Optimization Techniques for an Ergonomic Human-Robot Interaction

Baptiste Busch

université de BORDEAUX

Inria
INVENTEURS DU MONDE NUMÉRIQUE

THÈSE PRÉSENTÉE
POUR OBTENIR LE GRADE DE

DOCTEUR DE
L'UNIVERSITÉ DE BORDEAUX

ÉCOLE DOCTORALE DE
MATHÉMATIQUES ET INFORMATIQUE

SPÉCIALITÉ : INFORMATIQUE

Par Baptiste BUSCH

Optimization Techniques for an Ergonomic Human-Robot Interaction

---

Sous la direction de : Manuel LOPES

Date de soutenance : 27 février 2018

Membres du jury :

Olivier STASSE Rapporteur, Président du jury
Bram VANDERBORGHT Rapporteur
David DANEY Examinateur
Pierre-Yves OUDEYER Co-Directeur

Remerciements

Ces années de thèse ont été, pour moi, une excellente expérience, pleines de défis et de rencontres passionnantes.

Je souhaiterais avant tout remercier Manuel Lopes, pour m'avoir permis de vivre cette aventure, pour son soutien, sa confiance et l'appui qu'il m'a apporté tout au long de cette thèse. Je remercie aussi Pierre-Yves Oudeyer pour sa bienveillance et sa volonté de toujours faire ressortir le meilleur de chacun d'entre nous.

Au cours de cette thèse j'ai eu la chance d'être entouré par de brillants chercheurs qui m'ont aidé à mettre au point mes travaux et dans l'écriture de mes publications, Freek Stulp, Jonathan Grizou, Guilherme Maeda et Marc Toussaint qui, bien que surbookés, ont toujours pris le temps d'apporter leur aide et soutien.

Mes pensées sont aussi pour feu P^{r} Wisama Khalil, en charge du Master EMARO lors de ma graduation, qui m'a donné l'envie et les bases requises pour préparer cette thèse de doctorat dans les meilleures conditions. Je remercie aussi Sophie Sakka, ma directrice de thèse de Master, pour m'avoir formé à la recherche. Je n'oublie pas non plus mes professeurs de l'EISTI, particulièrement Peïo Loubières, Yannick Le Nir, Laureance Lamoulie, Nisrine Fortin, et Pierre Daugé, pour m'avoir formé à la programmation et fourni les bases en mathématiques et physiques qui m'ont servi tout au long de cette recherche.

Sans mes collègues, il est évident que cette thèse n'aurait pu être ce qu'elle est aujourd'hui. Tout d'abord ceux qui ont travaillé avec moi sur le projet 3rdHand, Thibaut Munzer qui a toujours su trouver du temps pour m'aider dans les moments de doutes, et Yoan Mollard à qui je souhaite une très belle carrière dans l'entrepreneuriat. Ceux qui ont partagé mon bureau, Benjamin Clément, pour nos discussions animées, et Didier Roy, qui m'a aussi partagé ses réunions téléphoniques, mais surtout sa bonne humeur. Puis il y a les anciens, Mattieu Lapeyre, Nicolas Rabault et Pierre Rouanet, pour cet attrait pour l'humour que nous partageons, Fabien Bénureau pour nos débats qui m'ont toujours permis de corriger bon nombre de mes faiblesses. Et le reste de l'équipe FLOWERS, Damien Caselli pour nos pauses cafés philosophiques, Théo Segonds, Florian Golemo, et Sébastien Forestier pour l'intérêt que nous partageons envers les nouvelles technologies, Stéphanie Noirpoudre pour sa gaieté contagieuse et Alexandra Delmas pour son courage et sa détermination exemplaire. Je n'oublierai bien sur pas Nicolas Jahier qui a toujours été présent pour nous simplifier les démarches administratives. Enfin je souhaiterais finir par William Schueller dont l'humour aura toujours, pour moi, le dernier mot.

En dernier, je remercie ma famille, ma mère qui a presque toujours compris l'ensemble de mes travaux, Fleur Vannière, ma future femme, qui m'a apporté son soutien indéfectible. Mes amis d'EMARO, Kateryna Parfeniuk, Eloise Matheson et Tilman Schmidt. Zait, Gasai, Swad et Tiny pour nos (trop) longues soirées jeux-vidéos. Et je terminerai par mes amis fidèles, Eric Hostalery, Gauthier Parent, Teddy Hourclats, et Nicolas Chauchat, qui, non content de m'avoir supporté pendant cette thèse, m'accompagnent aussi pour mon mariage.

Acknowledgments

Those years of thesis have been a wonderful experience, full of challenges and enriching encounters.

I would like first to thanks Manuel Lopes whith whom I was able to live this adventure, for his support, his trust, and the help he has provided throughout this thesis. I thanks also Pierre-Yves Oudeyer for his kindness, and his desire to bring out the best in every one of us.

During this thesis I had the chance to be surrounded with brilliant researchers who have helped me to devise my work and in the writing of my publications, Freek Stulp, Jonathan Grizou, Guilherme Maeda, and Marc Toussaint who despite being overbooked were always available to provide their help and support.

My thoughts are also with late Professor Wisama Khalil, in charge of the master EMARO by the time of my graduation, who gave me the desire and the necessary basis to prepare this doctoral thesis in the best conditions. I also thanks Sophie Sakka, my master thesis director, who taught me the basis of research. I don't forget to mention my teachers in EISTI, especially Peïo Loubières, Yannick Le nir, Laurence Lamoulie, Nisrine Fortin et Pierre Daugé, who have taught me the basis of programming, mathematics, and physics that have served me well throughout this research.

Without my colleagues, this thesis would have obviously not been the way it is today. First those who have worked with me on the 3rdHand project, Thibaut Munzer who made himself always available to help me during the hard times, and Yoan Mollard to whom I wish a great career in the entrepreneurship. Those who have shared my office, Benjamin Clément for our vivid discussions, and Didier Roy who has also shared his phone meetings but most importantly his cheerfulness. Then there are the former members of the FLOWERS team, Matthieu Lapeyre, Nicolas Rabault, and Pierre Rouanet for this attraction for humor that we share, Fabien Bénureau for our debates that have always helped me to correct my weaknesses. And the rest of team, Damien Caselli for our philosophical coffees, Théo Segonds, Florian Golemo and Sébastien Forestier for our shared attraction towards new technologies, Stéphanie Noirpoudre for her contagious joy, and Alexandra Delmas for her courage and her exemplary determination. I also don't forget Nicolas Jahier who have always been available to simplify administration as far as possible. Finally I would like to finish with William Schueller whose humor will always, to me, have the last word.

Last but not least, I would like to thanks my family, my mum who has almost always understand all my work, Fleur Vannière, my bride to be, who has provided me with unwavering support. My friends from EMARO, Kateryna Parfeniuk, Eloise Matheson and Tilman Schmidt. Zait, Gasai, Swad, and Tiny for our (too) long video-games evening sessions. And I will finish with my faithful friends, Eric Hostalery, Gauthier Parent, Teddy Hourclats, and Nicolas Chauchat, who were not only bearing with me during this thesis, but are also accompanying me with my wedding.

Abstract

L'interaction Humain-Robot est un domaine de recherche en pleine expansion parmi la communauté robotique. De par sa nature il réunit des chercheurs venant de domaines variés, tels que psychologie, sociologie et, bien entendu, robotique. Ensemble, ils définissent et dessinent les robots avec lesquels nous interagirons dans notre quotidien.

Comme humains et robots commencent à travailler en environnement partagés, la diversité des tâches qu'ils peuvent accomplir augmente drastiquement. Cela créé de nombreux défis et questions qu'il nous faut adresser, en terme de sécurité et d'acceptation des systèmes robotiques. L'être humain a des besoins et attentes bien spécifiques qui ne peuvent être occultés lors de la conception des interactions robotiques. D'une certaine manière, il existe un besoin fort pour l'émergence d'une véritable interaction humain-robot ergonomique.

Au cours de cette thèse, nous avons mis en place des méthodes pour inclure des critères ergonomiques et humains dans les algorithmes de prise de décisions, afin d'automatiser le processus de génération d'une interaction ergonomique. Les solutions que nous proposons se basent sur l'utilisation de fonctions de coût encapsulant les besoins humains et permettent d'optimiser les mouvements du robot et le choix des actions. Nous avons ensuite appliqué cette méthode à deux problèmes courants d'interaction humain-robot.

Dans un premier temps, nous avons proposé une technique pour améliorer la lisibilité des mouvements du robot afin d'arriver à une meilleure compréhension des ses intentions. Notre approche ne requiert pas de modéliser le concept de lisibilité de mouvements mais pénalise les trajectoires qui amènent à une interprétation erronée ou tardive des intentions du robot durant l'accomplissement d'une tâche partagée. Au cours de plusieurs études utilisateurs nous avons observé un gain substantiel en terme de temps de prédiction et une réduction des erreurs d'interprétation.

