Advancing the Acceptance and Use of Wheelchair-mounted Robotic Manipulators

by

Laura C. Petrich

A thesis submitted in partial fulfillment of the requirements for the degree of

Master of Science

Department of Computing Science

University of Alberta

Abstract

Wheelchair-mounted robotic manipulators have the potential to help the elderly and individuals living with disabilities carry out their activities of daily living independently. While robotics researchers focus on assistive tasks from the perspective of various control schemes and motion types, health research tends to concentrate on clinical assessment and rehabilitation. This difference in perspective often leads to the design and evaluation of experimental tasks that are tailored to specific robotic capabilities rather than solving tasks that would support independent living.

In addition, there are many studies in healthcare on which activities are relevant to functional independence, but little is known about how often these activities occur. Understanding which activities are frequently carried out during the day can help guide the development and prioritization of assistive robotic technology. By leveraging the strength of robotics (i.e., performing well on repeated tasks) these activities can be automated, significantly improving the quality of life for our target population.

Our first research goal is to investigate daily task frequency data in order to provide deeper insights and meaningful guidelines for future research developments in the field of assistive robotic manipulation. These guidelines are meant to shift focus towards better supporting the needs and performance requirements of the target population. While we have established that assistive robotic manipulators can help individuals regain functional independence, their performance is restricted by the underlying control system. Aside from

having direct control over the robots motion, each manipulator comes with a limited preprogrammed set of modes (e.g., drinking mode) that might not fully encompass the individual's needs. For mainstream use and acceptance, it is paramount to provide the target population with a way to refine the control and customize it to their personal needs. One method for learning new behaviours and skills on the fly is interactive shaping.

Interactive shaping is a method of transferring knowledge from a human to a learning agent by having the human teacher provide signals of approval (or disapproval) from instantaneous observations of the robots behaviour. This could be of great value in an assistive setting as it does not require the teacher to have expert domain knowledge. The teacher only needs to give feedback on whether the robot's previous action was "good" or "bad" in order to teach it new behaviours and skills, regardless of task difficulty. Our second research goal is to investigate whether interactive shaping can be used to teach robotic manipulators new autonomous tasks. To this end, we adapt TAMER, a framework for learning from human reward signals, to a seven degree of freedom robotic manipulator and carry out a proof-of-concept user study.

The work in this thesis is meant to open an avenue to better align research with the needs of individuals that will eventually leverage the technology in their daily life. To do so, we first bridge the gap between robotics and healthcare research to align both with respect to the target population's needs. Taking it one step further, we then adapt TAMER to allow the users themselves to add new autonomous behaviours to their wheelchair-mounted robotic manipulator. Together, these results introduce a new level of long-term autonomy for individuals living with disabilities.

Preface

In January 2018 I began working with the Computer Vision and Robotics Research Group led by Professor Martin Jägersand. For the first few months Masood Dehghan, Vincent Zhang, Menna Siam, Jun Jin and I worked on the Kuka Innovation award, building a robotic system capable of learning new objects and associated motions incrementally on the fly via human robot interaction. This work was published and presented as Dehghan, M., Zhang, Z., Siam, M., Jin, J., Petrich, L., and Jägersand, M., "Online Object and Task Learning via Human Robot Interaction," *in 2019 IEEE International Conference on Robotics and Automation (ICRA).* I contributed the graphical user interface that was employed to connect the deep learning based localization and recognition system with the hybrid force-vision robot control module.

In the summer of 2018 I conceived the idea of joining healthcare activities of daily living research with lifelogging data to gain insight into what functionalities would be required of assistive robotic manipulation systems with the goal of improving functional independence. This research serves as the foundation of Chapter 3. I organized and carried out the literature and dataset review, as well as led the task annotation and analysis of the NTCIR video dataset. Eisha Ahmed (an undergraduate student I was mentoring at the time) helped with annotating the video dataset with high level activity labels. Prof. Jägersand worked closely with me on this project and led the arm and hand motion analysis presented in Section 3.3. Prof. Jägersand and I wrote the first draft of this project in September 2018 for submission to ICRA. All authors helped edit the manuscript through a number of submissions, and it was published and presented as Petrich, L., Jin, J., Dehghan, M., and Jägersand, M., "A Quantitative Analysis of Activities of Daily Living: Insights into Improving

Functional Independence with Assistive Robotics," *in 2022 IEEE International Conference on Robotics and Automation (ICRA). Copyright © 2022, IEEE.* Chapter 3 is based on this work; for inclusion in this thesis changes have been made including: expansion of the literature review to include more recent works, addition of new insights and discussion to unify with Chapter 4, and general rewording and revision of the text.

During this time I also contributed to various projects in the lab (not included in this thesis) leading to two works published and presented at the *2019 IEEE International Conference on Robotics and Automation (ICRA)* as Siam, M., Jiang, C., Lu, S.W., Petrich, L., Gamal, M., Elhoseiny, M., and Jägersand, M., "Video Object Segmentation using Teacher-Student Adaptation in a Human Robot Interaction (HRI) Setting," and Jin, J., Petrich, L., Dehghan, M., Zhang, Z., and Jägersand, M., "Robot eye-hand coordination learning by watching human demonstrations: a task function approximation approach." In the first work I aided in setting up and collecting a video dataset of everyday kitchen tasks being carried out with a robotic arm; specifically I ran the underlying robot control during dataset collection. In the latter work with Jin et al. I developed the adaptive uncalibrated visual servoing module used for robot control and helped to carry out all robotic experiments. I was also significantly involved in manuscript writing and editing.

From 2018 to 2021 I worked closely with Jun Jin on his PhD research "Learning Geometry from Vision for Robotic Manipulation." This led to two additional publications (not included in this thesis): Jin, J., Petrich, L., Zhang, Z., Dehghan, M., and Jägersand, M., "Visual Geometric Skill Inference by Watching Human Demonstration," *in 2020 IEEE International Conference on Robotics and Automation (ICRA)*, and Jin, J., Petrich, L., Dehghan, M., and Jägersand, M., "A Geometric Perspective on Visual Imitation Learning," *in 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. I implemented the baselines, helped with the experiments, prepared the data, and participated in draft manuscript writing.

Chapter 4 of this thesis is based on a report submitted for credit during this MSc in 2020 for CMPUT 656: Interactive Machine Learning under the tutelage of Dr. Matthew Taylor. It was submitted as Petrich, L., and Przystupa, M., "Training a Robotic Arm Manipulator and Intent Inference from Human Reinforcement Feedback: A Case Study." Michael Przystupa and I both contributed equally to conceiving the idea of extending TAMER to the realm of robotic manipulation. I designed and developed the baseline control software for the robotic arm, took lead on carrying out the experiments, as well as manuscript composition. Przystupa implemented the human reward neural network model and helped integrate it with my robot control framework. He also generated figures, contributed helpful discussions, and was heavily involved in manuscript writing and editing. Note that for future publication the experiments in this chapter will need to be redone in a controlled setting as the ones included were carried out in my home during COVID lockdown. For inclusion in this thesis I have written new introduction and discussion chapters that place the TAMER project within the larger context of assistive technology. Chapters 3 and 4 are now part of an overarching narrative that adds additional insights that could not be drawn from each work individually.

Although not included in this thesis as a contribution, the underlying control software I developed during my MSc studies for the Kinova Gen3 ultra lightweight robot continues to be of use in the lab. This code base was used for the project Przystupa, M., Johnstonbaugh, K., Zhang, Z., Petrich, L., Dehghan, M., Haghverd, F., and Jägersand, M., "Learning State Conditioned Linear Mappings for Low-Dimensional Control of Robotic Manipulators," *submitted to 2023 IEEE International Conference on Robotics and Automation (ICRA).*

To Shaylee and Charlotte

For your unconditional support and inspiration.

What you do makes a difference, and you have to decide what kind of difference you want to make.

– Dr. Jane Goodall

Acknowledgements

First and foremost, I would like to acknowledge the support and guidance of my supervisor Prof. Martin Jägersand. He introduced me to the world of robotics and taught me the importance of hands-on learning and real-world applications. I value our many conversations over the years, they helped shape me into the researcher I am today.

Secondly, I owe a lot of gratitude to my labmates for their invaluable help, guidance, and support over the years. I would like to express special thanks to Dr. Masood Dehghan who took me under his wing and not only taught me everything he could about robotic control, but became a steadfast friend and constant pillar of support. I am grateful to Dr. Jun Jin for the numerous valuable discussions and continuous encouragement to pursue my interests. I would also like to thank Michael Przystupa for all his help, advice, and many laughs during our projects together; I look forward to many more in the coming years. Finally, I would like to thank everyone I had the honor of working with in the vision and robotics group: Dr. Xuebin Qin, Zichen Zhang, Connor Stephens, Dr. Mennatullah Siam, Jakub Piwowarczyk, Chen Jiang, Dhruv Sharma, Steven Lu, Junaid Ahmad, and Kerrick Johnstonbaugh.

I would like to acknowledge Dr. Matthew E. Taylor for his enthusiastic encouragement and advice when extending TAMER to the realm of robotic manipulation. I would also like to thank Dr. Pierre Lemelin for his thought-provoking discussions and teachings on human anatomy, they made me think about robotic manipulation in a new light. I am grateful to the Natural Sciences and Engineering Research Council of Canada (NSERC) and Alberta Innovates for their financial support during my studies.

I would like to thank my family for their unconditional support and love, none of this would have been possible without you here with me. Last but not least, I would like to thank my daughters, Shaylee and Charlotte, for putting up with all my shenanigans over the years and motivating me to keep going.

Contents

List of Tables

List of Figures

Acronyms

ADL Activities of Daily Living

ALS Amyotrophic Lateral Sclerosis

CV Computer Vision

DoC Degrees of Control

DoF Degrees of Freedom

GTEA Georgia Tech Egocentric Activity

HRI Human Robot Interaction

IADL Instrumental Activities of Daily Living

ICF International Classification of Functioning, Disability and Health

IoT Internet of Things

IRL Inverse Reinforcement Learning

LfD Learning from Demonstration

MDP Markov Decision Process

MIT Massachusetts Institute of Technology

NTCIR NII Testbeds and Community for Information access Research

RL Reinforcement Learning

SD Standard Deviation

TAMER Training an Agent Manually via Evaluative Reinforcement

WHODAS World Health Organization Disability Assessment Schedule

WMRM Wheelchair-mounted Robotic Manipulator

Chapter 1

Introduction

1.1 Motivation

Robotic assistive technologies focus on using robots to maintain and improve functional capabilities by assisting individuals living with disabilities and those who need extra help due to aging [21]. The overarching goal of this thesis is to promote the acceptance and use of wheelchair-mounted assistive robotic manipulators (WMRM) in order to improve functional independence. Here assistive robotics is defined as robotic arm systems designed to help a human achieve their manipulation goals by physically interacting with both the human and the surrounding environment. This falls within the broader field of physical human robot interaction (HRI) where robots physically interact with the environment while working alongside humans. In this thesis we focus on assistive robotic manipulators, where the goal is to assist end-users (i.e., individuals living with disabilities and the elderly) during the completion of their activities of daily living (ADL).

There have been rapid advances in the field of assistive robotics in recent years where robots work to improve the independence and quality of life of persons living with disabilities. Individuals with diminished arm and/or hand function may be completely dependent on the help of carers or relatives for carrying out daily tasks such as eating, drinking, opening doors, and other day-to-day activities. There are a wide range of robotic assistive devices available, such as single-use devices (e.g., robotic feeders), socially assistive robots, and robotic manipulators [9]. These systems are highly variable in the tasks they are capable of performing and the level of autonomy they provide. Wheelchair-mounted robotic manipulators (WMRMs) like the iARM produced by Exact Dynamics [19] or JACO by Kinova [48] have the potential to be of assistance in situations where individuals require caregiver help with their day-to-day activities, promoting functional independence.

Prior works have established five key factors that influence an end-users acceptance of a rehabilitation robot; thus determining market success. These factors are *cost*, *reliability*, *appearance*, *ease of use* (i.e., *usability*), and *function* [65]. However, to the best of our knowledge, no research has investigated how these factors contribute quantitatively towards acceptance. One could easily imagine that an assistive system that switches on a light (*functionality*) but takes fifteen minutes to do so (*usability*) would not be acceptable. In this work we focus on quantifying *usability* and *functionality* in order to increase user acceptance of assistive robotic systems. While *cost* and *appearance* are also of importance, they are beyond the scope of this thesis. We set out to gain insight into what is hindering the acceptance and use of assistive manipulation technologies and provide future directions to mitigate known challenges. Specifically, we are interested in influencing the design of control schematics for assistive robotics towards systems that are *intuitive*, *functional*, *reliable*, and *safe to use* in the real world.

Another key hindrance for the usage of assistive robotic devices is their limited customizability. Engineers cannot possibly anticipate all different real-world environments and situations these devices have to function well in. This will naturally result in a degraded performance when the system is used under real-world conditions. It is therefore essential to allow the user to *customize* the robotic arm to their individual needs — a key factor to success that has been neglected so far. We therefore propose and investigate the usage of TAMER, a machine-learning based approach to add new autonomous skills to the assistive device.

1.2 Research Questions

In this thesis we focus on two key research questions with the overarching goal of advancing the acceptance and use of wheelchair-mounted robotic manipulators. The first research question we aim to answer is — *what is hindering the acceptance and use of wheelchair-mounted robotic arms in the real-world?* To answer this question we look at existing literature and begin to quantify *functionality* and *usability* as key factors to the acceptance of assistive robotic systems. Our first research question is — *which daily living tasks should these devices be capable of carrying out?* To being answering this question we first look at existing literature to gain insight into what is hindering the acceptance and use of wheelchair-mounted robotic arms in the real-world; we identify *functionality* and *usability* as key factors to the acceptance of assistive robotic systems. We then carry out an analysis of existing lifelogging datasets collected with able-bodied participants. By looking at the lifelogging data, we identify a disconnect between the robotics community and the rehabilitation community. So far, development is focused on the assistive device rather than the needs of the intended user. Using the lifelogging data, we identify frequently executed tasks that should be automated in order to decrease the cognitive and physical load placed on the end user. We then deduct task guidelines to augment future robotics research in order to better align with the end user.

Our second research question is — *how can individuals easily teach their wheelchair-mounted robotic manipulator new tasks?* To this end, we introduce the usage of TAMER to teach robotic manipulators new skills. By providing simple signals of approval or disapproval, the end user can shape the behaviour of its robot to better fit their needs without the help of an expert. This will allow the user to customize the assistive system to their personal needs, significantly improving their quality of life. We carry out a proof-of-concept user study showcasing how TAMER can be used to teach a robot arm how to reach — a task that was highlighted in our lifelogging analysis to occur frequently.

1.3 Thesis Outline

In order to understand the rapidly changing field of assistive robotic manipulation and reduce the chance of technology abandonment, it is necessary to gain insight into the needs of end-users [54]. Chapter 2 will first provide background on activities of daily living and functional independence. This serves to motivate why it is important to understand and consider the needs of the target population while designing new assistive robotic systems. We then touch on robot arm control, including degrees of control and mode switching. Finally, this chapter covers decision-making techniques in assistive robotic manipulation that are built on in Chapter 4.

Chapter 3 presents an analysis of life-logging datasets in order to provide insight into what tasks would be of high priority for robotics researchers to concentrate efforts on. This section can serve as a guide to designing research protocols and robots focused on meeting the needs of the target population.

