative strategy and executed it consistently across conditions. Others, however, could not adapt successfully to the increased uncertainty of the eyes-free conditions and could not establish a working strategy, leading to disengaging from interaction and an imminent breakdown of collaboration. The biggest challenge was

to negotiate a strategy given the limited means for communication, as subjects could only perceive, but not drag or control each other directly. The efficiency was significantly higher in the visual than in the eyes-free conditions, together with a strong learning effect. Eyes-free interaction increased cognitive load, decreased performance and in the spatial-audio case only showed signs of a slower learning effect. As subjects usually covered the capacitive surface with their thumb they reported difficulties in controlling the discrete transitions at this specific resolution level. The spatial audio was more beneficial at the two extremes, left and right, than in the middle. The range of pitch-tones was considered too narrow and the mapping of high-low pitch to left-right direction not obvious and difficult to learn.

## 6.6 Conclusion

This chapter introduced a novel method for negotiated interaction by building on the abstract membrane metaphor as a proxy conveying the remote sense of touch and using restricted nonverbal means for communication between the partners. The feasibility of the proposed multimodal interaction method was evaluated in two collaborative target acquisition user studies using a shared calendar-like paradigm. The studies provided an insight into the dynamics of this social coordination task by revealing the emergence of specific collaborative strategies like tracking and dragging, as well as by highlighting certain patterns of human behaviour corresponding to lead-lag and turn-taking. Although the performance in the eyes-free conditions dropped significantly compared to the visual-enabled conditions, few pairs sustained their strategies across conditions and executed them relatively well even when eyes-free. The spatial audio presentation showed certain performance benefits, while the tactile feedback emphasised the sense of togetherness and social presence.

These results provide a better understanding of the potential different feedback mechanisms have for the creation of richer interpersonal communication systems and demonstrate that although an exact translation of a visual interface using other modalities such as audio or haptics is not plausible, it is possible to design equivalent eyes-free interfaces. This is becoming increasingly important for building more robust mobile interfaces, which enable users to maintain interaction by shifting their attention between modalities when vision is compromised.

The collected empirical data is used in the development of information-theoretic models and measures presented in the following chapters – characterising empowerment (Chapter 7) and quantifying interpersonal coordination (Chapter 8).

# Chapter 7

## Empowerment in Eyes-free Interaction

*‘Essentially, all models are wrong, but some are useful.’*

---

George Box

### 7.1 Introduction

This chapter presents an application of the empowerment formalism for modelling tracking behaviour building on the hypothesis that skilled human behaviour is driven by greedy empowerment maximisation, a criterion shown in recent work [19] to motivate guided self-organisation of collective complex systems. The model is based on the minimalist interactive paradigm corresponding to the Tactile condition of Experiment II (Section 6.4), which provides a clear distinction between the various feedback events. The essential performance properties of the proposed model are explored in series of simulated experiments and the resulting patterns of tracking behaviour are validated on data collected in Experiment II. Despite the simplicity of the proposed stochastic model parameterised by only two controlled variables – the levels of noise and delay respectively – the results suggest the plausibility of this approach for modelling human behaviour. Building on this model I compute the average empowerment levels from the empirical data of Experiments II and III (Chapter 6), and elaborate on their relation to perceived performance measures, however a part of these results is of rather speculative nature as it rests on the assumption that the model holds for all conditions in both experiments. This proof-of-concept example expands the scope of empowerment beyond providing only theoretical bounds, but also inferring the actual level of control experienced by subjects in the course of interaction.

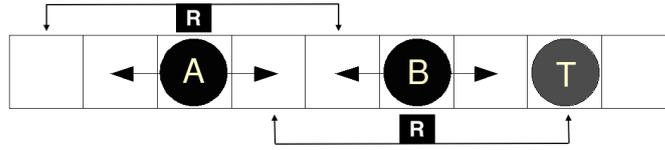


Fig. 7.1 Two agents (A and B) moving horizontally in a section of a line world aiming to jointly acquire a target. T denotes the target and R the range of contact, which specifies the sensing areas for both agents.

## 7.2 Two Agents in a Line World

This section will introduce a model of a dyad interacting in a discrete line world later used to generate specific patterns of human tracking behaviour in simulated experiments. Consider a 1-D grid (line) world of infinite size, in which two agents interact. Each agent has an absolute location in the grid and a rigid body which can occupy one tile. At each time-step the agents are allowed to move rigidly by a distance of one tile either back or forth or stay static. Both agents have a sensor, an actuator and a controller, which has a direct access to the state space at every step. The goal of both agents is to jointly capture a static target through collaborative interaction. The agents can sense whether they are within certain distance of each other and whether they have located the target (stepped on it). At every step, for simplicity, I will consider one of the agents as part of the environment and thus decrease the state space dimensionality when computing the empowerment for the other agent. This will be possible since the model assumes that the agents alternate in taking actions on even and odd steps. I will denote the agents with A and B, the target with T and the range of contact with R, within which both agents can sense each other and the target (see Figure 7.1).

### 7.2.1 Model

Building on empowerment maximisation I introduce a minimalist model reflecting human behaviour in the Tactile condition of Experiment II (see Chapter 6). As the Tactile condition lacks visual feedback I assume that the agents cannot observe their current position in space and are only guided by tactile events, meaning that the system state is a latent variable and the transitions are driven only by impact events. Each agent can sense the proximity of the other agent (within/out of range) and its own distance from the target (on/off target). Hence, the essential information driving the system's dynamics is the agent's relative position to the target and to the other agent. For every state  $x \in \mathcal{X}$  empowerment is defined as

$$\mathfrak{E}(x) = \max_{p(\vec{a})} I(A_t, \dots, A_{t+n-1}; S_{t+n}|x), \quad (7.1)$$

where the action space  $\mathcal{A}$  consists of three 1-step actions

$$\mathcal{A} = \{idle, left, right\},$$

and the sensor (observation) space  $\mathcal{S}$  is defined by two binary values

$$\mathcal{S} = \{0, 1\} \times \{0, 1\},$$

representing the impact events with the partner and the target. The system state space  $\mathcal{X}$  represents both agents' and the target's positions on the grid and the range of contact

$$\mathcal{X} = \mathbb{Z} \times \mathbb{Z} \times \mathbb{Z} \times \mathbb{Z}.$$

Since for each agent, the target, the other agent's position and the range of contact are fixed at every step in the empowerment computation I will use the following lower dimensional state space representing the position of the agent in question

$$\mathcal{X} = \mathbb{Z}.$$

Defining the impact events with the partner and the target by the random variables  $C$  and  $G$  respectively as follows

$$C(x) = \begin{cases} 1, x \in [b - r_t, b + r_t] \\ 0, x \notin [b - r_t, b + r_t] \end{cases} \quad G(x) = \begin{cases} 1, x = T \\ 0, x \neq T \end{cases},$$

where  $b$  is the current (fixed) position of the other agent and  $r_t$  is an instance of the range of contact Equation 7.1 takes the form of

$$\mathfrak{E}(x) = \max_{p(\vec{a})} I(A_t, \dots, A_{t+n-1}; C_{t+n}, G_{t+n} | x).$$

For both agents the transition probabilities  $p(x'|x, a)$  in state  $x \in \mathcal{X}$  are represented by a multivariate Gaussian distribution with  $\mu = (x, x-1, x+1)^T$  mean vector and  $\Sigma = \text{diag}(\sigma_0^2, \sigma_1^2, \sigma_1^2)$  covariance matrix, which in short will be written as

$$x'|x, a \sim \mathcal{N}(\mu, \Sigma).$$

The Gaussian mean vector is a function of the action  $a$  and the state  $x$ , while the covariance matrix depends only on the action.

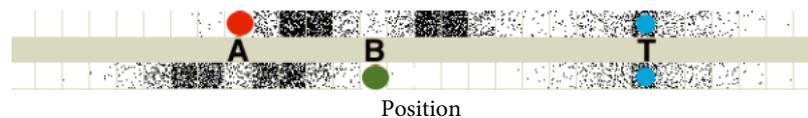


Fig. 7.2 An example visualising the empowerment levels of both agents (A and B). A, B and the target T are depicted with red, green and blue dots respectively. The empowerment levels are represented by density clouds (black corresponds to high and white to zero, i.e. the denser the cloud the higher the empowerment). For both agents the empowerment is higher at the edges of their contact regions and at the target.

## 7.2.2 Dyad Empowerment

Using the Blahut-Arimoto algorithm (Section 3.4) and the above model I compute the 1-step and 2-step empowerment for both agents separately while keeping the position of the other agent fixed. Following the discrete representation of the line world model in Figure 7.1 I split the presentation of results into two blocks reflecting the empowerment of agent A (top) and agent B (bottom) simultaneously. Figure 7.2 shows an example, in which the positions of the agents A and B are depicted with red and green dots respectively, while their corresponding empowerment levels are represented by density clouds (black corresponds to high and white to zero, i.e. the denser the cloud the higher the empowerment). The target T is depicted with a blue dot. For both agents the empowerment is higher at the target and at the edges of their respective contact region, i.e. the places where they can perceive the transition between sensing and not sensing events. Larger step horizons would require proper propagation of the position of the other agent, which necessitates a higher level model of decision making and is a potential topic for future research. Figure 7.3 presents the results in the deterministic case ( $\sigma_0^2 = 0$  and  $\sigma_1^2 = 0$ ) using two ranges of contact –  $r_t^A = 3$  for agent A and  $r_t^B = 1$  for

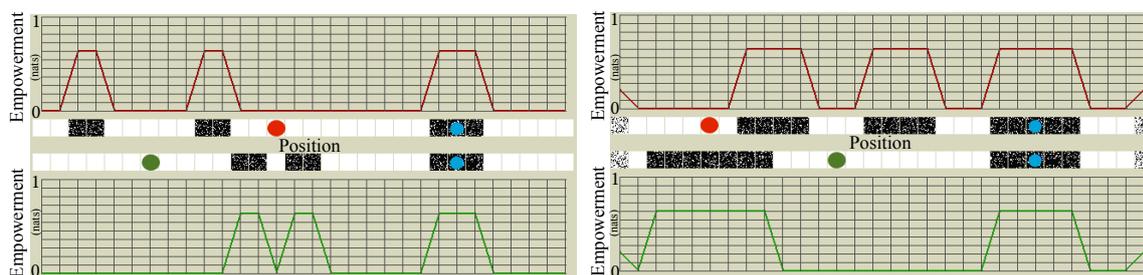


Fig. 7.3 Visualisation of 1-step (left) and 2-step (right) empowerment in the case of deterministic state transitions ( $\sigma_0^2 = 0, \sigma_1^2 = 0$ ) using two ranges of contact –  $r_t^A = 3$  for agent A and  $r_t^B = 1$  for agent B. The empowerment levels are represented by green and red curves as well as by density clouds. Inside larger contact regions empowerment is zero (top). 2-step action sequences inflict the merger of the two basins of high empowerment (bottom/right).

agent B. It shows that the places of high empowerment are located around the target and at the edges of the contact region, locations where both agents can influence the environment most by their actions and observe it. Figure 7.3/left shows that inside the contact regions empowerment is zero since any 1-step action keeps the sensor inputs constant, while using 2-step action sequences extends the basins of high empowerment (see Figure 7.3/right), which for  $r_t^A = 1$  merge with one another and fill the empowerment gap between the edges of the contact region (green curve).

Figure 7.4 presents the results in the stochastic case ( $\sigma_0^2 = 0.7, \sigma_1^2 = 1.5$ ) using various ranges of contact, in which the empowerment level on the outer edge of the contact region is slightly higher (0.54 nats for agent A and 0.63 nats for agent B) than on the inner edge (0.5 nats for agent A and 0.62 nats for agent B) reflecting more options when moving out of the contact region as opposed to moving in. The reason for this is the clearer distinction in the outputs of the moving actions when measured outside than inside the edge. Empowerment around the target peaks exactly at  $T$  (0.47 nats) due to the options provided by the ‘idle’ action, however that value is lower than the one on the edge due to the limited variability of the ‘idle’ action alone (moving actions have no influence in that location). Figure 7.4/bottom/right shows for  $r_t^A = 0$  the contraction of the contact region into the same shape as the target region.

As the agent approaches the target and the edge of the contact region reaches  $T$  the two basins of high empowerment merge as shown in Figure 7.5/top/left. The highest empowerment level reaches 0.94 nats, which reflects the increased number of options created by the combination of these two properties of the system’s dynamics. When both agents reach the vicinity  $r_t$  of  $T$  both contact ranges are reduced to zero ( $r_t^A = r_t^B = 0$ ) in order to

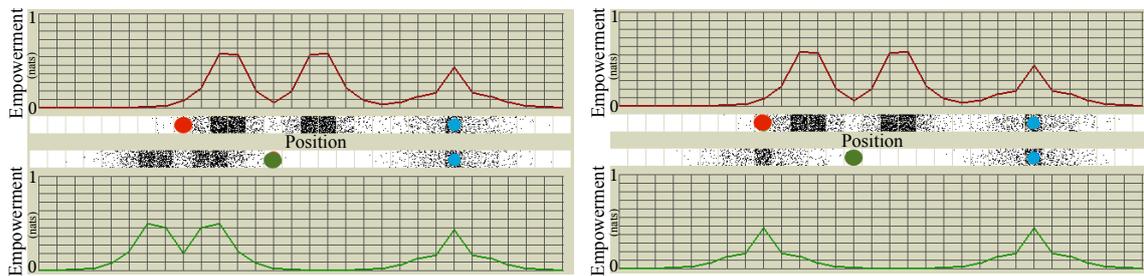


Fig. 7.4 Visualisation of 1-step empowerment in the case of stochastic state transitions ( $\sigma_0^2 = 0.7, \sigma_1^2 = 1.5$ ) using various ranges of contact –  $r_t^A = 1, r_t^B = 2$  (left) and  $r_t^A = 0, r_t^B = 2$  (right). The empowerment levels are represented by green and red curves as well as by density clouds. On the outer edge of the contact region empowerment is slightly higher than on the inner edge (left). For  $r_t^A = 0$  the contact region contracts and assumes the shape of the target region with a single peak (bottom/right).

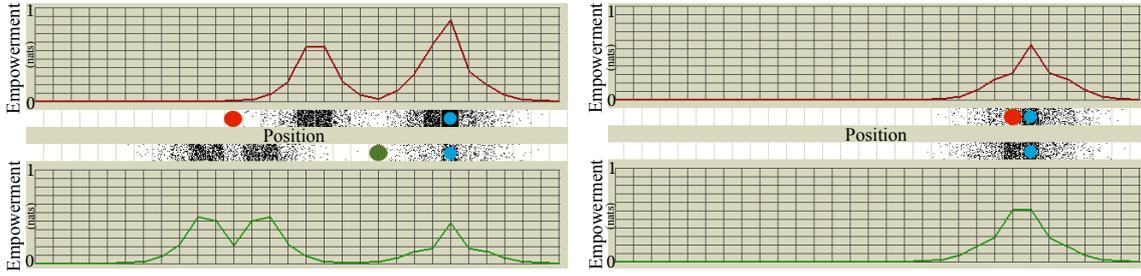


Fig. 7.5 Visualisation of 1-step empowerment in the case of stochastic state transitions ( $\sigma_0^2 = 0.7, \sigma_1^2 = 1.5$ ) using various ranges of contact –  $r_t^A = 1, r_t^B = 3$  (left) and  $r_t^A = 0, r_t^B = 0$  (right). The empowerment levels are represented by green and red curves as well as by density clouds. The two basins of high empowerment merge at the target providing the highest peak (top/left). The ranges of contact are reduced to zero at the target for both A and B resulting in a single peak (right).

facilitate the joint acquisition of  $T$ , since at that point the main driver for the agents is not to track each other, but to acquire the target. This yields a single peak of 0.64 nats in the empowerment curve, which is higher than previous values at  $T$  (see Figure 7.5/right).