Puis, nous nous sommes attelés au problème du choix des actions et des mouvements qui vont maximiser l'ergonomie physique du partenaire humain. En utilisant une mesure d'ergonomie des postures humaines, nous simulons les actions et mouvements du robot et de l'humain pour accomplir une tâche donnée, tout en évitant les situations où l'humain serait dans une posture de travail à risque. Les études utilisateurs menées montrent que notre méthode conduit à des postures de travail plus sûr et à une interaction perçue comme étant meilleure.

Mots-Clés: Interaction Humain-Robot, Ergonomie, Facteurs Humains, Mouvements Lisibles, Apprentissage par l'Interaction, Planification de Tâches et de Mouvements.

Ce travail a été supporté par des fonds nationaux au travers de Fundação para a Ciência e a Tecnologia (FCT), référence UID/CEC/50021/2013, et par EU FP7-ICT projet 3rdHand sous l'autorisation n^{o} 610878.

Abstract

Human-Robot Interaction (HRI) is a growing field in the robotic community. By its very nature it brings together researchers from various domains including psychology, sociology and obviously robotics who are shaping and designing the robots people will interact with on a daily basis.

As human and robots starts working in a shared environment, the diversity of tasks they can accomplish together is rapidly increasing. This creates challenges and raises concerns to be addressed in terms of safety and acceptance of the robotic systems. Human beings have specific needs and expectations that have to be taken into account when designing robotic interactions. In a sense, there is a strong need for a truly ergonomic human-robot interaction.

In this thesis, we propose methods to include ergonomics and human factors in the motions and decisions planning algorithms, to automatize this process of generating an ergonomic interaction. The solutions we propose make use of cost functions that encapsulate the human needs and enable the optimization of the robot's motions and choices of actions. We have applied our method to two common problems of human-robot interaction.

First, we propose a method to increase the legibility of the robot motions to achieve a better understanding of its intentions. Our approach does not require modeling the concept of legible motions but penalizes the trajectories that leads to late or mispredictions of the robot's intentions during a live execution of a shared task. In several user studies we achieve substantial gains in terms of prediction time and reduced interpretation errors.

Second, we tackle the problem of choosing actions and planning motions that maximize the physical ergonomics on the human side. Using a well-accepted ergonomic evaluation function of human postures, we simulate the actions and motions of both the human and the robot, to accomplish a specific task, while avoiding situations where the human could be at risk in terms of working posture. The conducted user studies show that our method leads to safer working postures and a better perceived interaction.

Keywords: Human-Robot Interaction, Ergonomics, Human Factors, Legible Motions, Learning from Interaction, Task and Motion Planning.

This work was supported by national funds through Fundação para a Ciência e a Tecnologia (FCT) with reference UID/CEC/50021/2013 and by the EU FP7-ICT project 3rdHand under grant agreement Nº 610878.

Contents

List of Figures

List of Tables

Résumé en Français

Le domaine de l'Interaction Homme-Robot (IHR) a vu son essor se renforcer, ces dernières années, par l'émergence de systèmes robotiques complexes et d'une intelligence artificielle de plus en plus poussée. Les robots ne sont plus simplement des machines industrielles cloisonnées afin d'éviter tout contact avec les ouvriers. De nombreuses industries investissent dans des solutions robotiques où ouvriers et machines travaillent en contact direct afin de bénéficier des forces et avantages des deux parties, i.e puissance et précision de la machine combinées aux capacités d'adaptation et de la dextérité accrue de l'être humain. On parle alors de *cobots* ou *cobotique*, un néologisme issu de la contraction des mots *coopération* et *robotique* proposé par Colgate, Peshkin, and Wannasuphoprasit (1996), tel que celui présenté en Figure 1.

Figure 1: Cobot de la firme ABB. Photo par Marco Verch.

Cette situation crée une opportunité pour l'émergence d'une robotique sociale où les robots assisteront l'Homme dans ses tâches quotidiennes, au travail mais aussi chez lui. De ce fait, elle force aussi à l'étude poussée des comportements humains afin de faciliter l'intégration des systèmes robotiques dans la société. L'être humain est un système complexe dont les comportements sociaux et les attentes qui en découlent en terme d'interactions ont été forgés par plusieurs milliers d'années d'évolution. À contrario, les systèmes robotiques, à l'échelle de l'humanité, n'en sont qu'à leurs balbutiements. Il est donc crucial de bien comprendre et analyser les comportements sociaux formant les interactions avec nos pairs afin de les retranscrire dans les robots que nous créons et déployons. Tout robot qui ne satisfait pas les attentes sociales les plus basiques sera simplement oublié et finira par être inutilisé (De Graaf et al. 2017). C'est en ce sens que l'on peut parler de la nécessité d'une véritable *robotique ergonomique*[1].

Améliorer l'ergonomie des systèmes robotiques passe bien évidemment par l'étude et la conception de robots plus attrayants et inspirant une plus forte confiance (Breazeal 2004). Néanmoins, le comportement et les actions choisies par le robot jouent aussi un rôle crucial dans son acceptation et le confort de son partenaire humain. Le projet européen 3rdHand, dans le cadre duquel s'inscrit cette thèse, vise à la création d'un assistant robotique pour les ouvriers sur chaîne de montage. Cet assistant se veut autonome et capable

[1] "L'ergonomie est l'étude scientifique de la relation entre l'homme et ses moyens, méthodes et milieux de travail. Son objectif est d'élaborer, avec le concours des diverses disciplines scientifiques qui la composent, un corps de connaissances qui dans une perspective d'application, doit aboutir à une meilleure adaptation à l'homme des moyens technologiques de production, et des milieux de travail et de vie." (Mosar et al. 1969)

de répondre aux besoins de l'ouvrier pour l'accompagner dans son travail. Les solutions envisagées doivent donc être le plus génériques possible et, de ce fait, il est nécessaire qu'elles ne soient pas spécifique à un modèle de robot en particulier. La cible principale du projet étant les petites et moyennes entreprises, qui possèdent des chaînes de montage variables où le changement de produits peut être fréquent, il est aussi requis de proposer des solutions s'adaptant automatiquement et demandant le minimum de programmation possible. C'est dans cette optique que nous avons choisi, au cours de cette thèse, d'étudier des méthodes permettant de générer une interaction plus ergonomique, sans apporter de modifications matérielles à un robot déjà existant. Notre choix s'est donc naturellement porté sur l'étude de l'impact des mouvements et des actions choisies par la machine sur le partenaire humain, tant d'un point de vue physique que psychologique.

Deux cas d'études ont alors été considérés. Le premier est la génération de mouvements plus prévisibles, permettant au collaborateur de lire les intentions du robot sans passer par une communication verbale[2]. En terme robotique, on parle de *lisibilité* des trajectoires (Dragan, K. C. Lee, et al. 2013). Cette lisibilité accrue permet d'améliorer la productivité de l'équipe humain-robot en réduisant les erreurs dues à une mauvaise compréhension des intentions de la machine. Si, en observant les mouvements du robot, je suis persuadé qu'il se dirige vers le tournevis je peux alors commencer à préparer les vis pour la tache suivante. Si, par malheur, il se dirigeait en fait vers un autre outil je vais devoir changer mon action complémentaire, entraînant un ralentissement du rythme de production et une certaine frustration.

Dans la littérature, la plupart des méthodes cherchant à améliorer la lisibilité des trajectoires reposent sur la création d'un modèle probabiliste de ce que représente une trajectoire prévisible. Cette démarche, longuement étudiée par Dragan and Siddhartha S. Srinivasa (2013) et illustré en Figure 2, suppose que la lisibilité de trajectoire est un concept universel, i.e. qui ne dépend pas d'un contexte culturel ou social particulier. Par ailleurs, cela demande la création d'un nouveau modèle pour chaque changement de taches, ce qui va à l'encontre d'un robot adaptatif et simple à programmer.

[2] Cette démarche s'inscrit dans le constat que la plupart des communications entre partenaires humains est non verbal (Harrison 1965). De ce fait il est nécessaire de considérer une approche similaire afin de générer une interaction plus naturelle.

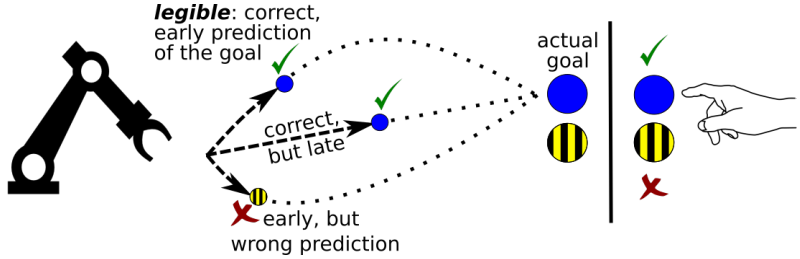

Figure 2: Illustration de trajectoires lisibles. Lorsque le robot est confronté à deux cibles possibles, une exagération de la trajectoire sur le côté entraîne une prédiction plus rapide. A l'inverse, une exagération du mauvais côté peut générer des erreurs d'interprétations.