In Chapter 4 we present a proof-of-concept system for teaching new manipulation behaviours and skills through interactive shaping. Our proposed system adapts TAMER, a framework for learning from human reward signals, to a seven degree of freedom robotic manipulator and highlights the trade-off between ease of use, task completion time, and varying levels of control.

Lastly, Chapter 5 summarizes key concepts presented in this thesis and provides directions for future research.

Chapter 2

Background

In this thesis we investigate what is impeding the acceptance and use of WMRMs in society. We present guidelines on where to focus future research efforts and developments in the field of assistive robotic manipulation and motivate why it is important to understand and consider the needs of the target population (i.e., individuals that would benefit from having a WMRM available for everyday use) during the experimental design process.

This chapter will provide the necessary background material for this thesis. It starts by describing what defines assistive robotic manipulation, functional independence, and providing motivation for improving this technology. We highlight factors imperative to the successful deployment of assistive technology, as well as what is impeding its acceptance and use in the real-world. This chapter also begins to answer our first research question — *which daily living tasks should these devices be capable of carrying out?* — by looking into previous studies on task importance and priorities in the target population.

We then introduce how robotic manipulators are commonly controlled and various methods for learning how to carry out tasks autonomously. The problem of learning from human reward, or interactive shaping, is described in detail. A framework for interactive shaping is extended in Chapter 4 of this thesis to learn new autonomous behaviours.

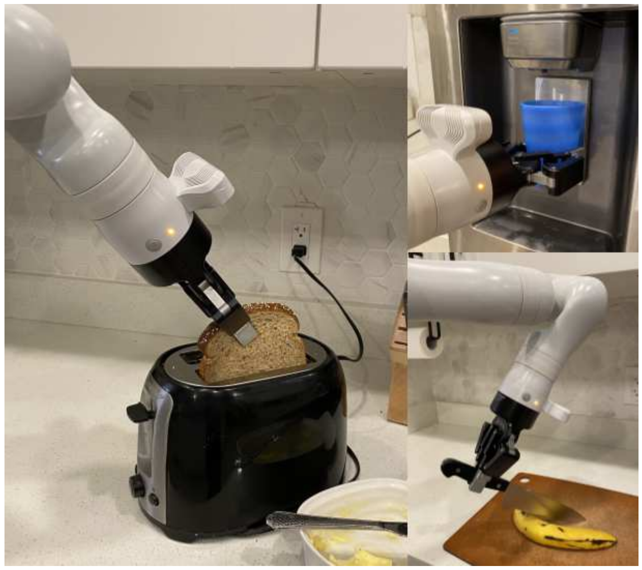

Figure 2.1: Wheelchair-mounted robotic manipulators can help promote independent living by providing individuals living with upper-body disabilities with the means to carry out their activities of daily living on their own. *Copyright © 2022, IEEE.*

2.1 Assistive Robotics

Activities of daily living (ADLs) can be a challenge for individuals living with upper-body disabilities and assistive robotic arms have the potential to help increase functional independence [63]. Wheelchair-mounted robotic manipulators (WMRMs), such as the Kinova Jaco [3] and Manus/iArm [19], have been commercially available for over a decade (see Figure 2.1). These devices can help to increase independence while decreasing the caregiver load and reducing healthcare costs [59]. WMRMs have the potential to be as important to individuals living with upper-body disabilities as power wheelchairs have become to those with lower-body disabilities. However, outside of research purposes, only a few hundred assistive arms, primarily in Europe and North America, are practically deployed and regularly used.

For the successful commercialization and acceptance of WMRMs in the general population it is imperative to understand what factors are related to assistive technology abandonment or disuse. In a survey study on device selection, acquisition, performance and use, four factors were found to be significantly related to device abandonment: *lack of consideration of user opinion*, *ease of device procurement*, *poor device performance*, and a *change in user needs or priorities* [54]. Chung et al. (2013) further claim that *reliability*, *cost-efficiency*, *appearance*, *functionality*, and *usability* are key factors necessary for the successful deployment of assistive robotic manipulators [13]. In this chapter we will focus on gaining insights into *functionality* and *usability* metrics while Chapter 4 proposes a proof-of-concept system that could adapt to the end-users *changing needs or priorities*.

The gap between robotic research and healthcare needs impedes the adoption of assistive devices. Healthcare professionals, assistive technology users, and researchers have differing biases as to which daily living tasks effort should be focused on. For assistive robotics research, knowing which ADLs are of high importance in the target population, as well as the necessary performance parameters for those high-priority tasks, will be crucial for real-world usability and deployment. In this thesis, we focus on establishing an understanding of what daily tasks are a priority (i.e., *functionality*) and what are acceptable time limits for accomplishing a specific task (i.e., *usability*). This would provide insight into what the target population requires of WMRM systems in order to decrease the risk of device abandonment and enhance user satisfaction. To build a task priority guide with a focus on functional independence, it is important to understand what defines independence and what is required to live independently; to this end we briefly review the World Health Organization Disability Assessment Schedule (WHODAS) in Section 2.1.1 [73], [77]. It should be noted that this classification was developed to determine an individual's level of disability and design appropriate rehabilitation plans, not to guide assistive robotics research.

2.1.1 Functional Independence

Real-world acceptance and use of WMRMs depends on user satisfaction in the technology. If a user is not satisfied then they are more likely to abandon the assistive technology. To this end, we argue that understanding what *functionality* the target population requires of the system is crucial and that the first

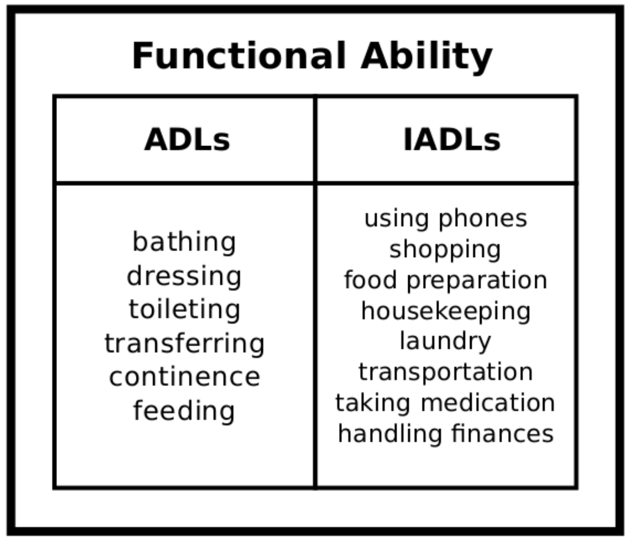

Figure 2.2: Measures of functional ability involve assessing the potential capacity of a person to perform the tasks and activities normally expected to be carried out everyday: Activities of Daily Living and Instrumental Activities of Daily Living.

step towards this is considering what defines functional independence. Information about an individual's ability is an important indicator of a population's health status, as it shows the impact functional limitations have on independence. This concept is known as functional disability, or the limitations one may experience in performing independent living tasks [64]. Measures of functional ability involve assessing the potential capacity of a person to perform the tasks and activities normally expected to be carried out everyday; these tasks are referred to as Activities of Daily Living (ADL) [35] and Instrumental Activities of Daily Living (IADL) [46]. ADLs are basic self-care tasks essential for independent living: bathing, dressing, toileting, transferring, continence, and feeding. IADLs are more complex tasks that are still a necessary part of everyday life, but require a higher level of autonomy: using phones, shopping, food preparation, housekeeping, laundry, transportation, taking medication, and handling finances. Together they can be viewed as a high-level priority list of key life tasks to help guide the development of assistive robotic systems (see Figure 2.2); in this work we will refer to these collectively as ADLs.

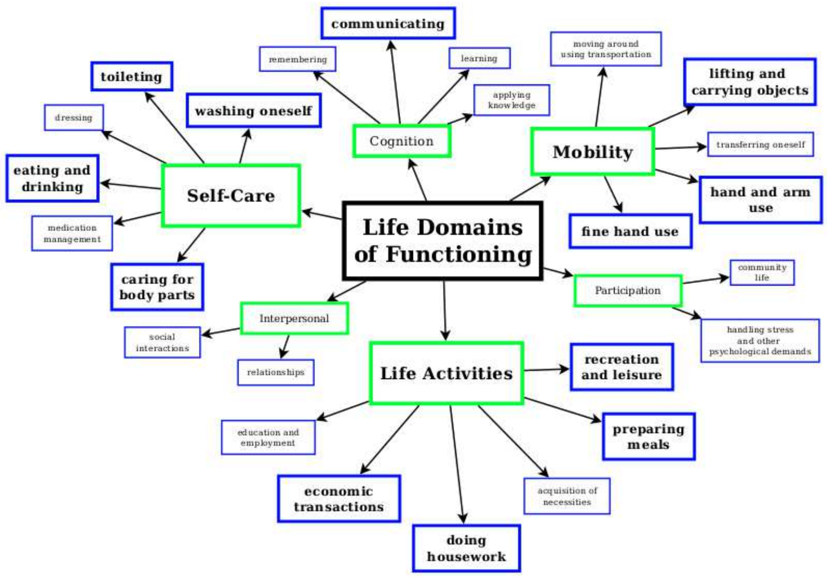

Figure 2.3: The major life domains of functioning and disability as set out in WHODAS2.0; physical manipulation activities relevant to robotics are highlighted in bold. *Copyright © 2022, IEEE.*

The International Classification of Functioning, Disability and Health (ICF) provides a framework for determining the overall health of individuals and populations [77]. The World Health Organization further developed the World Health Organization Disability Assessment Schedule (WHODAS2.0) from ICF as a standardized, cross-cultural measure of functioning and disability across all life domains [24], [73]. WHODAS2.0 is used to measure the impact of health conditions, monitor intervention effectiveness, and estimate the burden of physical and mental disorders across all major life domains. Figure 2.3 highlights these major life domains with associated tasks; the tasks most relevant to robotics research are emphasized in bold.

There have been many works that aim to inform the design and evaluation of assistive robotics that effectively meet the needs and preferences of individuals living with disabilities. This involves understanding how individuals use robotic arms in their daily activities, what types of tasks they are able to perform, and how to appropriately assess their impact on quality of life [5].

Tanaka et al. propose a framework of evaluation and design of assistive

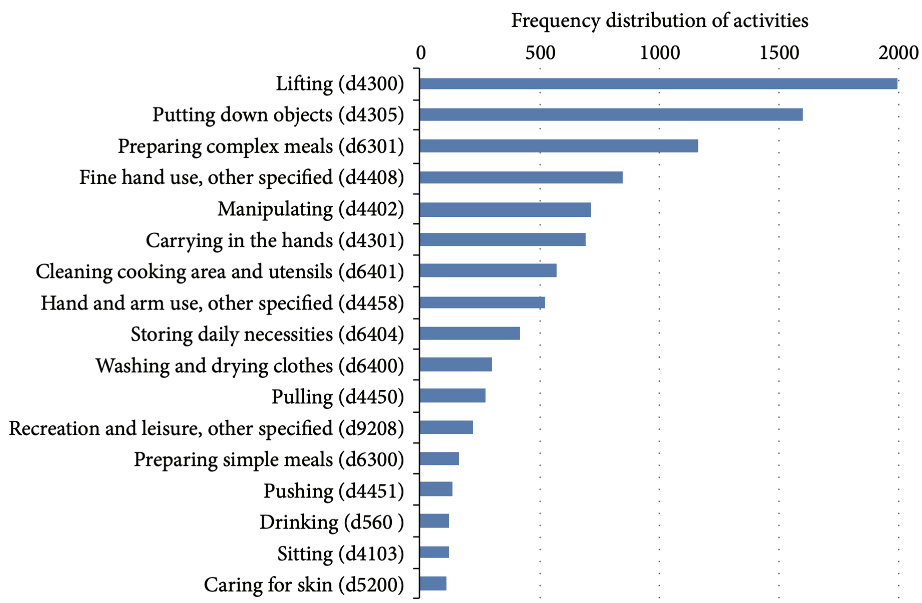

Figure 2.4: Frequency of distribution of activities extracted from life-log data captured for one person over five days [50]. *Note.* Reprinted from "A concept of needs-oriented design and evaluation of assistive robots based on ICF," by Y. Matsumoto, Y. Nishida, Y. Motomura and Y. Okawa, 2011, 2011 IEEE International Conference on Rehabilitation Robotics, p. 689. *Copyright © 2011, IEEE.*

robots based on the ICF. They claim that a cost-benefit and risk-benefit evaluation of assistive technologies will require estimating the objective and quantitative contribution of the device to improving the end users quality of life [69]. They recorded a voice recording life-log (action, time, duration, place, target object, and purpose) of a healthy individual capturing 3964 activities over the course of five days. Figure 2.4 shows the histogram of all recorded activities. Key findings of their study show that "pick and place" tasks - which are heavily studied in robotics - are the most frequently performed activities of daily life (corresponding to "Lifting" and "Putting down objects" in ICF). Through further analysis of the life-log they discovered that 90% of the objects manipulated during "lifting" activities were less than 300g. This finding is of importance to note as it implies that WMRM do not need to be capable of handling high payloads (which would make them heavier and more expensive) in order to have a large impact on quality of life.

Langdon et al. also propose a task taxonomy based on the ICF in order to analyze a robotic manipulators performance. They postulate that understanding the performance on a set list of tasks could help predict future performance on similar tasks, categorize different difficulties that arise during task completion, and help develop training strategies for effective operation [44].

With the goal of better understanding how assistive robotics can meet the needs and preferences of people with amyotrophic lateral sclerosis (ALS), the Healthcare Robotics Lab at Georgia Tech conducted a study to identify a prioritized list of 43 object classes for robotic retrieval [12]. They note that, especially in robotic manipulation, there is a lack of agreed upon benchmarks for evaluating the performance of assistive robotics systems. Existing research evaluation methods for object retrieval suffer from one or more of following drawbacks: (1) an insufficient number of objects (loss of generality); (2) insufficient variation in object type; and (3) objects without justification (objects cherry-picked to match robot capabilities). The ranked list of 43 everyday objects that Choi et al. recommend using in robotics research can be found in Table 2.1.