## 7.3 Experiment

### 7.3.1 Methodology

In order to find a good fit for the tracking patterns observed in Experiment II (see Chapter 6) I performed series of simulations using the stochastic model introduced in the previous section and the transfer functions defined by Equations 7.2 and 7.3. Assuming that the action  $A_t$  depends only on the sensor input  $S_t$  and the agent's decision making (the probabilistic mapping  $S_t \rightarrow A_t$ ) is time invariant the following transfer functions completely describe the state of the system at any given time

- Agent A

$$x_{t+1} = f(S_{x_t}, \mathcal{A}), \quad (7.2)$$

where

$$S_{x_t} = g_a(x_t, y_t, T)$$

and

- Agent B

$$y_{t+1} = f(S_{y_t}, \mathcal{A}), \quad (7.3)$$

where

$$S_{y_t} = g_b(x_t, y_t, T),$$

in which  $f$  is an empowerment maximisation over  $\mathcal{A}$  and the feedback display functions  $g_a$  and  $g_b$  represent the limited observational capabilities of the agents. Both agents' controllers take sensor input and select the next action based on greedy empowerment maximisation. The same stochastic parameters  $\sigma_0^2 = 0.7$  and  $\sigma_1^2 = 1.5$  are used for both agents in all trials while varying the range of contact  $r_t \in \{1, 2, 3, 4\}$  and the level of uncertainty injected into the perception–action loop, which is represented by a Gaussian state transition noise  $\mathcal{N}(0, \sigma^2 = 0, 1, 2)$  and an actuation delay (0, 2 or 4 steps). Twenty simulations were executed with each set of parameters while varying the initial positions of the agents according to a Gaussian distribution  $\mathcal{N}(\mu = 5, \sigma^2 = 3)$  and keeping the target position fixed. At every step both agents were driven by 1-step empowerment maximisation over the entire state space, i.e. by applying the stochastic transitions to move towards the place of the highest empowerment. Since the focus of this experiment was on tracking behaviour I assumed that both agents were aware of the direction towards the target location and always picked the edge of the contact region closer to  $T$  whenever the highest empowerment occurred at multiple places simultaneously. When the agents reached the close vicinity of  $T$  simultaneously both ranges of contact were reduced to zero in order to ensure quick convergence onto the target and avoid oscillating behaviour around it.

### 7.3.2 Results

In all simulation trials the agents succeeded to jointly acquire the target with certain performance variability, which is described below in more detail. The average number of steps<sup>1</sup> taken until successful target acquisition derived for twenty trials per condition and four contact range levels tested in the simulations (see Figure 7.6) is a strictly monotonically increasing function of uncertainty. The figure reveals that increasing the contact range helps stabilise the system and makes it more efficient and robust to uncertainty. While for  $r_t = 1$  (see Figure 7.6a) the average number of steps increases fast with uncertainty for  $r_t = 4$  (see Figure 7.6d) the growth is significantly slower, which suggests that the larger contact range has a damping effect on the disturbances. Figure 7.7 presents the mean distance between the agents A and B averaged over the course of interaction for twenty trials per condition, which also increases with uncertainty. The damping effect of larger contact regions appears at the level of maximal uncertainty (delay = 4,  $\sigma^2 = 2$ ) where this indicator is lower for  $r_t = 4$  than for  $r_t < 4$ . An interesting trade-off between the effects of noise and delay on the average

<sup>1</sup>Agents' steps are alternating, A makes a move on even and B on odd steps.

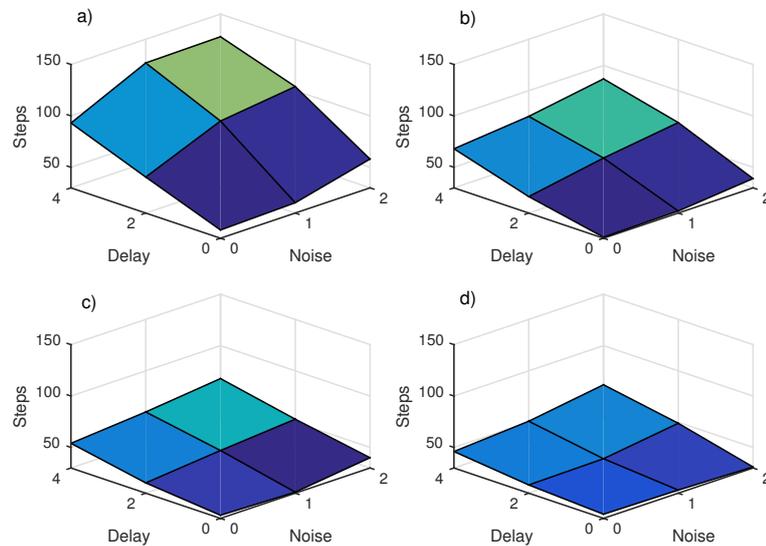


Fig. 7.6 Average number of steps until successful target acquisition per twenty simulated trials in each condition for four levels of the contact range: (a)  $r_t = 1$ , (b)  $r_t = 2$ , (c)  $r_t = 3$ , (d)  $r_t = 4$ . Performance drops with increased uncertainty, however incrementing the contact range introduces a damping effect on disturbances and improving efficiency.

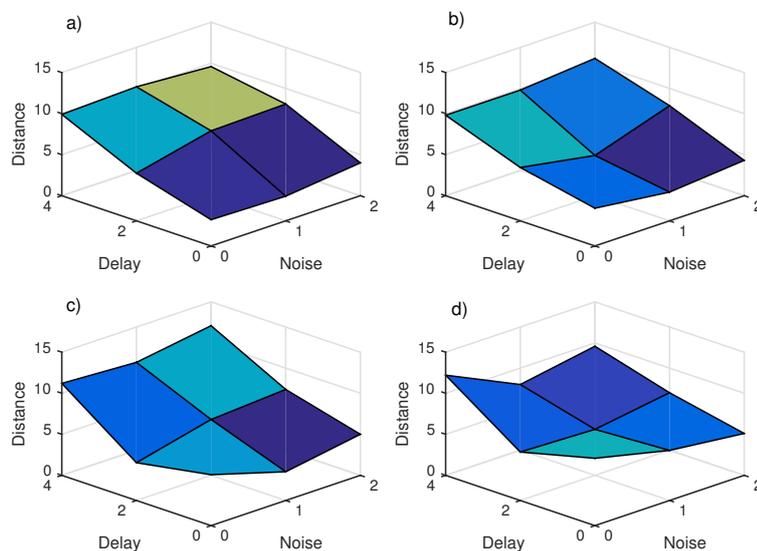


Fig. 7.7 Average distance between agents A and B per twenty simulated trials in each condition for four levels of the contact range: (a)  $r_t = 1$ , (b)  $r_t = 2$ , (c)  $r_t = 3$ , (d)  $r_t = 4$ . Performance generally drops with increased uncertainty, however a specific interplay between the highest delay and increasing noise levels improves efficiency indicating the adoption of a more cautious strategy by the agent.

distance can be observed in Figure 7.7d where at the highest delay level, surprisingly, the introduction of noise results in a drop in the average distance suggesting that the increased uncertainty makes the agent more cautious. Moderate levels of noise seem to improve this performance indicator, however, further increasing the noise level proves to have an adverse effect as agents struggle to cope with excessive uncertainty.

Figures 7.8 and 7.9 present the simulated trajectories of both agents' positions derived from a single trial in four extreme conditions – two deterministic (delay = 0,  $\sigma^2 = 0$ ) and two stochastic (delay = 4,  $\sigma^2 = 2$ ) using the smallest ( $r_t = 1$ ) and the largest ( $r_t = 4$ ) levels of the contact region.

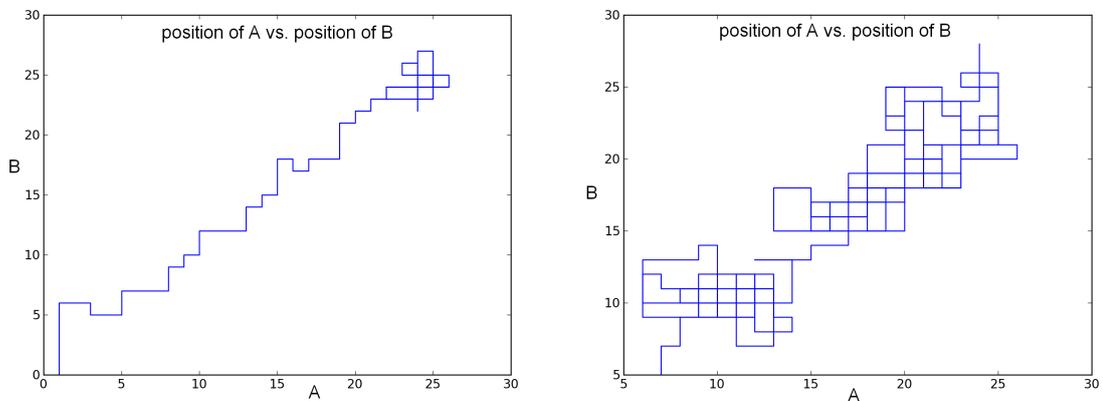


Fig. 7.8 Simulated trajectories of agent A vs. agent B for  $r_t = 1$  range of contact and (a) delay = 0,  $\sigma^2 = 0$ , (b) delay = 4,  $\sigma^2 = 2$  reflecting patterns of a tight (left) and a chaotic (right) tracking behaviour.

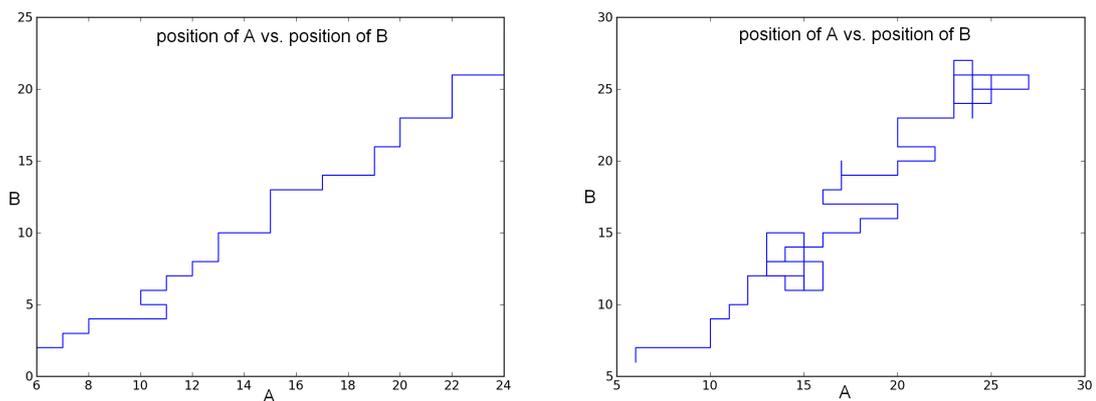


Fig. 7.9 Simulated trajectories of agent A vs. agent B for  $r_t = 4$  range of contact and (a) delay = 0,  $\sigma^2 = 0$ , (b) delay = 4,  $\sigma^2 = 2$  reflecting patterns of a tight (left) and a loose (right) tracking behaviour.

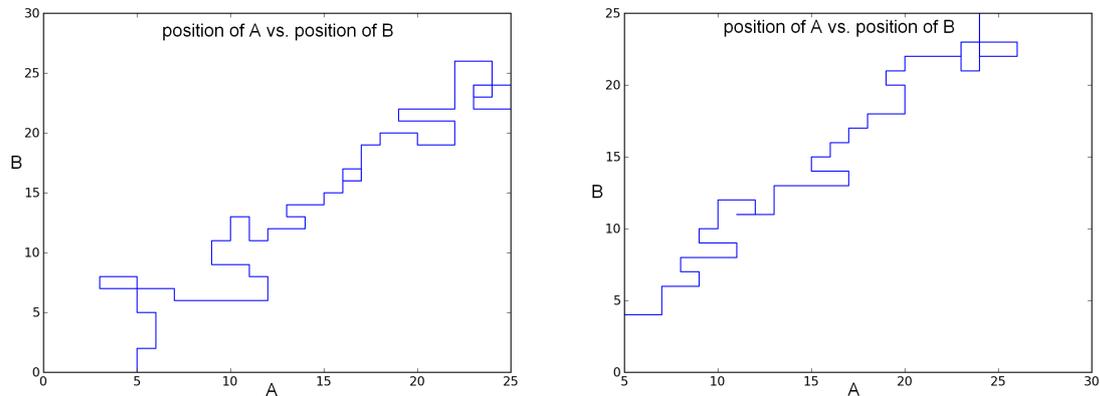


Fig. 7.10 Simulated trajectories of agent A vs. agent B for delay = 2,  $\sigma^2 = 1$  and (a)  $r_t = 2$ , (b)  $r_t = 3$  reflecting patterns similar to human tracking behaviour recorded in the user study.

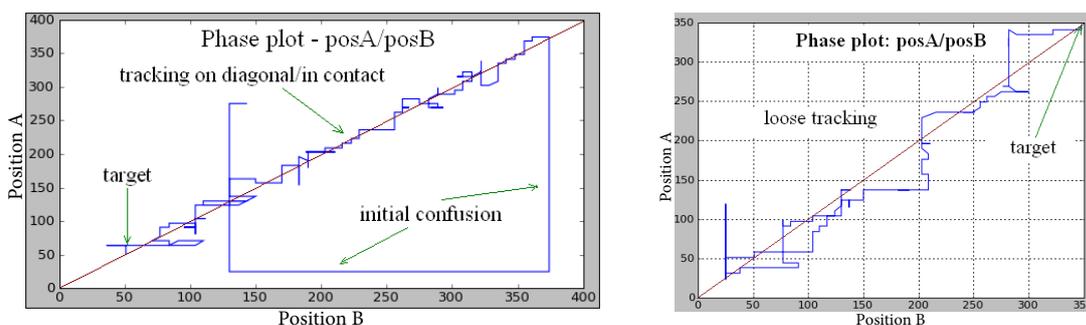


Fig. 7.11 Examples of tight (left) and loose (right) tracking performance provided by the experimental trajectories of subjects A and B in the Combined condition of Experiment II.

They show either a very tight (7.8a and 7.9a) or a too loose (7.9b) and sometimes chaotic tracking (7.8b). A qualitative examination of the trajectories in Figures 7.8b and 7.9b reveals the damping effect the large contact region has on uncertainty. Trajectories derived from simulations with moderate levels of noise and delay are presented in Figure 7.10a ( $r_t = 2$ ) and 7.10b ( $r_t = 3$ ). Qualitative analysis of Figures 7.8, 7.9 and 7.10 suggests that while the former two figures show only limited similarity to the empirical human data, the latter reveals more moderate behaviour resulting in patterns similar to human trajectories recorded in the user study, two characteristic examples of which are presented in Figure 7.11. These initial results provide qualitative evidence for the capacity of empowerment maximisation alone to generate patterns of collective human behaviour, which supports the hypothesis of empowerment as a guiding principle in self-organising biological systems. However, more rigorous quantitative analysis, based on suitable distance measure, is required to validate the proposed information-theoretic approach to model fitting.

## 7.4 Empowerment as a Measure of Control

The empowerment framework provides a theoretical measure quantifying the potential level of control, which serves as an upper bound characterising the limits of human performance. The inference of the actual level of empowerment experienced by subjects in the course of interaction requires the application of a theoretical model on empirical data.

To evaluate data recorded in Experiment II (Section 6.4) I applied the minimalist line world model introduced earlier in this chapter and computed the average levels of empowerment by counting the events of crossing places of high empowerment in the data as defined by this model. It must be noted that only the Tactile condition of Experiment II is described appropriately by this model, since the tactile display provides a binary (on/off) feedback. The visual display is not a true translation of the tactile one as it introduces an artefact providing a directional cue, which is not accounted for by this minimalist model. Nevertheless, I present the evaluation of all three experimental conditions for the sake of completeness.

As described in Section 7.2 there are two types of high empowerment basins – one at the edge of the contact zone and another one at the target. I computed the average levels of empowerment for each pair and both types of basins separately. Figure 7.12 presents the empirical evaluation of data recorded in the Combined, Visual and Tactile conditions of Experiment II. The between-pair variability of the relative time spent in places of high empowerment (see Figure 7.12a) reveals a trade-off between the two basins of high empowerment associated with the partner and the target respectively, and provides implications for modality selection. Furthermore, the between-pair variability of the subjective ratings reflecting how easy it was to locate the partner and the target (see Figure 7.12b) suggests a potential link between the two indicators.

