Mälardalen University Licentiate Thesis 271

Collaborative Adaptive Autonomous Agents

Mirgita Frasheri

MÄLARDALEN UNIVERSITY
SWEDEN

Mälardalen University Press Licentiate Theses
No. 271

COLLABORATIVE ADAPTIVE AUTONOMOUS AGENTS

Mirgita Frasheri

2018

School of Innovation, Design and Engineering

Copyright © Mirgita Frasheri, 2018
ISBN 978-91-7485-399-5
ISSN 1651-9256
Printed by E-Print, Stockholm, Sweden

Për mamin, babin, dajën, nënën, dhe gjyshin.

Acknowledgements

To my mom and dad, for your unbreakable spirits, courage and guts to take whatever life throws and defying odds. For holding my hand when I was learning how to write, or when I was struggling with some hard math problem. For looking after me, always. This life, and everything in it, I owe to you.

To my uncle. For walking me to school, German class, and wherever I had to go. For recounting the battles of Alexander the Great before I'd go to an exam, just in case I'd forgotten how to be brave. Thank you for showing me the good music, the kind that gets one through anything, discussing with me the beauty of Van Gogh's yellow, evolution, partisan battles and everything you thought I needed in order to grow.

To my grandparents, my heroes, for my joyous childhood, for showing me what it means to love unconditionally, being there for people, in the best and worst of times. Some of the best times in life I've had with you. You are home.

To Ina. It's been now more than 15 years when the weirdo kid met the stuck-up straight A geek. Despite that awkward beginning, I cannot imagine growing up, going through all the different phases in the years after, without you. You have always been in my corner, through thick and thin, through bad moods, annoying habits ... we have laughed a lot, taking on life those Friday nights at Olldi, or some terrace somewhere, singing 'ja maji erdhi...', 'një ditë të bukur maji...'. It has meant everything.

To Eda, Nilda, Pami, Laura, and Alba. As years pass, we are dispersed a bit everywhere, and sometimes it is tough to keep up with those everyday stories and little moments. No matter. We always go back, and it is as if nothing has changed. Thank you for your understanding and friendship through all this time.

To my math teacher, Mondi, for simply being a teacher in a million, paying attention to every single kid in class, believing that anyone could make it, and expecting us to use our heads.

To Elena. You are always there for me, bringing me chocolate when I am struggling with some bug, listening to my complaining and all I come up with, understanding, supporting, and helping, even when I've given up on asking for any of it. You're one of the kindest and bravest people I've met in life.

To Sveta. Thank you for noticing and giving value to the little things. For answering the questions that are there, but somehow remain unasked. For meticulously lecturing me on the different types of skis, boots, and goggles, so that I make an informed decision. You may think I don't take it seriously, but not really.

To Aida. It is amazing to me how generous you are with everyone. You look after all us, seeing that we do our best and always move forward. Thank you for that.

To my supervisors, Micke and Baran. Thank you for bearing with me all this time, supporting my ideas, and staying strong when those first drafts were coming.

To Alessandro. Thank you for all the discussions we've had. I have learned and keep learning so much from you. Most importantly, thank you for being inspired in what you do.

To the officemates that had (and still have) to bear with me. To Lana and Branko (current sufferers, be strong!), Elaine, Francisco, Alessio, Pablo and others, for sharing the little moments of everyday. To Afshin, for coming by, throwing some joke, and saving a lost day.

To the robotics crew, Fredrik, Carl, Lennie, Miguel and others, for the fun trips, hikes, naiad adventures, and good road-trip music.

To the guys of the mdh gang. To Abhi, Adnan, Andreas G, Asha, Cristina S, Dagge, Damir B, Eddie, Filip, Hossein, Irfan, Jakob, Jean, Leo, Lex, Marina, Martin, Maryam, Matthias, Mehrdad, Meng, Momo, Nabar, Nes, Nils, Nitin, Omar, Patrick, Per ES, Per H, Predrag, Raluca, Rob, Rong, Saad, Sahar, Sara A, Sara Ab, Sara Ab(x), Séve, Simin, Václav, Voica. I always felt like I was adopted into this big family when I came three years ago. Thank you for the lunches and fikas, and late fikas, and weekend fikas, movie nights, (video)game nights, charades, barbecues, billiards, ping-pong, badminton, skiing trips, discussions about home and country, and more. You all have given life and color to my journey.

To the administration crew. To Carola, Jenny, Malin, Maria, Michaela and others, for your continuous help and support, and for those Swedish fikas.

To Susanne and Radu, thank you for all your help, from the very beginning when everything started. Finally, thank you to the Euroweb+ project for funding this work.

Sammanfattning

Forskningen inom området autonoma agenter och fordon har verkligen tagit fart under de senaste åren. Detta syns i en omfattande mängd litteratur och den ökande investeringen av de stora industriaktörerna i utvecklingen av produkter såsom självkörande bilar. Ett olöst problem inom området är modelleringen av interaktionen mellan olika de aktörerna (exempelvis mjukvaruagenter och människor) på så sätt att det skapas ett smidigt samarbete och kontrollöverföring mellan de berörda parterna. Den övergripande ambitionen är att uppnå ett lagarbete som liknar det mänskliga. Ett sätt att närma sig detta problem är att använda sig av begreppet adaptiv autonomi. Agenter som uppvisar adaptivt autonomt beteende kan ändra sin nivå av autonomi som ett svar på förändrade omständigheter under drift, genom att fatta beslut om huruvida agenten ska vara beroende av andra eller låta de andra vara beroende av den. Det betyder att autonomi kan uttryckas i termer av beroende mellan agenter och eftersom dessa beroenden förändras över tid gör också självständigheten det.

Forskningen i denna avhandling föreslår en agentarkitektur som modellerar en agents interna verksamhet och använder konceptet att vilja interagera som grund för adaptivt autonomt beteende. Genom sina två komponenter, viljan att be om hjälp sannolikheten för att en agent kommer att be om hjälp och viljan att ge hjälp sannolikheten för att en agent kommer att hjälpa till om tillfrågad, innefattar viljan att interagera båda aspekterna av interaktion. Ett matematiskt ramverk har utvecklats som uppskattar komponenterna av viljan att interagera genom att ta med effekten av flera faktorer som hävdas vara relevanta för att skapa ett adaptivt autonomt beteende. Denna föreslagna agentarkitektur utvärderas genom att jämföra prestanda hos en population av agenter med adaptiv autonomi med an arkitektur där agenterna har tilldelats en statisk autonomi eller strategi. Dessutom undersöks hur systemet påverkas av olika initialkonfigurationer för villighet att interagera, olika system för uppdatering av

dess komponenter samt problem relaterade till agentutnyttjande. Resultaten indikerar påpotentiella fördelar med den föreslagna lösningen jämfört med agenter med statiska strategier.

Abstract

Research on autonomous agents and vehicles has gained momentum in the past years, which is reflected in the extensive body of literature and the investment of big players of the industry in the development of products such as self-driving cars. One open problem in the area is the modelling of interaction between different actors (software agents, humans), such that there is smooth collaboration and transfers of control between the involved parties. The overall ambition is the achievement of human-like teamwork. One way to approach this problem employs the concept of adaptive autonomy. Agents that display adaptive autonomous behaviour are able to change their autonomy levels as a response to changing circumstances during their operation, by making decisions about whether to depend on one another or allow others to depend on them. Thus, autonomy is expressed in terms of dependencies between agents. As these dependencies change in time, so does autonomy.

This research proposes an agent architecture that models the internal operation of an agent, and uses the concept willingness to interact as the backbone of adaptive autonomous behaviour. The willingness to interact captures both aspects of interaction, through its two components, the willingness to ask for help – the likelihood that an agent will ask for assistance, and the willingness to give help – the likelihood that an agent will help upon being requested. A mathematical framework is developed which estimates the components of the willingness to interact by integrating the impact of several factors that are argued to be relevant for shaping adaptive autonomous behaviour. This proposal is evaluated by comparing the performance of a population of agents with adaptive autonomy to agents with static autonomy or strategies. Moreover, the role of initial configurations for the willingness to interact, different schemes for the update of its components, and issues related to agent exploitation are investigated. Results indicate the potential benefit of the proposed solution with respect to agents with static strategies.

List of Publications

Papers Included in the Licentiate Thesis

Paper A *Towards Collaborative Adaptive Autonomous Agents*, Mirgita Frasheri, Baran Cürüklü, Mikael Ekström, 9^{th} International Conference on Agents and Artificial Intelligence (ICAART), 2017.

Paper B *Analysis of Perceived Helpfulness in Adaptive Autonomous Agent Populations*, Mirgita Frasheri, Baran Cürüklü, Mikael Ekström, Journal of Transactions on Computational Collective Intelligence, 2018.
Invited paper. Extension of Paper *A*.

Paper C *Comparison between Static and Dynamic Willingness to Interact in Adaptive Autonomous Agents*, Mirgita Frasheri, Baran Cürüklü, Mikael Ekström, 10^{th} International Conference on Agents and Artificial Intelligence (ICA ART), 2018.

Paper D *Adaptive Autonomy in a Search and Rescue Scenario*, Mirgita Frasheri, Baran Cürüklü, Mikael Ekström, Alessandro Vittorio Papadopoulos, accepted at the 12th IEEE International Conference on Self-Adaptive and Self-Organizing Systems, SASO 2018.

Additional Peer-Reviewed Publications, not Included in the Licentiate Thesis

1. *Collaborative Adaptive Autonomous Agents*, Mirgita Frasheri. Accepted at the 17^{th} International Conference on Autonomous Agents and Multi-agent Systems, Doctoral Symposium Track, 2018.

2. *Failure Analysis for Adaptive Autonomous Agents using Petri Nets*, Mirgita Frasheri, Lan Anh Trinh, Baran Cürüklü, Mikael Ekström. In 11th Joint Agent-oriented Workshops in Synergy (FedCSIS Conference) JAW S'17, pp. 293-297. 2017.

3. *An Optimized, Data Distribution Service-Based Solution for Reliable Data Exchange Among Autonomous Underwater Vehicles*, Jesús Rodrígu ez-Molina, Sonia Bilbao, Belén Martínez, Mirgita Frasheri, and Baran Cürüklü. Sensors 17, no. 8 (2017): 1802.

4. *Algorithms for the detection of first bottom returns and objects in the water column in sidescan sonar images*, Mohammed Al-Rawi, Fredrik Elmgren, Mirgita Frasheri, Baran Cürüklü, Xin Yuan, José-Fernán Martínez, Joaquim Bastos, Jonathan Rodriguez, and Marc Pinto. In OCEANS 2017-Aberdeen, pp. 1-5. IEEE, 2017.

Contents

I Thesis **1**

1 Introduction **3**
1.1 Thesis Overview 5

2 Background **9**
2.1 Intelligent Agents 9
2.2 Agent Autonomy 11
2.3 Directions of Research 12

3 Research Overview **17**
3.1 Research Problem 17
3.1.1 Problem Formulation 17
3.1.2 Hypothesis and Research Questions 18
3.2 Research Process 19
3.3 Thesis Contributions 20
3.3.1 C1 - Agent Architecture 20
3.3.2 C2 - The Willingness to Interact 22
3.3.3 C3 - Performance Analysis 28

4 Related Work **43**

5 Conclusions and future work **47**
5.1 Conclusions 47
5.2 Future Work 48

Bibliography **51**

II Included Papers 57

6 Paper A:
Towards Collaborative Adaptive Autonomous Agents 59
6.1 Introduction . . . 61
6.2 Related Work . . . 62
6.3 The Agent Model . . . 64
6.3.1 Interaction between Agents . . . 67
6.3.2 Agent Organization and Autonomy . . . 68
6.4 Experiment . . . 69
6.4.1 Setup . . . 69
6.4.2 Results . . . 71
6.5 Discussion . . . 74
6.6 Acknowledgements . . . 75
Bibliography . . . 75

7 Paper B:
Analysis of Perceived Helpfulness in Adaptive Autonomous Agent Populations 79
7.1 Introduction . . . 81
7.2 Related Work . . . 83
7.3 Agent Model . . . 86
7.3.1 Early Work (C1) . . . 87
7.3.2 Interactions between Agents . . . 89
7.3.3 Agent Organization and Autonomy . . . 91
7.3.4 The AA Agent (C2) . . . 92
7.3.5 ROS . . . 95
7.4 Perceived Helpfulness . . . 95
7.5 Simulations . . . 97
7.5.1 Setup and Results for C1 . . . 97
7.5.2 Setup and Results for C2 . . . 100
7.6 Discussion . . . 104
7.7 Conclusion . . . 112
7.8 Appendix . . . 112
Bibliography . . . 113

8 Paper C:
Comparison Between Static and Dynamic Willingness to Interact in Adaptive Autonomous Agents **119**
8.1 Introduction . . . 121
8.2 Related Work . . . 122
8.3 The Agent Model . . . 124
8.3.1 Willingness to Interact . . . 125
8.4 Simulations . . . 128
8.5 Results . . . 131
8.5.1 Static vs. Dynamic Willingness . . . 131
8.5.2 Base-line Update vs. Continuous . . . 134
8.6 Discussion . . . 138
8.7 Acknowledgments . . . 139
8.8 APPENDIX . . . 139
Bibliography . . . 139

9 Paper D:
Adaptive Autonomy in a Search and Rescue Scenario **143**
9.1 Introduction . . . 145
9.2 Background and Related Work . . . 145
9.2.1 The SAR Domain . . . 145
9.2.2 Agent Cooperation and Coordination . . . 146
9.3 The Adaptation Strategy . . . 147
9.3.1 The Agent Model . . . 147
9.3.2 Willingness to Interact . . . 148
9.3.3 Factor description . . . 150
9.4 Simulation setup . . . 151
9.4.1 Scenario instantiation . . . 152
9.4.2 Agent instantiation . . . 152
9.5 Results . . . 156
9.5.1 Evaluation metrics . . . 156
9.5.2 Numerical results . . . 157
9.6 Conclusion . . . 158
Bibliography . . . 158

I

Thesis

Chapter 1

Introduction

Recent years have seen an increase of interest in autonomous software and hardware, both from an academic and industrial perspective. A considerable body of research tackles issues related to autonomous behaviour from several viewpoints. On one hand, many efforts target the definition of autonomous behaviour itself. Aspects such as what it means for autonomy levels to change, how to achieve human-like teamwork and collaboration, how to specify a human' role in the interplay with agents and robots, and the role of ethical principles, are gaining considerable attention as well. On the other hand, companies such as Google and Tesla, among others, are moulding the state-of-art, by focusing on the realization of self-driving cars, that are a form of autonomous systems. A crucial component in these vehicles is the software which has to continuously adapt to a dynamic environment, and make decisions about what to do next. These decisions can be made autonomously, or with different degrees of feedback that can come from human operators or other systems. Five autonomy levels have been identified [1], from no automation, where the driver is the only actor that controls the vehicle, to full self-driving automation, where the driver is not expected to retain control of the vehicle, and the load is on the system.