À contrario, la solution que nous avons proposée est l'optimisation des trajectoires en se basant sur une méthode *essai-erreur* par interaction directe avec l'ouvrier sur la tâche à accomplir (Busch, Grizou, et al. 2017; Stulp, Grizou, et al. 2015). Cette approche présente l'intérêt de ne pas émettre d'hypothèses, potentiellement biaisées, sur le concept même de lisibilité. Par ailleurs, tout changement dans la tache à accomplir requiert simplement une nouvelle période d'adaptation.

La méthode d'optimisation que nous avons choisie repose sur le principe de *boite noire*[3] et est dénommée PI^{BB} (Stulp and Sigaud 2012). Nous avons testé cette approche via une expérience basée sur un système similaire à l'illustration en Figure2, i.e le robot décide d'un bouton à presser et le sujet doit presser le même bouton le plus rapidement possible. Pour évaluer les trajectoires générées par l'algorithme d'optimisation nous avons repris les fondements même du concept de lisibilité, tels que décrit par Dragan, K. C. Lee, et al. (2013), i.e. une trajectoire est lisible lorsqu'elle entraîne une prédiction plus rapide et sans erreur des intentions du robot. De ce fait, la fonction d'évaluation que nous avons considérée comporte trois éléments:

[3] L'optimisation en boite noire est utilisée lorsque l'on possède peu d'informations sur la fonction à optimiser. L'algorithme se base simplement sur une méthode d'évaluation des tirages aléatoires sans connaissance du modèle sous-jacent.

$$J = \underbrace{T_{\text{robot}} + T_{\text{subject}}}_{\text{Éfficacité}} + \underbrace{\gamma \delta_{\text{buttons}}}_{\text{Robustesse}} + \underbrace{\alpha \left| \dddot{\mathbf{q}}_{1\ldots N, 1\ldots T} \right|}_{\text{Énergie}} \tag{0.1}$$

Éfficacité: Le temps entre le début du mouvement du bras du robot (t_0) et la pression sur le bouton par l'humain (T_{subject}) et le robot (T_{robot}).

Robustesse: Erreur de prédiction. Une valeur est donnée si la personne a pressé le bon bouton (δ_{buttons} =0) ou non (δ_{buttons} =1). γ est un fort coût déterminé arbitrairement et fixé à 20 dans nos expériences, ce qui équivaut à une pénalité de 20*s* en termes d'efficacité.

Énergie: La somme du jerk, i.e. la dérivée troisième de la valeur angulaire des joints du robot ($\dddot{\mathbf{q}}_{t_i}$) à chaque pas i de la trajectoire. Le pas Δ_t utilisé pour calculer les dérivées a été fixé à 0.2. Le facteur d'échelle α est choisi de manière à ce que le coût en jerk représente 1/20 du coût total de la trajectoire initiale.

Les expériences menées auprès de 30 volontaires ont montré une nette diminution d'environ 20% du temps de prédiction lors de l'observation des trajectoires optimisées comparé aux trajectoires initiales correspondant à des lignes droites vers les boutons à presser. Au vu de ces résultats, nous avons effectué une étude plus approfondie afin de vérifier l'influence de la politique représentant les trajectoires sur la lisibilité finale et évaluer à quelle point les trajectoires optimisées peuvent se transférer d'un utilisateur à l'autre (Busch, Grizou, et al. 2017).

Le second cas étudié lors de cette thèse concerne une ergonomie plus physique. Selon Punnett et al. (2004) les Troubles Musculo-Squelettiques (TMS) représentent la majeure partie des accidents du travail dans les pays industrialisés. Ils sont souvent induits par une mauvaise posture ou de mauvaises habitudes de travail, ainsi que par la répétition de tâches à risque (Gallagher et al. 2013). Si nous sommes amenés à introduire des robots sur les chaînes de montage, cela ne peut se faire au détriment de la santé des ouvriers. De ce fait, il est important de mettre en place des solutions, en amont, afin de limiter les risques de TMS.

Le choix du placement des objets par le robot et des action qu'il peut accomplir n'est pas anodin et peut avoir un impact négatif sur la santé de son partenaire humain. Prenons l'exemple illustré en Figure 3. Sur la photo la plus à gauche, il est clair que le choix du placement de l'objet sphérique n'est pas judicieux. Ce type d'interaction aura, pour sûr, de lourdes conséquences sur la santé de l'ouvrier sur le long terme. À l'inverse, la photo de droite présente une interaction plus ergonomique.

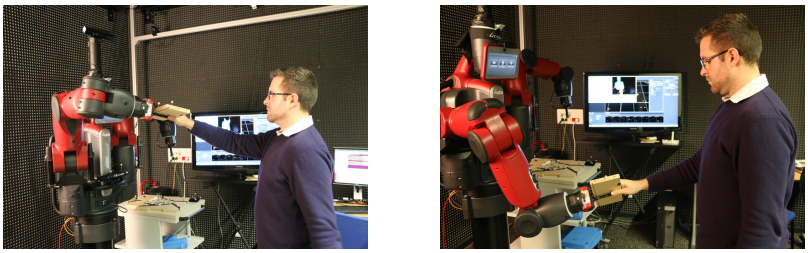

Figure 3: Différences entre deux types d'interaction. Une interaction ayant un impact négatif sur la santé de l'ouvrier (gauche) comparé à une interaction plus saine (droite)

La problématique est donc d'automatiser le choix des actions robotiques, d'un point de vue moteur comme d'un point de vue décisionnel, afin d'obtenir une interaction ergonomique et sécurisée pour l'ouvrier (Busch, Maeda, Mollard, et al. 2017; Busch, Toussaint, et al. 2018). Pour ce faire, nous proposons une approche d'optimisation à priori. À partir d'un modèle personnalisé de l'ouvrier nous déterminons la posture présentant le moins de risque de TMS lors de l'accomplissement d'une tâche puis établissons les mouvements du robot qui amèneront l'ouvrier à adopter cette posture optimal. Afin d'attribuer un score à la posture de l'ouvrier nous avons mis en place une évaluation automatique basée sur la méthode *Rapid Entire Body Assessment (REBA)* (Hignett et al. 2000). Le détail des calculs introduits par la méthode REBA est disponible en Annexe A. Une étude utilisateur sur 40 volontaires montre que l'interaction optimisée est largement préférée et les postures enregistrées présentent un risque réduit de TMS (Busch, Maeda, Mollard, et al. 2017).

Dans un second temps, nous avons introduis la méthode REBA dans un logiciel de planification développé par Toussaint and Lopes (2016). L'avantage de cette approche est de pouvoir résoudre les problèmes d'allocation de tâche et l'optimisation de mouvements simultanément[4]. De ce fait, cela nous permet de proposer une interaction ergonomique à la fois en terme de mouvements, mais aussi

[4] Ce type de problème est dénommé *Task and Motion Planning (TAMP)*. Il permet d'attribuer des tâches à valeurs sémantiques, e.g. *le robot prend le tournevis* et de résoudre leurs impacts logique et géométrique en terme d'effet sur la scène.

en terme du choix des actions à accomplir. Par exemple, les taches à faible valeur ajoutée, comme soulever des objets lourds, peuvent être automatiquement attribuées au robot, réduisant la fatigue physique des ouvriers (Busch, Toussaint, et al. 2018).

En conclusion, au travers de plusieurs expériences d'interaction, nous avons proposé des méthodes permettant de rendre l'interaction plus ergonomique. Dans un soucis de généralité et de réutilisabilité, ces méthodes sont indépendantes d'un robot particulier[5]. Cette philosophie transparaît aussi dans le choix des approches et leur implémentation. L'utilisation d'une approche non basée sur un modèle pour améliorer la lisibilité des trajectoires permet, par exemple, de s'affranchir de la programmation de spécificités qui peuvent être culturelles ou même dépendante d'un ouvrier en particulier. De ce fait, cela permet aussi de proposer une interaction unique et personnalisée. De la même manière, même si nos méthodes pour améliorer l'ergonomie physique, reposent sur l'utilisation d'une technique d'évaluation spécifique, la technique REBA, cette dernière peut être remplacée par une technique équivalente, potentiellement plus adaptée à certains besoins ou à certaines situations particulières.