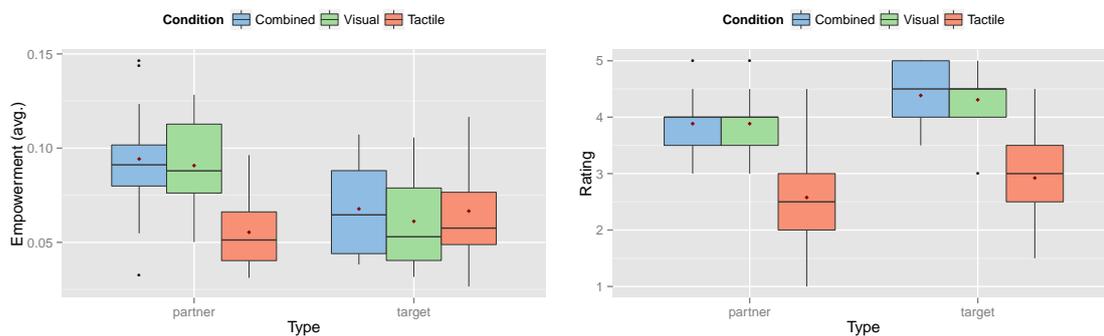


Fig. 7.12 Evaluation of data recorded in Experiment II presenting (a) the average empowerment, which reflects the relative time spent by both subjects in places of high empowerment at the edge of the partner's contact zone (left) or at the target (right) and (b) the subjective ratings indicating how easy it is to locate the partner (left) and the target (right).

A non-parametric Friedman test shows a significant effect ( $\chi^2 = 9.7$ ,  $df = 2$ ,  $p$ -value = .008) of the type of feedback on the time subjects spent at the edge of the contact region (see Figure 7.12a/left). Wilcoxon paired signed rank tests show a significant difference between the Combined and the Tactile and between the Visual and the Tactile conditions ( $p$ -value = .005 and .004 respectively with Bonferroni correction). This suggests that it is considerably easier to stay in contact with the partner in the visual-enabled conditions, which is influenced also by the directional information contained in the visual cue. The amount of time subjects spent at potential targets (see Figure 7.12a/right) is on the same level across conditions, however with a higher variability.

The subjective ratings (see Figure 7.12b) emphasise the higher difficulty level of the Tactile with respect to the other two conditions. A non-parametric Friedman test ( $\chi^2 = 20$ ,  $df = 2$ ,  $p$ -value = .000) shows a significant effect of the type of feedback on the difficulty level of locating the partner (see Figure 7.12b/left). Wilcoxon paired signed rank tests show significant difference between the Combined and the Tactile and between the Visual and the Tactile conditions ( $p$ -value = .007 and .005 respectively with Bonferroni correction). A non-parametric Friedman test ( $\chi^2 = 16$ ,  $df = 2$ ,  $p$ -value = .000) shows a significant effect of the type of feedback on the difficulty level of locating the targets (see Figure 7.12b/right). Wilcoxon paired signed rank tests show significant difference between the Combined and the Tactile and between the Visual and the Tactile conditions ( $p$ -value = .007 and .006 respectively with Bonferroni correction).

The results reveal an interesting trade-off. Subjects spent significantly more time in contact with the partner than at a potential target in the Combined and the Visual conditions (see Figure 7.12a) – indicated by Wilcoxon paired signed rank test ( $p$ -value = .068 and .008 respectively). However, in the Tactile condition the trend is the opposite (Wilcoxon paired signed rank test  $p$ -value = .068). This suggests that subjects spent considerably more time trying to control and track each other than searching for targets in the conditions, in which the visual feedback enhanced the interaction. The corresponding trends for the subjective ratings in the Combined and the Visual conditions are statistically significant (see Figure 7.12b), however reversed, indicated by Wilcoxon paired signed rank test ( $p$ -value = .025 and .049 respectively) suggesting that the static object is considered easier to localise as opposed to the unpredictable dynamic object.

Figure 7.13a presents the evaluation of empowerment in respect to the data recorded in all three conditions of Experiment III (Section 6.5). It shows that the empowerment levels around the partner and at the target are similar in the Visual-only, however they are significantly different in the Audio-tactile ( $p$ -value = .000) and in the Spatial audio-tactile ( $p$ -value = .001) conditions revealed by Wilcoxon paired signed rank tests. This suggests

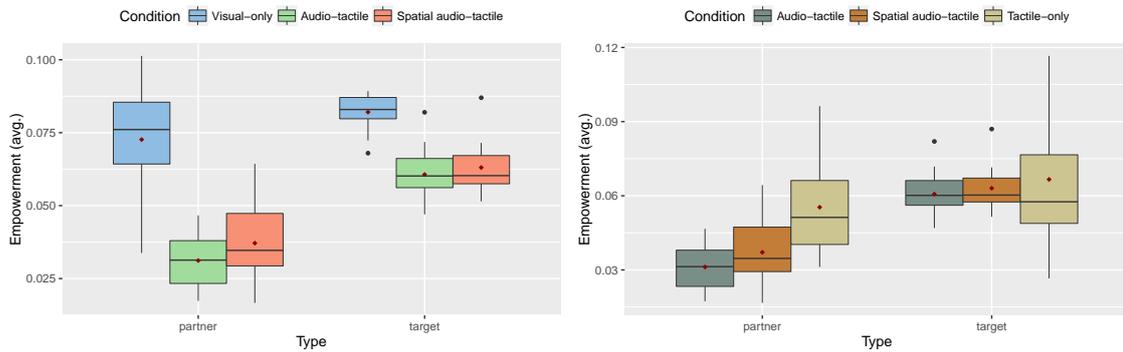


Fig. 7.13 Evaluation of data recorded in Experiments II and III presenting the average empowerment, which reflects the relative time spent by both subjects in places of high empowerment at the edge of the partner's contact zone or at the target in (a) all three conditions of Experiment III, and (b) eyes-free conditions of Experiments II and III.

that in the eyes-free conditions subjects spent more time exploring the static objects than interacting with the dynamic agent, which beside the difference in the difficulty level can also be attributed to the benefit of the dynamic audio feedback providing directional cue and helping follow the partner's movement.

Figure 7.13b reveals that the empowerment at the target in the eyes-free conditions of Experiments II and III is on the same level. Furthermore, the empowerment around the partner in the Tactile condition of Experiment II is significantly higher than in the Audio-tactile ( $p$ -value = .005) and Spatial audio-tactile ( $p$ -value = .03) conditions of Experiment III, whereas the latter two are not significantly different revealed by Wilcoxon paired signed rank tests. This confirms the earlier observation that the audio cue relieved the cognitive burden of tracking the other in the eyes-free conditions and decreased the time spent in close proximity of the partner, which provides further evidence for its benefits.

## 7.5 Discussion

The simulated experiments revealed a certain damping effect larger contact ranges have on disturbances. Larger contact ranges imply more spread basins of high empowerment, which seem to facilitate interaction between the agents and lead to improved performance. Furthermore, the results reveal the favourable impact of noise on the average distance between the agents for large levels of delay and contact range suggesting that increased uncertainty makes the agent more cautious. Simulated trajectories for moderate levels of disturbances reveal patterns similar in shape to human performance corresponding to close tracking behaviour, which confirms the observation that human interpersonal coordination is

modulated by a certain level of reaction delay and randomness in response due to inherent sensorimotor and decision-making constraints.

One conceptual disparity between the visual and the tactile feedback by design is that while the partner's direction of movement during contact is obvious from the visual cue the tactile cue does not provide such gradient information, which makes the prediction of the partner's direction harder and frequently leads to losing contact.

The trend in the Tactile condition is consistent for both the objective and the subjective measures, which can be linked to how empowered participants felt by this minimalist eyes-free interface. As noted earlier, the design of the theoretical model aims to fit primarily the agent and the environment in the Tactile condition, which helps understand the link between the empirical and the analytical results. When faced with this extreme condition subjects spent more time in the basins of high empowerment at the static targets than around the moving partners and their self-reported user experience is in line with that. This suggests a certain relationship between the two types of empowerment basins, which is not captured nor explained by the model and is an interesting topic for future research.

In the eyes-free conditions the audio cue proved beneficial in relieving the cognitive burden of tracking the partner revealed by the decreased time spent in his/her close proximity.

## 7.6 Conclusion

This chapter presented an application of the empowerment formalism for generating goal-oriented behaviour without specifying an explicit reward function, which highlights its benefit as a universal task-independent utility function identifying salient places in the environment and driving agents towards them through empowerment maximisation. The proposed minimalist model of tracking behaviour provided patterns similar to human performance based only on the system dynamics and driven by empowerment maximisation, which suggests the plausibility of this approach for modelling collective human behaviour. Furthermore, the evaluation of the user studies highlighted the notion of empowerment as instrumental in explaining experimental data and in providing a method for inferring the actual level of control experienced by participants in the course of interaction.

This initial study demonstrated the potential of the proposed model, which could be improved in future work in several aspects by 1) developing continuous non-linear dynamics and a predictive mechanism, 2) splitting the uncertainty treatment between the reference point and the internal states, 3) incorporating variable time delays in order to provide more realistic models and 4) defining a suitable distance measure for model fitting.

# Chapter 8

## Interpersonal Coordination in Collaborative Environments

*'It is a capital mistake to theorise before one has data. Insensibly one begins to twist facts to suit theories, instead of theories to suit facts.'*

---

Arthur Conan Doyle

### 8.1 Introduction

Social interaction is a high bandwidth multimodal and highly complex process and computer-mediated environments further increase its complexity by imposing various types of uncertainty, as for example transmission delays and measurement noise, on the perception–action loop (see Figure 8.1). Network switching, traffic peaks and processing bottlenecks create variability in transmission delays. The longer the delays and higher the noise, the higher the uncertainty [54]. Human sensorimotor constraints further contribute to uncertainty by adding up to few hundreds milliseconds of reactive delay, which is further magnified by a time-variant decision making. This mostly unpredictable variability characteristic for continuous dynamic human–human interaction poses a great challenge for the coupling of corresponding user actions. The aim of this chapter is to propose a formal measure quantifying the level of interpersonal coordination in computer-mediated interaction and to emphasise the main impediments arising from the inherent uncertainty in the form of noise and delays. The information-theoretic model builds on the interactive paradigm introduced in Chapter 6 and the corresponding measure is validated on empirical data of Experiment II (Section 6.4).

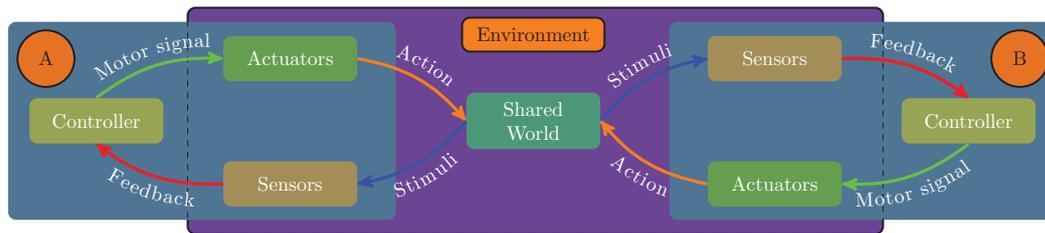


Fig. 8.1 Schematic diagram of a dyadic cognitive sensorimotor loop in a shared computer-mediated environment.

## 8.2 Background

Methods studying the dynamics of information shared between processes aim to detect the directionality of coupling and quantify the degree of asymmetry by assessing the interaction between two (sub-) systems at their outputs. Granger [32] determines the direction of causality between two variables by measures of causal lag and causal strength. Information theory provides a variety of concepts for the characterisation of information exchange between individual components in a system and the use of Shannon's mutual information to measure the overlap of information content between two systems is ubiquitous in this context [6, 114]. A particular interest lies in the identification of the 'flow of information' in a given system, for which typically variants of mutual information measures of correlative character are used [65, 86]. A measure of causal flow of information [6] captures the strength of a causal effect and realises a flow-like philosophy. The benefits of having such a measure are emphasised in [5, 49, 50, 113]. Transfer entropy [86] aims to quantify the information transfer between two systems by characterising the statistical coherence of the systems evolving in time. A non-parametric generalisation of Granger causality introduced in [64] is known as directed information.

Such measures could enable the quantification of a number of phenomena in the areas of synchronisation, game dynamics and the perception–action loop. For example, cooperative behaviour of coupled systems is related to a synchronisation phenomenon. When two players adapt their strategies over time the game dynamics could move towards cooperative or antagonistic behaviour and the flow of information could reveal a given player's contribution for the emergence of a particular cooperative or antagonistic strategy. Quantifying the correlation between the actions of interacting agents could serve as a measure of individuality or autonomy with respect to other agents, which is called intrinsic coordination in [36]. In a cooperative task under information processing constraints two agents arrive at intrinsic coordination in order to overcome limitations of their environment [36].

Without simplifying assumptions, however, the robust estimation of entropy-based functionals is notoriously difficult and usually requires a large number of data samples. Every method has its own free parameters and there is no consensus on an optimal way of estimating from a data set. Furthermore, measuring information-transfer-type quantities requires prior establishing the presence of and quantifying causal relationships as argued in [60], by applying, for example, information flow techniques as suggested in [6]. Revealing causal effects necessitates some type of perturbation or intervention of the source so as to detect the effect of that intervention on the destination.

The loose coupling of the agents in the collaborative scenarios introduced in Chapter 6 creates an asynchrony in the time series, which poses a challenge for a non-parametric information-theoretic analysis. As the emphasis of this chapter is on addressing this asynchrony, for simplicity I will use Shannon's mutual information to measure correlation, instead of the more complex methods cited above. The conditional variant of mutual information (Section 3.2) is defined for random variables  $X_A$  and  $X_B$  given  $X_S$  as follows

$$I_p(X_A : X_B | X_S) = \sum_{x_S} p(x_S) I_p(X_A : X_B | x_S), \quad (8.1)$$

where

$$I_p(X_A : X_B | x_S) = \sum_{x_A} p(x_A | x_S) \sum_{x_B} p(x_B | x_A, x_S) \log \frac{p(x_B | x_A, x_S)}{p(x_B | x_S)}.$$

Furthermore

$$I_p(X_A : X_B | X_S) \leq \min(H_p(X_A), H_p(X_B)), \quad (8.2)$$

providing an upper bound limited by the two entropies.

### 8.3 Model of Coordination

In order to get a deeper insight into the underlying human behaviour I applied a model-based approach for the characterisation of interpersonal coordination observed in the collaborative studies. The approach consists of defining a stochastic model, populating the model densities with content from the experimental data while making certain assumptions about how people coordinate, and computing the conditional mutual information from those densities. Figure 8.2 depicts a section of the causal Bayesian network representation of the perception-action loop unrolled over time, which is used in the formalisation of the model. This network specifies the causal relationships between both subjects' sensor states (visual and/or tactile stimuli  $s^a$  and  $s^b$ ), which are influenced by the respective actions (finger movements  $a$  and  $b$ ) through the environment (virtual membrane  $R$ ). Both subjects observe their current position,

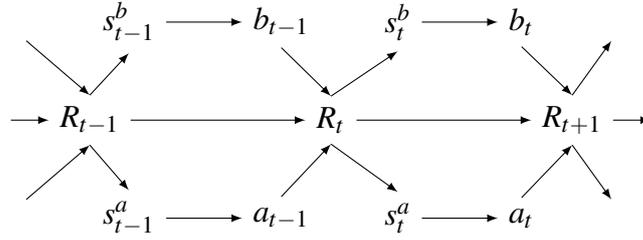


Fig. 8.2 Section of the causal Bayesian network representing the perception–action loop of a dyad ( $A - B$ ) interacting through the environment ( $R$ ) by applying actions ( $a$  and  $b$ ) in response to sensor stimuli ( $s^a$  and  $s^b$ ) unrolled over time.

keep an estimate of the distance from their partner (hereby called the error), and make action selection decisions among three discrete actions – stay in the current position or move up or down. The application of Equation 8.1 for computing the level of coordination requires a reliable estimate of the empirical probability density function from the collected experimental data. In order to ensure a consistent approximation of the conditional densities given the limited amount of data I defined low resolution action and state spaces of three elements each, derived from the sign functions below. Higher resolution spaces would require larger amounts of data to avoid sparsity and to provide a reliable empirical density. Assuming for simplicity that actions are influenced only by the error (i.e. distance between the pointers) denoted with the random variable  $X_S$  and the direction of motion is denoted with the random variables  $X_A$  and  $X_B$  respectively, let

$$X_S^t = \begin{cases} -1, & p_A^t \gg p_B^t \\ 0, & p_A^t \approx p_B^t \\ 1, & p_A^t \ll p_B^t \end{cases} \quad X_A^t = \begin{cases} -1, & p_A^t \gg p_A^{t+1} \\ 0, & p_A^t \approx p_A^{t+1} \\ 1, & p_A^t \ll p_A^{t+1} \end{cases} \quad X_B^t = \begin{cases} -1, & p_B^t \gg p_B^{t+1} \\ 0, & p_B^t \approx p_B^{t+1} \\ 1, & p_B^t \ll p_B^{t+1} \end{cases} \quad (8.3)$$

The random variables representing the actions and the error are defined with the sign functions specified by the Equations 8.3. The relation operators  $\approx$ ,  $\ll$  and  $\gg$  reflect the close proximity range used in the experiment, within which the subjects can perceive each other and refer to the relations between both subjects' positions  $p_A^t$  and  $p_B^t$  with a distance threshold of 20 pixels. The aim of this stochastic model is to capture the relationship between the random variables while keeping the complexity low. The joint probability distribution of  $X_A$ ,  $X_B$  and  $X_S$   $p(x_A, x_B, x_S)$  is estimated by counting the occurrences of the joint events in the data set.  $p(x_A)$ ,  $p(x_B)$  and  $p(x_S)$  are derived from  $p(x_A, x_B, x_S)$  by marginalisation and the conditionals  $p(x_A|x_S)$ ,  $p(x_B|x_S)$  and  $p(x_B|x_A, x_S)$  are obtained by applying the chain rule, which completes the list of densities required by Equation 8.1.