Rank	Object Class	Weight (grams)	Max Size (cm)
1	TV Remote	90	18
2	Medicine Pill	1	2.2
3	Cordless Phone	117	15
4	Prescription Bottle	25	7
5	Fork	39	18
6	Glasses	23	14
7	Toothbrush	15	19
8	Spoon	38	17
9	Cell Phone	76	9
10	Toothpaste	160	20
11	Book	532	24
12	Hand Towel	65	58
13	Mail	22	24
14	Cup / Mug	267	12
15	Soap	116	9.5
16	Disposable bottle	500	13
17	Shoe	372	30
18	Dish Bowl	154	13
19	Keys	24	8.5
20	Dish Plate	182	18
21	Pen / Pencil	3	14
22	Table Knife	76	24
23	Credit Card	5	8.5
24	Medicine Box	25	10
25	Bill (Money)	1	13.5
26	Straw	1	20
27	Magazine	206	27.5
28	Plastic container	49	13
29	Newspaper	247	31
30	Non-disposable bottle	709	20
31	Pants	539	100
32	Shirt	229	66
33	Wallet	116	100
34	Small Pillow	240	38
35	Socks	41	23
36	Hairbrush	100	24
37	Soda Can	350	6.4
38	Coin	6	2.5
39	Walking Cane	1140	94
40	Wrist Watch	86	10
41	Scissors	25	14
42	Purse / Handbag	380	24
43	Lighter	91	6

Table 2.1: A prioritized list of objects [12]. *Note.* Adapted from "A list of household objects for robotic retrieval prioritized by people with ALS," by Y. S. Choi, T. Deyle, T. Chen, J. D. Glass, and C. C. Kemp, 2009, 2009 IEEE International Conference on Rehabilitation Robotics, p. 515. *Copyright © 2009, IEEE.*

A common approach that drives research is to ask patients and caregivers for their preferences when it comes to robotic assistance [6], [65]. Notably, preferences vary and user opinions shift over time. Stanger et al. review nine task priority surveys reflecting the views of over 200 potential users of WMRMs [65]. The surveys reviewed by Stanger et al. include four pre-development questionnaires focused on what tasks they anticipate using an assistive robot for as well as five post-development questionnaires on their actual functional use of a specific robot. Pre-development participants favor picking up dropped objects and leisure-related tasks, with a shift more towards work-related tasks post-development [13]. The authors note various contributing factors which influence an end user's task priorities including, but not limited to, the user's age, disability, familiarity with technology, whether their disability has been present from birth or from an abrupt injury, and living situation (in an institution, with a family member, or independently with an in-home caregiver). Although survey results show differences between participants, overall task preference results imply that the ability to manipulate a large range of everyday objects within an unstructured environment is of the utmost importance. A summary of the pre- and post-development survey results (adapted from Chung et al. [13]) are shown in Table 2.2. In this thesis we combine user preferences obtained from these pre-existing survey results with quantitative ADL data from our lifelogging analysis in order to provide guidance to the robotic community on which daily life activities would make a meaningful impact in the target population.

2.1.2 Societal and Economic Impacts

The use of robotics to help increase functional independence in individuals living with upper-limb disabilities has been studied since the 1960's. With

Rank	Pre-development Survey Task	Post-development Survey Task
1	Cooking, fixing food, drinks	Work/school fetch and carry objects
2	Reaching, stretching, gripping, picking up objects	Personal hygiene
3	Gardening/hobbies and crafts/leisure	Feeding or eat/drink, preparing meals
4	Reach or pick up from the floor	Communication/phone
5	Personal hygiene, Dressing	Domestic opening doors, drawers, windows, closets, etc.

Table 2.2: A list of the top five preferred tasks from pre- and post-development user surveys [13]. *Note.* Adapted from "Functional assessment and performance evaluation for assistive robotic manipulators: Literature review," by C.-S. Chung, H. Wang, and R. A. Cooper, 2013, The Journal of Spinal Cord Medicine, vol. 36, no. 4, p. 278.

improved system *functionality*, *reliability*, and *ease of use*, more individuals in need of help with their daily living tasks could be reached. The United States Veterans Affairs estimate that approximately 150,000 Americans could benefit from currently commercially available wheelchair-mounted robot arms [13]. Many countries in the west and Asia have aging populations, and disabilities can affect anyone, regardless of age. Canada has a multi-ethnic population and characteristics similar to other industrialized nations. The proportion of seniors (age 65+) in Canada is steadily increasing, with seniors comprising a projected 23.1% of the population by 2031 [72]. In 2014, seniors constituted only 14% of the population, but consumed 46% of provincial public health care dollars [29].

Power wheelchairs allow individuals with reduced mobility move around independently. As the *reliability*, *functionality*, and *usability* of WMRMs improves, they could help increase independence and lower care needs for those living with reduced upper-limb function. Statistics Canada found from 2001 to 2006 there was a 20.5% increase in those identifying as having a disability, corresponding to over 2.4 million people in Canada [30]. One in twenty Canadians living with disabilities regularly receive assistance with at least one ADL on a daily basis, although not all of which will require the use of WMRMs. This suggests that there is a significant need for robotic solutions in Canada and similar countries world-wide. Some individuals may prefer automation integration with their smart homes, and some may require both cognitive and physical assistance. While artificial intelligence might provide some basic cognitive support, such as planning of the days tasks and reminders, it cannot eliminate the need for human contact and support. Robotic assistance can help free up humans from mundane, repetitive chores, allowing more time for caregivers to focus on high-quality help and personal interaction.

With an increasing portion of the population requiring help with ADLs additional pressure is placed on government budgets and healthcare personnel. A four year study of WMRM users found that a robot reduced the nursing assistance needed from 3.7 to 2.8 hours per day [59]. While cost savings from reduced nursing care are significant (approximately 20,000 USD per year), further savings and increased independence came from half of the users being able to move out of assistive living with one quarter being able to rejoin the workforce. Furthermore, a key advantage of WMRMs is that they are with the person at all times; this provides individuals with the freedom to independently complete their daily tasks whenever and however they desire.

2.2 Robot Control

A robot arm, or serial-link manipulator, is a chain of rigid links and joints [14]. A Degree of Freedom (DoF) is a joint that provides independent movement to the system, changing the relative pose of the connected links. A revolute joint provides one degree of rotational freedom. A prismatic or sliding joint provides one degree of linear freedom. The base of the robot arm is generally fixed (for the purpose of this thesis we assume it is mounted to the arm rest of the user's wheelchair) and the other end holds the end-effector or tool that is used to manipulate objects in the surrounding environment. Robotic manipulators are often referred to in terms of the total degrees of freedom they have. The

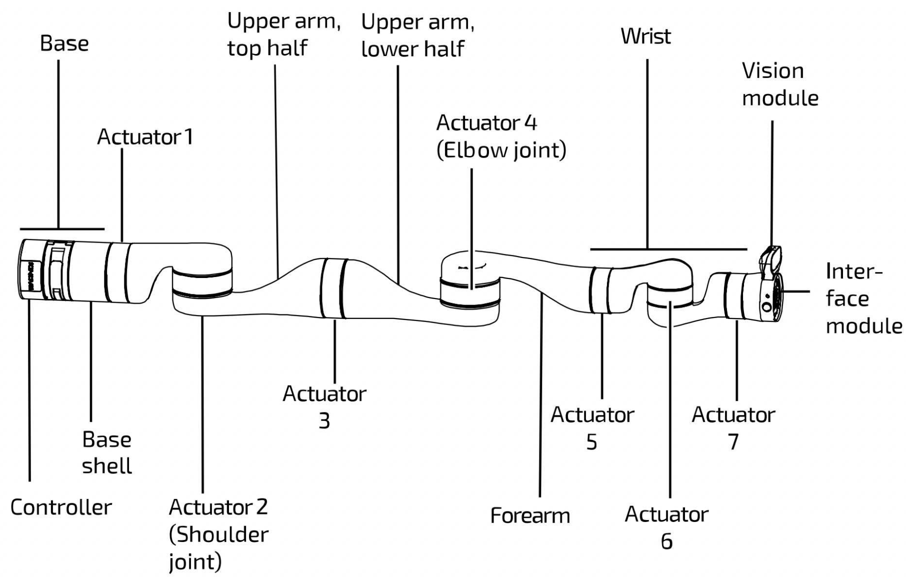

Figure 2.5: The main components of the Kinova Gen3 Ultra lightweight robot. Each actuator corresponds to one DoF. *Note.* Reprinted from “Kinova Gen3 Ultra lightweight robot User Guide”, by Kinova inc., 2022, p. 13. *© 2022 Kinova inc.*

robotic manipulator used in this work is the Kinova Gen3 Ultra lightweight robot; it has seven revolute joints (or actuators) corresponding to a total of seven DoF (see Figure 2.5) [39].

2.2.1 Degrees of Control

While DoF refers to the total number of joints that can be operated independently, it is often the case that only a subset is actually available to the human user, this is referred to as Degrees of Control (DoC). In the context of wheelchair-mounted robotic manipulation, input signals from the users existing device for controlling wheelchair motion (i.e., joystick) are mapped to control the robot arm. The number of input channels from these preexisting controller devices are often lower than the total DoF of the system. In order to increase the number of DoF that the human user can control, a switch or button is often used to allow for cycling through the available DoF. This leads to the problem of how to effectively aggregate the system's DoF into different *control modes* that can be easily manipulated by the human user. Each control mode corresponds to one DoC.

2.2.2 Mode Switching

Once the robot's DoF are grouped into different control modes, there remains a problem of how the human user should switch between the modes. Furthermore, how should the modes be presented to the user, how will they know what mode they are currently in, and how do we trigger mode switching. For example, control signals from a joystick need to be switched from moving the end-effector in free space to opening and closing of the gripper in order to pick up an object. One standard technique to address this problem is to have a pre-designed, optimized order for presenting the different modes and then have the human user switch between them in a cyclic manner. The number of available control options could also be selectively reduced by an expert, so that the user has access to only a subset of the systems DoF. While these standard switching-based methods allow for increased functionality and control, they are reported to be slow and difficult to use [61]. Edwards et al. show that adaptive switching, where mode switching predictions are learned real-time, can be used to significantly decrease the number of mode switches and total switching time for a modified box-and-blocks task [20].

2.3 Decision-making in Robotic Manipulation

In this section we provide an overview of the background material and related work necessary for Chapter 4. We focus on methods for learning robotic control signals where the general goal is to learn a policy as defined as the appropriate control action to take in response to a perceived state that serves to steer a dynamic system towards goal completion. Specifically, we are interested in human-teachable agents where the aim is to allow for more natural methods of imparting knowledge from non-expert humans to the agent. Our intent is not to emphasize that one method is better than the other, but rather that each complementary method demonstrates strengths and weaknesses under different situations. We postulate that the key to creating a robotic system capable of *continuous adaptation and learning in natural environments* will involve a combination of different techniques.

2.3.1 Markov Decision Processes

Sequential decision-making tasks are often modeled as Markov Decision Processes (MDPs) [67]. MDPs are a mathematical framework where the goal is to maximize the expected discounted returns $E[\sum_{i=1}^{T} \gamma^i r_i]$, where r_i represents the reward and $\gamma \in [0, 1]$ is the discount factor. To achieve this goal, an agent takes actions $a \in A$ following a policy $\pi(a|s)$ in a particular state $s \in S$. Typically, the sets of states S and actions A are assumed finite and a policy is a conditional probability distribution. When an action is taken, the transition between states is dictated by the transition dynamics distribution $p(s'|s, a)$ which, in the case of finite states and actions, can be represented as a table. As the agent transitions between states, a reward signal received and is modeled as a random variable $r_i \sim R(s, a)$ which constitute part of the objective in the MDP.

Many robotic tasks can be intuitively reduced to sequential decision-making, and can thus be modelled as Markov Decision Processes (MDPs) [43]. For example, the task of using an WMRM to get a drink of water can be broken down into sequential sub-tasks of retrieving a cup, putting down the cup, grasping the water pitcher handle, pouring the water into the cup, placing the pitcher back down, retrieving a straw, placing the straw in the cup, grabbing the now full cup and bringing it within reach of the users mouth. At the low level, each of these sub-tasks can be broken down into motion primitives [53] and would involve decisions on which specific control signals should be sent in order to produce the necessary movements. As the task is carried out, the user is required to choose new control signals (i.e., actions) to send according to the robots current trajectory and goal (i.e., state).

This naturally lends WMRM tasks to be framed as reinforcement learning (RL) problems, where the robot autonomously learns an optimal behaviour via trial and error interactions with the surrounding environment [68]. This helps to take the onus off of the user to directly control the anthropomorphic robotic manipulator which can be particularly challenging due to the many joints involved. The application of reinforcement learning to physical robotic systems has gained traction in recent years, but is not yet a straightforward undertaking and requires a certain level of skill to hand pick algorithmic parameters [43]. Furthermore, problem representations can be high-dimensional with continuous actions and states that are often partially observable and noisy. Since the true state must be estimated, it is important to take into consideration uncertainty in these estimates. Exploration with physical robotic systems comes at a high-cost, in terms of both computational time and physical resources, since each trial needs to run in real-time and must be closely monitored by an expert to ensure no harm comes to the robot or surrounding environment.

2.3.2 Reward Specification Problem

An important problem in reinforcement learning is how to specify a reward function that promotes quick learning in order to counteract the high cost of real-world exploration; this is known as *reward shaping* [45]. The choice of how to define a reward function to use is challenging in practice as it requires expert domain knowledge and has a direct impact on overall system feasibility. A poorly defined reward function can lead to undesired behaviors which could negatively impact a user's experience (and from Section 2.1 we know that *poor device performance* contributes towards device abandonment). Inverse reinforcement learning (IRL), where a caregiver manually performs the desired action and the reward function is subsequently learned from the demonstration is one solution to this problem. However, IRL can be restrictive since the learned reward function may not be representative of a patient's preferences and can be difficult to modify upon deployment.

2.3.3 Learning from Demonstration

Learning from demonstration (LfD) is an alternative approach to policy learning (i.e., learning a mapping between world state and actions) where a policy is learned from examples provided by a teacher [4], [60]. The LfD problem is split into two phases; first, examples (defined as recorded sequences of state-action pairs) are gathered; and second, a policy is derived from the gathered dataset of examples. This is in contrast with policy learning based on data acquired through exploration of the environment as is typically seen in reinforcement learning. A key limitation of this technique is that a LfD policy is defined only in the states encountered and actions taken during the expert demonstrations; in other words, it has difficulty generalizing to unseen states and actions. Apprenticeship learning is a type of LfD where the algorithm begins with an MDP$\backslash R$ (i.e., an MDP without a specified reward function), learns a reward function R, and then trains an agent on the new MDP from a single expert training session [2].

2.3.4 Learning from Implicit Human Feedback

Another source of reward signal can be from implicit human cues, which differ from previous methods that have explicitly incorporated human feedback. The problem of *learning from implicit human feedback* seeks to understand how an agent can learn a task from a human's involuntary or implicit reactions (e.g., emotions) to the agent's behaviour [16]. This requires the agent to accurately sense and interpret the meaning behind the humans implicit reactions. Cui et al. [16] define the general problem of learning from implicit human feedback and present the EMPATHIC framework to demonstrate the potential of using human facial reactions to improve learning.

2.3.5 Intent Inference and Shared Autonomy

Learning the human's desired goal, target, action, or behaviour is known as *intent inference* [33]. Accurate intent inference is fundamental to the field of assistive robotics where the set of potential task-related goals could be large and inferring the wrong goal, and thus providing the wrong assistance, is often worse than providing no assistance at all. For example, imagine the robot has grasped a glass of water and should bring it towards the users mouth to take a drink but instead places it down on the table. This problem is compounded by user concerns over perceived limitations in the robot's capabilities. In a study on user acceptance of robotic aid in their homes and workplace, Correal et al. [15] found that one of the main barriers to acceptance lies within concern regarding the robot's *lack of capabilities*. These challenges represent some of the difficulties with more widespread acceptance or trust in robotic systems for use in everyday life.

Given these current limitations of trust in the robotic system to do what the user expects, intent inference is usually employed in the form of confidence in predictions based on instantaneous observations. These confidence measures are used to determine how much control the autonomous system should take over and how much control the human should retain [25]. This is known as *shared autonomy*, where various methods are employed to decide how much of the control burden can be offloaded to autonomous systems. This reduces the cognitive load on the end user by providing an easier means of operation. Common approaches for shared autonomy include: control blending paradigms, control partitioning schemes, and having the user handle high-level goal planning while the the low-level control is handled autonomously [25], [33].