[5] Bien que vérifiées et implémentées sur un seul robot, le robot Baxter du projet 3rdHand, les méthodes présentées dans cette thèse sont indépendantes des caractéristiques physiques du robot et peuvent donc être implantées sur d'autres robots.

1
Introduction

We, human beings, are very social creatures. We interact with each other–and also other animals–on daily basis. Evolution has favored the development of complex social behaviors in humans, along with the brain architecture that supports them (Bjorklund et al. 1995). Compared to other mammals, humans have the largest neocortex, a brain area responsible for language acquisition, conscious thought and emotion regulation. The capacity for attributing mental states, i.e. intentions, beliefs, and desires, to others has been defined Theory of Mind (Premack et al. 1978) and is also thought to be regulated by the neocortex. We are, in a certain manner, hardwired to be sociable and to express social behaviors.

Figure 1.1: Animal used to help farming in Indonesia. Photo by Jan-Pieter Nap.

We also communicate, and cooperate, to achieve our goals. None of our greatest achievements or constructions were made by a single person but rather groups of people, often at the cost of their own life. This behavior is quite unique in the animal kingdom and probably originate from genetic and cultural evolution that has produced a species in which substantial numbers make sacrifices to uphold ethical norms and to help even total strangers (Bowles et al. 2011). Darwin proposed "social and moral faculties" as a key factor for the survival of human tribes in the early stage of their development (Darwin 1888):

> It must not be forgotten that, although a high standard of morality gives but a slight or no advantage to each individual man and his children over the other men of the same tribe, yet that an advancement in the standard of morality and an increase in the number of well-endowed men will certainly give an immense advantage to one tribe over another. There can be no doubt that a tribe including many members who, from possessing in a high degree the spirit of patriotism, fidelity, obedience, courage and sympathy, were always ready to give aid to each other and to sacrifice themselves for the common good, would be victorious over most other tribes; and this would be natural selection.

Not only do we collaborate with each other, but we also develop tools and make use of other animals to help us in our labor. Beeves and horses have been extensively used throughout history, and still nowadays in developing countries as illustrated in Figure 1.1, for their strength, compensating our own relative weakness. Tool creation and usage has greatly contributed to our evolution and is strongly linked with increased brain size, population size, and geographical range (Ambrose 2001). Our most complex robotic systems are, therefore, nothing more than an improvement of the first bone tool depicted by Kubrick in his movie *2001: A Space Odissey*.

Figure 1.2: Probably the most famous couple of robots in science fiction. C3PO and R2D2 have inspired many researches on social robots and human-robot interaction (HRI). Credit: Gordon Tarpley

Relying on this, roboticists are trying to enhance the human capabilities by developing machines, more and more capable, but also more and more complex. However, for a long time robots and humans had very few social interactions. For security reasons, norms in industry have created barriers between them, limiting the interactions to their strict minimum, most of them to program or repair the expensive machineries (*Robots and Robotic Devices – Safety Requirements for Industrial Robots – Part 1: Robots* 2011).

Science fiction, on the other hand, is full of example of social robots, interacting with human beings, for better or for worse. In its own manner, it has shape our ideas on robotics, both in terms of designs and capabilities[1]. Interestingly, it also highlights, and often relies on, the human capacity to recognize living patterns and social behaviors in other lifeforms, even artificial mechanism or machines. Referred as anthropomorphism, it seems to be an innate tendency of the human psychology. From animals we observe, to spiritual deities we worship, we tends to explain their behavior with human-like goals and intentions (Epley et al. 2007). This capacity allows us, for example, to interpret the language of R2D2, the can-like robot of Figure 1.2, which is only composed of a succession of *beeps* and whistles. Although this "language" is artificial[2], we seem to understand what the robot "feels" on the moment.

Over the past few years, the extensive progresses in terms of safety and sensing capacities opened up challenges and opportunities to make those fictional robots a reality. The neologism *cobot*, or *co-robot*, refers to a robot that shares its workspace with humans and physically interacts with them. Invented in 1996 by Colgate, Peshkin, and Wannasuphoprasit (1996) it is described in a US patent (Colgate and Peshkin 1999) as "an apparatus and method for direct physical interaction between a person and a general purpose manipulator controlled by a computer". Nowadays, many industrial companies have designed their own cobots. Some of them, like Baxter robot from Rethink Robotics are safe by design, i.e velocity of the robot motion and deployed forces are limited by the actuator design (Robotics 2017). Lots of effort are made to create robots that are safe and will be well accepted by industrial workers.

[1] Throughout history, myths and fictional arts have modeled our views on artificial machines like robots. It might explain why Western people tend to be afraid by most advanced robots while the same robots are well integrated in the Japanese society (Kaplan 2004).

[2] The sound designer of both movies, Ben Burtt, has pointed out that some of these sounds were purely made on a synthesizer, while others origin from recorded mechanical and motor sounds (Rinzler et al. 2010)

As robots become more and more social, new challenges arise, with the increase of end users expectations. Extensively studied by Reeves et al. (1996) our relations with computer and television might seem irrational and might look like interactions we have with our peers. This effect also appears in our relation with robotic system (Meerbeek et al. 2008). This type of relation is a great opportunity for roboticists as it facilitates the acceptance of their robotic systems and design. But it is, in fact, a double-edged sword as people tend to overestimate the capacities and intelligence of the robots they are interacting with. The disillusion when they realize its limitations is one of the first reason they simply drop using it (De Graaf et al. 2017).

To overcome this, robotic designers must understand human sociology to facilitate the acceptance of the their systems. For example, any robot that do not respect some simple codes of interaction will simply fail to complete its purpose. As such, it is not surprising to see why so many researches rely on "Wizard of Oz" which refers to social robots partially or fully remotely controlled. It has been named after the character in Baum (1900)'s fictional book. This solution aims at analyzing how people respond to their robotic system, mainly its design, before developing complex autonomous systems. In a way we can see this as the need for a truly *ergonomic robotics*[3]

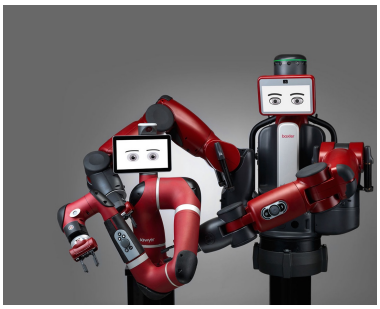

Figure 1.3: Baxter and Sawyer collaborative robots from Rethink Robotics.

[3] Ergonomics (or human factors) is the scientific discipline concerned with the understanding of interactions among humans and other elements of a system, and the profession that applies theory, principles, data and methods to design to optimize human well-being and overall system performance—International Ergonomics Association

The 3rdHand European Project, in which this thesis takes part, aims at developing a semi-autonomous robot assistant that acts as the third hand of a human worker. To facilitate the integration of such a technology in industry it is necessary to develop robotic behaviors that will ease the collaboration. Research areas leaded by the institutions of the 3rdHand committee, cover a large variety of topic directly linked to Human-Robot Interaction (HRI). One aspect, which motivated the researches in this thesis, is to improve the comfort of the human coworker during the interaction, i.e. to propose a more ergonomic interaction.

New robotic assistants should be able to prevent worker's discomfort as well as more aggravated work-related illness. By contrast with a human-human interaction, where both parties try to maximize their own comfort, sometimes at the expense of the other, an assistive robot should always maximize the comfort of its coworker. On the other hand, humans are complex systems and the notion of human comfort is very difficult to assess. It might be linked to the notion of acceptability, which is usually divided into social acceptability: *How does the society perceive robotic systems?*, and practical acceptability: *How do people perceive the robot when interacting with it?*

Aspect of the robot is an important factor of its acceptance (Breazeal 2004). For example a spider-like robot could be perfectly fit to operate in uneven terrain but probably not to interact with humans due to its dreadful design. Kajita et al. (2009) suggested this as a reason why humanoid robotics has undergone such a remarkable expansion in the last few years. Humanoid design and especially bipedal locomotion is probably not the most optimal but it greatly

facilitates the acceptance of the robotic system by the human workers[4].

However, when working with already designed robotic systems, such as Baxter robot which was used in all our experiments and is illustrated in Figure 1.3, the only left levering approach relies on improving the acceptance of the robot motions and behaviors. Choice of actions, robot motions, proximity with the human coworker[5], might lead to uncomfortable situations. Assessing those situations is not trivial. It usually relies on questionnaire answered after the experiment, recorded videos analyzed by experts, or physiological data such as cardiac or ocular activity (Dehais et al. 2011).