Collaboration between agents and robots is relevant in those scenarios in which they should work together to achieve some goal or task. An open issue remains the modelling of the interactions between agents, such that smooth collaboration between them is achieved. These interactions and collaborations impact the autonomy of the entities involved, when they depend on one another, or are subject to external influences (e.g. receiving help requests), with respect

to specific goals and associated tasks. Therefore, one definition of agent autonomy has been expressed in terms of dependence theory [2] i.e., an agent is non autonomous from another with respect to some task when it lacks means for its achievement and as a result decides to depend on the other agent. Whereas, adaptive autonomy (AA) generally assumes that the agent itself can change its autonomy levels as a response to different circumstances it finds itself in [3]. In this thesis, based on dependence theory, AA is further defined as the behaviour of changing dependence relations with other agents during operation. Thus, adaptive autonomous agents continuously make decisions on whether to depend on others, and on whether to allow others to depend on themselves. Therefore, AA can be used to model how agents can interact and collaborate with one another.

The research presented in this thesis tackles the modelling of adaptive autonomous behaviour, which enables agents to change their autonomy depending on the context, and thus collaborate with each other in order to complete tasks and improve their performance. To this end, an agent architecture is proposed which models the operation of an agent as a finite state machine, with different states capturing different aspects of operation. The criteria on the architecture are simplicity, and generality, such that it can be used in different domains. Furthermore, the concept of willingness to interact is proposed to be employed as the driver of adaptive autonomous behaviour. The willingness to ask for and give help are the two components that make up the concept of willingness to interact. Each of these elements model an aspect of interaction, and can be influenced by several factors. These factors are integrated in the proposed mathematical framework that estimates the willingness to interact.

Specific application domains for which this research is relevant include, without being limited to, search and rescue, agriculture, and generally areas in which it is envisioned to have autonomous systems work and collaborate with one another, possibly with human operators in the loop, and even cases in which human labour needs to be replaced due to extreme work conditions whilst still allowing for human or agent feedback or interference. Nevertheless, it might not always be possible to actively have the human in the loop or a central coordinator, due for instance to unreliable communication channels. This can be the case where agents and robots are located remotely. Therefore, it is of interest to allow agents to reason about when to interact and collaborate with each other. The research discussed in this thesis proposes one possible approach that could be applied and be useful in such scenarios.

1.1 Thesis Overview

This thesis is composed of two parts. The first part describes the content and obtained results of the research in adaptive autonomous agents. Initially Chapter 2 provides the necessary background that places the work in perspective with respect to the literature. Chapter 3 focuses on (i) describing the research problem by formulating the hypothesis and research questions, (ii) reporting on the research process, and (iii) detailing on the contributions of this work by simultaneously mapping to the relevant publications. Thereafter, Chapter 4 discusses in-depth relevant related work, and makes a generic comparison with the proposed approach. Finally Chapter 5 draws the conclusions and sketches potential directions for future work.

The second part of the thesis is composed of four publications (Papers *A-D*) that address the research problem, hypothesis and research questions. The included papers are:

Towards Collaborative Adaptive Autonomous Agents. Adaptive autonomy enables agents operating in an environment to change, or adapt, their autonomy levels by relying on tasks executed by others. Moreover, tasks could be delegated between agents, and as a result decision-making concerning them could also be delegated. In this work, adaptive autonomy is modelled through the willingness of agents to cooperate in order to complete abstract tasks, the latter with varying levels of dependencies between them. Furthermore, it is sustained that adaptive autonomy should be considered at an agents architectural level. Thus the aim of this paper is two-fold. Firstly, the initial concept of an agent architecture is proposed and discussed from an agent interaction perspective. Secondly, the relations between static values of willingness to help, dependencies between tasks and overall usefulness of the agents' population are analysed. The results show that a unselfish population will complete more tasks than a selfish one for low dependency degrees. However, as the latter increases more tasks are dropped, and consequently the utility of the population degrades. Utility is measured by the number of tasks that the population completes during run-time. Finally, it is shown that agents are able to finish more tasks by dynamically changing their willingness to cooperate.

Analysis of Perceived Helpfulness in Adaptive Autonomous Agent Populations. Adaptive autonomy allows agents to change their autonomy levels based on circumstances, e.g. when they decide to rely upon one another for completing tasks. In this paper, two configurations of agent models for adaptive

autonomy are discussed. In the former configuration, the adaptive autonomous behaviour is modelled through the willingness of an agent to assist others in the population. An agent that completes a high number of tasks, with respect to a predefined threshold, increases its willingness, and vice-versa. Results show that, agents complete more tasks when they are willing to give help, however the need for such help needs to be low. Agents configured to be helpful will perform well among alike agents. The second configuration extends the first by adding the willingness to ask for help. Furthermore, the perceived helpfulness of the population and of the agent asking for help are used as input in the calculation of the willingness to give help. Simulations were run for three different scenarios. (i) A helpful agent which operates among an unhelpful population, (ii) an unhelpful agent which operates in a helpful populations, and (iii) a population split in half between helpful and unhelpful agents. Results for all scenarios show that, by using such trait of the population in the calculation of willingness and given enough interactions, helpful agents can control the degree of exploitation by unhelpful agents.

Comparison Between Static And Dynamic Willingness To Interact in Adaptive Autonomous Agents. Adaptive autonomy (AA) is a behavior that allows agents to change their autonomy levels by reasoning on their circumstances. Previous work has modeled AA through the willingness to interact, composed of willingness to ask and give assistance. The aim of this paper is to investigate, through computer simulations, the behavior of agents given the proposed computational model with respect to different initial configurations, and level of dependencies between agents. Dependency refers to the need for help that one agent has. Such need can be fulfilled by deciding to depend on other agents. Results show that, firstly, agents whose willingness to interact changes during run-time perform better compared to those with static willingness parameters, i.e. willingness with fixed values. Secondly, two strategies for updating the willingness are compared, (i) the same fixed value is updated on each interaction, (ii) update is done on the previous calculated value. The maximum number of completed tasks which need assistance is achieved for (i), given specific initial configurations.

Adaptive Autonomy In A Search And Rescue Scenario. Adaptive autonomy plays a major role in the design of multi-robots and multi-agent systems, where the need of collaboration for achieving a common goal is of primary importance. In particular, adaptation becomes necessary to deal with dynamic

environments, and scarce available resources. In this paper, a mathematical framework for modelling the agents' willingness to interact and collaborate, and a dynamic adaptation strategy for controlling the agents' behavior, which accounts for factors such as progress toward a goal and available resources for completing a task among others, are proposed. The performance of the proposed strategy is evaluated through a fire rescue scenario, where a team of simulated mobile robots need to extinguish all the detected fires and save the individuals at risk, while having limited resources. The simulations are implemented as a ROS-based multi agent system, and results show that the proposed adaptation strategy provides a more stable performance than a static collaboration policy.

Chapter 2

Background

Concepts and theories that lie at the foundation of the work presented in this thesis are described in this chapter. Furthermore, a broader picture of the research area in autonomous agents is presented.

The chapter starts with a brief discussion on the notion of intelligent agent, which continues with a description on a prominent characteristic of such agents, that of autonomy. Thereafter, adaptive autonomy is defined, as used in this work, based on dependence theory [2]. Due to the fact that there is a plethora of works and models in agent autonomy, what follows is a description of related autonomy models, and six directions of research identified in this area. The chapter concludes by placing the work presented in this thesis in perspective with existing research.

2.1 Intelligent Agents

An agent is commonly defined as a computer system able to perceive and act upon an environment in which it is situated, through its sensors and actuators respectively, and capable of autonomous actions [4] [5]. Wooldridge *et al.* extend this definition further such that it covers intelligent agents, by specifying that these autonomous actions need to be flexible [4]. Flexibility refers to (i) social ability defined by interactions between other agents, (ii) reactivity to events that occur in the environment, and (iii) pro-activity i.e., ability to decide on one's own goals. This definition is considered as the weak notion of agency [6]. The strong definition further assumes that an agent can incorporate aspects

that are human-like, such as mental states, beliefs, desires, intentions and so on.

Wooldridge and Jennings consider three different lines of work in agent research that are agent theories, agent architectures, and agent languages [7]. Agent theories deal with defining the concept of an agent and can employ formal models to describe an agent and its properties. Whereas agent architectures focus on building these conceptualized agents using software and/or hardware. Agent languages consist of the software tools with which these agents are to be implemented.

There are three major paradigms for designing agent architectures, (i) symbolic (classical view), (ii) reactive, and (iii) and hybrid, achieved by a combination of the symbolic and reactive approaches [6]. The symbolic, or deliberative approach, considers agents, which have symbolic representations of their world, and are able to logically reason on the wanted course of action by running processes on these symbols (e.g. theorem proving). The reactive class of architectures assumes that intelligence comes from interaction through the environment, and not by processing symbols. A typical example of such architecture is the Subsumption by Rodney Brooks [8], in which the agent is a composition of competing behaviours, each aimed at achieving some task. The third class, that of hybrid architectures, consists of models that merge the two paradigms.

A similar taxonomy can be obtained by viewing agent architectures through the lens of cognition theories. There are two broad paradigms of cognition, (i) cognitivist (symbolic), and (ii) emergent [9]. Cognitivism assumes that cognition is a result of explicit manipulation over symbolic representations, and some models based on this theory are: Soar, EPIC, ACT-R, etc. Whereas, emergent systems are effective by self-organizing themselves, and maintaining their autonomy in the process. Examples of models based on this theory are: AAR, SASE, DARWIN etc. Furthermore, these two paradigms can be combined in a hybrid approach, examples include Cerebeus, and Cog among others.

The work presented in this thesis focuses on the autonomous behaviour and social aspects of intelligent agents, specifically on how social behaviour affects agent autonomy. Furthermore, the focus is on the design of a minimalistic agent architecture that supports autonomy adaptation. No assumptions are made with respect to cognitive theories. Based on the aforementioned definitions, in this thesis and included papers (*Paper B-D*), the agent is assumed to have a battery level, (ii) equipment (sensors, motors, manipulators, and actuators), (iii) skill/ability set, (iv) and knowledge. Knowledge is abstractly defined

Table 2.1: The 10 levels of autonomy proposed by Parasuraman *et al.* [11]

HIGH	10. The computer decides everything, acts autonomously, ignoring the human
	9. informs the human only if it, the computer, decides to
	8. informs the human only if asked, or
	7. executes automatically, then necessarily informs the human, and
	6. allows the human a restricted time to veto before automatic execution, or
	5. executes that suggestion if the human approves, or
	4. suggest an alternative
	3. narrows the selection down to a few, or
	2. The computer offers a complete set of decision/action alternatives, or
LOW	1. The computer offers no assistance: human must take all the decisions/actions

through what an agent knows about itself, others, the environment, and it can be used to assess or estimate variables of interest for the agent, among which it is possible to mention: outcomes of the interactions with other agents, estimate the likelihood of success in the environment, evaluate its own performance, its own progress on tasks, and even trade-offs between different tasks.