In assistive robotics, it is common for the user to specify the task goal and then have the WMRM autonomously carry out the specified goal. The problem then becomes, *how to have the human specify the goal task with enough information for the autonomous system to understand and carry it out*. One promising approach is to use geometric primitives to parameterize a task and then use geometric association constraints (e.g., point-to-point or co-linearity constraints) to build a controller that aims to minimize the task error [34].

2.3.6 Interactive Shaping

An alternative solution for the reward specification problem is *interactive shaping*. With interactive shaping, knowledge is transferred from a human to a learning agent (in this case the robotic manipulator) by having the human trainer provide signals of approval or disapproval (i.e., positive or negative reinforcement) in response to an observed behaviour. This is denoted as the *human reward* [41]. This concept is borrowed from animal learning literature where shaping is defined as training by reinforcing successively improving approximations of the target behaviour [8]. The *Training an Agent Manually via Evaluative Reinforcement* (TAMER) framework proposed by Knox et al. [41] allows a human trainer to interactively shape a learning agents policy via reinforcement signals. In the TAMER framework, supervised learning is used to model a hypothetical reward function and predict a numeric reward based on the systems state and action values. This is similar to apprenticeship learning, which also starts with an MDP$\backslash R$. With TAMER, a model of the human's reinforcement function is learned and dynamically adjusted over time [41] whereas in apprenticeship leaerning the reward function itself is directly learned from a single expert training session.

Interactive shaping (i.e., the problem of learning from human reward) involves understanding how reinforcement learning can be adapted to learn from rewards generated by a human trainer [41]. In other words, can the reward signals obtained from a live human observing states and actions be used to teach an agent new behaviours quickly? The TAMER framework offers a means to evaluate this question by mapping states and actions directly to a human's signals of approval or disapproval and deriving a policy greedy with respect to the predicted human reinforcement function.

In a case study using the MDS robot Nexi, a two DoF mobile robot platform, Knox et al. demonstrate the first successful teaching of a robot from pure human reward feedback using TAMER [42]. In this work the authors teach Nexi five different behaviours: go to, keep conversational distance, look away, toy tantrum, and magnetic control. They found that almost all unsuccessful experimental trials failed due to issues of transparency (i.e., a mismatch between what the trainer's belief is about what was occurring and what state-action pair was actually occurring). This is supported by common-ground research in human-robot interaction where it has been shown that as the level of autonomy increases issues related to a lack of transparency about the robots' decisions and logic tend to dominate [66].

Chapter 3

A Quantitative Analysis of Activities of Daily Living

In this chapter we focus on answering our first research question — *which daily living tasks should wheelchair-mounted robotic manipulators be capable of carrying out?*

Health care and robotic domains use different taxonomies to classify everyday activity tasks and motions [18], [35], [44], [46]. By merging these taxonomies and connecting health care needs with robotic capabilities we seek to bridge the two, often separate, communities. This would provide the robotics community with guidance as to which tasks have the potential to make a large impact (i.e., greatest increase in functional ability) on the target population if implemented. In the field of computer vision, recent interest in video object and activity recognition [74], [75] along with life-logging capture has resulted in numerous public data-sets [27]. In this chapter we aim to mitigate the gap dividing the health care and robotics communities; contributions include:

1. An analysis of long term video-recordings from publicly available life-logging data to extract quantitative measures of task priority;

2. From higher frame-rate video recordings of human kitchen activities, we analyze human arm and hand motion data to quantify the speed and variability of human movement; and

3. We discuss how understanding what tasks are of high importance, both quantitatively and qualitatively, will impact the acceptance and use of

assistive robotic technology in the real-world.

Extracted task frequencies of everyday activities provide insight into what tasks would be of high priority for robotics researchers to focus efforts on while analyses of human motions during task execution provides a gold standard for robotic manipulation to compare against.

3.1 Activity Analysis from Lifelogging Data

Lifelogging data is a valuable source of quantitative ADL and human motion information. Lifelogging involves long-term recording of all activities performed by an individual throughout the course of a day, usually through a video camera, and occasionally using other types of sensors [27]. While lifelogging research has been published on for over two decades [49], recent hardware and method advancements have allowed the field to grow substantially within the past few years [7]. Small, wearable cameras, such as the Microsoft Lifecam [76], with a longer recording duration has made it more practical when compared with the analog video cameras and recorders initially used. New methods for recognizing objects and actions has driven Computer Vision (CV) research interests to explore lifelogging data, which has been found to be a source of more realistic "in-the-wild"-type data than typical CV benchmarks [23], [55].

In this work we evaluated over 30 lifelogging datasets, most of which targeted the performance of a particular algorithm (e.g., object recognition in home environments) and therefore did not encompass the full day. As a result, these datasets did not typically have a statistically sound sampling over all objects and tasks performed in a day in order to meet our analysis inclusion criteria for this work. We found that video recordings taken over several days were done at one to two frames per minute, making the data useful for gross ADL task frequency analyses, but unsuitable for capturing detailed timings of individual arm and hand motions. An additional downfall of the low frame rate video datasets is that they fail to capture daily tasks repeated with high frequency but performed quickly, such as the opening doors or flipping a light switch. A higher frame rate (e.g., 30 frames per second) is required to cap-

Figure 3.1: In the NCTIR Lifelog Dataset [28] three people wore lifelogging cameras for a total of 79 days. These are sample images of the subjects egocentric environment collected at a rate of two frames per minute. *Copyright © 2020, IEEE.*

ture detailed timings of individual arm and hand motions. In this work, three sources of data were selected from the set of lifelogging datasets for further analysis: two from full day recordings to extract ADL task frequency [27], [70], and one from short-term recordings of individual tasks for individual arm and hand motion data [22]. For a detailed table of all datasets considered for inclusion in this study please visit our companion website[1].

3.2 Daily Living Activities Frequency Analysis

To compute quantitative data on ADL task frequency we analyzed both first-person (egocentric) lifelogging videos (referred to as 'NTCIR' [27], [28]), and third-person (exocentric) data from Internet-of-Things (IoT) type sensing built into home objects (referred to as 'MIT') [70]. Example lifelogging images from the NTCIR dataset are shown in Figure 3.1. Since the lifelogging video data was collected at only one to two frames per minute, the use of complementary sensing data turned out to be important for capturing a broader set of tasks. Tasks were inferred by manually labelling high-level actions in each image

[1] `http://webdocs.cs.ualberta.ca/~vis/ADL`

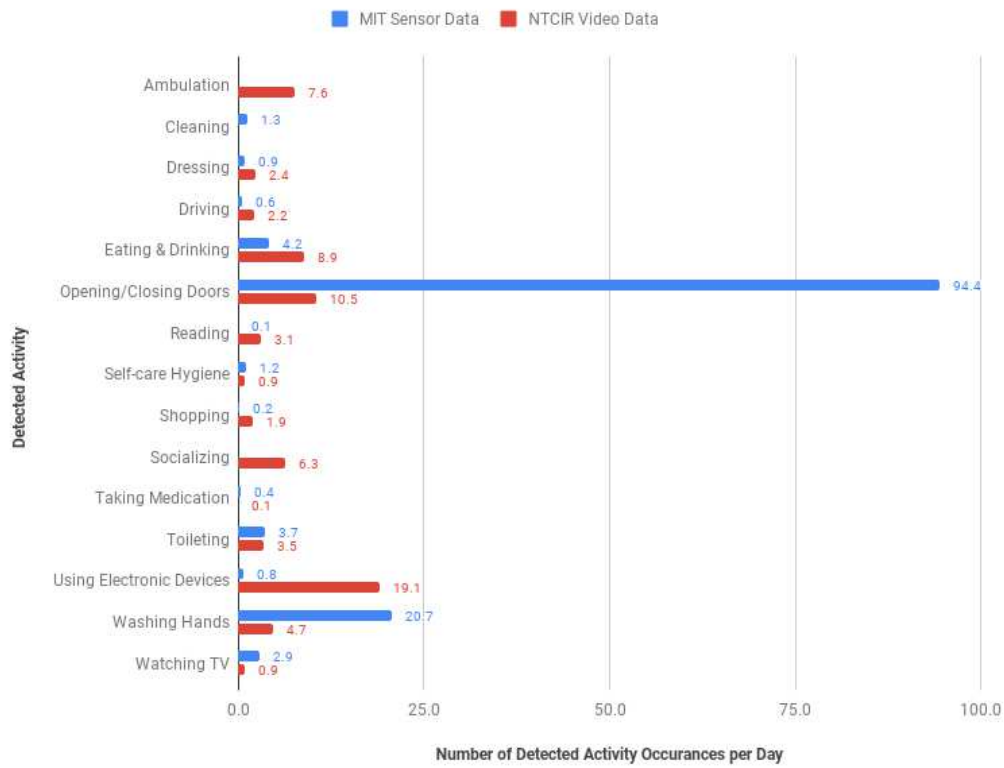

Figure 3.2: Frequencies of detected daily living activities from the MIT internet-of-things type sensor data (blue bars) and NTCIR lifelogging video data (red bars).

for a subset of the data and mapping them to automatically computed visual concepts provided with the NTCIR dataset. Our companion website[1] contains the visual context to actions inference bindings, so readers can replicate results or add other rules and actions to classify. This enabled us to label in-home data sequences spanning multiple days according to which ADLs were carried out at particular times and compute corresponding statistics.

Figure 3.2 illustrates the frequency of the most common daily living tasks found in these datasets, with the NTCIR video data results shown in red and MIT IoT sensor results in blue. Tasks corresponding to potential robot skills were grouped together during the detected activity labelling process. Some events are detected more reliably by the embedded sensors used in MIT, while others show up only in the image data. For example, sensors detect quick events more reliably than the low frame rate lifelogging video data. In contrast, the MIT sensor dataset is limited to detecting daily tasks only where sensors are placed, therefore missing events such as the outdoor activities that are captured in the video data. By combining results from both datasets, we were able to obtain a more accurate quantitative measure of daily task significance.

It should be noted that this work is limited by the small sample size of people represented in the three datasets as well as the collection methods. The NTCIR Lifelog Dataset was collected over the course of a month for each of the three subjects with a camera worn on a lanyard around their neck passively taking images every 30 seconds. The MIT IoT dataset was collected over 14 days for two subjects with 77 state-change sensors installed in the first subject's apartment and 84 in the second subject's apartment. Given the small sample size we do not assume our results can be generalized across all populations, rather we aim to capture an ordering of task importance and magnitude of difference in frequency. Furthermore, the subjects of both studies are all able-bodied and their set of daily activities may not be representative of the target populations needs. A potential avenue for future work to help mitigate these limitations would be to capture full day video recordings of a diverse set of individuals at a higher frame rate, with a focus on data collection within the end-user population.

Our results reflect that opening and closing doors is the most frequent task at 94 times per day; this category includes room doors, cabinet doors and drawers as they require similar robotic manipulation capabilities to carry out. We believe the MIT data was more accurate in this category since the data was obtained from built in door sensors, whereas the low video frequency of the NTCIR dataset missed quick openings, particularly of cabinet doors and drawers for object retrieval. Using electronics is the second most frequent task performed; referring to the use of electronic handheld devices and was dominated by smart phone use. These devices were mostly not covered by the MIT sensors, but were detected in the NTCIR video data. Drinking and eating were found to be essential tasks in both studies, with a frequency of 8.8/day from NTCIR and 4.4/day from MIT. MIT captured hand washing every time the faucet was turned on or off, which resulted in an overestimation of hand washing frequency. For inclusion in the graph we divided the total amount by two to designate one hand washing event resulting in an average of 20.7 per day for the MIT sensor data and 4.7/day for the NTCIR video data.

Task execution with a WMRM depends on the physical capabilities of the robot, as well as the time and cognitive load it takes the user to send control commands through the human-robot interface. Door opening/closing are often studied in robotics research [32], and robot feeding has been studied for over 30 years, with some promising recent results [6], [26]. In contrast, hand washing, which is also high-priority, has been studied in assistive CV [31], to prompt Alzheimer patients though the steps but, to the best of our knowledge, has yet to be studied in assistive robotic research.

These results capture the actions of able-bodied adults and can help guide robotics researchers as to what *functionalities* should be available in WMRM systems. A key finding in our work that was not mentioned by subjective surveys on user preferences (2.2) is that activities carried out on a frequent basis throughout the day should be easy and quick to carry out with a WMRM. We argue that activities that occur with high frequency but are not deemed as important or not mentioned by the target population should be carried out autonomously or semi-autonomously by the robot, as these tasks may not be worth the extra cognitive and physical effort required to manually control the robot during task execution.

3.3 Arm and Hand Motion Analysis.

The successful deployment and acceptance of assistive robotic manipulators also depends on the *usability* of the system. The goal of *usability* is to facilitate the user in accomplishing tasks within an acceptable time period; *usability* is influenced by the user interface and level of automation. In order to increase *usability* of WMRM systems, it is important to understand what an acceptable

Figure 3.3: Sample images from the GTEA Gaze+ dataset showing the top four most frequent kitchen motion tasks. *Copyright © 2020, IEEE.*

time period is for different tasks. Humans represent the gold standard for manipulation that WMRM systems should strive for. By analysing human arm and hand motion frequency and duration during task execution we can begin to understand where potential frustrations during the control of WMRMs stem from. *Poor device performance* is one of the four key factors related to device abandonment, therefore improving the *usability* by removing and/or reducing potential points of frustration should have a positive impact on the acceptance and use of WMRMs.

From high frame rate video datasets we were able to extract the number and timings of individual arm and hand motions required to perform a particular ADL and, for a few tasks, similar timings for robot execution. The Georgia Tech Egocentric Activity Datasets (GTEA Gaze+) [2] contains full frame rate (30 frames per second) video recordings of humans performing domestic tasks [22]. We analyzed the annotated GTEA Gaze+ dataset, containing ∼25GB of

[2] http://www.cbi.gatech.edu/fpv/

annotated kitchen activity videos, to extract individual human motion timings performed during task execution (sample images from this dataset are shown in Figure 3.3).

In Table 3.1 the frequency (occurrences per hour) and mean execution time for different kitchen motions captured in the GTEA Gaze+ dataset are presented. It is notable how quickly a human moves and how many movements we make while carrying out high-level activities (e.g., food preparation). Replicating human speed and agility should be the gold standard to benchmark robotic manipulation skills against. Previous works show that human motions are far faster than typical assistive robot motions. For example, reach motions that take us one second, can take anywhere from ten seconds to several minutes when carried out with a robotic arm [52]. This has implications for how many tasks a robot system can practically substitute in a day without taking up an excessive amount of time (leading to increased user frustration). For other motions, such as pouring liquids, the task itself constrains the human to proceed slowly, therefore human and robot task execution times are more similar. It is important to consider that the GTEA Gaze+ dataset is not a representative sampling of all human activities that occur through the day as it was collected solely within a kitchen setting. However, it is still notable that the number of reaches captured in the dataset is three times the number of door openings over the same 11 hours of video data. This implies that finding solutions for reliably and easily carrying out reaching motions with WMRMs would be beneficial for the target population in terms of reducing the overall task execution time as well as the associated cognitive load.