Although the acceptance is strongly linked to the worker's comfort, it only covers the psychological aspects of the comfort notion. More physical constraints need also to be considered. Work related injuries cost between 13$ to 20$ billions annually to US industries and cobots researches were essentially funded to address this issue (Akella et al. 1999). Some of those injuries are linked to traumatic accidents due to physical contact with the machineries but Musculoskeletal Disorders (MSDs)[6] are the largest category of work related illness (Punnett et al. 2004). Risk-factors for MSDs are work or individual related (Putz-Anderson et al. 1997). Work-related factors include high task repetition and wrong body posture, while individual-related factor comprise poor work practice and poor rest recovery. In both cases, a robotic approach could help reduce those factors by assigning the most burdensome tasks to the robot and assist the worker to help him or her keep good working posture and practices.

Both the psychological and physical impact of the interaction were motivating the research conducted during this thesis. The first situation we consider concerns the understanding of the robot actions. In a joint collaboration scenario, it is important that the human coworker understands what the robot plans to do or achieve. Failing to achieve such understanding could lead to stressful situations or unnecessary human actions. For example, if I believe the robot next action is to reach for the screws I will probably move toward the screwdriver. However, if I realize later on that it was aiming for a different part of the assembly I will have to put back the screwdriver on the table before starting the correct complementary action. On the other hand, a fast understanding of the robot actions allows me to anticipate on my own corresponding actions and generate a faster and smoother interaction. In an industrial setup where timings is often crucial this is a non negligible improvement to consider. To achieve such understanding, two approaches could be envisaged. The robot could explicitly details its next action, using for example signals on its screen or natural language. However, a coworker that says out loud all the actions he is planning of performing would slowly become annoying and the interaction would feel unnatural[7].

[4] Another reason is that most of the tools and workstations are designed for humans. Therefore, introducing humanoid robots in industry would not require a complete reshape of the working environment.

[5] People have special concerns about their personal space either in terms of intimacy and safety. We consider proxemics–physical and psychological distancing from othersEdward Twitchell Hall 1966–to address such concerns.

[6] MSDs include a wide range of inflammatory and degenerative conditions affecting the muscles, tendons, ligaments, joints, peripheral nerves, and supporting blood vessels.

[7] Actually, most of the communication between humans is non-verbal and uses what we refer as "body language" (Harrison 1965).

Another option is to share intentions directly from the robot motions. This concept, referred as legibility in the robotic community (Dragan, K. C. Lee, et al. 2013), mainly relies on the fact that living organisms tends to modify the way they move or accomplish actions to express their intentions. If you look at a dog willing to play, the way it moves informs you about its true intentions. The same approach can be applied to robotic systems to modify the robot motions to share intentionality (Dragan, K. C. Lee, et al. 2013; Sisbot and Alami 2012). Doing so will greatly benefit the productivity, which is always a concern for industrial applications, as it would reduce both the timings to accomplish a specific task and the coworker's cognitive load.

The second situation concerns the human comfort at a more physical level. In a physical human-robot interaction, where robot and humans cooperate to accomplish a specific task, the robot can provide support actions such as holding parts or handing-over objects. Such simple actions can negatively impact the posture of the human coworker. Imagine a situation where the robot holds a piece of furniture with the human to screw but the screw hole is facing the ground. The human would have to bend or squat to look for it and screw correctly. Over the years, such a repeated wrong body posture could lead to a high physical discomfort and/or world-related illness such as Musculoskeletal disorders (MSDs). The question that arises is how can we ensure that the worker will stay away from those high risk postures. In assembly-line production works, the entire workstation is adapted to the worker to prevent those situations (Das et al. 1996). However, in small production series, adapting the entire workstation introduce a non negligible cost which make it not suitable. As human-robot collaboration specifically targets this type of industry, there seems to be a great opportunity for the robot to adapt itself to the worker specificities, like height or eventual injuries, and preferences. Even simple types of interaction, such as handing-over an object, can be stressful if not performed correctly[8]. This creates a desire to form highly effective human-robot teams that combines strengths and abilities of both the robot and its human partner (Gombolay et al. 2015; Ogorodnikova 2008), while at the same time caring for the well-being of the people working with robots. Compared to a fully automated assembly line, a robot and human worker team also offers flexibility and adaptability to changing tasks (Bley et al. 2004; Krüger et al. 2009). This last point is particularly important for small assembly lines of customized products, where changes made in the workstations are frequent.

[8] As an example, consider a left-handed worker that has to interact on an object positioned and orientated for a right-handed interaction. Either it will force him or her to adopt a wrong body posture or to use the right hand, limiting his or her efficiency on the task.

Another interesting aspect lies on the self-awareness of risk associated to a given posture. The human body has an extensive amount of degrees of freedom which allows to accomplish tasks in various postures. Obviously, not all of those postures are ergonomically safe and one might perform a task in a wrong posture without even being aware of it.

In summary, our contributions for this thesis are twofold. First, we propose a solution to improve the legibility of the robot motions. Second, we study the ergonomics of the interaction and propose a method to reduce the potential risks of MSDs, by ensuring that the robot motions and actions will not force the user to adopt a high-risk posture.

To achieve our goals, we use, for both situations, a similar approach. We define a cost function that leads toward the desired behavior and use optimization algorithms to derive the robot motions. For both cases, we find a simple cost function and validate its usage with human subjects interacting with our Baxter robot. In the following paragraphs, we summarize our main contributions with their associated scientific publications.

Learning legible motions from interaction: In the first part of the thesis we describe our work on learning legible robotic motions. In Chapter 2 we explain the notion of legibility and provide some related work on improving the legibility of robotic motion. Our work on this question has leaded to two main scientific contributions.

First, a publication in the International Conference on Intelligent Robots and Systems (IROS 2015) (Stulp, Grizou, et al. 2015) in which we introduce a cost function based on human feedback to learn legible motions while interacting with human subjects. Coupled to an optimization algorithm we apply this cost function to two different tasks and analyze the gains in term of the subjects' prediction times and potential errors. We detail the experiment protocol and results in Chapter 6 of the thesis.

Our second contributions is an analysis on the universality of the learned legible motions. By comparing two different policies for representing the robot motions, we realize that we can achieve more universal legibility that require less adaptation. This extensive study has leaded to a publication in the International Journal of Social Robotics (Busch, Grizou, et al. 2017) and is detailed in Chapter 3.

Ergonomic human-robot interaction: In the second part of the thesis we focus on more physical aspects of the human comfort and risks associated with wrong body postures. We start by introducing the notion of ergonomics in Chapter 5 and review the researches made toward a more ergonomic human-robot interaction.

In Chapter 6 we present an optimization algorithm to optimize human postures to fulfill a specific task under ergonomic constraints. From the optimized posture, we derive a robot motion that leads the real user toward the safest posture he or she should adopt. In a user study we verify that our solution leads indeed to safer posture and is perceived as more ergonomic. This work has been published in IROS 2017 (Busch, Maeda, Mollard, et al. 2017).

Reusing the same ergonomic cost function, we have extended our study to also plan high-level sequences of actions[9]. We have based our work on the Logic-Geometric Programming (LGP) framework introduced by Toussaint and Lopes (2016) that allows to solve Task and Motion Planning (TAMP) problems simultaneously. Therefore, we propose an ergonomic interaction at both the low-level motor commands and the high-level actions. The integration also offers an interesting aspect of allocating task to either the robot or human agent based on ergonomic considerations. We provide details of the approach in Chapter 7 along with the results of the conducted user study. This work has been accepted for publication to the International Conference on Robotics and Automation (ICRA 2018) (Busch, Toussaint, et al. 2018).

On a last note, the images and pictures, used to illustrate the discussions, are either extracted from our publications, made specifically for the manuscript or under creative comment licensing.

[9] We oppose high-level actions to low level motor commands. A high-level action often has a semantic meaning such as "grasps the screwdriver".

2
Introduction to Legible Robotic Motions

In everyday life, people interact with each other using verbal communication or by means other than explicit natural language statements. Our body gestures and choice of actions communicate a lot our intentions. In 1974, Duke reported a story, attributed to Ann Landers[1]. She was responding to a query from a girl on how was the best way to meet a shy boy who rode on the same bus with her each day.

> Landers suggested to the girl that the important thing was to get the man to talk first because that was really the way he wanted it. To accomplish this, Landers recommended that the girl get off the bus with a heavy packages such as several dozen bricks well wrapped. Presumably the hero would be unable to avoid seeing a lady in distress and would immediately appear at the lady's side to offer assistance.
>
> If he didn't, Landers urged the girl to drop the package or twist her ankle; the male could never overlook this opportunity for demonstrating his gallantry, and he would be certain to carry both the girl and the package to their destination; once there, the girl could offer him refreshments, thank him and settle down for a get-acquainted chat.

In this story, most of the communication is non-verbal. It has been estimated that only 35% of the communication among people happens verbally (Harrison 1965). An important thing to note is that, here, the girl's real intention is hidden and her choice of actions do not clearly signal the goal she wish to achieve. In contrast with this anecdote, our body gestures and choice of actions might also reflect upon our true intentions.