2.2 Agent Autonomy

Defining autonomy has been approached from different angles in the past years. It is possible to express autonomy in terms of two dimensions (i) self-sufficiency or the descriptive dimension, which represents the capability to conduct an activity without external help, and (ii) self-directedness or prescriptive dimension, which represents the capability to generate one's own goals [10]. Johnson *et al.* consider interdependencies between different agents, which can be used in their understanding and design for supporting joint-activity [10]. Autonomy can represent a relational concept i.e., it can be described in relation to another entity such as the environment, where it refers to independence from outside stimuli, or other agents, where it refers to independence from the influence of others (social relation). The ten levels of autonomy scheme has been

proposed [11], the aim of which was to serve as a guideline for understanding and building systems that showcase different autonomy levels (Table 2.1). It can be observed, that at the lower levels the machine is merely a slave, and as the autonomy level increases the machine retains more and more initiative, and finally in the top level is a completely independent entity. Many theories on autonomy have been developed, and different definitions have been brought forward. Among these theories belong the following: sliding autonomy [12], adjustable autonomy, adaptive autonomy, mixed-initiative interaction [10] [3], collaborative control [13] and so on. Differences generally lie on the actor that has ability and authority to make changes in the levels of autonomy.

Castelfranchi uses dependence theory to define the concept of agent autonomy [2]. Specifically, dependencies between agents can shape their autonomy levels with respect to one another. The granularity of the dependences can vary, and depends on the particular agent architecture. Assume an agent A committed to achieve a task. Moreover, A may lack some means that prevents the successful completion of this task. Therefore, agent A may decide to rely on another agent B in order to compensate for what it lacks. In that sense, A is non autonomous from B with respect to the particular task, and the means it is missing for its fulfilment.

Adaptive autonomy, in a general sense, refers to an agent's ability to change its own levels of autonomy [14]. In this thesis, using the autonomy definition based on dependencies, adaptive autonomy is defined in terms of changing dependencies between agents. In fact, given different circumstances an agent can either achieve a task by itself or needs to depend on another agent for its completion. When the agent decides to ask for assistance during the execution of some task, then it will not be autonomous from the agent that is helping, with respect to that particular task.

2.3 Directions of Research

Six general directions of research have been identified in the course of the research conducted and described in this thesis. These are: (i) design methodologies for building systems that change their autonomy levels, (ii) regulatory systems that use policies and societal norms for shaping agent behaviour, (iii) agent architectures, (iv) research that aims at investigating the benefits of such behaviour and comparing different schemes for changing autonomy levels, (v) human-machine interfaces, and (vi) algorithms for changing autonomy levels.

In (i) Johnson et *al.* argue for the need of considering interdependencies

early on in the design of systems that should support joint-activity [10]. An agent A is considered to be interdependent on an agent B, if it relies on B during the execution of its own task. Moreover interdependencies could be soft i.e., help from B is not necessary for A to succeed, and hard i.e., help from B is necessary for A to succeed. Subsequently, a design method is proposed which gives system designers concrete tools for implementing systems that incorporate adjustable autonomy [15]. The method covers the following steps. First, possible interdependencies in the system are identified. Second, different mechanisms are designed in order to support them. Lastly, it is analyzed how these mechanisms affect the existing interdependent relationships.

In (ii) the role of regulatory systems (e.g. policy systems) in bringing about predictability and therefore coordination between agents has been discussed [16]. Systems such as Kaa [17], which builds on top of an existing policy system (KAoS) allow a central agent to override or change policies involving particular agents at runtime depending on the circumstances. The human is involved in the loop if a decision cannot be made by Kaa.

In (iii) belong agent architectures that have been developed in order to provide support teamwork and autonomy level adjustment such as, STEAM, DEFACTO, and THOR. STEAM [18] is an agent architecture which extends the Soar [19] by adding support for teamwork. Team operators i.e., reactive team plans, are introduced. The agents also have their individual plans which do not require teamwork. The approach is based on the joint intention theory. The proposal offers a synchronization protocol shared between agents for team plans. Moreover, the agents decide whether a specific communication is necessary, based on the likelihood that other agents have it already, the cost for miss-coordination, and the threat that some event poses to the joint plan. In this thesis, agents keep track of whom they are helping and with respect to what task. Moreover, the main assumption is that an activity that may start as an individual one, could turn out to need support, thus it could become a team-task or task that needs assistance from other agents. The DEFACTO framework [20] aims at providing support for transfers of control in continuous time, resolving human-agent inconsistencies, and making actions interruptible for real-time systems. Team THOR's entry in the DARPA Robotics Challenge [21] brings forward a motion framework for a humanoid robot which allows for low-level control, scripted autonomy i.e. invocation of robot movements by calling predefined scripts, and enables issuing high-level commands.

In (iv) are included studies, which aim at comparing scenarios involving: (a) static autonomy versus dynamic autonomy [22], (b) different implementations of dynamic autonomy, such as AA where agents change their own au-

tonomy, mixed-initiative interaction where both human and agent are able to change autonomy, and adjustable autonomy where the human is able to change the autonomy [3]. In the case of (a) decision frameworks such as master-slave, peer to peer and locally autonomous are dynamically switched in a series of experiments in order to show the benefits of dynamic versus static autonomy. Previous studies are used to set the right decision framework based on the environmental conditions. Their results provide ground and motivation for the actual usefulness of the ability to change autonomy. As for (b), the implementation of mixed-initiative interaction performs better than the other two, measured by the number of victims found in a search and rescue simulation environment.

In (v) research is conducted on interfaces between humans and machines, that allow actors to track and learn from one another, and to initiate changes in autonomy when necessary. A robot that perceives the human's controlling skills can adjust its autonomy accordingly [23]. The decision is based on navigational, manipulation such as gripping faculties, and multi-robot coordination skills of the human operator. Through custom interfaces a user can control a robot at different autonomy levels such as: low-level body control, setting navigation way-points or object selection, setting high-level goals such as "bring the can with a soft drink" [24]. Others target the issue of one human having to coordinate a team of robots which need only occasional help from the operator [25]. Furthermore, such interfaces need to allow for a common picture to be available between the involved parties at all times [26] [27] and their design is depended on how much autonomy a system is supposed to have [28].

In (vi) reasoning mechanisms which impact how much an agent should be influenced by outside stimuli have been devised, which use task urgency and dedication level to the organization as guidelines [29]. Changes in autonomy can also occur at the task level [12] i.e., the agent might execute a task autonomously, whilst needing tele-operation for another one. Others categorize tasks beforehand in two groups i.e., tasks which the robot is able to perform by itself, and tasks that need human supervision/assistance [30]. The design of these algorithms follow from such classification. Due to unexpected events, the operation of specific agents/robots might need to be interrupted by the human operator. Colored Petri-Nets are used to formalize team plans and adding interrupt mechanisms to handle such interventions [31].

Given this background, it is possible to place the research described in thesis in the intersection of architecture development and realization of reasoning mechanisms for adaptive autonomy. Furthermore, the concept of static autonomy has been used to serve as a base-line for comparison with the proposed

adaptive autonomous model.

Chapter 3

Research Overview

3.1 Research Problem

This Chapter provides a formulation of the problem which is tackled in this thesis. Furthermore, the hypothesis and research questions are derived from the research problem, and described in detail.

3.1.1 Problem Formulation

Nowadays, a considerable amount of research is being conducted that aims at the development of autonomous systems i.e., systems capable of making decisions and fulfilling tasks on their own. Nevertheless, these systems are envisioned to be working in some environment (real or virtual), in collaboration with each other, possibly with humans as well. To this end, it is of interest to equip them with mechanisms that allow them to change their levels of autonomy and collaborate as a response to dynamic circumstances such as changes in the environment, or capabilities and conditions for performing a task. In fact, one of the ambitions in the field is to achieve a point in which agents are able to modify their autonomy, transfer control to each other, and collaborate smoothly, with or without human operators in the loop, in order to achieve efficient human-like teamwork.

This thesis focuses on the problem of how agents should reason about interactions and collaborations with one another i.e., how they should decide whether to change their own autonomy levels to accommodate particular needs that may arise during operation. Moreover, these decisions are taken by agents

autonomously. Assume such a group of agents (defined as in Section 2.1) distributed in some environment, each with its own set of goals, no global order is assumed, and tasks that need to be fulfilled. Furthermore, assume that it is not possible to rely on solutions and recommendations coming from central entities (either human or software), due for instance to unreliable communication channels. In such conditions, during the execution of a task, agents need to reason on whether help is needed. It may be that at some point in time the agent can perform a task by itself, but at a later point circumstances might be such that autonomous operation is not possible (e.g. the agent is low on battery levels). Thus, such reasoning is a continuous process, and cannot be fully pre-determined. On the other hand, agents need to also reason about what to do with incoming requests, which similarly is a continuous process. In order to address this problem, this research investigates how a minimal agent for AA behaviour can be defined, and how dynamic interaction can be modelled, which allows agents to collaborate with each other during their operation such that their performance is optimized with respect to a specific parameter, e.g. with respect to the number of tasks completed.

3.1.2 Hypothesis and Research Questions

In accordance with the problem described in Section 3.1.1, the following hypothesis is investigated in the thesis.

Hypothesis. *A group of agents that displays adaptive autonomous behaviour has better performance than agents that use static strategies.*

The main motivation for achieving adaptive autonomous behaviour lies in the premise that it can allow agents to optimize their strategies whilst executing any task. The notion of performance has been extended throughout the work, and it is clearly specified how it is assessed in each paper in Section 3.3.3.

Two research questions have been formulated in order to address the hypothesis.

RQ 1. *How to design an agent architecture that is able to exhibit adaptive autonomous behaviour?*

The intent behind this research goes beyond the narrow scope of developing particular mechanisms and algorithms for enabling adaptive autonomous behaviour. The main interest lies in considering and achieving this behaviour in the context of a generic agent architecture. Specifically, how can such agent

be designed by using a minimalist approach, i.e. in the simplest terms, such that it can be flexible and applicable in different contexts.

RQ 2. *How does the willingness to interact impact agent performance and how can it be regulated during agent operation in order to optimize performance?*

The willingness to interact is the fundamental concept used for the realization of adaptive autonomous behaviour i.e., the collaborative behaviour of a population of agents. Therefore, its impact on the performance of the population as a whole is investigated. Moreover, the willingness is an individual characteristic of every agent and needs to be evaluated continuously as a response to changing circumstances. Practically the problem lies in understanding how behavioural changes at the agent level influence those at the population level.

3.2 Research Process

Research processes and methods in computer science are topics tackled by several studies, which have proposed road-maps and guidelines for conducting research in the field. Methods for experimental validation in the sub-field of software engineering have been identified as part of a twelve step model [32]. These include observational methods (e.g. project monitoring, case-study, assertion, and field study), historical (literature search, legacy data, static analysis and others), and controlled (replicated, synthetic environment, dynamic analysis, and simulation). Others have identified the three following methods: (i) theoretical computer science – close to logics and mathematics and therefore consists of classical processes of building theories, (ii) experimental computer science – involves the creation of prototypes in order to solve a particular problem, and (iii) computer simulation – used to investigate phenomena which cannot be replicated in laboratories [33].

Taxonomies have been produced by investigating the research methods adopted by researchers, mostly in software engineering [34] [35]. Shaw classifies the types of research questions, results, and validation techniques that can be found in such literature [34]. Generally, research questions are related to either methods of development, methods for analysis, design/evaluation/analysis of an instance, generalization, and feasibility. Whereas, obtained results can consist of techniques, qualitative models, empirical models, analytic models, tools/notations, specific solutions, judgement, or report. Validation might mean either analysis, experience, example, evaluation, and others. Combinations between the different categories produce what is called a research strategy.

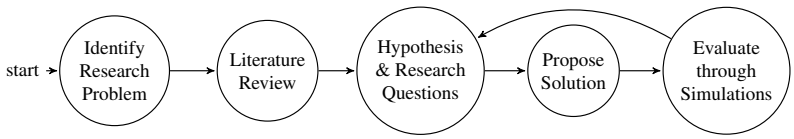

Figure 3.1: The research process

Whereas, later work categorizes research methods in computer science in nineteen groups, among which action research, analysis of literature, mathematical proof, simulation and so on [35].

In this thesis an empirical research process was followed (Figure 3.1). The literature review produced the necessary background for the formulation of early research questions and hypothesis, and the selection of existing conceptual frameworks which underlie this research and relate it to existing knowledge. Thereafter, a solution is proposed for the realization of the adaptive autonomous agent. In order to evaluate the behaviour and performance of such agents, computer simulations were used for the generation of relevant data. The latter was then quantitatively analysed through statistical analysis. Each paper included in this thesis comprises the same steps, thus uses computer simulation for data generation, and statistical analysis on such data.

3.3 Thesis Contributions

This thesis contributes three-fold to the knowledge regarding adaptive autonomous agents, as described in the following paragraphs (C1-3).

3.3.1 C1 - Agent Architecture

A high-level agent architecture has been proposed which models the internal operation of an agent. The concept of the willingness to interact, composed of the willingness to ask for and give help, is integrated in this agent model and gives rise to the adaptive autonomous behaviour.