3.4 Discussion

Understanding which daily living tasks are of high priority to the target population as well as which tasks occur at a high frequency throughout the day can provide insight and guidance in the field of assistive robotic manipulation. Table 3.2 lists these high priority tasks; the qualitative column consolidates results from prior end-user surveys [6], [65], while the quantitative column

Kitchen Motion Task	**Frequency**	**Duration (sec)**
Reach and pick item	88	1.5
Reach and place item	84	1.2
Turn switch on or off	10	2.1
Transfer food	6	8.6
Wash hands or items	3	6.7
Flip food in pan	2	4.9

Table 3.1: Frequency (occurrences per hour) and mean execution time (seconds) for various kitchen tasks captured in the GTEA Gaze+ dataset. *Copyright © 2020, IEEE.*

highlights the key findings of this work. The quantitative results calls attention to common activities not mentioned in the surveys, namely the frequent openings of cabinet doors and drawers, and the many switches and dials commonly found in homes. We believe that if the robotics community were to focus on solving these activities in the real-world it would have a large impact on the target populations quality of life. Tasks that appear only in the quantitative column and are feasible for a WMRM, such as opening/closing doors and flipping switches/buttons, should be offered as a preset automation as they occur frequently thoughout the day but are not deemed important by the end user.

Door opening and closing, eating and drinking, hand washing, and toileting would arguably be the most essential activities for assistive robot arm and hand systems to support out of all the ADL tasks analyzed in this work (3.2). The first three are relatively feasible to accomplish given the payload capacity of current robotic arms. Activities, including using electronics (primarily smartphones), socializing, and reading could be physically aided by WMRMs, but since these activities are not inherently physical, alternative solutions are possible and can be a simpler, more reliable solution (e.g., hands-free phone use or other computational automation). Toileting is a high priority task that involves person transfers. WMRMs do not generally support this, but there are specialized transfer devices that are used in health care, and can be easily installed in an individual's home [62]. Toilet lid opening/closing and flushing could be added as autonomous behaviours in a WMRM system as well as

Qualitative	**Quantitative**
picking up items	reach to pick/place
eating/drinking	eating/drinking
personal hygiene	hand washing
cleaning	toileting
leisure/recreation	using electronics
carrying objects	opening/closing doors
preparing food/drinks	switches/buttons

Table 3.2: High priority daily living activities that would have a large impact on the target community. The qualitative column reflects task priority preferences stated by the target population in surveys [6], [65]. The quantitative column highlights key results from our life-logging data analysis reflecting tasks that occur frequently throughout the day. *Copyright © 2020, IEEE.*

turning on/off sink faucets to assist with hand washing.

Overall, there is great potential for supporting ADLs for those living with reduced functional abilities. Over the past few decades there has been an increasing demand for health care services due to the rising elderly and disability populations [17]. Assistive robots can help bridge this gap by alleviating the labor burden for health care specialists and caregivers. Furthermore, an assistive robot could help one carry out daily living tasks they are otherwise incapable of managing on their own, thus increasing functional independence.

However, challenges remain before these robots will reach mainstream adoption, including but not limited to: *system costs*, *task completion times*, and *ease of use*. Currently costing around 30,000 USD, a robotic arm is a significant expense for an individual, who may already have a limited income. While western health insurance often covers expensive prostheses, only in the Netherlands does insurance cover the cost of a WMRM.

Speed of robot motion, which affects task completion time, is another challenge. While a human reach takes 1-2 seconds, assistive robots can take between 40-90 seconds, resulting in robot solutions that are magnitudes slower [36], [52], [57]. This results in decreased user satisfaction thus increasing the chance of device abandonment. In the GTEA Gaze+ kitchen tasks, humans performed 160 reaches per hour. Performing the same task with a WMRM would turn a 30 minute meal into a two hour ordeal. Anecdotal comments from users of assistive robot arms are that their morning kitchen and bathroom activities take them several hours to carry out.

Robots tend to solve tasks differently than humans as robots are often limited to grasping one item at a time, while humans can handle many. For example, when setting a table, humans pick several utensils at the same time from the drawer, whereas a robot would move each utensil individually. Analysing the publicly available TUM Kitchen Data Set of activity sequences recorded in a kitchen environment [71], we found that their robot strategy on average required 1.6 times more movements than a human. Users of assistive robots often adopt compromises to deal with the speed and accuracy of robots. For example, users are more likely to choose foods and drinks that can be held statically in front of the user by the robot (e.g., eating a snack bar or drinking with a straw) as they are far quicker and easier to consume than those requiring numerous robot reach motions, for instance, eating a bowl of cereal.

It has been shown that users prefer to have continuous in-the-loop control, especially when the robot will be interacting directly with the individual (e.g., while eating) [36], [57]. In recent work a low dimensional control space is learned from demonstrations. This allows a human user to have direct control over the robots motion in 3D space using a low DoF human-robot interaction interface, such as a joystick [47], [56]. Finding the proper balance between human interaction and semi-autonomous assistive systems is a challenge. Currently, most research is evaluated with a few participants testing out the system for about an hour each in a research lab setting. New human-robot interaction solutions will need to be deployed longer term in real end-users homes in order to properly evaluate *functionality*, *usability*, *reliability*, and *safety*.

3.5 Conclusions and Contributions of this Chapter

In this chapter we presented insights and meaningful guidelines that support the needs of the target population in order to help focus research and future developments in the field of assistive robotic manipulation. Understanding what tasks the target population expects to be able to carry out with an assistive robot arm (i.e., *functionality*) along with what the allowable time limits are for carrying out each task (i.e., *usability*) is critical for the acceptance and use of these systems in the real-world. We analyzed human task frequency from public life-logging datasets and computed motion timings from public CV data. A key finding of our quantitative analyses is that activities that occur at a high frequency but are not mentioned in user preference surveys would be ideal to be able to perform autonomously. Overall, reaching and door openings were the most frequently carried out motions. Drinking, eating, and hand washing are other high priority tasks that can be addressed by currently available assistive robot arms. Toileting and dressing, while ranking just below, are generally thought to be more challenging for robotics, since they require the transfer of body weight. However, wheelchair-mounted robotic manipulators can still be of assistance during the subtasks that make up toileting and dressing, including but not limited to: opening/closing toilet lids, flushing, turning faucets on/off, flipping light switches, and picking out clothes. Detailed data on frequency and duration information for all analyzed tasks and motions, as well as the analysis methods are available on our companion website: `http://webdocs.cs.ualberta.ca/~vis/ADL/`.

Chapter 4

Learning How to Reach from Human Reinforcement Feedback

In this chapter we focus on answering our second research question — *how can individuals easily teach their wheelchair-mounted robotic manipulator new tasks?*

Human robot collaboration is a cornerstone of assistive robotics as these devices are primarily used in an individual's natural environment, which is known to be highly unstructured and dynamic (e.g., the user may wish to pick up an object accidentally dropped on the floor). *Efficacy*, *safety*, and *reliability* of robotic systems in natural environments will need to be established in order for mainstream acceptance and use. As these systems become more common outside of research settings, it will be crucial for them to be able to learn how to carry out new tasks and manipulations as situations arise. A key factor found to be significantly related to device abandonment is a *change in user needs or priorities*, as discussed in Section 2.1 [54].

Ideally, the end user would be able to teach the learning system new behaviours and tasks without outside assistance from robotic experts or caregivers. Since the goal of assistive robotics is to improve quality of life and promote independent living, it is important to provide users with a way of teaching and customizing robotic behaviour to their liking that is simple, safe, and straightforward *as their needs change*. In this chapter we present a proof-of-concept system intended for teaching new autonomous behaviours to a robotic arm in the real-world.

4.1 Methodology

In this proof of concept work, we are interested in investigating whether the TAMER framework can be extended to high-dimensional robot manipulators. This is an alternative solution to the reward shaping problem prevalent in robotic reinforcement learning. The goal is to provide end users with a simple, yet reliable, method of teaching new behaviours to learning agents in the real world while carrying out tasks online. This is presented as a *learning mode* integrated into the existing control interface that can be activated as required. As discussed in Section 3.3, reaching to pick or place items is one of the most frequent motions made during task execution. In this proof-of-concept study, we demonstrate the feasibility of this system by teaching a seven DoF robotic manipulator how to reach for an object desired by the end user. In this section we frame the problem within the MDP framework and detail our adaptation of the TAMER algorithm.

4.1.1 Problem Formulation

The reaching behaviour can be formulated as an episodic task that ends once the system is confident enough in which object the trainer desires, at which point the system moves autonomously into the final grasping pose. All positions are defined relative to the robot's base frame coordinate system. Our state representation is a stacked vector of the robot's joint configuration $q \in \mathbb{R}^7$ and the difference in distance from the desirable objects $x^* \in \mathbb{R}^3$ and the robot's end-effector $x \in \mathbb{R}^3$ at each time step. The set G of potential goal objects $g_i \in G$ is provided as input to the system. The end-effector location is extracted from the robots pose $p \in \mathbb{R}^6$ which is continuously calculated from the robot's forward kinematics. At each time step, the agent can choose from six possible actions that correspond with the end effectors linear velocity along one of the three Cartesian axes: forward (+x), backward (-x), left (+y), right (-y), up (+z), or down (-z).

We present our implementation of the TAMER framework in Algorithms 1, 2, and 3 where the agent seeks to learn the human trainer's reinforcement

Algorithm 1

Input: α, ϵ

procedure TamerAgent3D
    $t \leftarrow 0$
    $s_t \leftarrow$ *GetState()*
    InitializeHumanModel(s_t)
    InitializeReinforcementListener()
    $a_t \leftarrow SelectRandomAction()$
    TakeAction(a)
    while $PredictIntent() < \alpha$ **do**
        $x \leftarrow random()$
        if $x < \epsilon$ **then**
            $a_t \leftarrow SelectRandomAction()$
        else
            $s_t \leftarrow GetState()$
            $f_t \leftarrow GetFeatures(s_t)$
            $p_t \leftarrow ModelPredict(f_t)$
            $a_t \leftarrow argmax(p_t)$
        $TakeAction(a_t)$
        $t \leftarrow t + 1$

Algorithm 2

procedure InitializeReinforcementListener
    while *true* **do**
        $r_t \leftarrow GetHumanSignal()$
        if $r_t \neq 0$ **then**
            $s_t \leftarrow GetState()$
            $a_t \leftarrow GetAction()$
            $UpdateModel(f_t, a_t, r_t)$

function $H : S \times A \to \mathbb{R}$. At each time step, which lasts 1.5 seconds to help avoid the credit assignment problem (discussed further in Section 4.4), the desired object confidence is calculated as

$$c_{g_i} = max(0, 1 - \frac{d + \alpha r_t \Delta d}{D}), \tag{4.1}$$

where d is the Euclidean distance to the goal, Δd is the change in distance since the last time step, α is a weighing term determined experimentally (set to 10 in our experiments), r_t is the human reinforcement signal received at the current time step, and D is the maximum reach of the robotic arm (1.0m). If an object is determined to be the one desired by the user (intent threshold was set to

Algorithm 3

Input: $r_t, s_t, \vec{\theta_{t-1}}, \alpha$ **Output:** $\vec{\theta_t}$

procedure UPDATEMODEL
  prediction $\leftarrow \hat{H}_{\theta_{t_1}}(s, a)$
  error $\leftarrow r_t -$ *prediction*
  $\vec{\theta_t} \leftarrow \vec{\theta_{t-1}} + \alpha * error * \nabla_{\theta_{t-1}} \hat{H}_{\theta_{t-1}}(s, a)$
  return $\vec{\theta_t}$

0.85 for our experiments) then the system moves autonomously into the final grasping pose. Otherwise, the agent selects and carries out a new action using a dynamic ϵ-greedy policy that more heavily weighs random actions as the number of negative reinforcement signals increases. During action selection, the learned model $\hat{H}$ is consulted to choose the action a that maximizes the predicted human reinforcement signal. By learning a humans reinforcement function instead of the reward directly, TAMER is capable of generalizing to unseen states. This generalization is especially important in a robotic setting where exploring the entire state space of the physical system would take too much time or is simply not possible due to kinematic constraints. TAMER also avoids the credit assignment problem inherent in RL by assuming the human reinforcement signal is fully informative about the quality of recent actions since the human would have intuitively taken the long-term impact of the behaviour into consideration [42]. This assumption is directly exploited by our confidence prediction algorithm as we assume the human trainer will only give positive reinforcement signals when the robot is moving towards the desired object.

4.1.2 Human Reward Function

For our model of the human reward function $\hat{H}$ we train a single hidden layer, fully connected, neural network. Mathematically, we represent this as

$$\hat{H}_\theta(s, a) = \sigma(\sigma(W^h[s] + b^h)W^o + b^o)[a], \tag{4.2}$$

where the learning parameters are $\theta = \{W^h, b^h, W^o, b^o\}$, and $\sigma(\cdot)$ is the activation function. In our experiments, we used the hyperbolic tangent function

Figure 4.1: The Kinova Gen3 Ultra robotic manipulator

and the hidden layer had 100 neurons. We do not include actions as input, but instead represent the reward signal for each action as a separate output. In Equation 4.2, our notation is for accessing the output dimension for the corresponding action index a which is a finite, discrete set. We found this helped speed up learning, and was motivated by work from deep reinforcement learning research [51]. We train our model with the Adam optimizer [37] and minimize the mean-squared error as our objective function $\frac{1}{2}||H(s,a) - \hat{H}_\theta(s,a)||_2^2$.

4.2 Experiments

Our experimental work aims to evaluate the performance of TAMER on the Gen3 robot manipulator, as well as the feasibility of using the same human reinforcement signals for inferring which object the participant would like to grasp. To establish a comparative baseline, we conduct a proof-of-concept

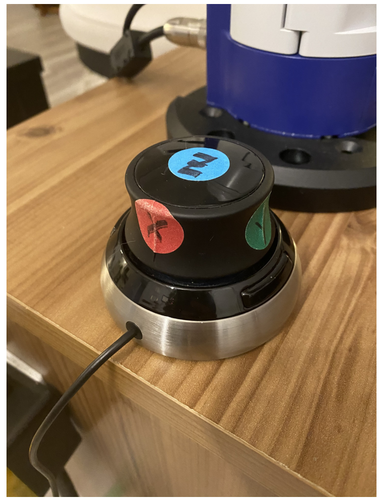

Figure 4.2: The SpaceMouse Compact used as the input device for the baseline experiments.

user study where six participants perform a simple reaching task. Input signals from three different control interfaces commonly used in the assistive domain are mapped to the robots motion (i.e., mimicking traditional teleoperation control). In this study we are interested in how the task execution time and cognitive load on the user are affected as the dimensionality of the input signal increases. We then compare the TAMER method for control with these baselines to gain insight into how a learning agent compares with the traditional teleoperation control schemes in the same task setting. In this section we describe the hardware used and experimental studies carried out.

4.2.1 Hardware

The Gen3 Ultra robotic platform, developed by Kinova, is a lightweight (8.2kg) robotic manipulator designed for assistive robotic research [38]. It has a maximum reach of 902 cm and is capable of handling payloads up to 4.0kg. The arm is outfitted with a two finger Robotiq adaptive robot gripper (model 2F-85) [58], for a total of seven DoF (one for each joint) plus the additional open/close position of the gripper (Figure 4.1). The open source API provided by Kinova was used for software development and all experiments were carried out with the Robot Operating System (ROS) on a system running Ubuntu 16.04. The

Figure 4.3: Experimental setup with the desired end-pose shown. Participants were instructed to move the robotic arm to the Tylenol bottle.