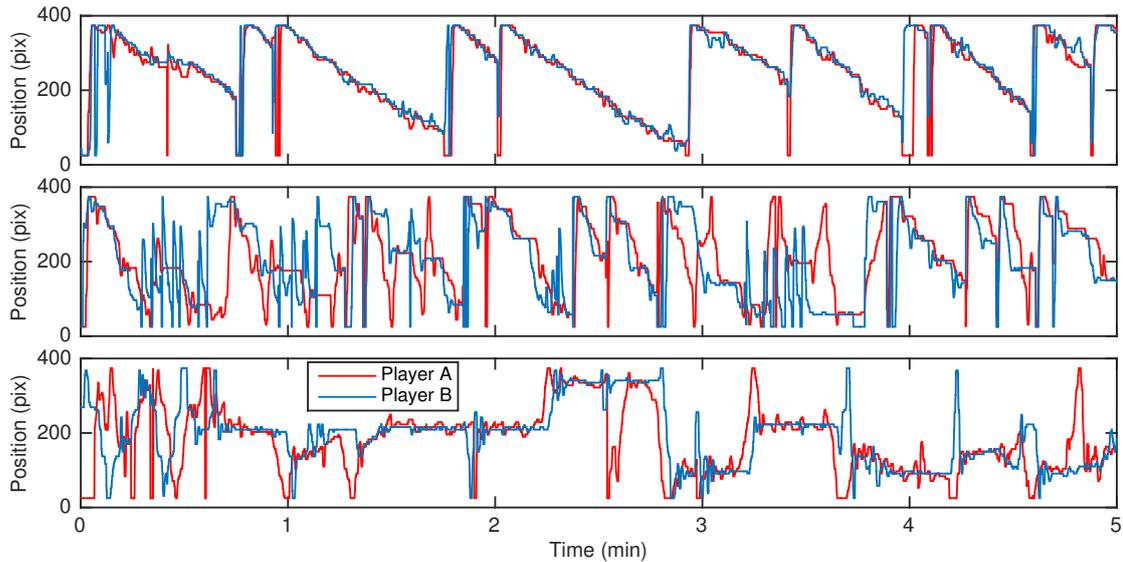


Fig. 8.3 The position time series recorded in the Combined condition of Experiment II representing three characteristic strategies – (top) tight tracking (pair F), (middle) loose tracking (pair G), and (bottom) random strategy (pair E).

## 8.4 Measure of Coordination

Qualitative analysis of data recorded in Experiment II, three characteristic examples of which are presented in Figure 8.3, suggests that the interaction consists of series of discrete messages, which poses a challenge for the stochastic model to infer from the ‘raw’ data set. To cancel the effect of noise I applied a moving average filter on the ‘raw’ data using a sliding window of two seconds, deriving a ‘filtered’ data set.

### 8.4.1 Time-shift Inference

In order to infer the time delay in reaction to the partner’s action attributed primarily to human decision making I reconstructed the velocity profiles from the ‘filtered’ data set in the form of sign functions reflecting the direction of motion – up, idle, down (i.e. -1,0,1 respectively), an example of which is presented in Figure 8.4/middle where one of the variables is scaled by half for improved visibility. I computed the standard cross correlation on these velocity profiles and present the resulting time-shift characteristics in Figure 8.4/bottom. The analysis reveals that the delay varies in the range of  $[-2.5, 2.5]$  seconds for this particular sample of close-tracking behaviour and rarely exceeds the five-second mark, which provides a confidence range for the window size to be used in the delay compensating dynamic time warping algorithm.

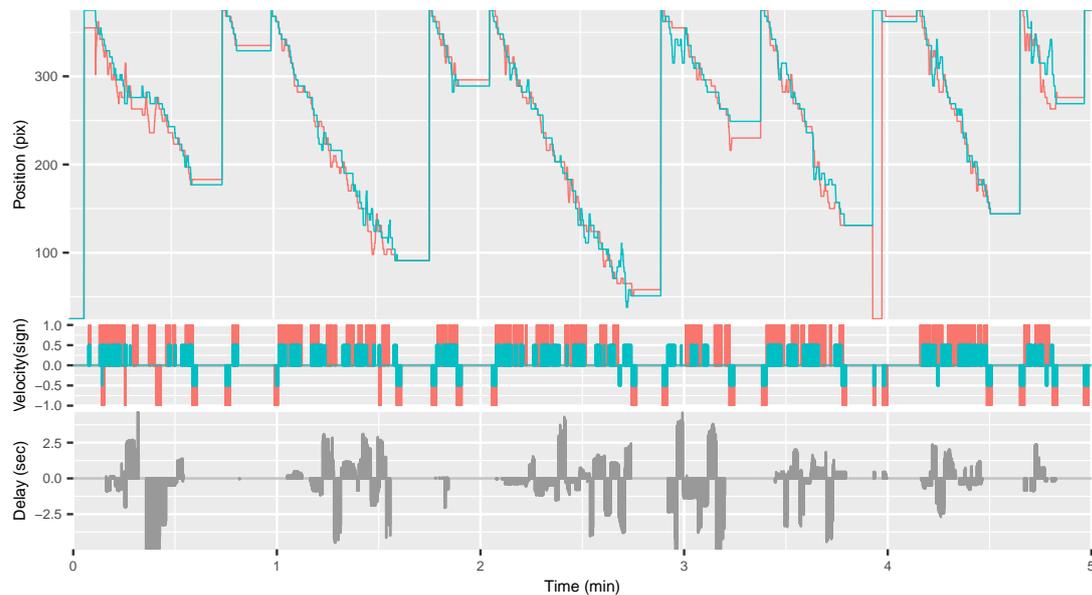


Fig. 8.4 Delay inference using cross correlation between the velocity profiles of both subjects. The raw time series of pair F (top), the velocity profiles (middle), and the inferred time-shift characteristics concentrated in the range of  $[-2.5, 2.5]$  seconds (bottom).

The densities of the inferred time-shift characteristics for all pairs in the Combined condition can be categorised in regard to close tracking into two categories as shown in Figure 8.5. This figure presents the distribution of delays over the  $[-5, 5]$  seconds range revealing short, regular delays on the left and longer, irregular delays on the right. The peak on the left represents close tracking performance, while on the right delays spread more uniformly. The leftmost bars in both charts represent the delays falling outside of the  $[-5, 5]$  seconds range. Seven out of 13 pairs fall into Category 1 and the rest into Category 2.

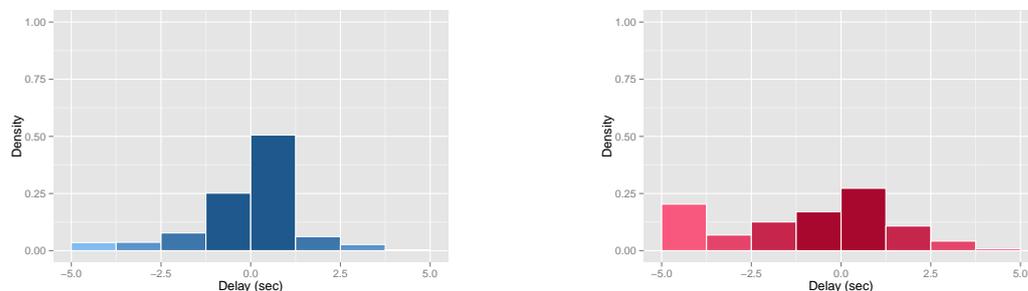


Fig. 8.5 Delay density profile implied performance categorisation of all 13 pairs in the Combined condition, of which seven fall into Category 1 – characterised by short, regular delays (left), and six fall into Category 2 – exhibiting long and irregular delays (right).

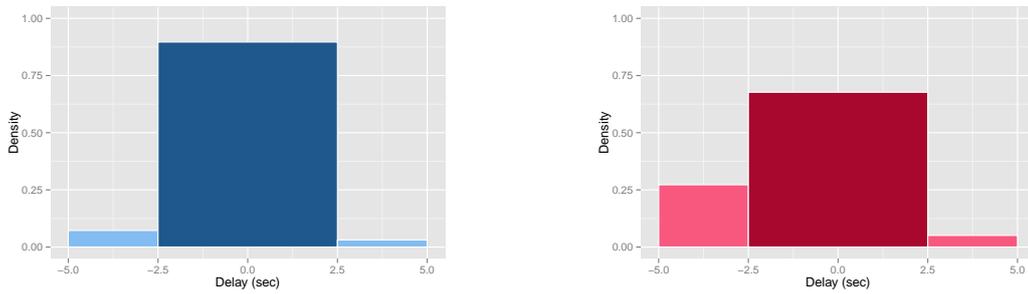


Fig. 8.6 Cumulative delay density profile implied performance categorisation of all 13 pairs in the Combined condition. Category 1 (left) consists of seven and Category 2 (right) of six pairs characterised by short, regular and long, irregular delays, respectively.

Figure 8.6 gives a further insight into the [-2.5,2.5] seconds range by aggregating densities from Figure 8.5. Figure 8.6/left provides a sufficient evidence for using the [-2.5, 2.5] seconds window range in the delay compensation algorithm.

Figure 8.7 presents the cumulative delay characteristics based on this categorisation for all three experimental conditions. The upper chart shows shorter and regular delays in the Visual and the Combined conditions and reveals that in the Tactile condition even good performers could not do better than the rest as the task difficulty increased significantly. The figures in the Combined condition show higher density of short delays than in the Visual condition, which is an evidence for the improved response time attributed to the additional (tactile) feedback modality. The bottom chart shows lower levels and a more uniform distribution across conditions, revealing that pairs from Category 2 have little variation in delay characteristics between the eyes-free and the visual-enabled conditions.

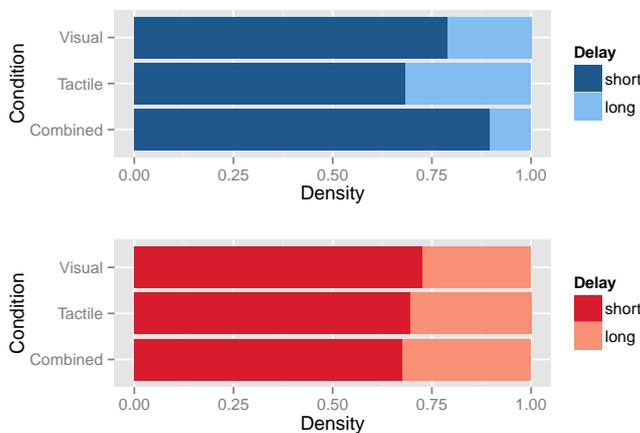


Fig. 8.7 Cumulative delay density profile for all three experimental conditions. Short delays refer to the [-2.5,2.5] seconds interval and long delays to larger levels. Category 1 (top) reveals higher density of short delays in the Combined than in the Visual condition, whereas Category 2 (bottom) shows lower levels and little variation across conditions.

### 8.4.2 Probability Distributions

The inferred delay threshold of 2.5 seconds between the correlated actions of both partners was used for delay-compensating the ‘filtered’ data set by applying a dynamic time warping algorithm and the resulting ‘delayed’ data set served as a basis for populating the empirical probability distribution functions  $p(x_A, x_B, x_S)$ . Figure 8.8 presents examples for both categories revealing the variability of such PDFs over the five-minute long sessions computed with a 60-seconds wide sliding window. The total number of random events according to the stochastic model is  $3^3 = 27$ , however only few of them show certain (limited) variability in the left chart (Category 1), which suggests consistent trends. The chart on the right shows less consistent and structured patterns as most of the probabilities vary significantly in a random fashion over time.

Figure 8.9 provides a further insight into the effect of the delay level used for dynamic time warping on the resulting empirical PDF. It presents stacked time-aligned series of raw data, empirical PDFs derived with delay levels of 2.5, 5 and 10 seconds respectively, and the corresponding levels of coordination (in this order from the top down). The figure reveals the effect the increase of delay has on the variability of the empirical PDFs, and how this translates into higher coordination levels. The three PDFs show similar patterns, however their structure gets sparser from the top down as the warping window (delay threshold) gets larger, which indicates the smoothing effect of the dynamic time warping algorithm. The patterns in the coordination levels show similar trends for different window sizes with quasi constant offsets, which suggests the consistency of the measure and offers the possibility for fine-tuning the appropriate value in the overall data analysis.

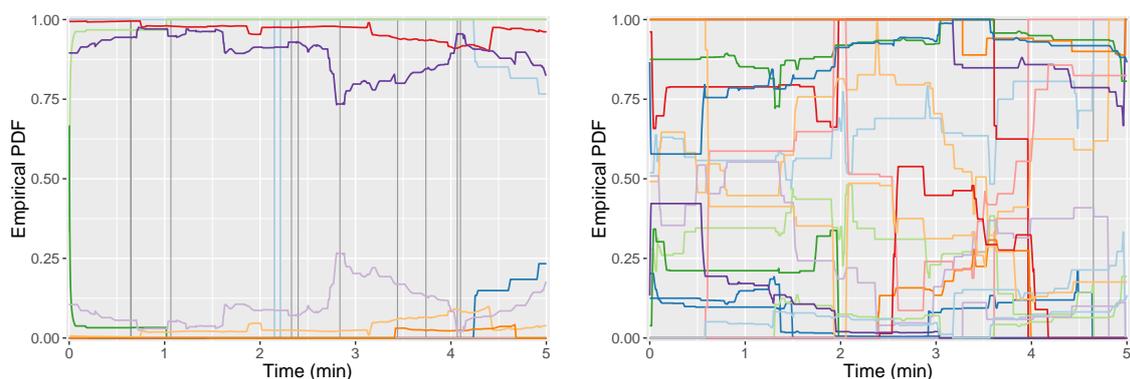


Fig. 8.8 Variability of the empirical probability distributions computed with a 60-seconds wide sliding window over the five-minute long sessions in the Combined condition – samples from Category 1 revealing minor fluctuations and consistent trends (left), and more chaotic and unstructured patterns for Category 2 (right).