The agent is expressed as a finite state machine with five states: *idle*, *execute*, *interact*, *regenerate*, and *out of order* (Figure 3.2). Agent operation always starts in *idle*, in which the agent is assumed not to be committed to any task. In this state, the agent can generate a task if certain conditions are met. In Papers *A*-*C*, tasks are abstract and task generation happens with probability p. In Paper *D*, a search and rescue scenario is implemented, in which tasks are

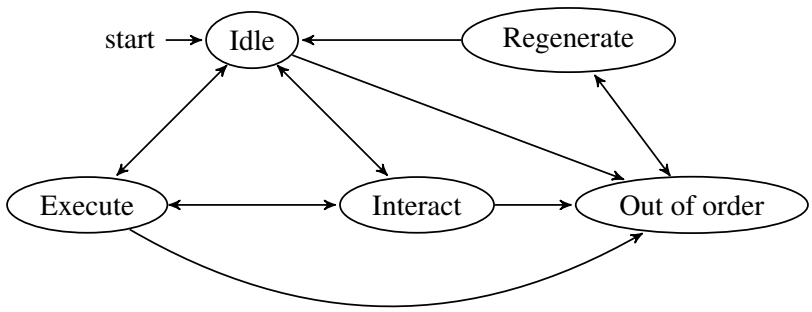

Figure 3.2: The proposed agent model composed of five states [36].

generated based on if a disaster site is within the visibility range of an agent. As soon as a task is generated, an agent shifts its state to *execute*, and the task execution takes place. In Papers *B*, and *C*, a task is an atomic step during execution, whereas in Papers *A*, and *D* a task is composed of iterations, thus is interruptible. Whilst in *execute* an agent decides on whether to make a request for assistance based on its willingness to ask for help. This is expressed as a likelihood to initiate an interaction. When an agent receives any such requests, it will shift its state into *interact*, where based on its willingness to give help will decide whether to accept or reject them. As in the case of the willingness to ask for help, the willingness to give help is expressed as a likelihood. Once a task is completed, successfully or otherwise, an agent will return to *idle*. It is possible to shift from any other state to *out of order* when some condition is met (e.g. battery level is below a certain threshold). The agent can shift to *regenerate* from *out of order*, where regeneration processes can take place (e.g. battery recharge), after which it shifts to *idle*.

C1 targets the first research question (RQ1) and provides the conceptual foundations on which the other contributions are built on. C1 is mainly addressed in Papers *A* and *B*. Paper *A* outlines the initial concept of the agent and models adaptive autonomous behaviour through the willingness to give help, whilst Paper *B* extends the agent itself and considers both parts of the willingness to interact (asking and giving help).

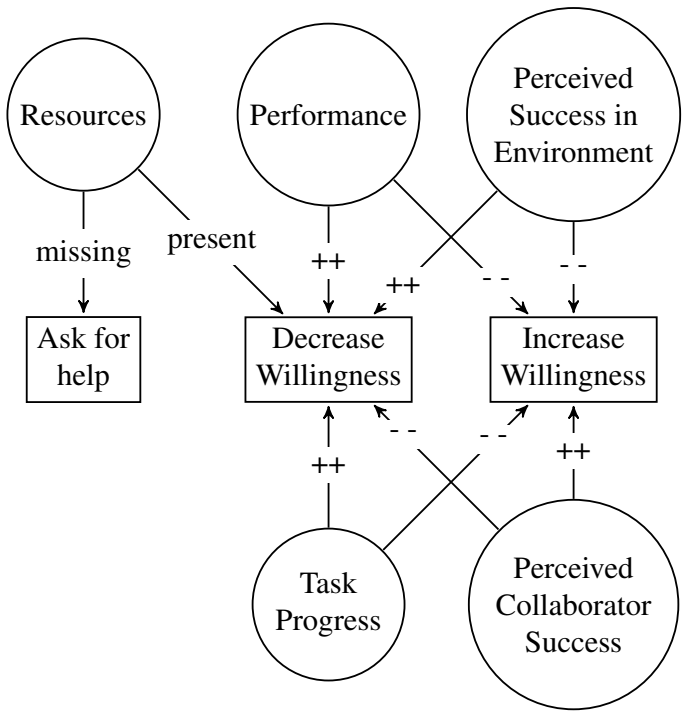

Figure 3.3: Factors that influence the willingness to ask for help.[1]

3.3.2 C2 - The Willingness to Interact

Given the agent architecture, the element of interest is the willingness to interact that drives adaptive autonomous behaviour. Initially, only the willingness to give help is considered (Paper *A*), to be afterwards extended with the willingness to ask for help (Papers *B*-*D*). The willingness to ask for help represents the likelihood that an agent will ask for help at some time t, whereas the willingness to give help represents the likelihood that an agent will give help upon receiving a request. Thus, a complete model of interaction is achieved on the conceptual level. The second contribution of the thesis comprises the analysis of the factors that can impact the willingness to interact. Furthermore, a mathematical framework is proposed such that the willingness to interact can be

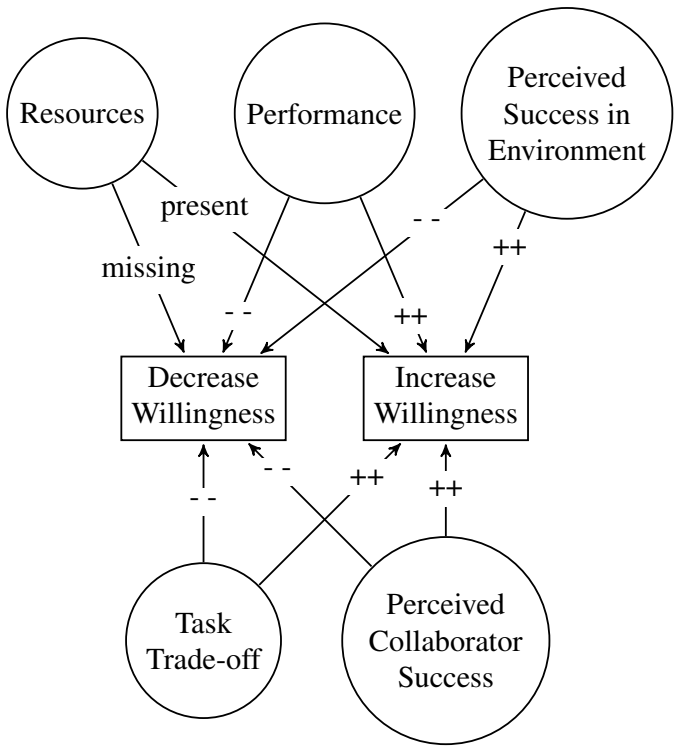

Figure 3.4: Factors that influence the willingness to give help.[2]

updated as circumstances of agents change. Contribution C2 addresses partly RQ1 – by elaborating on how the willingness to interact, and more precisely how its components can be modelled, in order to achieve adaptive autonomous behaviour (Papers *B* and *C*), and partly RQ2 – by proposing the mathematical framework used in the estimation of the willingness to interact (Paper *D*).

[1]The corresponding figure presented in Paper *C* depicts only internal resources such as the battery level, skill set, equipment, knowledge. However, other resources such as tools can determine whether asking for help is necessary. Therefore, the notation in the present figure is changed to resources.

[2]See footnote [1].

Factor Description

In the research presented in this thesis nine factors are identified for each component of the willingness to interact. The willingness to ask for help (Figure 3.3) is affected by battery level, ability/skill set efficiency, knowledge accuracy, equipment accuracy, resource quality (e.g. quality of external resources such as tools), performance, perceived success in the environment, task progress, and perceived success of collaborator. Similarly, the willingness to give help (Figure 3.4) is affected mostly by the same factors. However, task trade-off, between the current task and the requested task, is considered instead of task progress. Each factor affects the willingness to interact in a certain way, and these relations can be explained through the following example. Factors are derived from the agent definition and resource types generally needed for achieving a task.

Assume an agent *A* committed to a particular task. At the beginning and during the execution of a task *A* will reason about its willingness to ask for help (Figure 3.3). If the agent has the necessary battery levels required by the task, then it can proceed autonomously. As a result it will tend to decrease its likelihood to ask for assistance. Otherwise, it will ask for help. Note that this is deterministic, since the desired behaviour for an agent with low battery level is to delegate its task, and not attempt it on its own. Thus this factor dominates all others. In the same way it is argued for *A*'s skill set, knowledge, equipment, and other resources relevant for this particular task. If *A* lacks a certain skill, or the efficiency of the skill is not appropriate for the task, then it should ask for help and so on. The other factors do not deterministically affect the willingness. If performance is high, *A* is less inclined to ask for help. If the performance is low, then *A* is more inclined to ask for help. If the perceived success in the environment is high, then *A* will decrease its willingness to ask for help. Otherwise, if the perceived success is low, then *A* will increase its willingness to ask for help. If interaction with a potential collaborator is perceived as successful then *A* will increase its willingness to ask for help, otherwise it will decrease it. If *A* perceives a progress in the execution of the task then it will decrease its willingness to ask for help, otherwise it will increase it.

Assume that *A* makes the decision to ask for help, and forwards this request to another agent *B*. Upon receiving the request, *B* will reason on whether to accept it, determined by the willingness to give help (Figure 3.4). The process is analogous to reasoning on the willingness to ask for help. Consequently, *B* will increase its willingness to give help if it has the necessary battery level, and decrease it in the opposite case. Similarly it is argued for the skill set, knowledge,

equipment, and other resources i.e., if the agent has the required skill set then it will increase its willingness to give help, otherwise it will decrease it. Note that there is a difference to how the willingness to ask for help is estimated. There is a chance that even if the battery level is not sufficient, *B* will accept to help. There is no dominance between factors. In this case *B* will ask another agent *C* for help, thus triggering a chain reaction. It is also possible to enable *B* to reject the request altogether if it does not have the required battery level. It has not been investigated which scheme is preferable. If the performance of *B* is high, then the likelihood to give help increases. The opposite also holds. In case the perceived success in the environment is high, the likelihood of giving help increases, otherwise it decreases. The same reasoning holds for perceived success with a potential collaborator. When receiving the request, it is possible for *B* to be engaged in another task. As such if the new task is perceived as more rewarding, then the willingness to give help will increase, otherwise it will decrease.

The list of factors is not meant to be exhaustive i.e., depending on a particular scenario, or application domain some of these factors can be removed, or new ones can be added. Nevertheless, the rationale behind the choice of these particular factors is as follows. The first five factors express what is required by the agent in order to accomplish a particular task. The performance factor serves as a generic cue to the agent that can be used to determine whether operation is going well. Task-progress on the other hand is a specific metric related to performance with respect to the specific task currently being executed. Both the perceived success in environment and of the collaborator represent facets of reality that can have an impact on the agent's operation. Finally task trade-off captures the information which an agent can use when deciding what task to perform among several. A weighted average that integrates the effect of these factors and uses them for the calculation of the willingness to interact has been employed and is described in the following paragraphs. This choice was made due to its simplicity. A full list of the variables employed in the following are given in Table 3.1.

Adaptation Strategy

The willingness to ask for help at some time t, $\gamma_t \in [0, 1]$, is calculated by applying a correction to its initial value γ_0 (Equation 3.1).

$$\gamma_t = \mathrm{sat}(\gamma_0 + u_t), \tag{3.1}$$

where $\mathrm{sat}(x) := \min(\max(x, 0), 1)$, and u_t is the *adaptive correction* at time t, calculated as

$$u_t = \sum_{i=1}^{n} \phi_i f_t^{i\gamma}, \tag{3.2}$$

where n is the number of factors, ϕ_i, $i = 1, \ldots, n$, are weights (constant or calculated at runtime), such that $\phi_i \in [0, 1]$, $\sum_{i=1}^{n} \phi_i = 1$, and $f_t^{i\gamma}$, are the considered factors at time t.

Similarly, the willingness to give help at some time t, $\delta_t \in [0, 1]$, is calculated by applying a correction to its initial value δ_0 (Equation 3.3).

$$\delta_t = \mathrm{sat}(\delta_0 + v_t). \tag{3.3}$$

The correction for the willingness to give help is

$$v_t = \sum_{i=1}^{n} \zeta_i f_t^{i\delta} \tag{3.4}$$

where n is the number of factors, ζ_i, $i = 1, \ldots, n$, are weights (constant or calculated at runtime), such that $\zeta_i \in [0, 1]$, $\sum_{i=1}^{n} \zeta_i = 1$, and $f_t^{i\delta}$, are the considered factors at time t.

The k-th factor f^{k*} is calculated continuously based on the current measurements ψ_k of relevant quantities, e.g., the battery level, the accuracy in executing a specific task and so on, and on a minimal acceptable value $\psi_{k\,\min}$ for each. The update rules for the willingness to ask for and give help are the following:

$$f^{k\gamma} = \begin{cases} \psi_{k\,\min} - \psi_k, & \psi_k > \psi_{k\,\min} \\ (1-\alpha)(\psi_{k\,\min} - \psi_k) + \alpha, & \psi_k \leq \psi_{k\,\min} \end{cases} \tag{3.5}$$

$$f^{k\delta} = \begin{cases} \beta(\psi_k - \psi_{k\,\min}) & \psi_k > \psi_{k\,\min} \\ \beta(1-\alpha)(\psi_k - \psi_{k\,\min}) - \alpha, & \psi_k \leq \psi_{k\,\min} \end{cases} \tag{3.6}$$

where $\alpha \in \{0, 1\}$ $\beta \in \{-1, 1\}$ are *control parameters*, that influence the calculation of the factors to achieve the desired behavior, and $f^{k*} \in [-1, 1]$. In case $f^{k\gamma} \to 1$, then the likelihood of the agent to ask for help is high, whereas for $f^{k\gamma} \to -1$ the opposite holds. Similarly it can be stated for $f^{k\delta}$.