SpaceMouse Compact designed for 3D navigation was used as the control input device for all baseline experiments (Figure 4.2) [1]. The control interfaces emulated in this work were a one-axis Sip-N-Puff device, two-axis joystick, and three-axis joystick. Teleoperation with control interfaces that are lower dimensional than the control space of the robotic arm requires the user to switch between one of several control modes, known as *mode switching*. This maps the input dimensions to a subset of the arms controllable degrees of freedom. In our experiments, input signals from the SpaceMouse axes are mapped to robot motion commands in Cartesian space. The left button is used for mode switching and the right button to end the experiment session. As the DoF of the input device decreases, the number of modes required to represent the robotic control space increases. Direct teleoperation with the three-axis joystick requires switching between three different modes, teleoperation with the two-axis joystick requires four modes, and the one-axis Sip-N-Puff interface requires seven different modes. As the number of available modes increases, so does the cognitive load placed on the user as reflected as an increase in task completion time.

4.2.2 Experimental Design and Protocol

In this proof of concept study six participants (three male and three female, ages 11 to 61 years) participated in the baseline experiments. Due to COVID restrictions at the time, the participants were unable to return to participate in the TAMER experiments therefore the TAMER results presented were gathered by the authors. The experimental setup is shown in Figure 4.3. At the start of each session, the participant was instructed to reach for the Tylenol bottle from a top down orientation and had a 10-minute training period for each control interface to become familiar with the mapping between the input signals and robotic motion. After the training period, the reaching task was carried out three times for each baseline control interface for a total of nine trials per participant. During each trial, the subject provided input signals via the SpaceMouse and was able to toggle through the available modes using the left button; the current mode is printed to the terminal for transparency (Figure 4.4). The current mode, motion command, Cartesian pose, joint angles, and gripper position data are continuously logged during each trial as well as the total number of mode switches and task duration time at the end of each session.

For the TAMER experiments, the same environmental setup shown in Figure 4.3 was followed, except human feedback was captured from the keyboard instead of the SpaceMouse. Two studies were carried out; in the first study the authors were instructed to press the up arrow key when the robot moved towards the Tylenol bottle (positive reinforcement signal) and press the down arrow key (negative reinforcement signal) when the robot moved away from the Tylenol bottle. For this study, we were interested in seeing how the agent would perform in comparison with baseline methods without pre-training the agent for a single object goal. The aim of the second study was to test whether the object desired by the subject could be accurately inferred from the same reinforcement signals used for action selection. The only difference between this study and the first is that the authors were instructed to pick any object in the scene (instead of the Tylenol bottle) and provide reinforcement signals

```
******************** FILE INFO *************************************

Filename: 1D1reach12121224
Task: reach
Participant: 1

******************* DIRECTIONS *************************************

Please try to pick up the Tylenol bottle from the top

Press the LEFT button to switch modes
Press the RIGHT button to stop the experiment
Push FORWARD/BACKWARD on the SpaceMouse to move the robot

Modes Available for 1D Joystick:
        Forward/Backward
        Left/Right
        Up/Down
        Vertical Rotation
        Horizontal Rotation
        Wrist Rotation
        Open/Close Gripper

******************* EXPERIMENT *************************************

Starting timer, please try to pick up the Tylenol bottle from the top

Current Mode: FORWARD/BACKWARD
Current Mode: LEFT/RIGHT
Current Mode: UP/DOWN
Current Mode: VERTICAL ROTATION
Current Mode: LATERAL ROTATION
Current Mode: WRIST ROTATION
Current Mode: OPEN/CLOSE GRIPPER
Current Mode: FORWARD/BACKWARD

TOTAL EXECUTION TIME: 0 MINUTES 28 SECONDS
TOTAL MODE SWITCHES: 7
```

Figure 4.4: An example terminal output shown to participants during experimental trials that displays the current mode in operation. At the end of the trial, the total number of mode switches and task completion time are printed on the screen.

with respect to the robot's motion towards that object. Experimental results showed no change in system performance as the number of potential goal objects increased between the two studies, this implies that the reinforcement signals provided by the trainer could also be used to predict the users intent and is a promising avenue for future work.

4.3 Results

In this section we present the performance metrics used for evaluation, discuss our experimental results, as well as various challenges faced while adapting TAMER for 3D robotic arm motion.

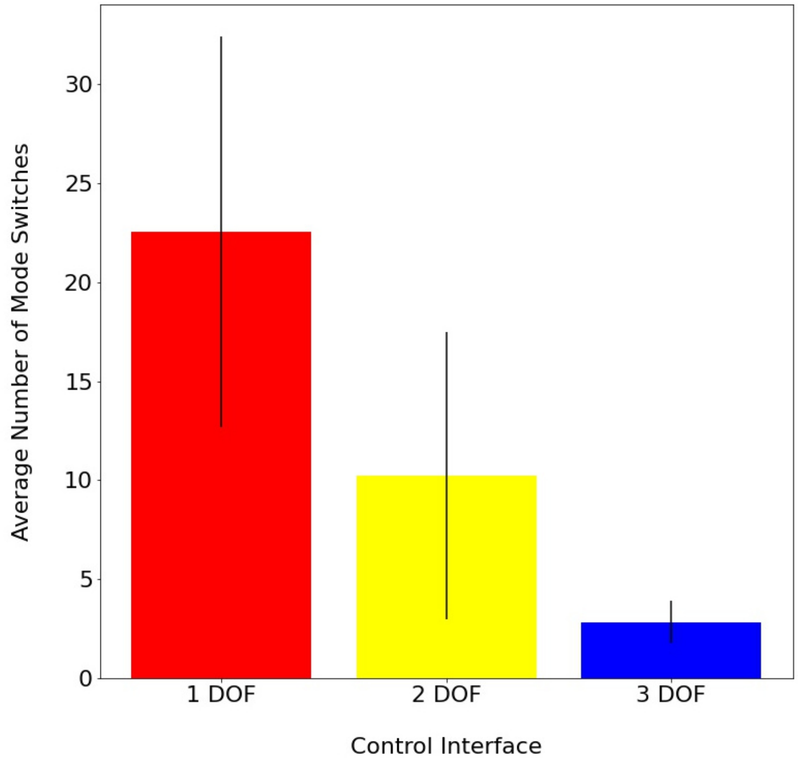

Figure 4.5: The average number of mode switches participants made during trials of the reaching task. The black bar is the standard deviation of the number of mode switches.

4.3.1 Performance Evaluation

In this work the performance of each control interface is evaluated on the total number of control mode switches, number of human input signals received, and the time taken to complete the reaching task (i.e., task completion time).

Mode Switching

Teleoperation involves mapping input signals from the end users existing control interface to robotic arm motion in 3D space. Since the dimensionality of the robot manipulator is generally higher than the input signal DoF, the end user must switch between various control modes that are then mapped to a subset of the robots control space. This can be particularly challenging in the assistive domain where individuals with motor impairments generally use input devices with only 1 or 2 DoF (e.g., the 1D Sip-n-Puff). In this work we are interested in investigating the relationship between the dimensionality of the input interface and the number of mode switches required to carry out tasks. We hypothesize that as the number of input dimensions decreases (and thus the number of available modes increases) the task will become more difficult for the user to carry out and this will be reflected as an increase in the number of mode switches required, input signals received, and task completion time.

Figure 4.5 shows the number of mode switches required averaged across all participants with standard deviations (SD) for each baseline control interface. The mean number of mode switches was 22.6 (SD 9.9), 10.2 (SD 7.3), and 2.8 (SD 1.1) for the one, two, and three DoF control interfaces, respectively. These results support our hypothesis that as the input dimensionality increases, the number of times the user needs to switch modes while controlling the robotic arm decreases. One of the benefits of the TAMER control interface is that it removes the need to switch between modes (i.e., the number of mode switches is zero across all trials) which is why there is no bar for TAMER in Figure 4.5. We hypothesize that this reduces the cognitive load placed the user during operation since they do not need to mentally keep track of the mapping between input signals and robotic motion along the various axes in 3D space. Future studies should measure the cognitive load placed on the participant to more concretely test this hypothesis.

Input Signals

In the assistive domain, it can be particularly challenging for individuals to send even one DoF signal due to motor impairments. Furthermore, since WMRMs are usually mapped to the user's pre-existing control interface, they are unable to use the interface for its original use (e.g., wheelchair motion) at the same time as operating the WMRM. Therefore, to encourage the acceptance and use of these systems in the real-world, it is of interest to develop control methods that reduce the number of input signals required from the user. For the baseline joystick control interfaces this is counted as the total input sig-

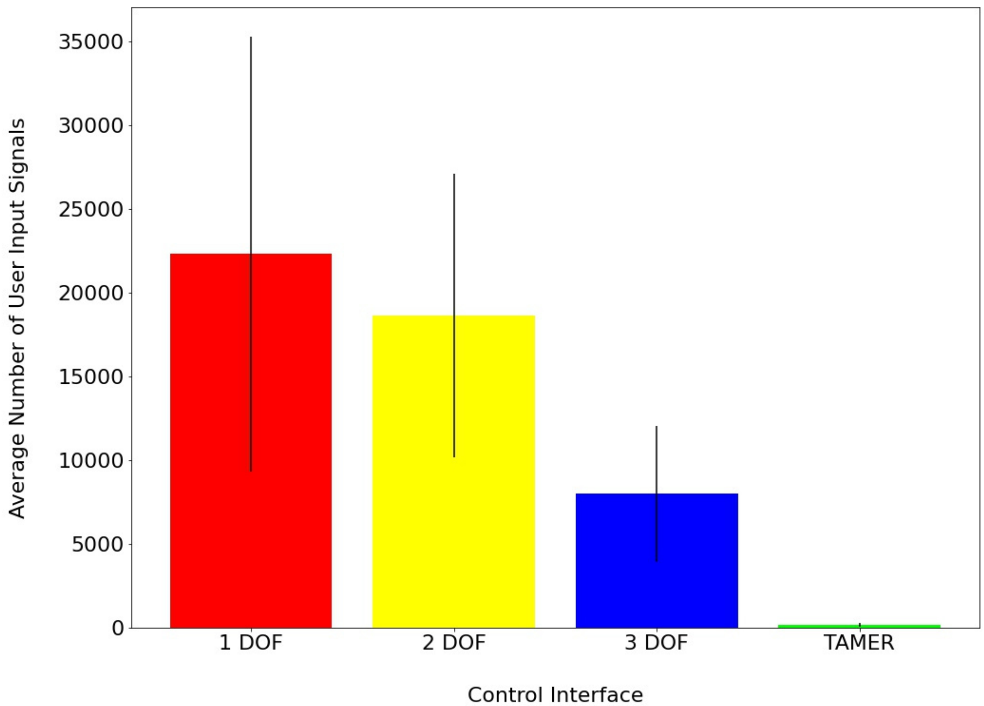

Figure 4.6: The number of user inputs required on average to complete the task. The black bar represents the standard deviation. Results suggest that training a TAMER agent requires far fewer user inputs in order to complete the task.

nals received from the SpaceMouse (motion commands and mode switches), and for our TAMER implementation it is the total number of key presses (positive and negative reinforcement signals) received during an episode. A decrease in the number of input signals received while carrying out a task reflects a reduced physical load on the user.

Figure 4.6 shows the average number of input signals received for each control scheme evaluated. The mean number of input signals received was 22320.4 (SD 12980.0), 18633.4 (SD 8444.0), and 7993.3 (SD 4063.4) for the one, two, and three DoF control interfaces, respectively. The TAMER trials show a substantial reduction in the number of input signals required to carry out the task with a mean of 96.7 (SD 32.1) signals. This result reflects the reduced physical and mental load placed on a user when compared to more traditional forms of teleoperation. Note that this is without pre-training, therefore we

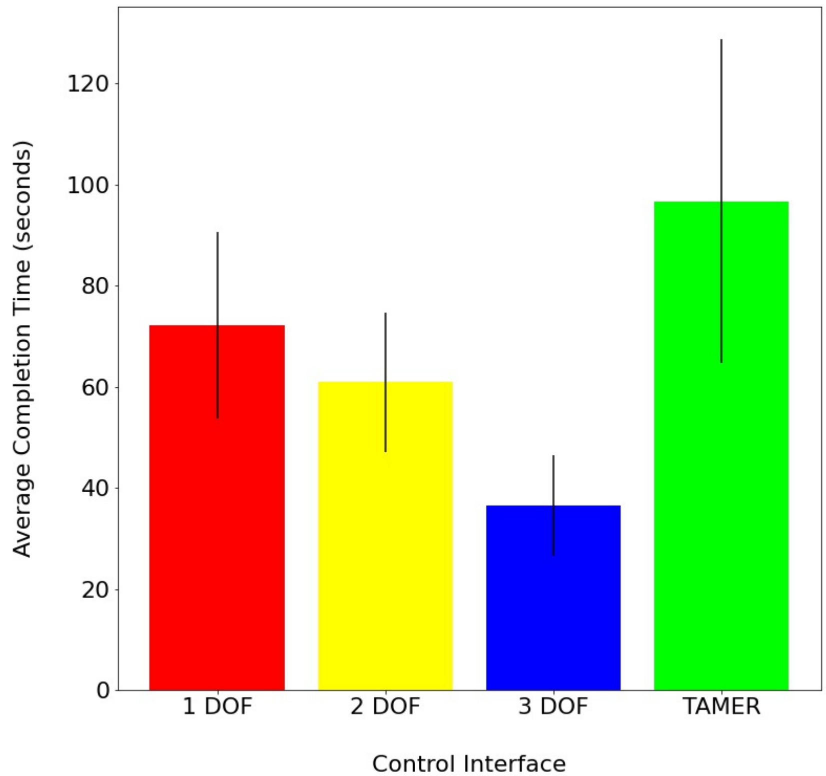

Figure 4.7: The average completion time for the reaching experiment. The black bar represents the standard deviation.

expect this to improve as the agent learns over time and will be explored in future work. For the baseline controllers, the general trend is that the number of commands sent by the subject decreases as the number of input dimensions increases, supporting our hypothesis.

Task Completion Time

A key hindrance to the acceptance of robotic arms in the real-world is the increased complexity and time taken to carry out simple tasks in comparison with having a human performing the same task. Our hypothesis is that as the number of input DoF increases, the user will be able to complete a task faster as they are not required to switch between as many modes. It is important to mention that, in the assistive domain, increasing the number of input dimensions is often not possible due to physical limitations of the end user. Shared autonomy and knowing when to switch to an autonomous system could help decrease the time it takes to carry out tasks by offloading a portion of complexity to the robot. In our experiments, we are interested in investigating how intent inference (i.e., predicting which object in the environment the user wishes to manipulate) can be used to decide when to switch to an autonomous agent. Our results show that the reinforcement signals received by the system can be successfully used to determine when an autonomous agent should take over to move the robotic arm to the final grasping pose. Furthermore, that the number of potential goal objects in the scene does not have a large affect on system performance emphasizing the potential of this type of learning agent in a real-world setting.

Figure 4.7 shows the task completion time averaged across all trials for each method of control. The mean time taken to reach for the Tylenol bottle was 72.1 (SD 18.5), 60.9 (SD 13.8), 36.6 (SD 10.0), and 96.7 (SD 32.1) seconds for the one DoF, two DoF, three DoF, and TAMER control interfaces, respectively. Results show that task completion time decreases as the input dimensionality increases, reflecting a difference in control difficulty. The TAMER agent takes longer to complete the reaching task, although some trials performed comparably with the one DoF control interface. This is likely due to the exploratory nature of the TAMER agent and we hypothesize that task completion time will improve over time as it learns more of the state space. This is left to future work.