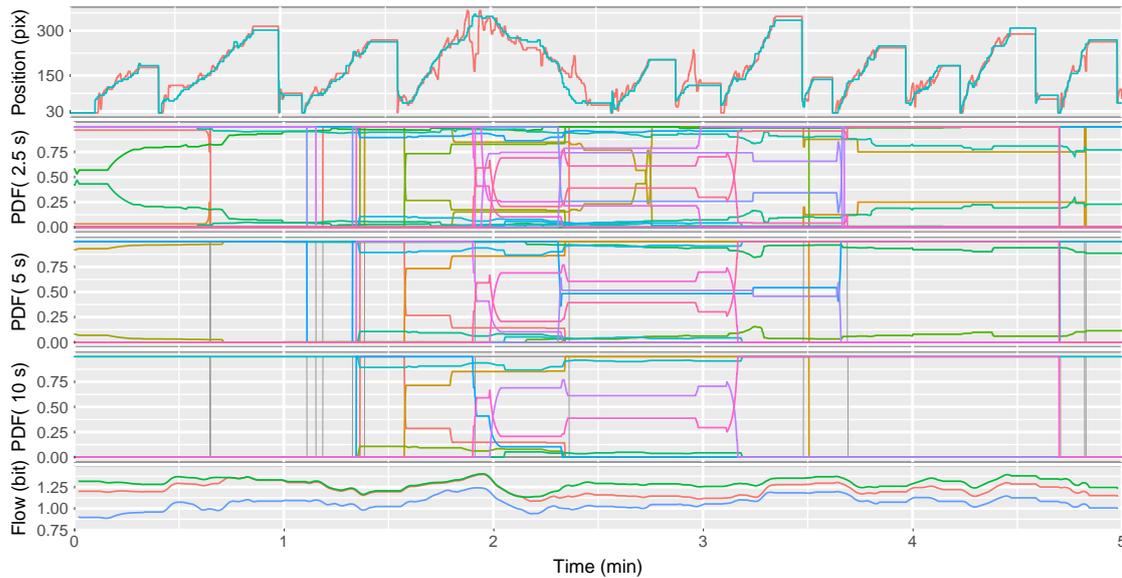


Fig. 8.9 An example of close tracking performance – raw time series (top), PDF variability corresponding to 2.5, 5 and 10 seconds of time-warping respectively, computed with a 60-seconds wide sliding window (middle) and the corresponding levels of coordination in blue, red and green respectively (bottom). The PDF structure becomes sparser from the top down reflecting the smoothing effect of the DTW algorithm. The coordination curves are similar with quasi constant offsets.

### 8.4.3 Delay Sensitivity

In order to explore the relation between coordination and delay compensation in more detail I generated the empirical PDFs and computed the corresponding coordination for several levels of delay and all pairs in the Combined condition. Figure 8.10 presents the resulting trade-off curves, which show an increasing monotonic trend and split the pairs into two distinct groups. The first group formed around the blue curve has higher levels and shows a steeper increase of coordination saturating at around five seconds delay, which provides a further evidence for the reasonable range of  $[-2.5, 2.5]$  seconds used for delay compensation in order to avoid over-smoothing of data for larger delays. Above that threshold the data variability is washed away and important information is lost. The second group formed around the green curve has lower levels and shows a slower increase lacking a clear sign of saturation even at the 25 seconds mark.

The range of  $[0, 5]$  seconds delay reveals for both groups consistent and compact curve segments. Larger delays introduce higher variability and lead to the emergence of a single outlier depicted in red. Delays longer than five seconds, however, do not fit the model representing tracking behaviour and the corresponding figures are not considered informative in this analysis. They are included only for the sake of completeness.

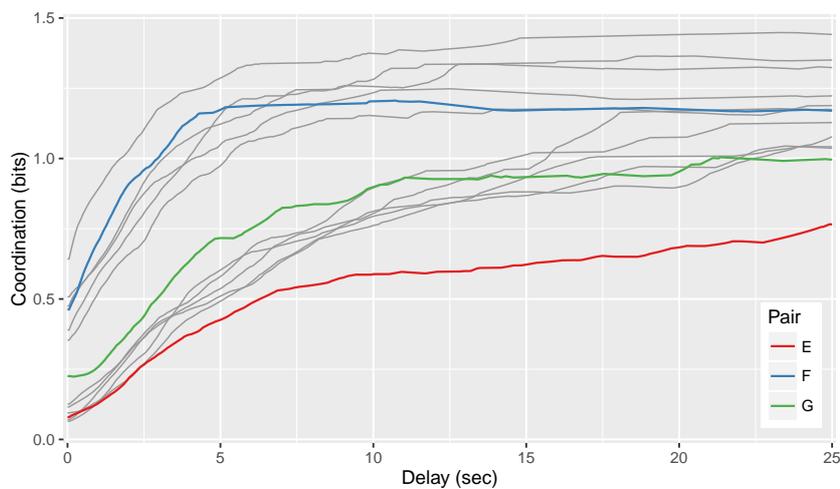


Fig. 8.10 Trade-off curves of the coordination level vs. delay compensation computed for all pairs in the Combined condition. The coordination levels and the saturation rate split the pairs into two distinct groups formed around the blue and the green curves. Larger delays introduce higher variability and a single outlier depicted in red.

#### 8.4.4 Noise Sensitivity

Further analysis reveals the sensitivity of the proposed measure to data quality. Figure 8.11 presents the coordination levels computed on the three data sets – ‘raw’, ‘filtered’ and ‘delayed’ – for three characteristic tracking patterns (delivered by the pairs F, G and E). As expected, the ‘raw’ data set provides low values of mutual information due to various sources of noise and delays diminishing the correlation in the time series (see Figure 8.3). Data smoothing provides higher levels of mutual information, calculated on the ‘filtered’ data set. Delay compensation yields a further increase computed on the ‘delayed’ data set, revealing the key role of delays associated with human motor control and decision making.

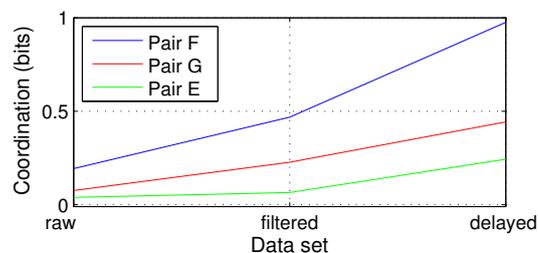


Fig. 8.11 Overall levels of coordination computed for pairs F, G and E (see Figure 8.3) from – (a) raw data, (b) noise-filtered data, and (c) delay-compensated noise-filtered data.

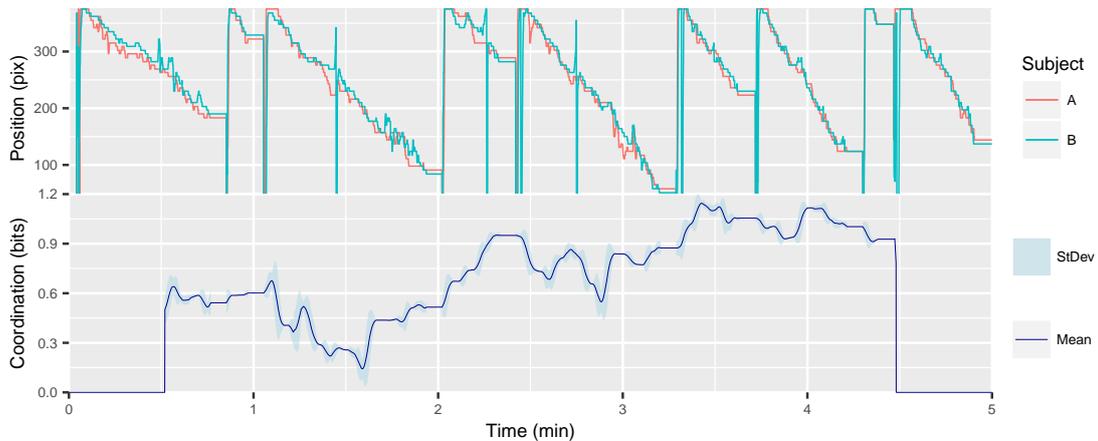


Fig. 8.12 Sensitivity of the coordination measure to data variability – raw data (top), corresponding coordination level mean and standard deviation over 50 consecutive values computed on 60-seconds wide sliding windows (bottom). Minor variations in action correlation can lead to a significant variability in the coordination levels.

These results show a direct correspondence to the three characteristic types of tracking behaviour presented in Figure 8.3. Pair F provides an example of a tight tracking performance throughout the trial session, which yields 0.98 bits of coordination. Pair G demonstrates loose tracking and sometimes erratic behaviour with longer and irregular delays, yielding 0.44 bits of coordination. Pair E executes a random strategy, which exhibits very little tracking behaviour and achieves 0.24 bits of coordination.

Figure 8.12/bottom reveals another aspect of the sensitivity of the measure to data variability by presenting the running means and standard deviations over 50 consecutive levels of coordination computed on 60-seconds wide sliding windows, while the corresponding raw data is time-aligned in the top. At this granularity level the analysis reveals that minor variations in the correlation of actions could result in a significant variability in the coordination levels. This reflects the fact that more consistent and reliable estimates of information-theoretic quantities usually require large amounts of data. The complete evaluation of coordination for all pairs across conditions is performed on the full sets of data corresponding to the five-minute long sessions and the results are presented in Section 8.5.

### 8.4.5 Simulations

In order to further explore the properties of this measure I investigated its behaviour on a range of simulated realisations of the proposed stochastic model. For this purpose I generated simulated PDFs as a combination of the three empirical F, G and E (noise-filtered and

delay-compensated) and five artificially created probability densities P4 – P8, defined on the same set of random variables  $\{X_A, X_B, X_S\}$ . The model of P5 corresponds to tightly-coupled controllers, P6 – P8 to different types of non-coordinated behaviour, and P4 to a mixture of the above. Using these eight models I defined a basis in  $(x, y)$  and generated 10000 alternating models in a two dimensional grid of  $100 \times 100$  resolution with the following linear interpolation. Given four joint densities  $q_1, q_2, q_3$  and  $q_4$  over  $\{X_A, X_B, X_S\}$  I defined a new density  $q(x_A, x_B, x_S)$  following Equations 8.4-8.6, and computed the corresponding level of coordination (see Figure 8.13).

$$r1 = x * q1 + (1 - x) * q2, \quad x \in [0, 1] \quad (8.4)$$

$$r2 = x * q3 + (1 - x) * q4, \quad x \in [0, 1] \quad (8.5)$$

$$q = y * r1 + (1 - y) * r2, \quad y \in [0, 1] \quad (8.6)$$

Figure 8.14 provides further insight into the sensitivity of the measure to changes in the underlying distributions by showing the Jensen-Shannon divergence between the three empirical (F, G and E) and all densities across the grid. In this figure the red lines connect

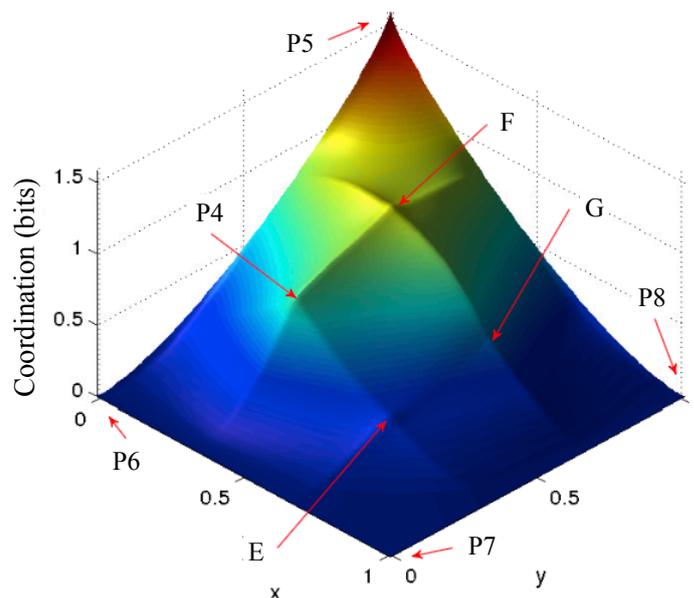


Fig. 8.13 The coordination levels computed on simulated PDFs generated with a linear interpolation on the three empirical (F, G, E) and five artificial (P4-P8) densities. P5 corresponds to tightly-coupled controllers yielding the theoretical maximum for this model of  $\log 3 = 1.58$  bits. P6 – P8 correspond to different types of non-coordinated behaviour.

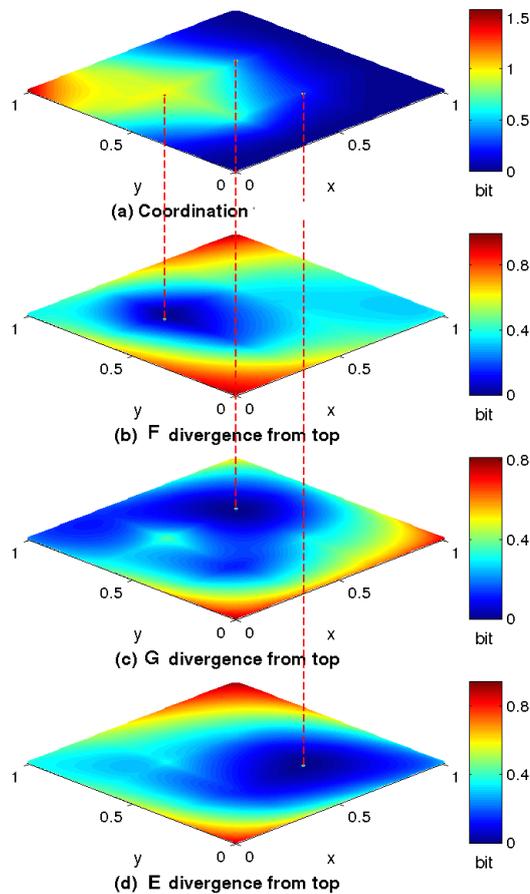


Fig. 8.14 Jensen–Shannon divergence between the empirical F, G and E and all densities in the grid. (a) Coordination levels (as in Figure 8.13); (b), (c) and (d) J–S divergence between F, G and E respectively and the corresponding densities in (a). The red lines connect F, G and E from (a) with the corresponding points of zero divergence on (b), (c) and (d) respectively.

the points of zero divergence on (b), (c) and (d) respectively to their corresponding empirical densities F, G and E in (a), providing a further perspective for the visualisation of the trends on the coordination curve.

## 8.5 Evaluation

This section provides the complete evaluation of coordination in all three conditions. The overall coordination levels computed on the full data sets per pair and condition (see Figure 8.15) are generally lower in the Tactile condition, particularly for pairs performing a close tracking otherwise. In the visual-enabled conditions good performers approach and one pair even exceeds the one bit mark, which can be considered a success given this short experiment and the theoretical upper bound of  $\log 3 = 1.58$  bits. Figure 8.16 presents the within-pair variability of the coordination levels computed on 60-seconds wide sliding windows with

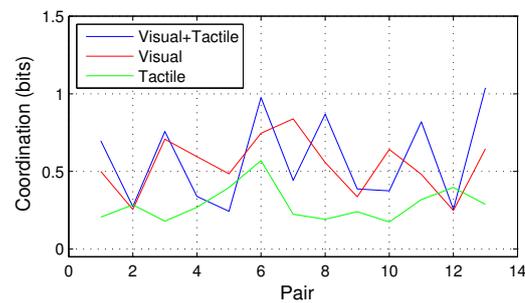


Fig. 8.15 The coordination levels of all pairs in all three conditions. The Visual and the Combined are generally higher than the Tactile condition.

mean (red dots) and median levels. The overall coordination levels computed on the full data sets are represented with yellow dots. A closer look at the correspondence between the yellow and the red dots leads to the following observations. In few cases the respective dots overlap, while in others they are apart from each other. All three relation types are theoretically possible as is reflected in this figure, nevertheless the range for both quantities is the same  $[0, \log 3]$  and the deviation depends on the data variability. A locally coherent data results in higher mean levels and lower global consistency decreases the overall coordination, which is generally the case in Figure 8.16. Nevertheless, the yellow dots are in the ballpark of the boxplots and are consistent with the trends. Figure 8.16 also shows that in the opposite

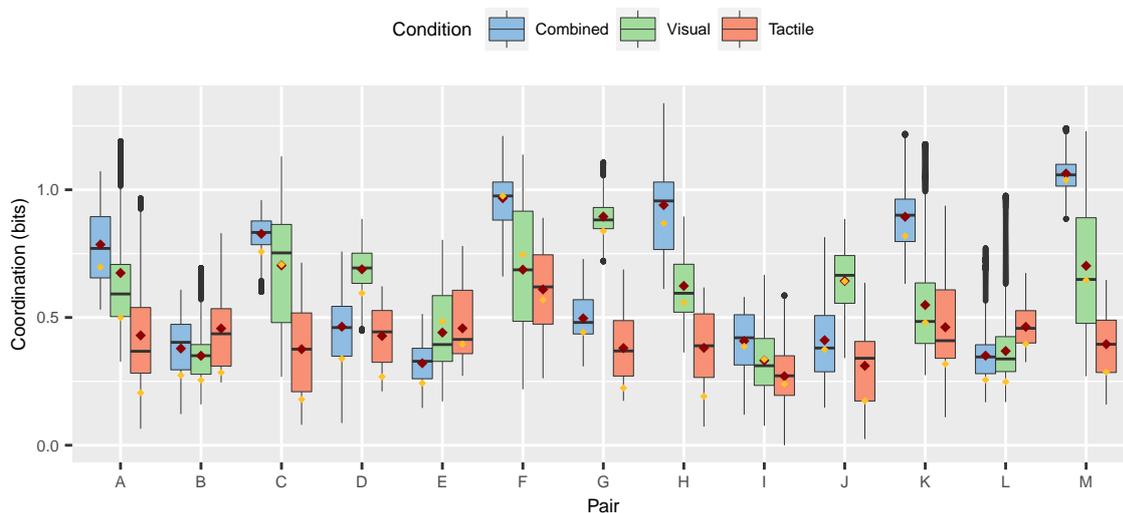


Fig. 8.16 Within-pair variability of the coordination levels for all pairs in the three experimental conditions computed on 60-seconds wide sliding windows – mean (red dots), and the overall coordination levels computed on the full data sets (yellow dots). A high level of coordination with a low variability reveals a consistently good performance.

case despite lower local coherency higher overall levels of coordination could be achieved (e.g. pair F). These deviations suggest that such measures must be investigated with care. In several cases high coordination is achieved with a relatively low variability (e.g. F, G, M), which is a sign of good performance throughout the session.