When doing so, it is important to select actions that will clearly convey this information to our interlocutor. The story of John Hinckley Jr. is an example of a poor choice of actions that failed this purpose. In 1981, he attempted to assassinate U.S. President Ronald Reagan with the mean of impressing the actress Jody Foster for whom he devoted a true obsession. It might be unclear how this specific action could achieve Hinckley's real goal[2]. Therefore, a wrong choice of actions might lead to the complete opposite of what you wished to accomplish.

[1] Ann Landers is a pen name created by Chicago Sun-Times advice columnist Ruth Crowley in 1943 and taken over by Esther Pauline "Eppie" Lederer in 1955. For 56 years, the "Ask Ann Landers" syndicated advice column was a regular feature in many newspapers across North America.

[2] Some extra information might help understanding his choice. During her career, Jodie Foster has played a role in the 1976 film Taxi Driver, in which disturbed protagonist Travis Bickle (Robert De Niro) plots to assassinate a presidential candidate. With that information in mind, we could see why Hinckley chose this action. A common knowledge might also help the interaction.

However, not only does our actions reflect upon our intentions but also the very way we accomplish them. Humans, and other animals, are exquisitely attuned to recognize goals and intentions from the motion of living organisms. Watching a dog making a circling motion leads us toward its intentionality to play. It is not even necessary to have an organism-shaped body to convey intention in the motion, as Heider et al. (1944) showed[3]. Their animations of moving triangles and circles elicited strong subjective impressions of goals and intentions, e.g., of one triangle trying to catch another.

[3] Several videos of this animation of moving mathematical shapes can be found on Youtube https://www.youtube.com/watch?v=sx7lBzHH7c8.

This non-verbal communication is also central in the context of joint task between humans to coordinate their actions (Sartori et al. 2011). If you move an open bottle in my direction, I will place my glass in an appropriate position for you to pour. If you reach for the screwdriver, I will lift the shelf to hold it in place for you. By monitoring the actions of others and inferring their intentions, a human can predict and preemptively initiate the appropriate complementary actions without the need for verbal communication (Sartori et al. 2011; Sebanz et al. 2006; Timmermans et al. 2012). Furthermore, it has been shown that humans unconsciously change their behavior, for instance the speed of task execution, to improve coordination (Vesper et al. 2011).

If we are to develop social robots that interact with people on a daily basis, it is important to understand this non-verbal communication mechanism, not only to infer people intentions, but also to communicate in a similar manner. The actions performed by the robot, and the way to achieve them, should reflect upon its goal and convey sufficient information.

Recently, Knepper et al. (2017) proposed a framework to formally model this implicit communication in human-robot interaction. This model relies on the notion of surprisal (Hohwy 2013)[4]. They state that actions or motions inducing a high surprisal is a mean to convey information. In other terms, if an action $\hat{a}$ seems improbable to an observer, it is chosen over a more probable action a^* because it has a specific meaning m^* that the observer should understand. This, obviously, requires some basic concepts to be respected:

[4] Any communicative action will be perceived by an observer with a certain level of surprisal, which is an encoding of how probable the observer believes the action to be, given the context. The higher an observer's surprisal, the more improbable the observer believes the action to be in the given context.

1. There exist multiple actions to achieve a specific goal ($a^* \neq \hat{a}$)
2. One action should be more probable than the others given the context
3. The meaning m^* should be an easy explanation of the observation of $\hat{a}$
4. There is no other meaning $\hat{m}$ that is an explanation of $\hat{a}$

Hinckley's story is a good example that violates the third rule. His true intention could not be easily derived from his choice of actions. The fourth rule could be culturally dependent and requires to know exactly how our actions are perceived by the observer. One example of action that might fall under this rule is nodding the head for acceptance[5]. Although, in most countries, it might be correctly understood, there is a few exceptions where the meaning is swapped and actually means a refusal[6]. The fourth rule is particularly important as it highlights the fact that an action that successfully convey some information to one person might fail to do so with another one. As we develop robotic behaviors to be very informative, this might question the universality of our findings. However, this could be overcome if both parties agree, beforehand, on the meaning of a specific action.

[5] A nod of the head is a gesture in which the head is tilted in alternating up and down arcs along the sagittal plane.

[6] In Greece, Iran, Lebanon, Syria, Turkey, Bulgaria, Albania, and Sicily a single nod of the head up (not down) indicates a "no".

We can apply this formalism, in human-robot interaction. First, the choice of individual or sequences of robot actions can highlight a greater goal and trigger a human response. As an example, if the robot moves to grasp the screwdriver, the human coworker should understand that he or she needs to reach for the screws. Second, the robot can signal its intention directly from its motion[7]. In the second case, the human would go for the screws before the robot has even reached the screwdriver. Achieving such a comprehension between the robot and the human would induce a gain in terms of productivity as it would reduce the time to accomplish a specific task.

[7] Robot motion might refer to the movement of its mobile base as it has been studied with mobile robots moving in a human environment (Sisbot and Alami 2012) or to the motion of its arms when working with fixed base robots like assembly robots (Dragan, K. C. Lee, et al. 2013). In this work, we only focused on the latest scenario.

Conveying information through motion is a concept referred as legibility, and efforts have been made to improve the legibility of the robot behaviors (Dragan and Siddhartha S. Srinivasa 2013; Sisbot and Alami 2012)[8]. The main idea is that one can infer the robot intention by only observing its motion. Let us take the following example:

[8] Some researchers prefer to use the term "readability" rather than "legibility" (Takayama, Dooley, et al. 2011).

> John Doe is on his bike, arriving at a zebra crossing. A car is coming quite fast and John wonders if it will stop to let him cross the road. The car does not seem to be slowing down, John set foot on the ground. But finally the car driver brakes and flashes the headlights, signaling John to pass.

In this situation, John could only rely on the car motion and speed to infer if it would stop on time to let him cross the road. However, the car driver failed to convey this information and had to use a more informative signals, a flashing of headlights. Should the car driver break earlier, John would have continue on its path without the need of setting foot on the ground.

The same situation arises in human-robot interaction. We could have a special signal for each of the robot actions, which akin to the flashing of headlights in the car bike situation. Or we could improve the robot legibility to smooth the interaction and trigger an appropriate human response.

In human-human interaction, this concept has been extensively studied and a key component that helps to facilitate social interaction is believed to be the action observation network. Motor processes underlie the execution of actions as well as the understanding of other's people intended actions (Decety et al. 2006; Gallese 2001; Gallese and Goldman 1998). This process relies on some specific neurons in our brain called *mirror neurons*[9] that allow us to understand a person intention by projecting his or her motor commands onto ourselves.

[9] A mirror neuron is a neuron that fires both when an animal acts and when the animal observes the same action performed by another.

Could the same process be used to recognize robot intentions from its motion? Cross et al. (2012) shows that it could, in fact, go beyond the simple imitation mechanism and it is also activated by the observation of non-human agents, specifically robotic systems. Moreover, if a robot mimic well a human movement, our anthropomorphism capacity might simply take over and facilitate the recognition process.

The question is then, *how can we improve the legibility of robot motions?* One way to achieve this can be to imitate the human motion in the same task context. The minimum jerk model (Flash et al. 1985) makes the assumption that human hand motion can be mathematically retrieved, by minimizing the jerk in Cartesian space, during a grasping task. On an industrial robot, however, trajectories generally follow a trapezoidal joint velocity profile (Craig 2005). Research has shown that predicting this type of motion is harder than a minimum jerk profile (Glasauer et al. 2010)[10].

[10] In the light of human psychology and neuroscience, those results suggest that, because trapezoidal joint velocity profile are "non-natural", our action observation network might be unable to map it onto ourself, which make the recognition process harder.

For specific tasks, it is possible to manually define motion that convey the desired intention. This can be made for different applications. For instance to facilitate handing over an object (Alami et al. 2006; Cakmak, Siddhartha S. Srinivasa, et al. 2011; M. K. Lee et al. 2011; Mainprice et al. 2010; K. W. Strabala et al. 2013; K. Strabala et al. 2012), or to coordinate robot soccer players (Pagello et al. 1999; Stulp, Isik, et al. 2006). This involves understanding and modeling how humans interpret actions, and implementing controllers based on this knowledge. Explicit task-specific encoding of intention prediction has also been used to coordinate robot soccer players (Pagello et al. 1999; Stulp, Isik, et al. 2006). The concept of legibility has also been studied in the context of safe navigation in the presence of humans (Lichtenthäler et al. 2012; Sisbot and Alami 2012).

Another approach, taken by Dragan and Siddhartha S. Srinivasa (2013) is to provide a *general-purpose* definition of legibility: how probable is a goal, given a partially observed trajectory? Higher legibility implies earlier divergence of probabilities for different goals. They also clearly formalize the difference to predictability: what is the most probable trajectory, given knowledge of the goal?