The reasoning behind (3.5) and (3.6) is that factors should be updated based on the distance of a current measurement ψ_k from the minimal allowed value $\psi_{k\,\min}$. Specifically, if $\psi_k \geq \psi_{k\,\min}$, then the agent does not necessarily need

Table 3.1: Variable Notation

Variable	Description
	General symbols
γ_t	**Willingness to ask for help** at time t
δ_t	**Willingness to give help** at time t
	Factors
f_t^{1*}	**Battery factor** – influences the willingness to interact based on a_i's battery level at time t
f_t^{2*}	**Knowledge factor** – influences the willingness to interact based on a_i's confidence on own knowledge at time t
f_t^{3*}	**Skill factor** – influences the willingness to interact based on a_i's efficiency of own skills at time t
f_t^{4*}	**Equipment factor** – influences the willingness to interact based on a_i's equipment accuracy at time t
f_t^{5*}	**Resource factor** – influences the willingness to interact based on a_i's resource quality at time t
f_t^{6*}	**Performance factor** – influences the willingness to interact based on a_i's own performance at time t
$f_t^{7\gamma}$	**Task progress factor** – influences the willingness to interact based on a_i's task progress at time t
$f_t^{7\delta}$	**Task trade-off factor** – influences the willingness to interact based on the trade-off between τ_j and τ_p at time t
f_t^{8*}	**Environment factor** – influences the willingness to interact based on a_i's perceived success in the environment at time t
f_t^{9*}	**Collaborator factor** – influences the willingness to interact based on a_i's perceived success with its collaborator at time t
	Measured quantities and parameters
ψ_1	Agent's battery level
ψ_2	Confidence on knowledge concerning τ_j
ψ_3	Efficiency of skills concerning τ_j
ψ_4	Efficiency of equipment related to τ_j
ψ_5	Quality of resources related to τ_j
ψ_6	Agent performance
$\psi_{7'}$	Task progress
$\psi_{7''}$	Task trade-off between τ_j and τ_p
ψ_8	Perceived success in environment
ψ_9	Perceived success of collaborating

help to complete its task, thus its willingness to ask for help will decrease. Otherwise, it is able to give help, and as a result its willingness to give help will increase.

It is important to note that a factor f^{k*} does not necessarily dominate the other factors. In this thesis, only factors $f^{1-5\gamma}$ are considered to be dominant. Therefore, the corresponding weights ϕ_{1-5} will be set to 1 when the value of the factor is 1. Moreover, the weights for the other factors $f^{6-9\gamma}$ are set to 0, because $\sum_{i=1}^{n} \phi_i = 1$. If two or more of the dominant factors is equal to 1, then only one of the corresponding weights need be changed. The desired behaviour of asking for help will be achieved. In all other conditions, the weights for the factors are the same and calculated as $1/n$.

3.3.3 C3 - Performance Analysis

The final contribution relates to the analysis of the impact of the willingness to interact on agent performance. The concept of performance has been extended throughout this thesis. Nevertheless the ground metric is the percentage of dependent tasks completed, defined below, which gives a measure of how successful the collaborations between agents are. The following paragraphs summarize the results from each paper included in the thesis.

Paper *A* outlines the initial concept of an agent for adaptive autonomy, and models such behaviour through the willingness to give help (δ). The aim of the simulations is to investigate the performance of homogeneous agent populations in different settings of task dependency. The latter refers to percentages of tasks that depend on other tasks in order to be completed (dependent tasks), in two modes with static and dynamic willingness to give help. Performance is measured with two metrics, the completion degree (CD) which measures the percentage of completed dependent tasks, and drop-out degree (DD) which measures the percentage of dropped tasks with respect to all attempted tasks.

In the static case, simulations are run for a population size of 10 and 30, where δ takes values in $[0.0, 0.25, 0.5, 0.75, 1.0]$[3], and the task dependency takes values in $[0.0, 0.25, 0.5, 0.75, 1.0]$. An additional simulation was run for population size equal to 10 with a finer resolution segment for δ equal to $[0.0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0]$[4]. In the dynamic case, sim-

[3] Note that the likelihoods for an agent to give or ask for help, given the values for the willingness to give or ask for help, are generated through the random function from python's random library.

[4] These ranges were chosen because they cover the whole spectrum from completely unhelpful to completely helpful agents.

ulations were run for population size equal to 10, for two different scenarios, one in which there is one dynamic agent among static agents, and another in which all agents adopt a dynamic strategy for their willingness to give help. During simulation runs agents ask each other for help in case the task currently being executed depends on some other task. Dependencies between tasks are given before run-time and do not change.

Results produced by simulations (Figure 3.5) indicate that when agents dynamically change their willingness to give help, then they have better performance (in terms of dependent completed tasks). Nevertheless, task dependency has a non-negligible effect on performance. In fact, for high percentages of task dependency, performance degrades among static and dynamic agent populations alike. Furthermore, it is possible to notice that static agents that are not willing to help perform better than their counterparts when task dependency is 1.0.

Paper *B* is a journal extension of Paper *A*. This paper outlines the full concept of the agent, and proposes the willingness to interact, which captures both aspects of adaptive autonomous behaviour, that of asking for (γ) and giving help (δ). Several factors that are assumed to influence the willingness to interact are discussed. Given the extended model, the aim of the paper is to investigate how the perceived helpfulness of surrounding peers should be used in determining the willingness to give help in order to control agent exploitation. Specifically three different scenarios are considered in which the population of agents is composed of (i) one unselfish ($\langle \delta_0 = 1.0, \gamma_0 = 0.0 \rangle$), the rest selfish ($\langle \delta_0 = 0.0, \gamma_0 = 1.0 \rangle$), (ii) one selfish, the rest unselfish, and (iii) half selfish, half unselfish. Furthermore, simulations were run for each scenario and population sizes, 10 and 30. Note that γ_0 and δ_0 refer to the initial values for the willingness to ask for and give help. In the first scenario (Figure 3.6), it is possible to observe, by inspecting the percentage of accepted requests with respect to all requests received, that an unselfish agent, by considering the perceived helpfulness of the population in calculating its willingness to give help, can adapt and control exploitation by others. In the second scenario (Figures 3.7 and 3.8), it is possible to observe that an unselfish population will adapt to a selfish individual by denying it help. Nevertheless, the size of the population will affect the adaptation process as it can be seen in Figure 3.8, the selfish agent achieves a higher percentage of dependent tasks. Moreover, another scenario is simulated in which the selfish agent adapts to the population, thus it becomes more helpful. It can be observed that this agent completes more dependent tasks than the original selfish version. In the third scenario, similar results are observed (Figure 3.9), in which the unselfish agents adapt to selfish

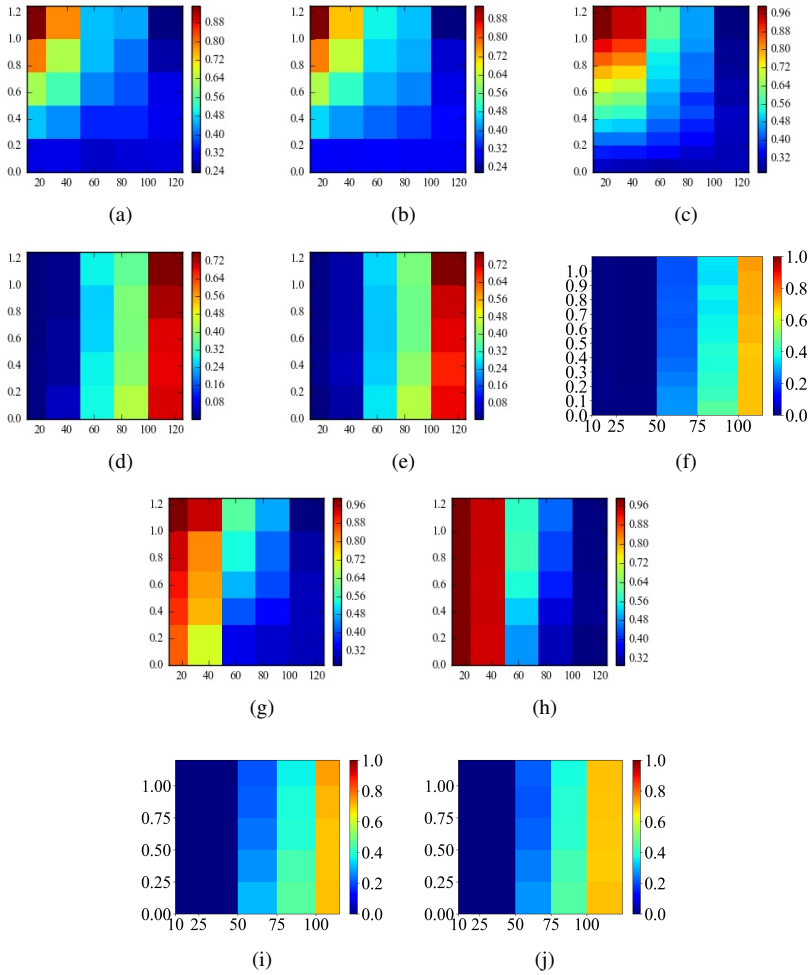

Figure 3.5: Heat maps of CD and DD utility measures, for simulations with static δ and dynamic δ, and different popsize [37]. (a) CD for popsize = 10 with static δ. (b) CD for popsize = 30 with static δ. (c) CD for popsize = 10 with finer resolution of static δ. (d) DD for popsize = 10 with static δ. (e) DD for popsize = 30 with static δ. (f) DD for popsize = 10 with finer resolution of static δ. (g) CD for popsize = 10, one agent with dynamic δ. (h) CD for popsize = 10, all agents with dynamic δ. (i) DD for popsize = 10, one agent with dynamic δ. (j) DD for popsize = 10, all agents with dynamic δ.

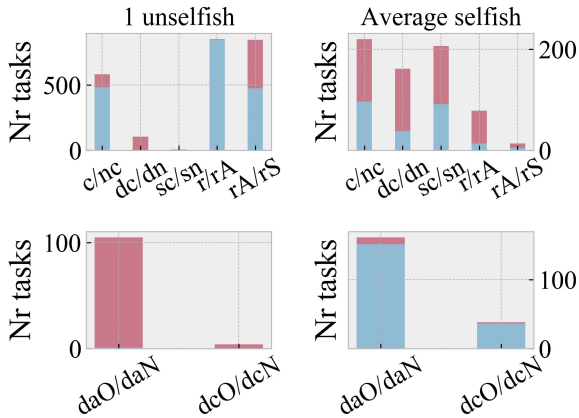

(a)

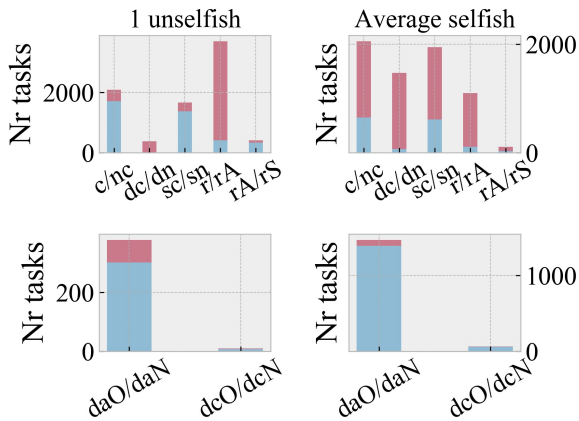

(b)

Figure 3.6: Results for S1, for $popsize = 10$. Notation: c/nc - completed/not completed tasks (c - blue bottom bar, nc - top red bar), dc/dn - dependent completed/dependent not completed tasks, sc/sn - completed self-generated/not completed self generated, r/rA - request received/requests accepted (rA - blue bottom bar), rA/rS - requests accepted /requests succeeded (rS - blue bottom bar), daO/daN - dependent self-generated tasks attempted/dependent not self-generated tasks attempted, dcO/dcN - dependent self-generated completed/ dependent not self-generated completed. (a) Selfless agent that does not adapt to a selfish population. (b) Selfless agent that adapts to a selfish population by becoming more selfish. Figures on the left correspond to the agent $\langle \delta_0 = 1.0, \gamma_0 = 0.0 \rangle$, while those on the right are averaged over the agents $\langle \delta_0 = 0.0, \gamma_0 = 1.0 \rangle$

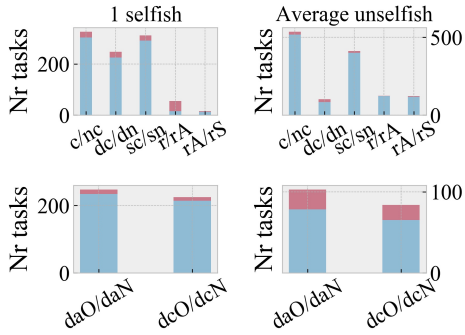

(a)