4.4 Challenges

In this work we were interested in creating an agent that learns online to carry out new tasks without any prior training. We do this to show that our adaptation of TAMER can be directly used in place of traditional teleoperation schemes. Future work will expand on this finding and aim to show that it can learn to adapt and perform better over time. In our work, we found the main source of failure was due to issues of transparency where the human's reinforcement signals were assigned to the wrong actions. As discussed in our

Figure 4.8: A singular configuration of the assistive robotic arm reached during a TAMER trial.

background section, this problem was also observed in the work of Knox et al. [42]. Assigning signals to the wrong actions was even more detrimental in our implementation because the action space is larger and it took longer for the model to recover. To reduce the number of wrong assignments, we experimented with varying the duration of actions and found that a minimum duration of 1.5 seconds per time step to be sufficient.

All prior works on TAMER have run under the implicit assumption that all actions were feasible in every state. A challenge we experienced in expanding TAMER to 3D robotic arm motion is that not every action is achievable in every state. This happens when the robotic arm reaches a singularity and loses one or more degrees of freedom (Figure 4.8). A robot singularity is a joint configuration where the end effector can no longer be displaced in certain directions [11]. When this happened in the baseline experiments the end user automatically compensated by moving in the opposite direction and moving the joints to a different configuration. On the other hand, the TAMER agent was unable to realize when an action was unachievable and the trainer needed to send negative reward signals to elicit a different action, which in turn corrupted the learned model. To accommodate for this we needed to make a number of adjustments to the algorithm. The joint space configurations were added as input to the supervised learning model to learn the association between the Cartesian action space and joint space. Furthermore, if the agent received multiple negative signals within one time step it would carry out a random action that is different than the last action taken. These additions were enough to help the agent learn how to move out of singular configurations without hard coding specific behaviours into the system.

Another problem we encountered is that the agent would occasionally get stuck in local minima where it continuously alternated between two actions when consulting the learned human reinforcement model regardless of the input provided by the trainer. This generally occurred when the robot reached the end of it's workspace and was close to a singular configuration (Figure 4.8). To counteract this and encourage exploration, we implemented an epsilon-greedy action selection policy instead of always choosing the greedy action.

4.5 Conclusions and Contributions of this Chapter

In this work, we presented an adaptation of the TAMER framework that allows for online teaching of new behaviours and control of assistive robotic arms. We implemented three baseline control interfaces commonly used in the assistive domain for teleoperation of robotic arms and undertook a comparative analysis of system performance. Our experimental results show that positive and negative reinforcement signals can be successfully used to carry out robotic reaching tasks and that these signals can also be used to predict the intended object a participant is reaching towards.

The main limitation of this study is that we work under the assumption that the goal locations of all desired objects are known. Future extensions of this work will incorporate methods, such as object detection and segmentation [10], to dynamically estimate the location and pose of objects in the robots coordinate system. If the goal locations can be estimated online then they can be used to interactively paint the state space with reinforcement signals as was done with the robot Nexi [40]. Future work will also include evaluating system performance as it learns over time, specifically whether the task completion time for TAMER can be reduced to achieve human-like task performance (i.e., task not performed with a robot).

Chapter 5

Conclusions and Future Work

The overall goal of this thesis has been to advance the acceptance and use of wheelchair-mounted robotic manipulators by emphasizing the importance of increasing *functionality* and *usability*. While our long-term goal is to build an adaptable control system where the user can dynamically choose the level of autonomy they desire, this thesis is primarily focused on gaining insight into the needs of the target population and highlighting the potential that learning agents have for providing a simple method of teaching new autonomous behaviours.

Current methods of controlling wheelchair-mounted robotic arms are tedious and frustrating, generally requiring non-intuitive mappings between the input device and robotic motion. Being able to offload a portion of the control burden to an autonomous agent would lead to decreased task execution time and improved user satisfaction. This begs the question – what part of the control process does the end-user retain control of and what is off-loaded to the autonomous agent? Many works have attempted to answer this question in order to design new systems of control that are easier to use, but none have yet to make the jump from research study to commercial implementation. We are interested in understanding why this is the case. What exactly is hindering the acceptance and use of wheelchair-mounted robotic arms in the real-world?

A common method of off-loading some of the control is to have the user designate the high-level goal while the system autonomously handles the low-level planning and execution. The problem with this is what if they change their mind on what they want after the robot is already moving? Furthermore, prior works have shown that individuals prefer to retain more control over the system. In this thesis we postulate that for real-world acceptance of these systems we need to provide a system with a sliding scale of autonomy where the user can choose the level of control they want at any given time, as well as a system capable of learning new autonomous behaviours on the fly by the end-user. This real-world adaptability will be crucial, as it is not reasonable to assume that an expert will always be available to teach the robot new behaviours. This thesis takes one step towards this goal of a truly adaptable real-world system for robotic arm control.

Chapter 3 investigates the frequency of activities that make up a person's daily life in order to gain insight into which activities occur at a high frequency. We propose that activities that occur frequently throughout the day but are not of high importance (i.e., are not mentioned in previous studies on user preferences) should be carried out autonomously at the users request. Furthermore, we study human motion timings to begin understanding usability metrics that should be standardized in robotic manipulation. One previous study of a voice-recorded lifelog found that pick and place tasks were the most frequently performed daily living tasks (corresponding with "Lifting" and "Putting down objects" in Figure 2.4) followed by "Preparing complex meals". Furthermore, 90% of the objects being manipulated during pick and place tasks were less than 300 grams. The Healthcare Robotics Lab at Georgia Tech also look at task importance in an object-centric manner. They developed a prioritized list of objects for robotic retrieval with the top three objects on their ranked list being: TV Remote, Medicine Pill, and Cordless Phone (see Table 2.1). Another common approach is to ask patients and caregivers which tasks they would prefer to be able to carry out with a robotic manipulator. Survey results reflecting the views of over 200 potential users show that preferences vary and user opinions on what functionalities are most useful shift over time as they use the robot in real-life. Before acquiring a WMRM, the top three activities users anticipate using the device for are: preparing food/drinks, pick and place tasks, leisure/hobby activities. After acquiring the assistive devices, users preferences shift to more work related pick and place tasks, personal hygiene tasks, and feeding/preparing meals. Two activities that were missing from pre-development survey results, but observed by participants after they had the robotic arm for everyday use were communication/phone use and the opening of different doors and drawers. These notable shifts in preferences led us to wonder whether there were other daily living activities that were missing from subjective survey results that would make a significant improvement in quality of life if the user had the option of carrying them out autonomously with a WMRM. Key results from our quantitative daily task analysis show that opening/closing doors and eating/drinking are two of the most frequent activities that make up a persons day (part of the top three in both the video and sensor datasets). The MIT sensor dataset also highlights washing hands as a frequently occurring task, emphasizing the importance of being able to detect and manipulate different knobs and handles with a robotic arm. The NTCIR video dataset included using electronic devices as one of the top three tasks, supporting the post-development preference survey and object ranking results noted in prior works [12], [65]. Our kitchen tasks motion analysis reveals that reaching to pick or place items, flipping switches, and transferring food are the most frequent motions carried out during food preparation. Overall task importance results from previous subjective works as well as our quantitative analysis imply that the ability to manipulate a large range of everyday objects within an unstructured environment is of the utmost importance.

A limitation of our work is that the underlying datasets analysed does not represent a diverse set of individuals and does not include the population of interest. The NTCIR video dataset is captured at a low frame rate, thus is susceptible to missing tasks that take less than 30 seconds to carry out. The MIT IoT sensor dataset is restricted to capturing tasks that have an associated sensor installed, thus misses tasks that occur outside of the home and tasks that involve objects without a sensor. Both datasets are limited by the low number of participants included in the studies. A promising avenue for future work is to collect a new lifelogging dataset at a higher frame rate with an emphasis on recruiting participants with a diverse background, taking care to include the population of interest as well as caregivers. This would help us to better understand the direct needs of those that would personally benefit from this technology. This new dataset could be used to extract not only task frequency and duration, but also capture the objects manipulated (and how often), tasks that are challenging to carry out (i.e., does task completion take longer than expected), tasks that individuals ask others to help them with, and whether or not a task requires one hand or two to properly carry out. Chapter 3 identifies tasks that would be beneficial to offer as autonomous high-level behaviours (e.g., flip light switch or open door) available for the end-user to choose as needed. Engineers cannot possibly anticipate all potential high-level behaviours that could be required for daily living. A natural way to solve this problem, would be to allow the end-user to teach their robotic system new behaviours as the need arises.

In Chapter 4 we propose a proof-of-concept system for teaching new autonomous behaviours, such as those identified in Chapter 3, through reinforcement signals of approval or disapproval. We demonstrate that it is possible to extend the TAMER framework to the realm of 3D robotic manipulation and show that this framework can be used to teach new reaching behaviours without an expert present. Such a system can help to reduce the risk of device abandonment by alleviating the need to reach out for technical support whenever the end-users needs change. A limitation of this work is that the goal location of the object being reached for was predefined. A natural extension of this is to use object detection and segmentation methods to automatically detect the goal objects location on the fly. Future work will also focus on using TAMER to teach the other high-priority behaviours highlighted in Chapter 3, such as automatically opening/closing doors, flipping switches, picking up and placing down objects. This could be incorporated into the end-users existing control system and offered as a *teaching mode* then the newly taught behaviours could be offered as new *autonomous modes* that are available. New experiments should be carried out to include participants from the population of interest.

References

[1] 3Dconnexion. “Spacemouse compact.” (2020), [Online]. Available: `https://3dconnexion.com/uk/product/spacemouse-compact/` (visited on 12/10/2020). 41

[2] P. Abbeel and A. Y. Ng, “Apprenticeship learning via inverse reinforcement learning,” *Proceedings, Twenty-First International Conference on Machine Learning, ICML 2004*, pp. 1–8, 2004. DOI: `10.1145/1015330.1015430`. 20

[3] P. Archambault and P. S. Archambault, “Evaluation of the JACO robotic arm: clinico- economic study for powered wheelchair users with upper-extremity disabilities Evaluation of the JACO robotic arm,” *Proceedings of the 2011 IEEE Int. Conf. on Rehabilitation Robotics (ICORR)*, 2011. DOI: `10.1109/ICORR.2011.5975397`. 6

[4] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, “A survey of robot learning from demonstration,” *Robotics and Autonomous Systems*, vol. 57, no. 5, pp. 469–483, 2009, ISSN: 09218890. DOI: `10.1016/j.robot.2008.10.024`. 19

[5] M. Beaudoin, J. Lettre, F. Routhier, P. S. Archambault, M. Lemay, and I. Gélinas, “Impacts of robotic arm use on individuals with upper extremity disabilities: A scoping review,” *Canadian Journal of Occupational Therapy*, vol. 85, no. 5, pp. 397–407, 2018. 9

[6] J. M. Beer, C.-A. Smarr, T. L. Chen, *et al.*, “The domesticated robot: design guidelines for assisting older adults to age in place,” *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction - HRI ’12*, pp. 335–342, 2012, ISSN: 2167-2121. DOI: `978-1-4503-1063-5`. [Online]. Available: `http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6249576`. 13, 28, 30, 32

[7] M. Bolanos, M. Dimiccoli, and P. Radeva, “Toward Storytelling from Visual Lifelogging: An Overview,” *IEEE Transactions on Human-Machine Systems*, vol. 47, no. 1, pp. 77–90, 2017, ISSN: 21682291. DOI: `10.1109/THMS.2016.2616296`. 24

[8] M. Bouton, *Learning and Behavior: A Contemporary Synthesis*. Sinauer Associates, 2007. 22

[9] S. W. Brose, D. J. Weber, B. A. Salatin, *et al.*, "The role of assistive robotics in the lives of persons with disability," *American Journal of Physical Medicine and Rehabilitation*, vol. 89, no. 6, pp. 509–521, 2010, ISSN: 08949115. DOI: 10.1097/PHM.0b013e3181cf569b. 1

[10] F. Chaumette, S. Hutchinson, P. Corke, *et al.*, "Visual Servoing," in *Handbook of Robotics*, B. Siciliano and O. Khatib, Eds., 2nd edition, Springer, 2016, pp. 841–866. 50

[11] S. Chiaverini, "Singularity-robust task-priority redundancy resolution for real-time kinematic control of robot manipulators," *IEEE Transactions on Robotics and Automation*, vol. 13, no. 3, pp. 398–410, 1997. DOI: 10.1109/70.585902. 49

[12] Y. S. Choi, T. Deyle, T. Chen, J. D. Glass, and C. C. Kemp, "A list of household objects for robotic retrieval prioritized by people with als," in *2009 IEEE International Conference on Rehabilitation Robotics*, IEEE, 2009, pp. 510–517. 11, 13, 54

[13] C.-S. Chung, H. Wang, and R. A. Cooper, "Functional assessment and performance evaluation for assistive robotic manipulators: Literature review," *The Journal of Spinal Cord Medicine*, vol. 36, no. 4, pp. 273–289, 2013, ISSN: 1079-0268. DOI: 10.1179/2045772313Y.0000000132. [Online]. Available: http://www.tandfonline.com/doi/full/10.1179/2045772313Y.0000000132. 7, 13, 14

[14] P. Corke, *Robotics, Vision and Control - Fundamental Algorithms in MATLAB®* (Springer Tracts in Advanced Robotics). Springer, 2011, vol. 73, pp. 1–495, ISBN: 978-3-642-20143-1. [Online]. Available: http://dblp.uni-trier.de/db/series/star/index.html#Corke11. 15

[15] R. Correal, A. Jardón, S. Martínez, R. Cabás, A. Giménez, and C. Balaguer, "Human-Robot Coexistance in Robot-Aided Apartment," *23rd International Symposium on Automation and Robotics in Construction*, pp. 170–175, 2006. 21

[16] Y. Cui, Q. Zhang, A. Allievi, P. Stone, S. Niekum, and W. B. Knox, "The EMPATHIC Framework for Task Learning from Implicit Human Feedback," pp. 1–23, 2020. arXiv: 2009.13649. [Online]. Available: http://arxiv.org/abs/2009.13649. 20

[17] Department of Economic and Social Affairs: Population Division, *World Population Ageing 2017 - Highlights*. 2017, p. 46, ISBN: 9789211515510. DOI: ST/ESA/SER.A/348. [Online]. Available: http://www.un.org/en/development/desa/population/publications. 32

[18] A. Dollar, "Classifying Human Hand Use and the Activities of Daily Living," *Springer Tracts in Advanced Robotics*, vol. 95, 2014, ISSN: 1610742X. DOI: 10.1007/978-3-319-03017-3. 23

[19] B. Driessen, H. Evers, and J. v Woerden, "Manus—a wheelchair-mounted rehabilitation robot," *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of engineering in medicine*, vol. 215, no. 3, pp. 285–290, 2001. 2, 6

[20] A. L. Edwards, M. R. Dawson, J. S. Hebert, *et al.*, "Application of real-time machine learning to myoelectric prosthesis control: A case series in adaptive switching," *Prosthetics and orthotics international*, vol. 40, no. 5, pp. 573–581, 2016. 17