Imagine a reaching task scenario where two objects are closed to each other. You have to guess which object the robot is going to pick only by looking at the robot arm motion. As illustrated in Figure 2.1, if you knew that the robot would pick the blue object, you would expect it to perform a straight line toward it. This is the most probable action $a*$ according to Knepper et al.'s formalism. Dragan, K. C. Lee, et al. refer to it as predictable motion.

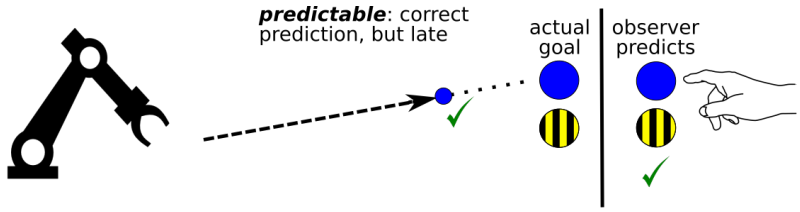

Figure 2.1: Predictable motion in a reaching task scenario. This is the most probable motion and the least informative.

The main drawback of this predictable motion is its lack of information. With two possible targets, very close to each other, you would need to observe a large snippet of the trajectory before being able to guess the robot target.

To convey more information, the robot trajectory would then need to diverge from the straight line. Now, observe the trajectory illustrated in Figure 2.2. This type of motion, an $\hat{a}$ action, is very informative. It allows you to correctly infer the robot target, even if you see only a short snippet of the trajectory.

Figure 2.2: Legible motion in a reaching task scenario. This is the most informative type of motion.

On the other hand, the robot could also diverge on the other side of the straight line. However, this type of motion depicted in Figure 2.3 would mislead the observer on the real target. It is breaking the third rule of Knepper et al.'s formalism. In fact, the yellow hashed target is an easier explanation of the observation of this trajectory. It is referred as deceptive motion (Dragan, Holladay, et al. 2014).

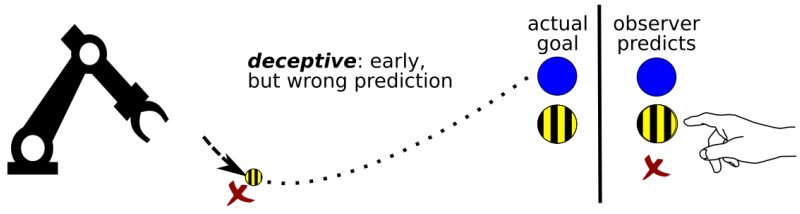

Figure 2.3: Deceptive motion in a reaching task scenario. This motion mislead the observer on its prediction.

One question is, *how can we generate those legible motions?*. Dragan and Siddhartha S. Srinivasa (2013) propose a model-based method to optimize the robot trajectories for legibility. Although they define legibility and predictability as general ideas, they are implemented as cost functions that might not apply to all task contexts. It is a non-trivial task to adapt this cost function to novel task contexts, and especially to different (classes of) users.

The fundamental problem of trying to model such a complex concept as legibility is being unable to capture its full scope. Each person might have his or her own idea on what is a legible motion, although probably unconsciously. The whole concept of legibility might be also culturally dependent, as the concept of nodding for acceptance presented earlier. And finally, it might be also task dependent. Therefore, accounting for those specificities is almost impossible with a model-based approach.

The question becomes then, *can a robot learn to generate legible motions, adapted to its human coworker?* This question is the motivation behind this first part of the thesis. Our goal is to learn how to generate legible motion in a model-free manner, to adapt the robot behavior to the observer's preferences and his or her perception of legibility.

Rather than defining legibility as an explicit property to be optimized, we investigate legibility as an emergent adaptive property of robots who are rewarded for efficiently cooperating with humans. Our approach is based on model-free optimization, where the robot iteratively improves its legibility through trial-and-error interaction with a human. By not providing a model we need to find other means to reward the robot which would lead to a more legible behavior. Dragan, K. C. Lee, et al. (2013) defines legible motions as "motion that enables an observer to quickly and confidently infer [a] correct goal G". Therefore, a simple cost function that captures the time an observer takes to infer a goal G and his or her eventual prediction errors seems to be a good reward candidate. It would force the emergence of legible behaviors without the need of providing a probabilistic definition of the whole concept of legibility. Moreover, this approach has the advantage that no assumptions about the task or the human must be made, and the robot automatically adapts its legibility to the user preferences during the interaction.

Since our approach does not require a model, it is applicable to different tasks without modification. However, it does require a training phase to learn to generate legible behavior, and the resulting behavior generalizes to different tasks. A novel task thus requires learning a new behavior. In contrast, previous model-based methods (Alami et al. 2006; Cakmak, Siddhartha S. Srinivasa, et al. 2011; M. K. Lee et al. 2011; Mainprice et al. 2010; K. W. Strabala et al. 2013; K. Strabala et al. 2012) are able to generate legible behavior on-the-fly, but require task-specific models of legibility. A novel task thus requires the design of a novel model by an expert.

Our approach is thus well suited for scenarios where not all tasks are known in advance, and where similar tasks are executed many times. In assembly lines where humans and cobots work together for instance, the resulting behavior is used thousands of times. The number of trials required to learn the behavior (<100) may thus well be worth the investment, and could also be performed on-the-job.

To conclude, the difference between model-based and model-free approaches can be summarized using the *bottom-up* and *top-down* designs[11]. Model-based approach is a *top-down* design as you start from a general concept of legibility and refines it if necessary. For example, if you model a robot to nod for acceptance, you will probably start by the most common design, i.e. an up and down head motion. If you wish to later deploy your robot in Greece or Lebanon you will have to add a specific condition as this head motion could be seen as a "no". The more exceptions and specificities you have to adapt to, the more complicated the model grows. Moreover, every refining require programing skills to rewrite a new adapted model. Conversely, our model-free approach is a *bottom-up* design. You start from a specific concept of legibility, adapted to a single user, and verify that what you have learned holds for other users. When you wish to adapt your learned behavior to a new user or specificities you use a trial and error approach, re-optimizing in the process. No programming is, therefore, required in this situation.

[11] *Top-down* and *bottom-up* are both strategies of information processing and knowledge ordering, used in a variety of fields including software, humanistic and scientific theories

One important question that we need to verify is whether the robot learns to generate universally legible behavior, or rather idiosyncratic behavior that the observer learns to interpret. The difference between the two is illustrated in Figure 2.4. To understand this difference we can come back to our biker example:

> John is riding his bike on a very busy road. At the next intersection, he is going to cross the road to turn left. To signal the cars behind him he raises his left arm before turning.

This is an example of "universal" and legible behavior. Even for cultures in which cycling is not widespread, an arm spread out to the left is likely to convey the intention that the cyclist will make a left turn.

> John is in the leading pack during a bicycle race. He wants to signal his team that he is going to attack[12]. He drops his left arm straight down and joins his thumb with his index.

[12] An attack is the action of quickly accelerating while riding in a pack, or in smaller numbers, with a view to create a gap between yourself and other riders.

And this is an idiosyncratic behavior. It is meant only for John's teammate to be understood. For any other observer, this is breaking the third rule of Knepper et al.'s formalism as there is no universal meaning that can easily explain this signal. Coming back to our robotic situation, the robot can learn to generate idiosyncratic behavior such as the loop in Figure 2.4 to signal which object it is going to grasp.

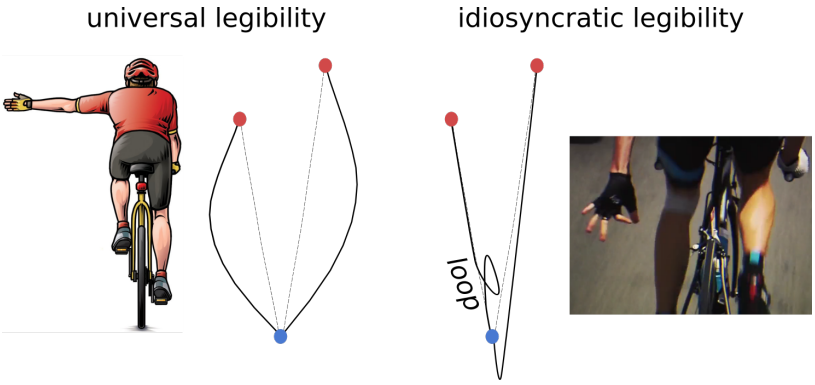

Figure 2.4: Distinction between universal and idiosyncratic legibility. The left graph with trajectories has been adapted from Dragan and Siddhartha S. Srinivasa (2013) work.