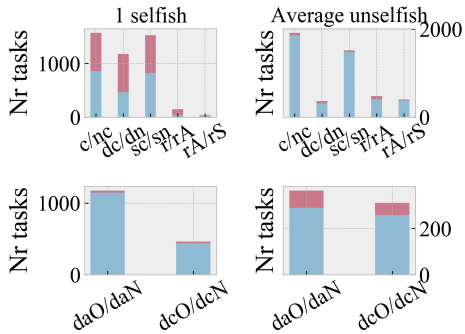

(b)

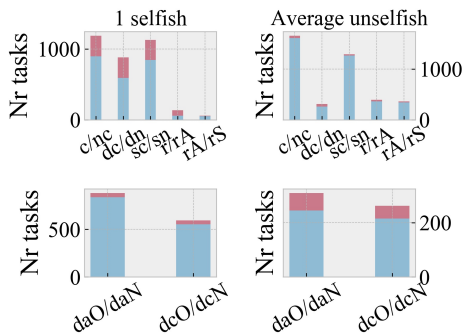

(c)

Figure 3.7: Results for S2, for $popsize = 10$. Notation same as Figure 3.6. (a) Agents that do not adapt to the selfish individual. (b) Agents that adapt to the selfish individual. Figures on the left correspond to the agent $\langle \delta_0 = 0.0, \gamma_0 = 1.0 \rangle$, while those on the right are averaged over the agents $\langle \delta_0 = 1.0, \gamma_0 = 0.0 \rangle$ and $k = 4$. (c) Selfish agent that adapts to the population by becoming more helpful. Figures on the left correspond to the agent $\langle \delta_0 = 0.0, \gamma_0 = 1.0 \rangle$, while those on the right are averaged over the agents

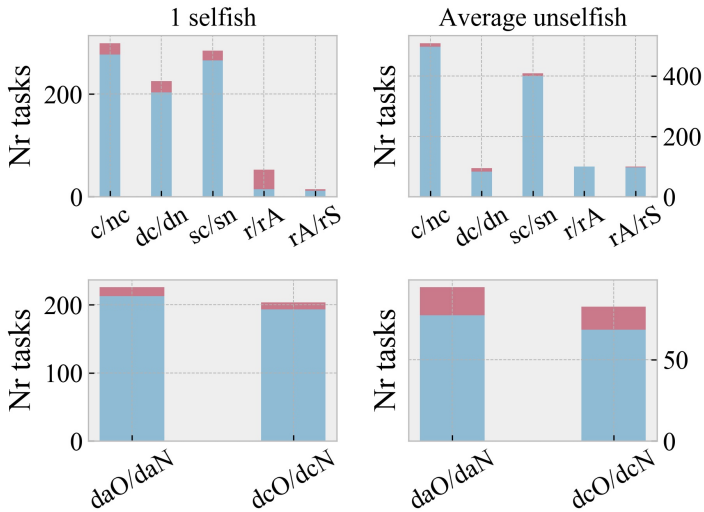

(a)

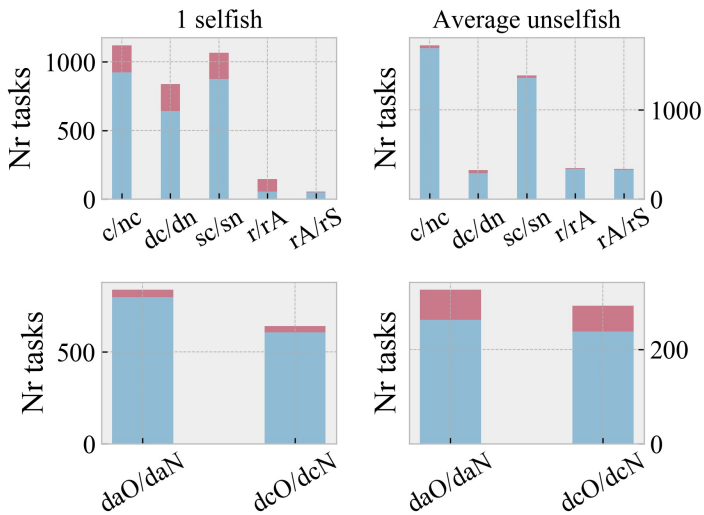

(b)

Figure 3.8: Results for S2, for $popsize = 30$. Notation same as Figure 7.7. (a) Agents that do not adapt to the selfish individual. (b) Agents that adapt to the selfish individual. Figures on the left correspond to the agent $\langle \delta_0 = 0.0, \gamma_0 = 1.0 \rangle$, while those on the right are averaged over the agents $\langle \delta_0 = 1.0, \gamma_0 = 0.0 \rangle$

agents. Nevertheless, it can be seen that less dependent tasks are completed even among selfless agents compared to the two other scenarios, which is due to the fact that the algorithm for the selection of the collaborator always picks considering the entire group. A finer mechanism could allow selfless agents only to ask for assistance from one another.

Paper *C* uses the extended model presented in paper *B*, and extends the scope of the simulations described in Paper *A*, by investigating how the willingness to interact affects agent performance (in homogeneous populations), in three different scenarios with increasing difficulty, for agents with dynamic and static strategies. Difficulty is expressed as a probability that an agent lacks either the skill set, equipment, knowledge, or tools required for a task, and is annotated as P_D[5]. This variable takes the following values $0.2, 0.5, 0.8$. Two different strategies for the update of the components of the willingness to interact are investigated, (i) the base-line update in which the components of willingness to interact are calculated based on the same initial values, and (ii) continuous update in which the components are calculated based on the previous calculated values. Results show that (Figure 3.10-3.12), overall dynamic agents with the base-line update perform better than their static counterparts with respect to the percentage of dependent tasks completed[6]. Nevertheless, there are exceptions in the higher difficulty scenarios (Figure 3.12) for $\langle \delta_0 = 1.0, \gamma_0 = 0.5 \rangle$, $\langle \delta_0 = 0.8, \gamma_0 = 0.2 \rangle$, and $\langle \delta_0 = 0.8, \gamma_0 = 0.5 \rangle$ for $P_D = 0.8$. Finally, it is also shown that the base-line update produces better results than the continuous update, and such results are stable across the three difficulty scenarios.

Paper *D* focuses on the formal description which underlies the computation of the willingness to interact, and aims to shift from an abstract concept of tasks into a real-world like scenario in which the adaptive autonomous behaviour can be useful. In these scenarios it is not always possible to rely on a central entity positioned remotely due to unreliable communication channels. One such scenario is taken from the search and rescue domain [38], in which three kinds of agents (fire fighter, ambulances, police) explore a 2D space and aim to extinguish all fires and save all trapped individuals. In order to simulate the

[5] Note that in the implementation, the probability that any of the four factors is lacking, is evaluated separately as $P_D/4$. Consequently, the likelihood that a task is a dependent task will be lower than P_D, because for a task to be dependent it is necessary that just one of any such resources lacks. This is not properly explained in Paper *C*.

[6] Note that, in the static case for $\gamma = 0.0,\ or\ 1.0$ and $\delta = 0.0$, no simulations were run. In these cases no dependent tasks can be completed because either agents never ask or give help, or all of them are always asking for help without completing any task. In these cases, values for the total number of completed tasks in the static scenario (s_A) are represented with 0.

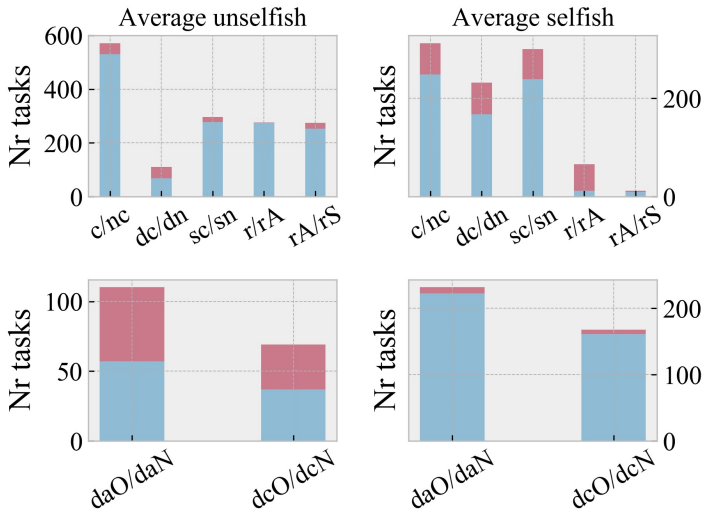

(a)

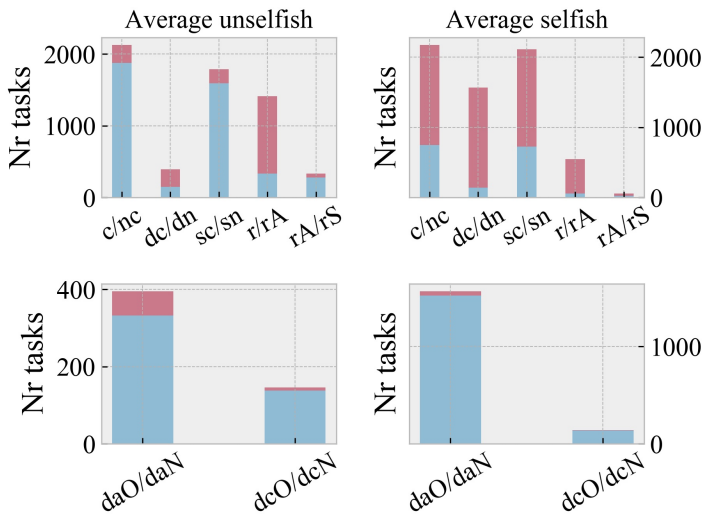

(b)

Figure 3.9: Results for S3, for $popsize = 10$. Notation same as Figure 7.7. (a) Selfless agents that do not adapt to selfish agents. (b) Selfless agents that adapt to selfish agents by becoming more selfish. Figures on the left correspond to the average of half the population of agents with $\langle \delta_0 = 1.0, \gamma_0 = 0.0 \rangle$, while those on the right are averaged over the other half $\langle \delta_0 = 0.0, \gamma_0 = 1.0 \rangle$

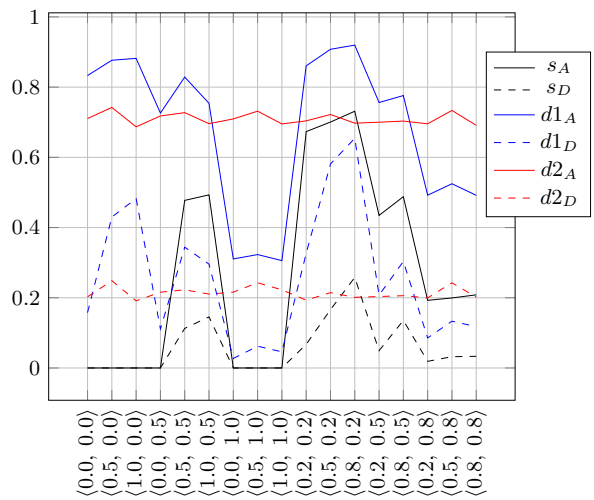

Figure 3.10: Task completion percentages for $P_D = 0.2$. s_A refers to the completion degree for all tasks in the static phase, whereas s_D corresponds to completion degree of dependent tasks. The interpretation of the subscript is the same for $d1$ and $d2$. Both refer to the dynamic case, $d1$ to the base-line update, and $d2$ to the continuous update respectively.

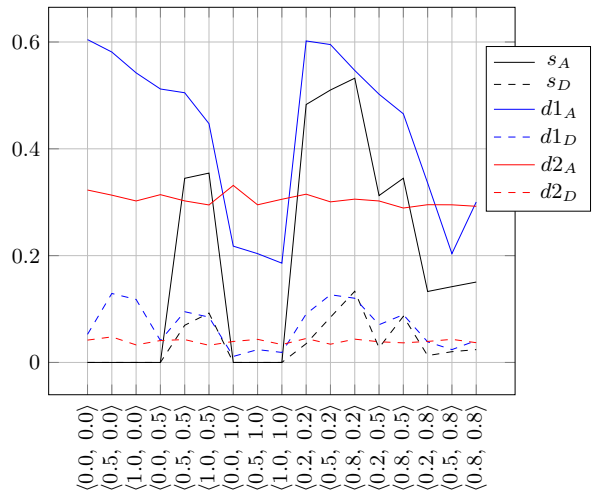

Figure 3.11: Task completion percentages for $P_D = 0.5$. Notation interpretation as in Figure 8.4.