### 8.5.1 Analysis of Variance

A non-parametric Friedman test shows a significant effect of the type of feedback on the mean, median and overall level of coordination (p-values .05, .04 and .02 respectively), where the mean and the median levels are computed from values corresponding to 60-seconds wide sliding windows over each session and the between-pair variability is presented in Figure 8.17. Wilcoxon paired signed rank tests show a significant difference between the Combined and the Tactile, and between the Visual and the Tactile conditions for the mean (p-value = .05 with Bonferroni correction), median (p-value = .04 and .05 respectively with Bonferroni correction) and overall (p-value = .024 and .007 respectively with Bonferroni correction) levels of coordination. This is consistent with measured and perceived performance metrics, which suggests the coherence of the applied information-theoretic measure.

### 8.5.2 Learning Effect

The separate analysis of the Visual and the Combined conditions shows evidence of a learning effect, when the execution order of these two conditions is considered and the effect of the Tactile condition and the tactile feedback is disregarded (i.e. the tactile modality is excluded from consideration). Figure 8.18 presents the respective target-acquisition scores (bottom) and the corresponding overall coordination levels (top) for all pairs in the order of execution of these two conditions. The total scores show more pronounced improvement over time, nevertheless the coordination levels also suggest similar trends for most pairs reflected in

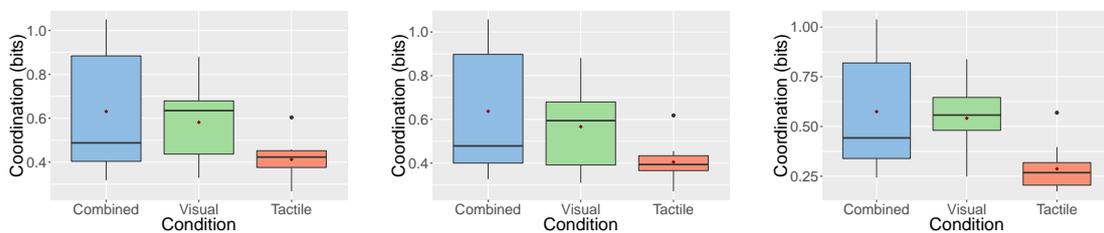


Fig. 8.17 Between-pair variability of the coordination levels for all pairs in the three experimental conditions computed on 60-seconds wide sliding windows – mean (left) and median (middle) levels. Overall coordination levels are computed on the full data sets (right). The Combined and the Visual are significantly higher than the Tactile condition.

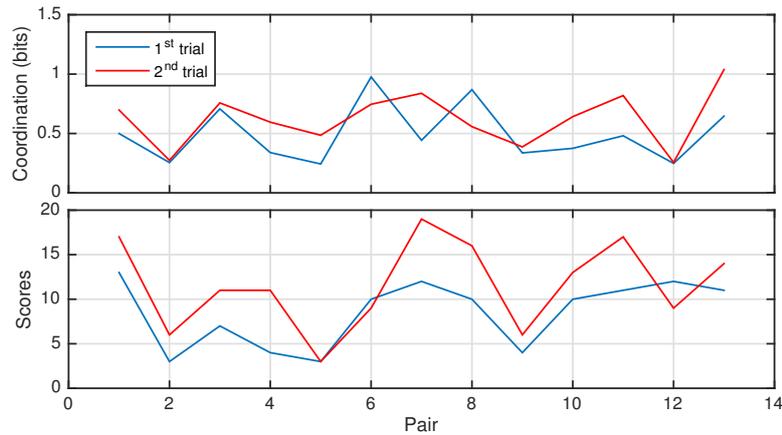


Fig. 8.18 Learning effect in the Visual and the Combined conditions in their respective trial order per pair – overall coordination levels (top) and numbers of targets acquired (bottom).

Table 8.1, which groups pairs by trends. Ten out of 13 pairs have kept or increased their levels with just one pair (H) achieving higher score with lower coordination, which may appear counter-intuitive, however is possible considering the respective strategy.

The learning trends are presented in more detail in Figure 8.19, which reveals the variability, the mean, the median and the overall coordination levels for all pairs. Wilcoxon paired signed rank tests show significant difference between the first and the second trial only for the overall coordination level (p-value = 0.04), presented in Figure 8.20/right, although

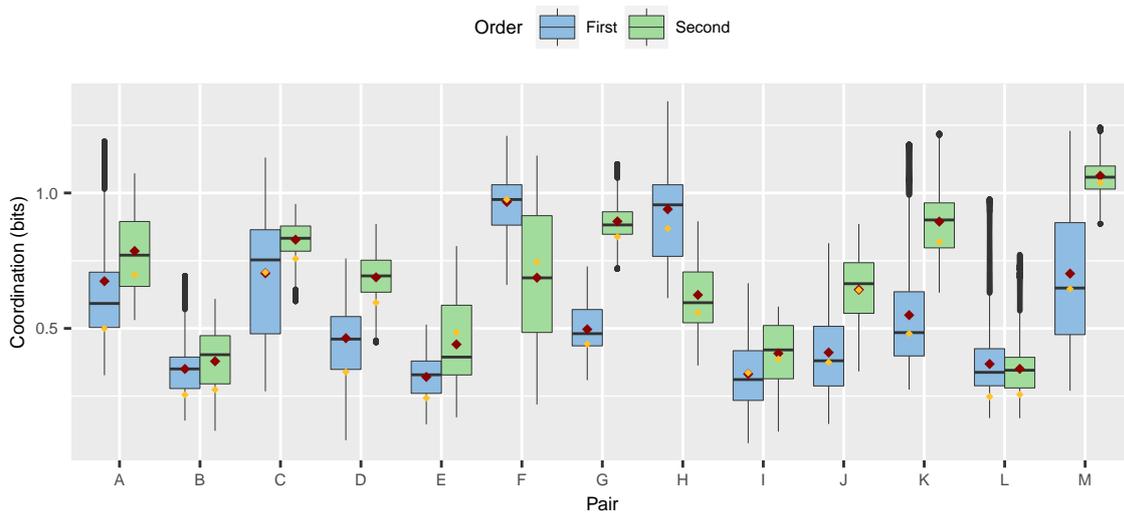


Fig. 8.19 Within-pair variability of the coordination levels in the Visual and the Combined conditions in their respective trial order computed on 60-seconds wide sliding windows – mean (red dots). Overall coordination levels are computed on the full data sets (yellow dots). A marked learning effect is visible for pairs C, D, G, K and M.

Table 8.1 Trends in the total scores and the overall coordination levels between the Visual and the Combined conditions in their respective trial order ('up' represents an increase from the 1<sup>st</sup> to the 2<sup>nd</sup> trial). Ten out of 13 pairs have kept or increased their levels.

Score \ Coordination	Coordination		
	up	level	down
up	A, D, G J, K, M	B, C, I	H
level	E	×	F
down	×	L	×

the mean – Figure 8.20/left (p-value = 0.08) and the median – Figure 8.20/middle (p-value = 0.07) levels have similar trends except for the higher variance. This suggests that the proposed model and coordination measure could be used also for evaluating the proficiency level of subjects in real time during the training phase preceding user studies.

## 8.6 Discussion

The in-depth analysis of the proposed approach demonstrated that the levels of coordination are higher for pairs engaged in a close tracking, e.g. pair F (see Figure 8.3a), and drop significantly for pairs who occasionally or more frequently disengage from tracking each other, e.g. pair G (see Figure 8.3b) suggesting the potential of this indicator to capture salient properties of the experimental data and infer the actual level of coordination. The applied measure reveals specific patterns of human behaviour and identifies structures in the interaction dynamics and as such it may not always have a direct correspondence to standard performance metrics as for example the total score. A very smooth and consistent tracking performance might be too slow in locating and acquiring targets as the partners are extremely careful to not lose contact with each other and thus achieve a lower total score (e.g. pair F).

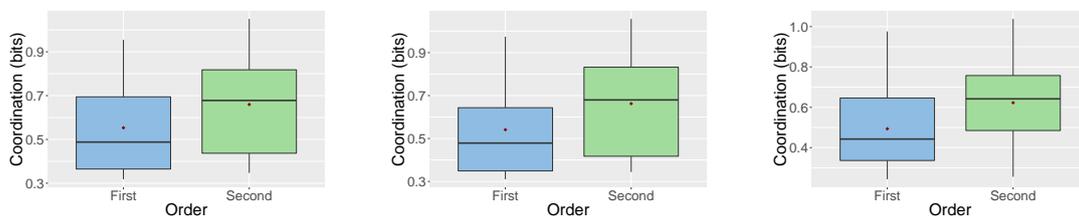


Fig. 8.20 Between-pair variability of the coordination levels in the Visual and the Combined conditions in their respective trial order computed on 60-seconds wide sliding windows – mean (left) and median (middle) levels. The first and the second trial are significantly different only for the overall coordination levels computed on the full data sets (right).

An alternative strategy might be to systematically jump from one target to the next without paying too much attention about the constant contact with the partner, which might result in a less consistent tracking, but a faster target acquisition leading to a higher total score (e.g. pair G). Having no strategy at all or at least not one qualifying as tracking could still achieve a relatively high total score by, for example, randomly jumping from one target to another. Future work could decouple the dyad and measure the flow of influence between the players. The trajectories of pair G (see Figure 8.3b) reveal patterns of leader–follower behaviour reflecting a command-and-control interaction style, while those of pair F (see Figure 8.3a) suggest a turn-taking leadership resembling a smooth dance, in which the leader and the follower roles are not clearly defined.

This work focused on developing a low-level perceptual model of tracking behaviour, which could be extended in the future with higher level mental models of coordination. Patterns of higher level decision making could be identified in the experimental data, e.g. when the subjects reach the end of the membrane and jump to the other end interrupting the continuous tracking.

## 8.7 Conclusion

This chapter demonstrated the potential of the proposed information-theoretic approach to characterise interpersonal coordination from empirical data. The results suggest that the utilised measure of coordination can be used for evaluating the proficiency level of subjects in real time, e.g. during the training preceding user studies. Such a measure could facilitate the assessment of cooperative behaviour in coupled systems and could support making inferences about transitions of strategies towards cooperative or antagonistic behaviour over time. Detailed sensitivity analysis subject to various factors is provided and specific key properties of the measure are highlighted, which require special attention in order to ensure reliable estimates.

Applying this information-theoretic approach, however, requires certain amount of prior modelling, which contributes to the evaluation cost. The cost of modelling is largely influenced by the quality of data, to which factors such as noise, delays, sensing and feedback resolution all contribute by implicating the involvement of advanced data enhancing methods in the modelling process.

The coordination levels of the analysed empirical data set reveal expected trends, which along with the theoretical coherence of the framework suggest its potential as part of a future toolset for understanding interactive systems. However, further work is required to expand and validate the approach in other domains.

## **Part IV**

# **Conclusions**



# Chapter 9

## Conclusions and Discussion

*‘Science never solves a problem without creating ten more.’*

---

George Bernard Shaw

This thesis presents an information-theoretic account of one specific aspect of human behaviour – stochastic manual control – by tapping into the perception–action loop of dynamic interactive systems. The premise of the thesis rests on the assumption that information theory could provide a calculus for identifying and formulating fundamental principles guiding biological perception–action loops. The information-theoretic perspective is universal, general and theoretically grounded and provides analytical tools for modelling the bounded-rational human agent in uncertain, dynamic and mobile contexts, which is a great challenge for user interface designers. The thesis provides proof-of-concept examples for the application of three information-theoretic utility measures characterising human behaviour in dynamic human–machine and human–human interactive scenarios in stochastic environments and reveals the relation of these measures to traditional metrics used in HCI research.

### 9.1 Contributions of the Thesis

The main contributions of this thesis are the adaptations of three generic information-theoretic concepts – empowerment, relevant information and mutual information (Chapters 4, 5 and 8 respectively) – to continuous manual control problems and their applications for the characterisation of specific aspects of human performance. The extensions of the original information-theoretic frameworks accommodate for the inherent lags in the perception–action loop characteristic for human sensorimotor control, decision making and other technology-related factors and provide more realistic models of human–machine couplings.

### 9.1.1 Empowerment in HCI

Building on the principle of empowerment I proposed a measure for quantifying the uncertainty in dynamic interactive systems and suggested its potential in making predictions and providing theoretical bounds on standard metrics of human performance based solely on properties of the environment. I developed a formal extension of the original empowerment framework, which captures uncertainty attributed to various factors common for interactive systems (e.g. noise, delays, errors, etc.) in a generic theoretical measure reflecting the level of control a user has over the environment in the course of interaction. This expands the capability of empowerment to model interactive tasks, in which delays are inevitable, and enables its application in the domain of human–computer interaction. I developed a parametric empowerment model characterising human performance in general in respect to a particular user interface and demonstrated how empowerment could facilitate the theoretical evaluation of system’s usability in various environmental conditions as it only requires the probabilistic model of the system’s feedback loop. Furthermore, I developed an empowerment model providing a particular type of natural human tracking behaviour in collaborative interaction through empowerment maximisation and derived empirical values of empowerment for individual users corroborating specific subjective metrics. In this thesis I presented empowerment both as (1) a theory about the user’s objective function and (2) a measure of the utility of interaction design.

### 9.1.2 Information Parsimony

The exploration of the informational cost in the perspective of a performance-oriented study by way of example for the application of relevant information to the domain of manual control suggested its potential to characterise and model human behaviour and provide an insight into the structure of behaviour and information processing. Applying the concept of relevant information, which provides bounds on the information required for decision making in order to achieve a certain level of utility I developed a reinforcement learning model of a dynamic interaction task and derived information-theoretic limits of human performance and a model of a particular pattern of human behaviour. This concept builds on the hypothesis that the ability to trade off value and informational cost lies at the core of natural behaviour. It adopts the view that biological systems implement an information parsimony principle, i.e. they achieve a given level of performance at the lowest informational cost possible. This hypothesis implies that humans attempt to realise valuable behaviours while minimising the cognitive cost and may resort to suboptimal solutions if they save a significant amount of cognitive resources.

### 9.1.3 Coordination of Behaviour

Building on a novel minimalist interaction paradigm for non-verbal remote collaboration I developed a framework for characterising dyad couplings in the scope of shared computer-mediated environments. The proposed stochastic model of interaction compensates for reaction delays and discrete patterns characteristic for human behaviour and facilitates the application of Shannon's mutual information as a measure of interpersonal coordination. An in-depth sensitivity analysis suggests the potential of the approach to characterise mutual entrainment and synchrony between participants and provide insights about the smoothness of interaction as a benchmark of a user interface.

## 9.2 Scope and Limitations

The studies presented in this thesis are of exploratory nature and provide conceptual examples, which could serve as a guideline for HCI practitioners. For this purpose the continuous interactive tasks were represented with low-granularity discrete models building on certain simplifying assumptions, which inevitably affect the model accuracy, however keep the complexity of the presentation low for readers not adept with information theory. Furthermore, the discrete versions of the applied information-theoretic concepts are more extensively studied and better understood than their continuous counterparts. Therefore, the extent to which these models make explain data is limited to the general trends and require further refinements to provide higher precision. While the relevant information model fits one particular type of human behaviour, the empowerment model aims to predict human performance in general.