[21] P. Encarnação and A. M. Cook, Eds., *Robotic Assistive Technologies: Principles and Practice*. 2018, ISBN: 9781498745727. 1

[22] A. Fathi, Y. Li, and J. M. Rehg, "Learning to recognize daily actions using gaze," *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 7572 LNCS, no. PART 1, pp. 314–327, 2012, ISSN: 03029743. DOI: 10.1007/978-3-642-33718-5{_}23. 25, 29

[23] A. Fathi, X. Ren, and J. M. Rehg, "Learning to recognize objects in egocentric activities," *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 3281–3288, 2011, ISSN: 10636919. DOI: 10.1109/CVPR.2011.5995444. 24

[24] S. Federici, M. Bracalenti, F. Meloni, and J. V. Luciano, "World Health Organization disability assessment schedule 2.0: An international systematic review," *Disability and Rehabilitation*, vol. 39, no. 23, pp. 2347–2380, 2017, ISSN: 14645165. DOI: 10.1080/09638288.2016.1223177. 9

[25] D. Gopinath, S. Jain, and B. D. Argall, "User-driven Optimization of Shared Autonomy in Assistive Robotics," *IEEE Robotics and Automation Letters*, vol. 2, no. 1, pp. 2–4, 2017. DOI: 10.1109/LRA.2016.2593928. 21

[26] E. K. Gordon, X. Meng, T. Bhattacharjee, M. Barnes, and S. S. Srinivasa, "Adaptive robot-assisted feeding: An online learning framework for acquiring previously unseen food items," *IEEE International Conference on Intelligent Robots and Systems*, no. 1, pp. 9659–9666, 2020, ISSN: 21530866. DOI: 10.1109/IROS45743.2020.9341359. arXiv: 1908.07088. 28

[27] C. Gurrin, H. Joho, F. Hopfgartner, L. Zhou, and R. Albatal, "NTCIR Lifelog: The First Test Collection for Lifelog Research," *39th International ACM SIGIR conference*, pp. 705–708, 2016. DOI: 10.1145/2911451.2914680. 23–25

[28] C. Gurrin, H. Joho, F. Hopfgartner, *et al.*, "Overview of NTCIR-13 Lifelog-2 Task," *Proceedings of the 13th NTCIR Conference on Evaluation of Information Access Technologies*, pp. 6–11, 2016. 25

[29] *Health care in canada, 2011. a focus on seniors and aging*, https://secure.cihi.ca/, Accessed: 2018-07-30. 14

[30] *Help with activities of daily living for people with a disability*, https://www150.statcan.gc.ca/n1/pub/11-008-x/2010002/article/11345-eng.htm1, Accessed: 2018-07-30. 14

[31] J. Hoey, P. Poupart, A. v. Bertoldi, T. Craig, C. Boutilier, and A. Mihailidis, "Automated handwashing assistance for persons with dementia using video and a partially observable markov decision process," *Comput. Vis. Image Underst.*, vol. 114, no. 5, pp. 503–519, May 2010, ISSN: 1077-3142. DOI: 10.1016/j.cviu.2009.06.008. [Online]. Available: https://doi.org/10.1016/j.cviu.2009.06.008. 28

[32] A. Jain and C. C. Kemp, "Pulling open doors and drawers: Coordinating an omni-directional base and a compliant arm with equilibrium point control," *Proceedings - IEEE International Conference on Robotics and Automation*, pp. 1807–1814, 2010, ISSN: 10504729. DOI: 10.1109/ROBOT.2010.5509445. 28

[33] S. Jain and B. Argall, "Probabilistic Human Intent Recognition for Shared Autonomy in Assistive Robotics," *ACM Transactions on Human-Robot Interaction*, vol. 9, no. 1, pp. 1–23, 2020, ISSN: 2573-9522. DOI: 10.1145/3359614. 20, 21

[34] J. Jin, L. Petrich, Z. Zhang, M. Dehghan, and M. Jagersand, "Visual Geometric Skill Inference by Watching Human Demonstration," *Proceedings - IEEE International Conference on Robotics and Automation*, pp. 8985–8991, 2020, ISSN: 10504729. DOI: 10.1109/ICRA40945.2020.9196570. arXiv: 1911.04418. 21

[35] S. Katz, "Assessing Self-maintenance: Activities of Daily Living , Mobility , and Instrumental Activities," *Journal of the American Geriatrics Society*, vol. 31, pp. 721–727, 12 1979. 8, 23

[36] D.-J. Kim, R. Hazlett-Knudsen, H. Culver-Godfrey, *et al.*, "How autonomy impacts performance and satisfaction: Results from a study with spinal cord injured subjects using an assistive robot," *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, vol. 42, no. 1, pp. 2–14, 2011. 32, 33

[37] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, Y. Bengio and Y. LeCun, Eds., 2015. [Online]. Available: http://arxiv.org/abs/1412.6980. 39

[38] Kinova. "Open and modular kinova gen3 robots." (2020), [Online]. Available: https://www.kinovarobotics.com/en/products/gen3-robot (visited on 12/10/2020). 40

[39] *Kinova gen3 ultra lightweight robot user guide*, Kinova inc., 2022. [Online]. Available: https://www.kinovarobotics.com/uploads/User-Guide-Gen3-R07.pdf. 16

[40] B. Knox. "Nexi: The teaching strategy of painting." (2020), [Online]. Available: http://www.bradknox.net/nexi/ (visited on 12/10/2020). 51

[41] W. B. Knox and P. Stone, "Interactively shaping agents via human reinforcement: The TAMER framework," *Proceedings of the 5th International Conference on Knowledge Capture*, pp. 9–16, 2009. DOI: 10.1145/1597735.1597738. 21, 22

[42] W. B. Knox, P. Stone, and C. Breazeal, "Training a robot via human feedback: A case study," *Lecture Notes in Computer Science*, vol. 8239 LNAI, pp. 460–470, 2013, ISSN: 03029743. DOI: 10.1007/978-3-319-02675-6_46. 22, 38, 49

[43] J. Kober, J. A. Bagnell, and J. Peters, "Reinforcement Learning in Robotics: A Survey," *Springer Tracts in Advanced Robotics*, vol. 97, pp. 9–67, 2013, ISSN: 1610742X. DOI: 10.1007/978-3-319-03194-1_2. 18, 19

[44] B. Langdon, H. Wang, and D. Ding, "Analysis of Assistive Robotic Manipulator (ARM) Performance Based on a Task Taxonomy," *RESNA Annual Conference*, 2016. [Online]. Available: https://www.resna.org/sites/default/files/conference/2016/emerging_technology/langdon.html. 11, 23

[45] a.D. Laud and A. Laud, "Theory and application of reward shaping in reinforcement learning," Ph.D. dissertation, University of Illinois at Urbana-Champaign, 2004. 19

[46] M. P. Lawton and E. M. Brody, "Assessment of Older People: Self-Maintaining and Instrumental Activities of Daily Living," *The Gerontologist*, vol. 9, no. 3, pp. 179–186, 1969, ISSN: 0016-9013. DOI: 10.1093/geront/9.3{_}Part{_}1.179. [Online]. Available: http://gerontologist.oxfordjournals.org/cgi/doi/10.1093/geront/9.3_Part_1.179. 8, 23

[47] D. P. Losey, K. Srinivasan, A. Mandlekar, A. Garg, and D. Sadigh, "Controlling Assistive Robots with Learned Latent Actions," *Proceedings - IEEE International Conference on Robotics and Automation*, pp. 378–384, 2020, ISSN: 10504729. DOI: 10.1109/ICRA40945.2020.9197197. arXiv: 1909.09674. 33

[48] V. Maheu, P. S. Archambault, J. Frappier, and F. Routhier, "Evaluation of the JACO robotic arm: Clinico-economic study for powered wheelchair users with upper-extremity disabilities," *IEEE International Conference on Rehabilitation Robotics*, pp. 4–6, 2011, ISSN: 19457898. DOI: 10.1109/ICORR.2011.5975397. 2

[49] S. Mann, "Wearable computing: A first step toward personal imaging," *Computer*, vol. 30, no. 2, pp. 25–32, 1997. DOI: 10.1109/2.566147. 24

[50] Y. Matsumoto, Y. Nishida, Y. Motomura, and Y. Okawa, "A concept of needs-oriented design and evaluation of assistive robots based on icf," in *2011 IEEE International Conference on Rehabilitation Robotics*, IEEE, 2011, pp. 1–6. 10

[51] V. Mnih, K. Kavukcuoglu, D. Silver, *et al.*, "Playing atari with deep reinforcement learning," *CoRR*, vol. abs/1312.5602, 2013. arXiv: `1312.5602`. [Online]. Available: `http://arxiv.org/abs/1312.5602`. 39

[52] K. Muelling, A. Venkatraman, J. Valois, *et al.*, "Autonomy infused teleoperation with application to brain computer interface controlled manipulation," *Autonomous Robots*, vol. 41, pp. 1401–1422, 2017. 30, 32

[53] J. Peters and S. Schaal, "Reinforcement learning of motor skills with policy gradients," *Neural Networks*, vol. 21, no. 4, pp. 682–697, 2008, ISSN: 08936080. DOI: `10.1016/j.neunet.2008.02.003`. 18

[54] B. Phillips and H. Zhao, "Predictors of Assistive Technology Abandonment," *Assistive Technology*, vol. 5, no. 1, pp. 36–45, 1993, ISSN: 19493614. DOI: `10.1080/10400435.1993.10132205`. 4, 7, 35

[55] H. Pirsiavash and D. Ramanan, "Detecting activities of daily living in first-person camera views," *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 2847–2854, 2012, ISSN: 10636919. DOI: `10.1109/CVPR.2012.6248010`. 24

[56] C. Quintero, M. Dehghan, O. Ramirez, M. Ang, and M. Jägersand, "Flexible virtual fixture interface for path specification in tele-manipulation," *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 5363–5368, 2017. 33

[57] C. Quintero, R. Tatsambon, M. Gridseth, and M. Jagersand, "Visual pointing gestures for bi-directional human robot interaction in a pick-and-place task," *Proceedings - IEEE International Workshop on Robot and Human Interactive Communication*, pp. 349–354, 2015. DOI: `10.1109/ROMAN.2015.7333604`. 32, 33

[58] Robotiq. "2f-85 and 2f-140 grippers." (2020), [Online]. Available: `https://robotiq.com/products/2f85-140-adaptive-robot-gripper` (visited on 12/10/2020). 40

[59] G. Romer, H. Stuyt, and A. Peters, "Cost-savings and economic benefits due to the assistive robotic manipulator (arm)," vol. 2005, Jan. 2005, pp. 201–204, ISBN: 0-7803-9003-2. DOI: `10.1109/ICORR.2005.1501084`. 6, 15

[60] S. Schaal, "Learning from demonstration," *Proceedings of the 9th International Conference on Neural Information Processing Systems*, pp. 1040–1046, 1996. DOI: `10.5555/2998981.2999127`. 19

[61] E. Scheme and K. Englehart, “Electromyogram pattern recognition for control of powered upper-limb prostheses: State of the art and challenges for clinical use.,” *Journal of Rehabilitation Research & Development*, vol. 48, no. 6, 2011. 17

[62] S. Sivakanthan, E. Blaauw, M. Greenhalgh, A. M. Koontz, R. Vegter, and R. A. Cooper, “Person transfer assist systems: A literature review,” *Disability and Rehabilitation: Assistive Technology*, vol. 16, no. 3, pp. 270–279, 2021. 31

[63] C. Smarr, C. B. Fausset, and W. A. Rogers, “Understanding the Potential for Robot Assistance for Older Adults in the Home Environment,” Atlanta, GA: Georgia Tech, School of Psychology, Tech. Rep., 2010, p. 36. 6

[64] W. Spector and J. Fleishman, “Combining activities of daily living with instrumental activities of daily living to measure functional disability,” *Journal of Gerontology*, vol. 53, no. I, pp. 46–56, 1998. [Online]. Available: `http://psychsocgerontology.oxfordjournals.org/content/53B/1/S46.short`. 8

[65] C. A. Stanger, C. Anglin, W. S. Harwin, and D. P. Romilly, “Devices for Assisting Manipulation: A Summary of User Task Priorities,” *IEEE Transactions on Rehabilitation Engineering*, vol. 2, no. 4, pp. 256–265, 1994. 2, 13, 30, 32, 54

[66] K. Stubbs, D. Wettergreen, and P. J. Hinds, “Autonomy and common ground in human-robot interaction: A field study,” *IEEE Intelligent Systems*, vol. 22, no. 2, pp. 42–50, 2007, ISSN: 15411672. DOI: `10.1109/MIS.2007.21`. 22

[67] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd. Cambridge, Massachusetts: The MIT Press, 2018, ISBN: 9788578110796. arXiv: `arXiv:1011.1669v3`. 18

[68] R. S. Sutton, D. Precup, and S. Singh, “Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning,” *Artificial Intelligence*, vol. 112, no. 1, pp. 181–211, 1999, ISSN: 00043702. DOI: `10.1016/S0004-3702(99)00052-1`. 18

[69] H. Tanaka, M. Yoshikawa, E. Oyama, Y. Wakita, and Y. Matsumoto, “Development of assistive robots using international classification of functioning, disability, and health: Concept, applications, and issues,” *Journal of Robotics*, vol. 2013, 2013. 11

[70] E. M. Tapia, S. S. Intille, and K. Larson, “Activity recognition in the home using simple and ubiquitous sensors,” in *International conference on pervasive computing*, Springer, 2004, pp. 158–175. 25

[71] M. Tenorth, J. Bandouch, and M. Beetz, “The tum kitchen data set of everyday manipulation activities for motion tracking and action recognition,” pp. 1089–1096, 2009. 33

[72] *The sustainability of health care spending in canada 2017*, `https://www.fraserinstitute.org/sites/default/files/sustainability-of-health-care-spending-in-canada-2017.pdf`, Accessed: 2018-07-30. 14

[73] T. B. Üstün, S. Chatterji, N. Kostanjsek, *et al.*, "Developing the world health organization disability assessment schedule 2.0," *Bulletin of the World Health Organization*, vol. 88, no. 11, pp. 815–823, 2010, ISSN: 00429686. DOI: `10.2471/BLT.09.067231`. 7, 9

[74] P. Wang, L. Sun, S. Yang, A. Smeaton, and C. Gurrin, "Characterizing everyday activities from visual lifelogs based on enhancing concept representation," *Comput. Vis. Image Underst.*, vol. 148, pp. 181–192, 2016. 23

[75] P. Wang and A. F. Smeaton, "Using visual lifelogs to automatically characterize everyday activities," *Information Sciences*, vol. 230, pp. 147–161, 2013, ISSN: 00200255. DOI: `10.1016/j.ins.2012.12.028`. [Online]. Available: `http://dx.doi.org/10.1016/j.ins.2012.12.028`. 23

[76] G. Wilson, D. Jones, P. Schofield, and D. Martin, "Experiences of using a wearable camera to record activity, participation and health-related behaviours: Qualitative reflections of using the sensecam," *Digital Health*, vol. 2, Nov. 2016. DOI: `10.1177/2055207616682628`. 24

[77] World Health Organization (WHO), "International Classification of Functioning, Disability and Health," *World Health Organization*, no. Version 2.1a, Clinician Form, pp. 1–15, 2003, ISSN: 08963207. [Online]. Available: `http://www.who.int/classifications/icf/icfchecklist.pdf?ua=1`. 7, 9