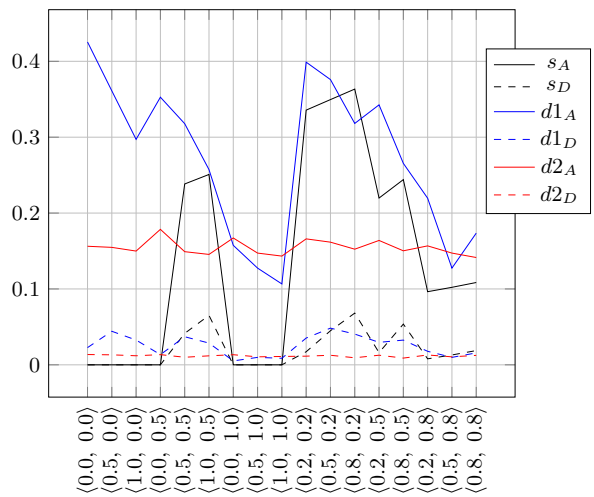

Figure 3.12: Task completion percentages for $P_D = 0.8$. Notation interpretation as in Figure 8.4.

behaviour of these agents in such context, the previously developed simulation set-up (Papers *A*-*C*) was extended. Furthermore, two other modifications in the behavior of agents ware made. (i) If an agent receives a help request, whilst itself is waiting upon the help of another agent, then it will immediately decline giving help. (ii) An agent will ask for help all agents it knows, one after the other, until there are no more agents, or its request has been successful. Agent performance is investigated across two scenarios that increase in difficulty. The difficulty is specified as the ratio between the maximum amount of resources an agent can have to perform a task over the amount of resources required initially at any site. The amount of resources for fire fighter agents is fixed at 25 units which refer to the substance used in extinguishing, whereas for ambulances is fixed at 5 and refers to available spots for individuals. The resources required at each site, depending on the difficulty, can be $\langle 75, 15\rangle, \langle 100, 20\rangle]$. These values are fixed and the same for all sites distributed in the 2D space. In total, thirty independent trials were conducted for each configuration and scenario difficulty. Performance is measured through four metrics: (i) percentage of individuals saved (% saved), (ii) percentage of fires extinguished (% fires), (iii) percentage of returns to the base with respect to tasks completed (task per return), (iv) percentage of completed dependent tasks (% dependent). Results are displayed as probability density functions (PDFs) in Figures 3.13 and 3.14. Each sub-figure contains the outcomes for the eight initial configurations and the two modes of execution (the static and dynamic strategies). The red and magenta lines refer to the static case (indicated with s in the legend), while the blue and black lines refer to the dynamic strategy (indicated with d in the legend). Results show the efficacy and stability of the proposed model across four scenarios increasing in difficulty, against the static counterparts. It can be observed that all agents, except static ones with $\gamma_0 = 0.5$ manage to extinguish all fires and save all individuals. Moreover, dynamic agents complete a high percentage of dependent tasks across difficulty scenarios and initial configurations, and have a higher ratio of completed tasks versus return to the base. Static agents with $\gamma = 0.0,\ 0.2$ show the best performance as compared to other static agents configured differently. In these cases the number of collaborations is minimal, in fact for $\gamma = 0.0$ no collaborations take place.

The produced results point to the effectiveness of a homogeneous society of agents, initiated with the same willingness to interact, where they are inclined to help one another, given that the need for such help is limited, as discussed in Papers *A* and *C*. Paper *B* considers three different scenarios in which the population of agents is not homogeneous. The issue of exploitation is investigated and it is shown how such phenomenon can be controlled by adjusting

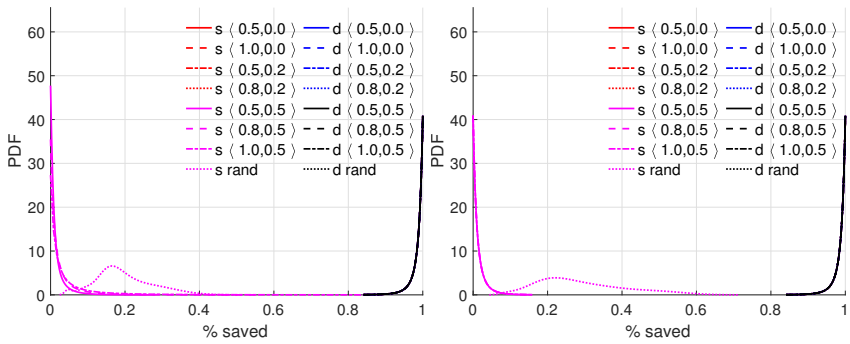

(a) % of saved individuals for SC1

(b) % of saved individuals for SC2

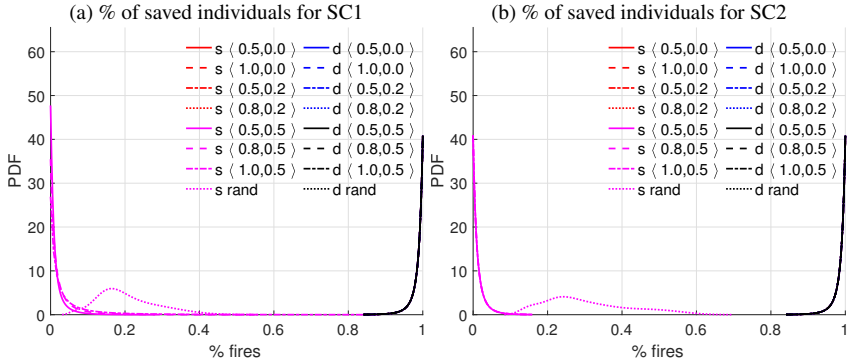

(c) % of extinguished fires for SC1

(d) % of extinguished fires for SC2

Figure 3.13: PDFs of the results for the *% saved* and *% fires* metrics, in SC1–2.

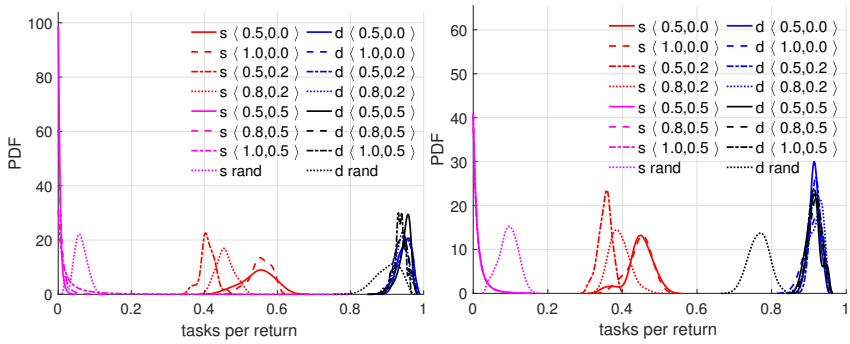

(a) Task per return to base for SC1

(b) Task per return to base for SC2

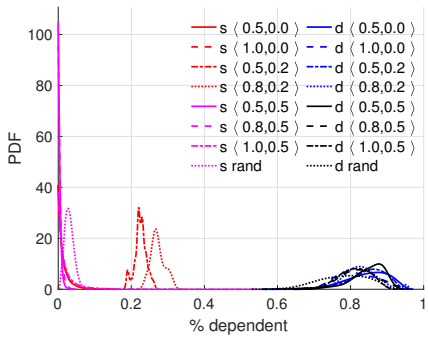

(c) % dependent for SC1

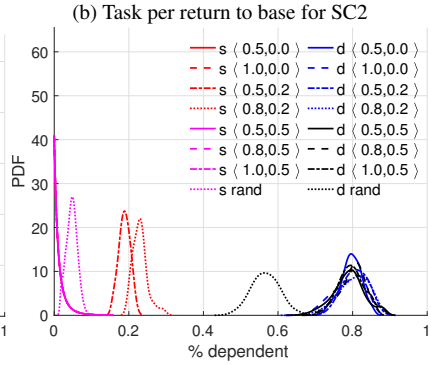

(d) % dependent for SC2

Figure 3.14: PDFs of the results for the *task per return* and *% dependent* metrics, in SC1–2.

accordingly the willingness to give help. Based on the results displayed in Papers *A*-*C*, the goal of an agent can be formulated as follows: the agent aims to be as helpful as it can, whilst not being inclined to ask for help. This goal is explicitly adopted in Paper *D*, where the results indicate the stability of the performance produced by the proposed model. Finally, the platform on which the simulation run can have an impact on the results. It is expected that in more powerful machines, the results for all agents will improve overall, whilst in less powerful ones, the results will be poorer.

Contribution C3 addresses research question RQ2. Agents with configurations that have produced higher results in the previous Papers *A* and *C* have been investigated further in Paper *D*. The results are used to shape the goal of agents with respect to the adaptation of the willingness to interact and have guided the development of the underlying equations.

Own contribution: I am the main driver of the four papers included in the thesis. Among my contributions are the modelling of adaptive autonomous behaviour through the willingness to interact, simulation design, implementation, and writing process. Nevertheless, my supervisors and co-authors have contributed with feedback with respect to different phases – from idea crystallization to results display – thus helping shape up the presented work.

Chapter 4

Related Work

This chapter discusses several works, in chronological order, that are considered to be closest to the research presented in this thesis, and draws comparisons to the proposed approach.

Static versus dynamic autonomy. Martin and Barber have investigated the benefits of adapting decision-making schemes – referred to as frameworks – as a result of changes in environmental conditions [22]. Three such frameworks are investigated, master-slave, peer to peer, and locally autonomous, and affect agents at the task level. Consequently, in the master-slave framework, an agent A with authority over another agent B, can assign tasks to B. Agent B is required to perform such task. In those scenarios when the communication is down, agents make local decisions without other considerations, thus they are locally autonomous. The third framework assumes that an agreement is achieved between agents, thus no agent has authority over others.

The scenario in which these three frameworks are compared is that of a group of three agents which manage radio frequencies on military ships. No involvement of human operators is present. Scripts are executed which map a specific framework to specific environmental conditions. This mapping is deduced by previous work from the authors. Therefore, the reasoning mechanisms that an agent might use to switch framework are not considered, as the aim of the authors is provide arguments in favour of adaptive autonomous behaviour. It is shown that by changing the decision framework during operation, agents perform better, then when under a single decision framework.

In the work presented in this thesis an agent model and specific reason-

ing mechanisms are developed which address exactly the issue of how agents should change their strategy, in the face of changing circumstances, which could refer to the environment, other agents, tasks, and the agent itself. The concept of adaptive autonomy itself is considered rather fluid, and it is not expressed in terms of switching a specific framework. The behaviour of the agent is shaped by its willingness to interact. Nevertheless, the performance of static agents is treated as a base-line for comparison with the dynamic ones. Furthermore, the benefit of adaptive agents over static ones, represents the main hypothesis of this work, and is investigated through different simulated scenarios.

Impact of external influences on decision-making. Van der Vecht *et al.* propose a meta-reasoning model – expressed through heuristics – that shapes how an agent is influenced by external factors [39]. These heuristics cover the relevance of received information with respect to the current goals, who is sending the message (is it a superior or a peer?), and how trustworthy is the sender of the message. Furthermore, the heuristic to adopt at some specific time is decided based on particular circumstances, thus the agent dynamically changes its behaviour in terms of how and by whom it is affected. These heuristics affect the priorities of the rules that shape the reasoning of the agent.

In their work, the authors adopt a fire-fighter scenario, composed of fire-fighter agents, and a superior coordinator who is able to send requests to the agents involved. These agents have two goals, extinguishing fires, and remaining healthy by keeping low levels of tiredness. In this setting, specific heuristics that agents can use include:

1. Task urgency. When this variable reaches high levels, then the agent will filter out any information that is not related to the task at hand. In case agents feel tired, and there is no time for a break, then the urgency of of the goal of taking a rest will increase.

2. Dedication level. Presumably, an agent part of an organization is dedicated and follows its norms. Moreover, dedication to the goals set by the organization is higher than those set by the self. Nevertheless, if wrong information keeps coming from the organization, then the dedication level will drop, and the agent will be more likely to take initiative in determining its own goals.

It is shown that adaptive behaviour can be achieved through this model, and that it performs better than the cases where no such meta-reasoning is used.

The work presented in this thesis relates to the concepts discussed by Van der Vecht *et al*. The ground for comparison is the consideration of specific factors, external or internal, that allow agents to display adaptive behaviour. Nevertheless, the model presented in this thesis builds on the willingness to interact, a rather generic concept that allows agents to display adaptive autonomous behaviour, and is not rule-based. Furthermore, a wider selection of factors is analysed, and a mathematical framework is developed that integrates their influences on the willingness to interact. Some of these factors are similar to task-urgency and dedication, such as task progress and success of collaboration, although they are not expressed in the same way.

Experiments in autonomy management. Hardin and Goodrich investigate, through a series of experiments, different schemes for changing autonomy within a team composed of a human operator and simulated robotic searchers, to the range of 200 [3]. Their interest lies in the relation between operator workload, prior knowledge, autonomy scheme, and the impact on performance.

Three different autonomy schemes are compared, adjustable autonomy, adaptive autonomy, and mixed-initiative interaction. Adjustable autonomy gives the operator the ability to select the correct level of autonomy for the searchers involved, whereas adaptive autonomous agents decide themselves about the desired level of autonomy. Mixed-initiative interaction allows both agent and operator to make decisions on the adequate autonomy level based on the circumstances.

Performance is measured through primary and secondary metrics, among which number of individuals found, and number of cues found in the environment, respectively. Their experiments indicate that mixed-initiative interaction achieves better results with respect to primary metrics, whilst adaptive autonomy achieves higher scores on the secondary metrics. This is a result of a better symbiosis of operator and searchers in mixed-initiative interaction. The latter will allow the operator more interferences, thus lowering a searcher's own autonomy, and is especially beneficial when the operator can provide relevant information. In the adaptive autonomy case, searchers attempt at all times to maintain a higher level of autonomy, and could not be directly given commands by the operators.