The stochastic model of coordination was informed by a qualitative analysis of empirical data and built on specific assumptions of how people interact with each other. This model could provide a formal account of negotiated strategies building on the notion of mutual engagement and could replace qualitative categorisation. The comprehensive analysis of the proposed measure of coordination revealed its sensitivity to various factors emphasising that such quantities require careful examination to mitigate the risk of misinterpretation of results.

## 9.3 Future Work

### 9.3.1 Continuous Models

Exploring the proposed empowerment model for systems with time-varying irregular delays in larger empirical studies is a direct extension of this work. Furthermore, more realistic models of dynamic interaction could benefit from the continuous variants of the empowerment

and the relevant information algorithms. However, since these approximation methods rely on Monte Carlo sampling their accuracy is a function of the computational complexity. A recently developed method provides fast continuous approximation of empowerment by making certain Gaussian and linearity assumptions.

### **9.3.2 Discrete Interaction**

An important extension of this work will be the application of empowerment to other domains of HCI, e.g. to the traditional command-and-control discrete interaction style. Empowerment has been extensively studied for discrete problems and the construction of such interactive models should be straightforward. This would provide more accurate models at a relatively lower cost compared to dynamic interactive scenarios.

### **9.3.3 Non-parametric Methods**

Various non-parametric estimators of information-theoretic functionals exist, providing numeric representation of empirical data, e.g. mutual information, transfer entropy or directed information. The application of such model-free methods on observed data is typically straightforward, assuming satisfactory quality of data, however their causal implications need a careful interpretation. When used properly such measures could provide an insight into the dynamics of information transfer in complex systems and could identify turn-taking and leader-follower patterns as well as discern cooperative from antagonistic behaviour.

## **9.4 Final Remarks**

This work provides an important bridge between theory and experiments. It highlights the notion of empowerment as instrumental in explaining empirical data and as a criterion in user interface design optimisation. One of the aims of the thesis is to raise the awareness of the research community about the potential empowerment has in contributing to the more solid theoretical foundations for the science of HCI. The results demonstrate how information-theoretic measures can be treated naturally side-by-side along traditional metrics used in HCI research. Employing entropy-based information-theoretic principles to address uncertainty describes more appropriately the variability in human behaviour. The adoption of such utility measures could foster the foundation of a more solid theoretical framework for the study of HCI as well as provide a number of practical benefits. They could help designers treat and evaluate interactive systems in a general fashion and could augment current usability studies improving quality of design while at the same time reducing risk and evaluation costs.

# References

- [1] Annett, J. (1969). *Feedback and human behavior*. Penguin Harmondsworth.
- [2] Anthony, T., Polani, D., and Nehaniv, C. (2008). On preferred states of agents: how global structure is reflected in local structure. In *Proceedings of the eleventh international conference on the simulation and synthesis of living systems*, pages 25–32. MIT Press.
- [3] Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall Ltd.
- [4] Auvray, M., Lenay, C., and Stewart, J. (2009). Perceptual interactions in a minimalist virtual environment. *New Ideas in Psychology*, 27:32–47.
- [5] Ay, N. and Krakauer, D. (2007). Geometric robustness theory and biological networks. *Theory in biosciences*, 125(2):93–121.
- [6] Ay, N. and Polani, D. (2008). Information flows in causal networks. *Advances in Complex Systems*, 11(1):17–41.
- [7] Bar-Hillel, Y. (1964). *Language and Information*. Addison-Wesley, Reading, Mass.
- [8] Baren, J., IJsselsteijn, W., Markopoulos, P., Romero, N., and Ruyter, B. (2004). Measuring affective benefits and costs of awareness systems supporting intimate social networks. In: *Nijholt, Nishida (Eds.), 2004. Social Intelligence Design*, pages 13–19.
- [9] Basdogan, C., Ho, C., Srinivasan, M., and Slater, M. (2000). An experimental study on the role of touch in shared virtual environments. *ACM TOCHI*, 7(4):443–460.
- [10] Bellman, R. E. (1957). *Dynamic Programming*. Princeton University Press.
- [11] Benford, S., Greenhalgh, C., Rodden, T., and Pycock, J. (2001). Collaborative virtual environments. *Commun. ACM*, 44(7):79–85.
- [12] Bertsekas, D. P. (1976). Dynamic programming and stochastic control. *Mathematics in Science and Engineering, Volume 125*.
- [13] Bialek, W., de Ruyter van Steveninck, R. R., and Tishby, N. (2007). Efficient representation as a design principle for neural coding and computation. *arXiv.org:0712.4381*.
- [14] Bialek, W., Nemenman, I., and Tishby, N. (2001). Predictability, complexity and learning. *Neural Computation*, 13:2409–2463.
- [15] Blahut, R. (1972). Computation of channel capacity and rate distortion functions. *IEEE Transactions on Information Theory*, 18(4):460–473.

- [16] Breazeal, C. (2002). *Designing sociable robots*. MIT Press, Cambridge, MA.
- [17] Brewster, S. A. (1994). Providing a structured method for integrating non-speech audio into human–computer interfaces. *Ph.D. thesis, Department of Computer Science, University of York*.
- [18] Cao, Y., Theune, M., and Nijholt, A. (2009). Modality effects on cognitive load and performance in high-load information presentation. *In Proceedings of the 13th international conference on Intelligent user interfaces*, pages 335–344.
- [19] Capdepuy, P., Polani, D., and Nehaniv, C. (2007). Maximization of potential information flow as a universal utility for collective behaviour. *In Proceedings of the First IEEE Symposium on Artificial Life*, pages 207–213.
- [20] Chen, X. and Howes, A. (2012). A reinforcement learning model of bounded optimal strategy learning. *In Proceedings of the 11th International Conference on Cognitive Modelling*. TU Berlin.
- [21] Clark, H. and Brennan, S. (1991). Grounding in communication. *In: Resnick, Levine, Teasley (Eds.), 1991. Perspectives on socially shared cognition*, pages 127–149.
- [22] Clark, H. H. (1996). *Using Language*. Cambridge University Press.
- [23] Cockton, G. (2004). From quality in use to value in the world. *In Proceedings of CHI*, pages 1287–1290.
- [24] Cover, T. M. and Thomas, J. A. (1991). *Elements of Information Theory*. Wiley, New York.
- [25] D’Este, C. (2004). Sharing meaning with machines. In Berthouze, L., Kozima, H., G.Prince, C., Sandini, G., Stojanov, G., Metta, G., and Balkenius, C., editors, *Proceedings of the Fourth International Workshop on Epigenetic Robotics*, pages 111–114. Lund University Cognitive Studies.
- [26] Dourish, P. and Bellotti, V. (1992). Awareness and coordination in shared workspaces. *In Proc. CSCW ’92*, pages 107–114.
- [27] Fitts, P. M. (1954). The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology*, 47(6):381–391.
- [28] Fuster, J. M. (2001). The prefrontal cortex - an update: Time is of the essence. *Neuron*, 30:319–333.
- [29] Fuster, J. M. (2006). The cognit: A network model of cortical representation. *International Journal of Psychophysiology*, 60(2):125–132.
- [30] Galantucci, B. (2005). An experimental study of the emergence of human communication systems. *Cognitive Science*, 29(5):737–767.
- [31] Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. Houghton Mifflin Company, Boston.

- [32] Granger, C. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3):424–438.
- [33] Greenberg, S. and Buxton, B. (2008). Usability evaluation considered harmful (some of the time). In *Proceedings of CHI*, pages 111–120.
- [34] Greenberg, S. and Thimbleby, H. (1992). The weak science of human-computer interaction. In *Proceedings of CHI*.
- [35] Guiard, Y. and Beaudouin-Lafon, M. (2004). Special issue: Fitts' law 50 years later: Applications and contributions from human-computer interaction. *International Journal of Human-Computer Studies*, 61(6).
- [36] Harder, M., Polani, D., and Nehaniv, C. L. (2010). Two agents acting as one. In *In Proc. Artificial Life*, pages 599–606.
- [37] Hart, S. G. and Staveland, L. E. (1988). Development of nasa-tlx: Results of empirical and theoretical research. *Hancock, Meshkati, (Eds.), Human Mental Workload*, pages 139–183.
- [38] Hebb, D. (1964). The american revolution. *The Cognitive Processes: Readings*.
- [39] Hobson, P. (2002). *The cradle of thought: Challenging the origins of thinking*. Macmillan, London.
- [40] Hoggan, E. and Brewster, S. (2007). Designing audio and tactile crossmodal icons for mobile devices. In: *Proceedings of 12th International Conference on Multimodal Interfaces*, pages 162–169.
- [41] Howard, R. A. (1966). Information value theory. *IEEE Transactions on Systems Science and Cybernetics*, 2(1):22–26.
- [42] Howes, A., Lewis, R. L., and Vera, A. (2007). Bounding rational analysis. *Integrated Models of Cognitive Systems*, page 403.
- [43] Jagacinski, R. J. and Flach, J. M. (2003). *Control theory for humans: quantitative approaches to modeling performance*. L. Erlbaum Associates, Mahwah, N.J.
- [44] Jordan, J., Mortensen, J., Oliveira, M., Slater, M., Tay, B., Kim, J., and Srinivasan, M. (2002). Collaboration in a mediated haptic environment. In *Proc. PRESENCE*.
- [45] Jung, T., Polani, D., and Stone, P. (2011). Empowerment for continuous agent-environment systems. *Adaptive Behavior*, 19(1):16–39.
- [46] Kaplan, F. and Hafner, V. (2006). The challenges of joint attention. *Interaction Studies*, 7(2):135–169.
- [47] Kelley, C. R. (1968). *Manual and automatic control*. Wiley, New York.
- [48] Kelso, S. (1997). *Dynamic Patterns: The Self-Organization of Brain and Behavior*. Complex Adaptive Systems.

- [49] Klyubin, A., Polani, D., and Nehaniv, C. (2005a). *Empowerment: A universal agent-centric measure of control*. In *Proc. IEEE congress on Evolutionary Computation*, pages 128–135.
- [50] Klyubin, A. S., Polani, D., and Nehaniv, C. L. (2004). Organization of the information flow in the perception-action loop of evolved agents. In *In Proceedings of 2004 NASA/DoD Conference on Evolvable Hardware*, pages 177–180. IEEE Computer Society.
- [51] Klyubin, A. S., Polani, D., and Nehaniv, C. L. (2005b). All else being equal be empowered. *Advances in artificial life*, 3630:744–753.
- [52] Klyubin, A. S., Polani, D., and Nehaniv, C. L. (2007). Representations of space and time in the maximization of information flow in the perception-action loop. *Neural Computation*, 19(9):2387–2432.
- [53] Klyubin, A. S., Polani, D., and Nehaniv, C. L. (2008). Keep your options open: An information-based driving principle for sensorimotor systems. *PLoS ONE*, 3(12).
- [54] Kohrs, C., Angenstein, N., and Brechmann, A. (2016). Delays in human-computer interaction and their effects on brain activity. *PLoS ONE*, 11(1).
- [55] Körding, K. P. and Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, 427:244–247.
- [56] Laughlin, S. B. (2001). Energy as a constraint on the coding and processing of sensory information. *Current Opinion in Neurobiology*, 11:475–480.
- [57] Lázaro-Gredilla, M., Quiñonero-Candela, J., Rasmussen, C. E., and Figueiras-Vidal, A. R. (2010). Sparse spectrum gaussian process regression. *Journal of Machine Learning Research*, 11:1865–1881.
- [58] Lemmelä, S., Vetek, A., Mäkelä, K., and Trendafilov, D. (2008). Designing and evaluating multimodal interaction for mobile contexts. In *Proceedings of the 10th International Conference on Multimodal Interfaces (Crete, Greece), ICMI’08*, pages 265–272.
- [59] Lewis, R. L., Howes, A., and Singh, S. (2014). Computational rationality: Linking mechanism and behavior through bounded utility maximization. *Topics in Cognitive Science*, 6(2):279–311.
- [60] Lizier, J. T. and Prokopenko, M. (2010). Differentiating information transfer and causal effect. *Eur. Phys. J. B*, 73(4):605–615.
- [61] MacKay, D. J. C. (2003). *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press.
- [62] MacKenzie, I. S. and Ware, C. (1993). Lag as a determinant of human performance in interactive systems. In *INTERCHI*, pages 488–493.
- [63] MacKenzie, S. I. (1992). Fitts’ law as a research and design tool in human-computer interaction. *Human-Computer Interaction*, 7:91–139.
- [64] Massey, J. L. (1990). Causality, feedback and directed information. In *Proceedings of the International Symposium on Information Theory and its Applications*, pages 303–305.

- [65] Matsumoto, K. and Tsuda, I. (1988). Calculation of information flow rate from mutual information. *J. Phys. A: Math. Gen.*, 21(6):1405.
- [66] McRuer, D. and Jex, H. (1967). A review of quasilinear pilot models. *IEEE Trans. Human Factors Elect.*, pages 231–249.
- [67] Murray-Smith, R. (2009). Empowering people rather than connecting them. *International Journal of Mobile HCI*, 1(3):18–28.
- [68] Norman, D. A. (1988). *The Design of Everyday Things*. Basic Books, New York, USA.
- [69] Oakley, I., Brewster, S., and Gray, P. (2001). Can you feel the force? an investigation of haptic collaboration in shared editors. *In Proceedings of Eurohaptics*, pages 54–59.
- [70] Oaksford, M. and Chater, N. (2007). *Bayesian Rationality: The Probabilistic Approach to Human Reasoning*. Oxford University Press.
- [71] O’Regan, J. K. and Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behav. Brain Sci.*, 24:939–973.
- [72] Oviatt, S. (2006). Human-centered design meets cognitive load theory: designing interfaces that help people think. *In Proceedings of the 14th annual ACM international conference on Multimedia, New York, NY, USA*, pages 871–880.
- [73] Pearl, J. (1995). Causal diagrams for empirical research. *Biometrika*, 82(4):669–710.
- [74] Pearl, J. (2000). *Causality: Models, Reasoning, and Inference*. Cambridge University Press.
- [75] Polani, D. (2009). Information: Currency of life? *HFSP Journal*, 3(5):307–316.
- [76] Polani, D., Nehaniv, C., Martinetz, T., and Kim, J. T. (2006). Relevant information in optimized persistence vs. progeny strategies. *In Artificial Life X: Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems*, pages 337–343. MIT Press.
- [77] Poulton, E. C. (1974). *Tracking Skill and Manual Control*. Academic Press, New York.
- [78] Powers, W. T. (1973). *Behavior: The control of Perception*. Aldine de Gruyter.
- [79] Rasmussen, C. E. and Williams, C. K. I. (2006). *Gaussian Processes for Machine Learning*. MIT Press.
- [80] Reeves, L. M., Lai, J., and Larson, J. A. (2004). Guidelines for multimodal user interface design. *Communications of the ACM*, 47(1):57–59.
- [81] Rieke, F., Warland, D., van Steveninck, R., and Bialek, W. (1997). *Spikes: Exploring the neural code*. MIT Press.
- [82] Robinson, S., Eslambolchilar, P., and Jones, M. (2008). Point-to-geoblog: gestures and sensors to support user generated content creation. *In: 10th international Conference on Human Computer interaction with Mobile Devices and Services*, pages 197–206.