In this thesis, the concept of adaptive autonomy is rather relaxed. At the core, it is still assumed that the agent itself decides when to change its autonomy, that is when to depend on others, or whether to accept requests, thus allowing others to influence own goals. Nevertheless, it is still possible for agents to influence each other, by requesting assistance for the achievement of

some task. Furthermore, the adaptive behaviour is considered to be resultant of deciding on whether to interact based on particular circumstances. Agents are wired such that they will tend to help as much as possible, whilst attempting tasks mostly on their own. As a result, complete independence is not particularly relevant to the definition of adaptive autonomy, as is shaping one's behaviour as a response to the change of what are considered relevant parameters. Finally, this thesis does not consider human operators explicitly. An agent can in principle be a human, but there is no modelling of human behaviour.

Chapter 5

Conclusions and future work

This chapter concludes on the contributions and obtained results presented in the previous chapters, and outlines potential directions of research for future work.

5.1 Conclusions

The main goal of the thesis was to develop an agent architecture that supports adaptive autonomous behaviour, such that the performance of such agents is better than agents with static strategies. To this end, first it is shown how such an architecture can be designed, by modelling in generic terms the internal operation of an agent, and using the willingness to interact as the one concept that underlies adaptive autonomous behaviour. Furthermore, the components of the willingness to interact are analysed. Several factors are argued to be relevant in their estimation. Thereafter, a mathematical framework is developed which calculates each component of the willingness to interact. Second, the agent architecture and model are evaluated through computer simulations in order to evaluate the main hypothesis of this work, which states that adaptive autonomous agents have higher performance than those agents which do not change their autonomy. Overall, the obtained results support the hypothesis. Agents that are more likely not to ask for help, but very likely to give help will perform better, by completing more dependent tasks, than agents configured differently. Nevertheless, there are cases in which agents with a static strategy perform just as good, or even better. It is also shown how in hetero-

geneous agent populations, where there is a mix of selfish and selfless agents, it is possible in some settings to control agent exploitation, by using the perceived helpfulness of the population as an additional cue in the calculation of the willingness to give help. Finally, the model has been applied to a search and rescue simulated scenario, borrowed from the search and rescue application domain. The results are consistent with the previous, in that agents with adaptive autonomy overall have better performance than the static one. Finally, the performance measure is stable not only across configurations but over different scenario difficulties as well.

5.2 Future Work

There are several directions for future work.

1. It is of interest to investigate how the adaptive autonomy agent model fits in a wider framework that also includes high-level planning [40] and human operators in the loop. Therein, the question lies in how to optimize the collaboration between the different components, considering that each represents a specific layer of abstraction, with the human being in the top level, and the agent being in the bottom one.

2. There are other dimensions of interaction such as the presence of hierarchy, authority, or societal norms. According to Castelfranchi [2], this aspect is orthogonal to all other possible dependencies that can arise between agents. Assuming that to be the case, it is interesting to explore how the model and the adaptive behaviour of the agent is affected, once deontic aspects are taken into account. This direction is in-line with the previous one.

3. The notion of trust, shallowly addressed in this work (Paper *B)*, remains relevant in multi-agent system domains, due to the importance that selecting a trustworthy collaborator has for a task to be successful. Trust influences hierarchies and norms as well [29]. It is of interest to explore these relations and their impact on agent autonomy.

4. Finally, interaction patterns between agents and emergent adaptive behaviour need to be studied further in the context of large agent populations, to the range of 100s and 1000s, to better understand the impact of the model. Furthermore, other scenarios from relevant application

domains can be used to evaluate the model, such that it is possible to generalize the current results.

Bibliography

[1] Christina Rödel, Susanne Stadler, Alexander Meschtscherjakov, and Manfred Tscheligi. Towards autonomous cars: the effect of autonomy levels on acceptance and user experience. In *Proceedings of the 6th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*, pages 1–8. ACM, 2014.

[2] Cristiano Castelfranchi. Founding agent's 'autonomy' on dependence theory. In *Proceedings of the 14th European Conference on Artificial Intelligence*, pages 353–357. IOS Press, 2000.

[3] Benjamin Hardin and Michael A Goodrich. On using mixed-initiative control: A perspective for managing large-scale robotic teams. In *Proceedings of the 4th ACM/IEEE international conference on Human robot interaction*, pages 165–172. ACM, 2009.

[4] Michael Wooldridge. Intelligent agents: The key concepts. In *ECCAI Advanced Course on Artificial Intelligence*, pages 3–43. Springer, 2001.

[5] Nicholas R Jennings, Katia Sycara, and Michael Wooldridge. A roadmap of agent research and development. *Autonomous agents and multi-agent systems*, 1(1):7–38, 1998.

[6] Michael Wooldridge and Nicholas R Jennings. Intelligent agents: Theory and practice. *The knowledge engineering review*, 10(2):115–152, 1995.

[7] Michael Wooldridge and Nicholas R Jennings. Agent theories, architectures, and languages: a survey. In *International Workshop on Agent Theories, Architectures, and Languages*, pages 1–39. Springer, 1994.

[8] Rodney Brooks. A robust layered control system for a mobile robot. *IEEE journal on robotics and automation*, 2(1):14–23, 1986.

[9] David Vernon, Giorgio Metta, and Giulio Sandini. A survey of artificial cognitive systems: Implications for the autonomous development of mental capabilities in computational agents. *IEEE transactions on evolutionary computation*, 11(2):151–180, 2007.

[10] Matthew Johnson, Jeffrey M Bradshaw, Paul J Feltovich, Catholijn M Jonker, Birna Van Riemsdijk, and Maarten Sierhuis. The fundamental principle of coactive design: Interdependence must shape autonomy. In *Coordination, organizations, institutions, and norms in agent systems VI*, pages 172–191. Springer, 2011.

[11] Raja Parasuraman, Thomas B Sheridan, and Christopher D Wickens. A model for types and levels of human interaction with automation. *IEEE Transactions on systems, man, and cybernetics-Part A: Systems and Humans*, 30(3):286–297, 2000.

[12] Jonathan Brookshire, Sanjiv Singh, and Reid Simmons. Preliminary results in sliding autonomy for assembly by coordinated teams. In *Intelligent Robots and Systems, 2004.(IROS 2004). Proceedings. 2004 IEEE/RSJ International Conference on*, volume 1, pages 706–711. IEEE, 2004.

[13] Terrence Fong, Charles Thorpe, and Charles Baur. *Collaborative control: A robot-centric model for vehicle teleoperation*, volume 1. Carnegie Mellon University, The Robotics Institute, 2001.

[14] Michael A Goodrich, Dan R Olsen, Jacob W Crandall, and Thomas J Palmer. Experiments in adjustable autonomy. In *Proceedings of IJCAI Workshop on Autonomy, Delegation and Control: Interacting with Intelligent Agents*, pages 1624–1629, 2001.

[15] Matthew Johnson, Jeffrey M Bradshaw, Paul J Feltovich, Catholijn M Jonker, M Birna Van Riemsdijk, and Maarten Sierhuis. Coactive design: Designing support for interdependence in joint activity. *Journal of Human-Robot Interaction, 3 (1), 2014*, 2014.

[16] Paul J Feltovich, Jeffrey M Bradshaw, William J Clancey, and Matthew Johnson. Toward an ontology of regulation: Socially-based support for coordination in human and machine joint activity. In *International Workshop on Engineering Societies in the Agents World*, pages 175–192. Springer, 2006.

[17] Jeffrey M Bradshaw, Hyuckchul Jung, Shri Kulkarni, Matthew Johnson, Paul Feltovich, James Allen, Larry Bunch, Nathanael Chambers, Lucian Galescu, Renia Jeffers, et al. Kaa: policy-based explorations of a richer model for adjustable autonomy. In *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, pages 214–221. ACM, 2005.

[18] Milind Tambe. Agent architectures for flexible. In *Proc. of the 14th National Conf. on AI, USA: AAAI press*, pages 22–28, 1997.

[19] John E Laird. *The Soar cognitive architecture*. MIT Press, 2012.

[20] Nathan Schurr, Janusz Marecki, and Milind Tambe. Improving adjustable autonomy strategies for time-critical domains. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, pages 353–360. International Foundation for Autonomous Agents and Multiagent Systems, 2009.

[21] Stephen G McGill, Seung-Joon Yi, Hak Yi, Min Sung Ahn, Sanghyun Cho, Kevin Liu, Daniel Sun, Bhoram Lee, Heejin Jeong, Jinwook Huh, et al. Team thor's entry in the darpa robotics challenge finals 2015. *Journal of Field Robotics*, 2016.

[22] Cheryl Martin and K Suzanne Barber. Adaptive decision-making frameworks for dynamic multi-agent organizational change. *Autonomous Agents and Multi-Agent Systems*, 13(3):391–428, 2006.

[23] Bennie Lewis, Bulent Tastan, and Gita Sukthankar. An adjustable autonomy paradigm for adapting to expert-novice differences. In *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*, pages 1656–1662. IEEE, 2013.

[24] Sebastian Muszynski, Jörg Stückler, and Sven Behnke. Adjustable autonomy for mobile teleoperation of personal service robots. In *RO-MAN, 2012 IEEE*, pages 933–940. IEEE, 2012.

[25] Andreas Birk and Max Pfingsthorn. A hmi supporting adjustable autonomy of rescue robots. In *Robot Soccer World Cup*, pages 255–266. Springer, 2005.

[26] Gloria L Calhoun, Michael A Goodrich, John R Dougherty, and Julie A Adams. Human-autonomy collaboration and coordination toward multi-rpa missions. *Remotely Piloted Aircraft Systems: A Human Systems Integration Perspective*, page 109, 2016.

[27] Michael J Barnes, Jessie YC Chen, and Florian Jentsch. Designing for mixed-initiative interactions between human and autonomous systems in complex environments. In *Systems, Man, and Cybernetics (SMC), 2015 IEEE International Conference on*, pages 1386–1390. IEEE, 2015.

[28] Kristen Stubbs, Pamela J Hinds, and David Wettergreen. Autonomy and common ground in human-robot interaction: A field study. *IEEE Intelligent Systems*, 22(2), 2007.

[29] Bob van der Vecht, Frank Dignum, and JJ Ch Meyer. Autonomy and coordination: Controlling external influences on decision making. In *Web Intelligence and Intelligent Agent Technologies, 2009. WI-IAT'09. IEEE/WIC/ACM International Joint Conferences on*, volume 2, pages 92–95. IEEE, 2009.

[30] Sanghyun Kim, Mingon Kim, Jimin Lee, Soonwook Hwang, Joonbo Chae, Beomyeong Park, Hyunbum Cho, Jaehoon Sim, Jaesug Jung, Hosang Lee, et al. Team snu's control strategies for enhancing a robot's capability: Lessons from the 2015 darpa robotics challenge finals. *Journal of Field Robotics*, 2016.

[31] Alessandro Farinelli, Masoume M Raeissi, Nathan Brooks, Paul Scerri, et al. Interacting with team oriented plans in multi-robot systems. *Autonomous Agents and Multi-Agent Systems*, pages 1–30, 2016.

[32] Marvin V Zelkowitz and Dolores Wallace. Experimental validation in software engineering. *Information and Software Technology*, 39(11):735–743, 1997.

[33] Gordana Dodig-Crnkovic. Scientific methods in computer science. In *Proceedings of the Conference for the Promotion of Research in IT at New Universities and at University Colleges in Sweden, Skövde, Suecia*, pages 126–130, 2002.

[34] Mary Shaw. What makes good research in software engineering? *International Journal on Software Tools for Technology Transfer*, 4(1):1–7, 2002.

[35] Hilary J Holz, Anne Applin, Bruria Haberman, Donald Joyce, Helen Purchase, and Catherine Reed. Research methods in computing: what are they, and how should we teach them? In *ACM SIGCSE Bulletin*, volume 38, pages 96–114. ACM, 2006.

[36] Mirgita Frasheri, Baran Çürüklü, and Mikael Ekström. Analysis of perceived helpfulness in adaptive autonomous agent populations. *LNCS Transactions on Computational Collective Intelligence*, 2017.

[37] Mirgita Frasheri, Baran Çürüklü, and Mikael Ekström. Towards collaborative adaptive autonomous agents. In *9th International Conference on Agents and Artificial Intelligence 2017 ICAART, 24 Feb 2017, Porto, Portugal*, 2016.

[38] Alexander Kleiner, Alessandro Farinelli, Sarvapali Ramchurn, Bing Shi, Fabio Maffioletti, and Riccardo Reffato. Rmasbench: benchmarking dynamic multi-agent coordination in urban search and rescue. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*, pages 1195–1196. International Foundation for Autonomous Agents and Multiagent Systems, 2013.

[39] B v Vecht, Frank Dignum, and JJ Meyer. Autonomy and coordination: Controlling external influences on decision making. In *Proceedings of the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology-Volume 02*, pages 92–95. IEEE Computer Society, 2009.

[40] Branko Miloradovic, Baran Çürüklü, and Mikael Ekström. A genetic mission planner for solving temporal multi-agent problems with concurrent tasks. In *International Conference on Swarms Intelligence*, August 2017.