- [83] Rosenblum, M. and Pikovsky, A. (2001). Detecting direction of coupling in interacting oscillators. *Physical Review Series E*, 64(4):045202.
- [84] Sallnäs, E. (2005). Effects of communication mode on social presence, virtual presence, and performance in collaborative virtual environments. *Presence*, 14(4):434–449.
- [85] Sallnäs, E. (2010). Haptic feedback increases perceived social presence. *LNCS*, 6192:178–185.
- [86] Schreiber, T. (2000). Measuring information transfer. *Phys. Rev. Lett.*, 85(2):461–464.
- [87] Shannon, C. E. (1949). The mathematical theory of communication. In Shannon, C. E. and Weaver, W., editors, *The Mathematical Theory of Communication*. The University of Illinois Press, Urbana.
- [88] Sheridan, T. (2000). Interaction, imagination and immersion: Some research needs. In *Proceedings of Virtual Reality Software and Technology*, pages 1–7.
- [89] Sheridan, T. (2002). Some musings on four ways humans couple: Implications for systems design. *IEEE Transactions On Systems, Man, And Cybernetics*, 32(1):5–10.
- [90] Spence, C. (2002). Multisensory attention and tactile information-processing. *Behavioral Brain Research*, 135, pages 57–64.
- [91] Steels, L. and Kaplan, F. (2001). Aibo’s first words. the social learning of language and meaning. *Evolution of Communication*, 4(1):3–32.
- [92] Steels, L., Kaplan, F., McIntyre, A., and Looveren, J. (2002). Crucial factors in the origins of word-meaning. In Wray, A., editor, *The Transition to Language*. Oxford University Press.
- [93] Stefanovska, A. and Palus, M. (2003). Direction of coupling from phases of interacting oscillators: An information-theoretic approach. *Physical Review E*, 67(5):055201.
- [94] Steuer, J. (1992). Defining virtual reality: Dimensions determining telepresence. *Journal of Communication*, 42(4):73–93.
- [95] Strachan, S. and Murray-Smith, R. (2008). Geopoke: Rotational mechanical systems metaphor for embodied geosocial interaction. In: *Proceedings of the fifth Nordic conference on Human-computer interaction*, pages 543–546.
- [96] Strachan, S. and Murray-Smith, R. (2009). Bearing-based selection in mobile spatial interaction. *Personal and Ubiquitous Computing*, 13(4):265–280.
- [97] Sutton, R. S. and Barto, A. G. (1998). *Reinforcement Learning*. MIT Press, Cambridge, Mass.
- [98] Taylor, S. F., Tishby, N., and Bialek, W. (2007). Information and fitness. *arXiv.org:0712.4382*.
- [99] Thelen, E. and Smith, L. B. (1996). *A Dynamic Systems Approach to the Development of Cognition and Action*. Bradford Book.

- [100] Thimbleby, H. (2007). *Press On Principles of Interaction Programming*. MIT Press.
- [101] Tishby, N. and Polani, D. (2011). Information theory of decisions and actions. In Cutsuridis, V., Hussain, A., and Taylor, J., editors, *Perception-Action Cycle: Models, Architecture and Hardware*, pages 601–636. Springer.
- [102] Todorov, E. and Jordan, M. (2002). Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, 5(11):1226–1235.
- [103] Tohidi, M., Buxton, W., Baecker, R., and Sellen, A. (2006). Getting the right design and the design right. In *Proceedings of CHI*, pages 1243–1252.
- [104] Tomasello, M. (1995). Joint attention as social cognition. In Moore, C. and Dunham, P., editors, *Joint attention: its origins and role in development*, pages 103–130. Lawrence Erlbaum Associates.
- [105] Tomasello, M. (1999). *The cultural origins of Human Cognition*. Harvard University Press.
- [106] Tomasello, M. (2009). *Why We Cooperate*. Boston Review Books.
- [107] Touchette, H. and Lloyd, S. (2000). Information-theoretic limits of control. *Phys. Rev. Lett.*, 84(6):1156–1159.
- [108] Touchette, H. and Lloyd, S. (2004). Information-theoretic approach to the study of control systems. *Physica A*, 331:140–172.
- [109] van Dijk, S. and Polani, D. (2011). Look-ahead relevant information: Reducing cognitive burden over prolonged tasks. In *IEEE Symposium Series in Computational Intelligence 2011 – Symposium on Artificial Life*, pages 46–53.
- [110] van Dijk, S. G. (2013). *Informational Constraints and Organisation of Behaviour*. PhD thesis, University of Hertfordshire.
- [111] Wasinger, R. and Krüger, A. (2006). Modality preferences in an instrumented environment. In *Proceedings of the 10th international conference on Intelligent User Interfaces*, pages 336–338.
- [112] Weiser, M. (1991). The computer for the twenty-first century. *Scientific American*, 265(3):94–104.
- [113] Wennekers, T. and Ay, N. (2005). Finite state automata resulting from temporal information maximization. *Neural computation*, 17(10):2258–2290.
- [114] Wheeler, J. A. (1990). Complexity, entropy and the physics of information. *Santa Fe Studies in the Sciences of Complexity*, pages 3–28.
- [115] Witmer, B. and Singer, M. (1998). Measuring presence in virtual environments: A presence questionnaire. *Presence: Teleoperators and Virtual Environments*, 7(3):225–240.
- [116] Wolpert, D., Doya, K., and Kawato, M. (2003). A unifying computational framework for motor control and social interaction. *Philosophical Transactions of the Royal Society of London Series B*, 358(1431):593–602.

- [117] Wolpert, D. M. (2007). Probabilistic models in human sensorimotor control. *Human Movement Science*, 26(4):511–524.
- [118] Wolpert, D. M. and Körding, K. P. (2006). Bayesian decision theory in sensorimotor control. *Trends in Cognitive Sciences*, 10(7):319–326.
- [119] Zimmerman, J., Forlizzi, J., and Evenson, S. (2007). Research through design as a method for interaction design research in hci. In *Proceedings of CHI*, pages 493–502.

# Appendix



## A Experiment II – User Experience

The results of the extended NASA-TLX questionnaire (see Figure 1) reveal that the Tactile is generally inferior to the Combined and the Visual conditions. Note that lower scores correspond to better ratings except for Perceived Performance and Overall Preference, which are reversed in this case. A non-parametric Friedman test shows a significant effect of the type of feedback on all measures presented in Table 1. Pair-wise Wilcoxon signed ranks tests (including Bonferroni correction) on all factors in the extended NASA-TLX reveal that the Tactile is significantly lower than the Combined and the Visual conditions (see Table 2). No significant difference was found between the Combined and the Visual conditions, perhaps due to the dominant effect of the visual modality feedback, which was present in both conditions.

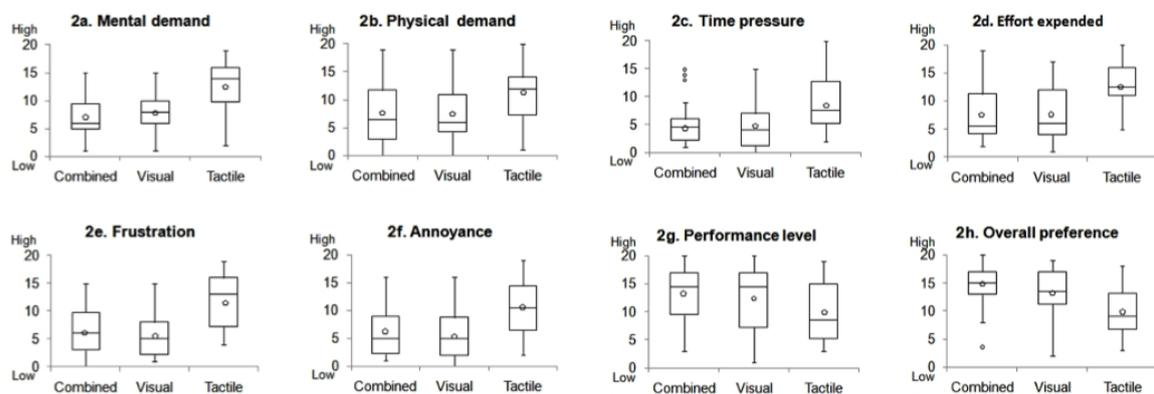


Fig. 1 Results of the extended NASA-TLX questionnaire measuring perceived workload on a 20-point scale. Lower scores correspond to better ratings except for Perceived Performance and Overall Preference, which are reversed. Tactile is generally inferior to the other two conditions.

Table 1 Friedman test statistics revealing a significant effect of the type of feedback on all measures of the extended NASA-TLX.

Factor	$\chi^2$	$df$	$p$	$N$
Mental demand	30.6	2	.000	26
Physical demand	17.4	2	.000	26
Time pressure	24.1	2	.000	26
Perceived effort	30.3	2	.000	26
Perceived performance	10.3	2	.006	26
Frustration	27.4	2	.000	26
Annoyance	17.1	2	.000	26
Overall preference	21.1	2	.000	26

Table 2 Pair-wise Wilcoxon signed ranks tests of the extended NASA-TLX (including Bonferroni correction) revealing that the Tactile is significantly lower than the Combined and the Visual conditions in all measures.

Factor	Tactile vs. Combined ( <i>p</i> )	Tactile vs. Visual ( <i>p</i> )
Mental demand	.000	.000
Physical demand	.045	.006
Time pressure	.015	.003
Perceived effort	.000	.000
Perceived performance	.01	.043
Frustration	.001	.000
Annoyance	.018	.001
Overall preference	.002	.015

## B Experiment II – Negotiation Strategies

Insets from the time series of various pairs reveal further details of different tracking behaviours. Figure 2 shows periods of almost perfect synchrony, during which the partners are constantly in contact (left) and patterns of loose tracking, where the partners occasionally lose contact, however regain it quickly (right). Other successful strategies, different from continuous tracking, are shown in Figure 3, in which the pairs quickly acquired the targets with minimal contact between each other. Figure 4 presents details of the negotiation process in search for a working strategy of a pair, who admittedly did not have a joint strategy. Although this pair found it easy to learn the interaction technique of finding the holes and the partner, they found it difficult to get the other to move to the same direction.

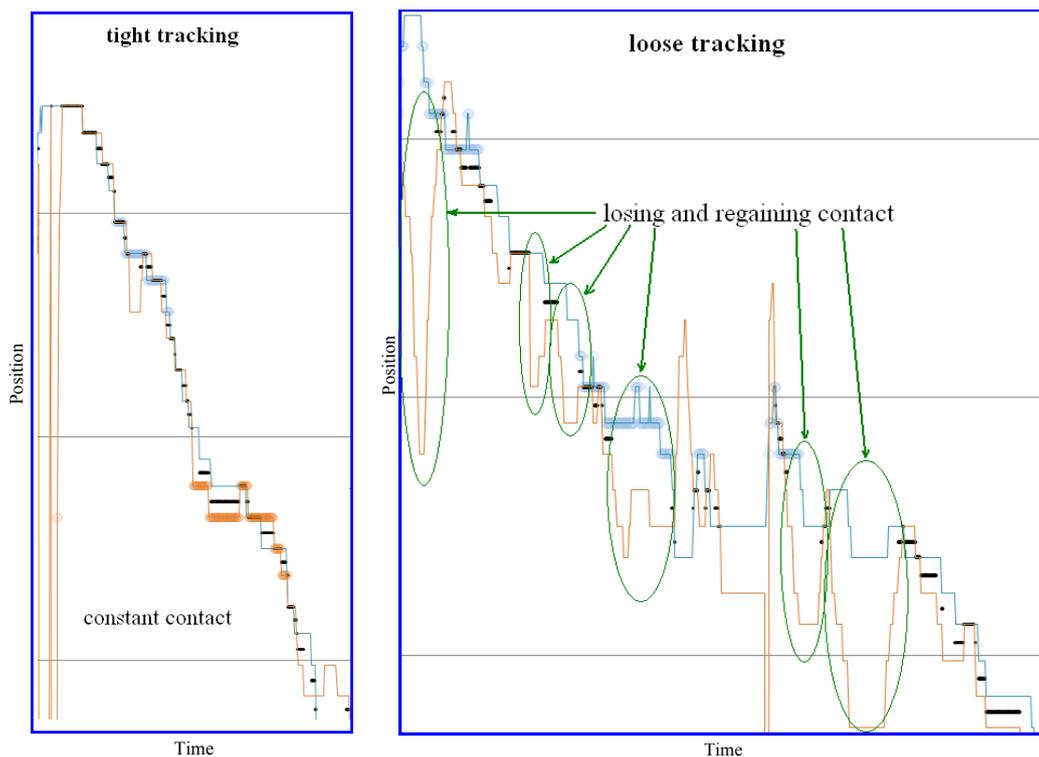


Fig. 2 Examples of tight (left) and loose (right) tracking behaviour in the Visual condition. The positions of subjects A and B are represented with blue and red curves, potential target localisation events with blue and red spheres, in-contact events with black dots, and successful target acquisitions with yellow spheres.

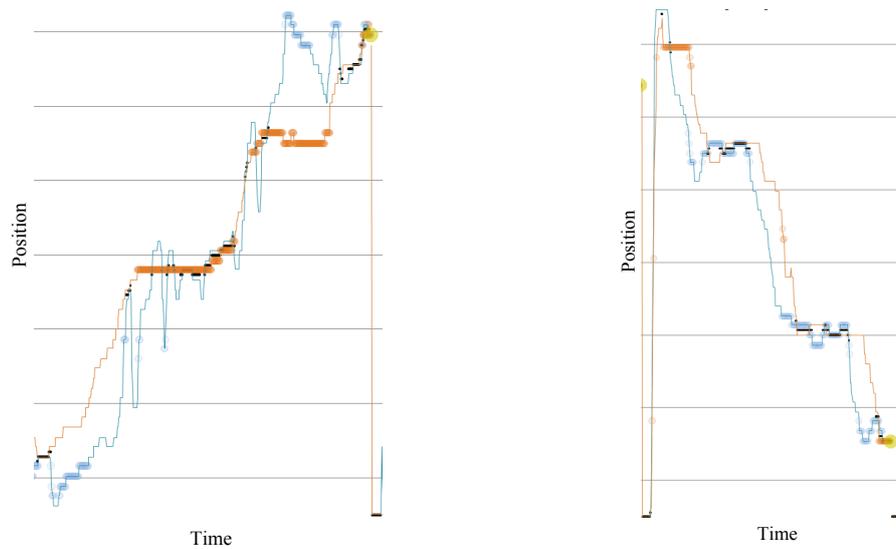


Fig. 3 Examples of target acquisitions – searching from the bottom up (left), and from the top down (right) in the Visual condition. The positions of subjects A and B are represented with blue and red curves, potential target localisation events with blue and red spheres, in-contact events with black dots, and successful target acquisitions with yellow spheres.

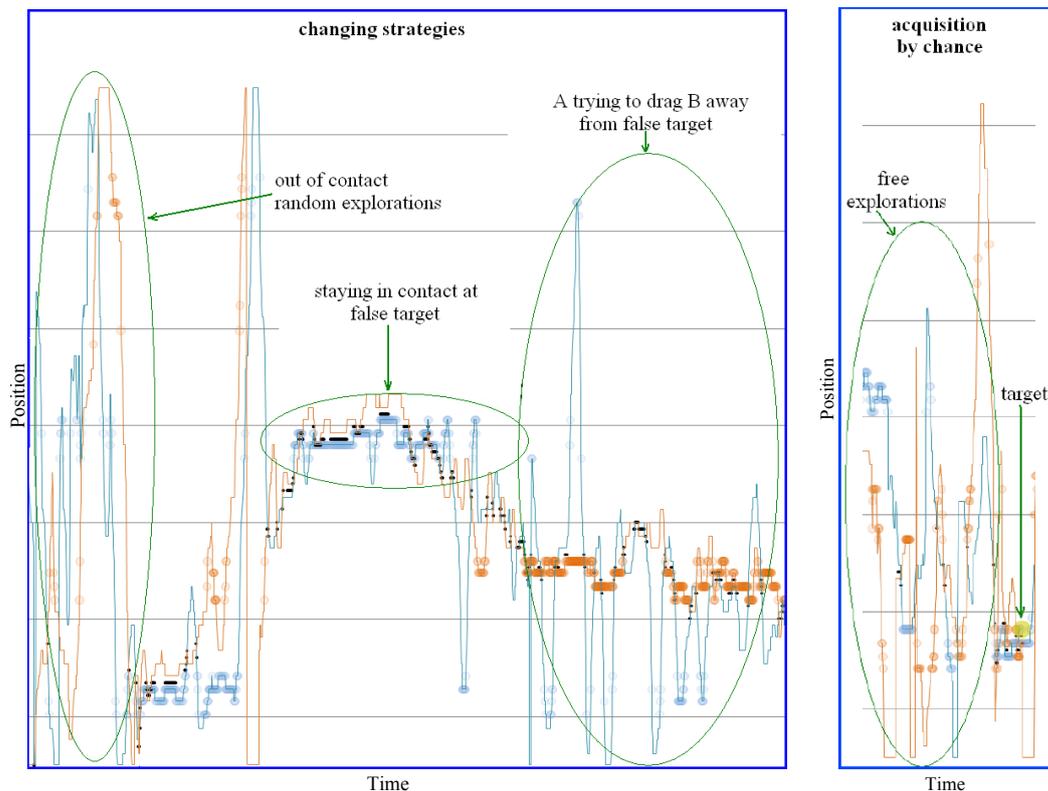


Fig. 4 Examples of random behaviour in the Visual condition. Positions of subjects A and B are depicted with blue and red curves, potential target localisations with blue and red spheres, in-contact events with black dots, and successful target acquisitions with yellow spheres.