In the next chapters we detail our works on learning legible motions from interaction. Chapter 3 introduces the experimental protocol used to learn legible motions and the user study we have conducted on two different tasks[13]. Following our initial work, we have studied the concept of *universal* legibility and how we can tend to achieve it. Details are provided in Chapter 4.

[13] The text of Chapter 3 is mainly extracted and adapted from our conference paper (Stulp, Grizou, et al. 2015) while most of the text of Chapter 4 comes from our journal article (Busch, Grizou, et al. 2017)

3
Model-Free Optimization for Legibility

Contents

THE KEY IDEA of our approach is that legibility of robot behavior does not needs to be defined and optimized explicitly, but that it arises automatically if joint task execution is penalized for not being efficient. In particular, this behavior should arise from direct interaction with users, so as to tailor the legibility to specific user preferences.

As we need to optimize a cost function, but do not have models of individual users, we formulate the problem as a model-free optimization. We describe the generic optimization algorithm and policy representation used in our contribution.

Policy Improvement through Black-Box optimization

Policy improvement is a form of model-free optimization technique, where the parameters ϑ of a parameterized policy π_ϑ are optimized through trial-and-error interaction with the environment. The optimization algorithm we use is PI^{BB}, short for "Policy Improvement through Black-Box optimization" (Stulp and Sigaud 2012). It optimizes the parameters ϑ with a two-step iterative procedure. The first step is to locally *explore* the policy parameter space by sampling K parameter vectors ϑ_k from the Gaussian distribution $\mathcal{N}(\vartheta, \mathbf{\Sigma})$, to execute the policy with each ϑ_k, and to determine the cost J_k of each execution. This exploration step is visualized in Figure 3.1, where $\mathcal{N}(\vartheta, \mathbf{\Sigma})$ is represented as the large (blue) circle, and the samples $J_{k=1...10}$ are small (blue) dots.

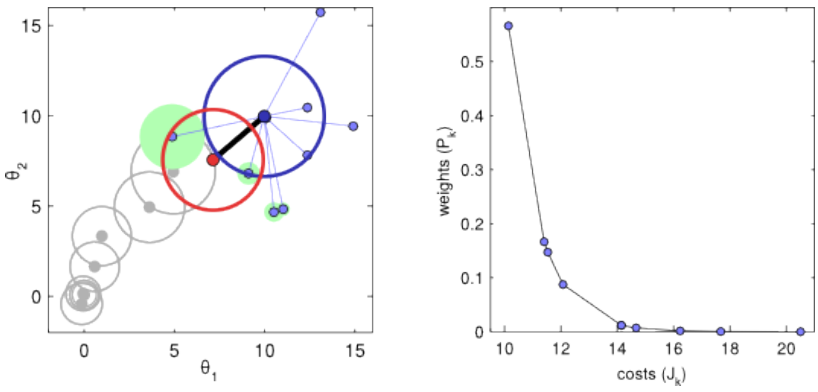

Figure 3.1: Illustration of the PI^{BB} algorithm on a simple cost function $J(\vartheta) = \|\vartheta\|$ (without policies). Left: iterative updating of the parameters, where the exploratory samples for the first iteration are shown. Right: mapping the costs J_k to weights P_k for the first iteration. The algorithmic parameters are $K = 10$, $\lambda = 0.7$.

The second step is to *update* the policy parameters ϑ. Here, the costs J_k are converted into weights P_k with

$$P_k = e^{\left(\frac{-h(J_k - \min(\mathbf{J}))}{\max(\mathbf{J}) - \min(\mathbf{J})}\right)} \tag{3.1}$$

where low-cost samples have higher weights. For the samples in Figure 3.1, this mapping is visualized (to the right). The weights are also represented in the left figure as filled (green) circles, where a larger circle implies a higher weights. The parameters ϑ are then updated with reward-weighted averaging

$$\vartheta \leftarrow \sum_{k=1}^{K} P_k \vartheta_k \tag{3.2}$$

Furthermore, exploration is decreased after each iteration $\mathbf{\Sigma} \leftarrow \lambda \mathbf{\Sigma}$ with a decay factor $0 < \lambda \leq 1$. The updated policy and exploration parameters (red circle in Figure 3.1) are then used for the next exploration/update step in the iteration. Despite its simplicity, PI^{BB} is able to learn robot skills efficiently and robustly (Stulp, Herlant, et al. 2014). Alternatively, algorithms such as PI^2, PoWeR, NES, PGPE, or CMA-ES could be used[1].

[1] see Kober et al. (2011) and Stulp and Sigaud (2012) for an overview and comparisons

Policy Representation

The policy π_ϑ itself is implemented as a Dynamical Movement Primitive (DMP) (Ijspeert, Nakanishi, and Schaal 2002). The authors define a control policy by a set of nonlinear differential equations with a well-defined attractor dynamics. For a single degree of freedom denoted by y, which, in our case, is one of the internal joint angles, the following system of linear differential equations with constant coefficients has been proposed as a basis for motion specification (Schaal et al. 2007),

$$\tau \dot{z} = \alpha_z(\beta_z(g - y) - z) - f \tag{3.3}$$

$$\tau \dot{y} = z. \tag{3.4}$$

where τ is a time constant and α_z and β_z are positive constants. Provided that the parameters α_z and β_z are selected appropriately, e.g., $\alpha_z = 4\beta_z$ and the forcing term $f = 0$, this system has a unique attractor point at $(y, z) = (g, 0)$.

Choosing the forcing function f to be phasic, i.e active in a finite time window, leads to a point attractive system. Therefore, f can be defined as

$$f(t) = \frac{\sum_{i=1}^{N} \Psi_i(t)\omega_i}{\sum_{i=1}^{N} \Psi_i(t)}, \tag{3.5}$$

where Ψ_i are fixed basis function and ω_i are adjustable weights[2]. However, the explicit time dependency creates a non-autonomous dynamical system, or more precisely a linear time-variant dynamical system. Therefore, Ijspeert, Nakanishi, Hoffmann, et al. (2013) introduce a time replacement x following first order linear dynamics

$$\tau \dot{x} = -\alpha_x x, \tag{3.6}$$

with α_x a constant. Starting from some arbitrarily chosen initial state x_0 such as $x_0 = 1$ the state x converges monotonically to zero. x can thus be conceived of as a phase variable, where $x = 1$ would indicate the start of the time evolution and x close to zero means that the goal g has essentially been achieved. Equation (3.6) is called the *canonical system*.

Following this, f can be rewritten as a function of the *canonical system*

$$f(x) = \frac{\sum_{i=1}^{N} \Psi_i(x)\omega_i}{\Psi_i(x)} x(g - y_0), \tag{3.7}$$

where $\Psi_i(x)$ are N Gaussian basis functions defined as

$$\Psi_i(x) = exp\Big(-\frac{1}{2\sigma_i^2}(x - c_i)^2 \Big) \tag{3.8}$$

with σ_i and c_i are constants that determine respectively the width and centers of the basis functions. y_0 is the initial state at time $t = 0$. With this formulation, the system is designed to have a unique equilibrium point at $(z, y, x) = (0, g, 0)$ with y evolving toward the goal g from any initial conditions.

Therefore DMPs are very convenient for our experiments, as they ensure convergence towards a goal, whilst allowing the trajectory towards this goal to be adapted by changing the parameters ω_i of the forcing term f.

In our experimental setup, we represent f with 3 basis functions, i.e. $N = 3$ and DMPs are fist initialized by recording straight lines toward the goal. Then, the weights ω_i of the basis functions are optimized using the black-box optimization introduced in Section 3.1.

[2] According to Ijspeert, Nakanishi, Hoffmann, et al. (2013) Representing arbitrary nonlinear function as a normalized linear combination of basis functions has been a well-established methodology in machine learning (Bishop 2006) and also has similarities with the idea of population coding in models of computational neuroscience (Dayan et al. 2001)

Experimental setup

In the joint human-robot task, depicted in Figure 3.2, the robot reaches for and presses one of two buttons. Subjects are given two goals:

- **Efficiency**: press the same button as you think the robot will, as quickly as possible,
- **Robustness**: avoid making mistakes, i.e. pressing a different button from the one the robot will.

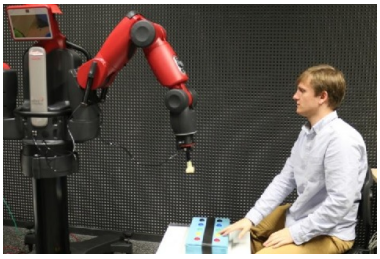

Figure 3.2: Button pressing experiment set-up with the Baxter robot, human subject, and the two rows of buttons that they will press. The two possible targets corresponds to the "red" and "yellow" button on the box, the two buttons on the left side of the subject.

The protocol of an experiment is as follows. The experiment starts with a *habituation phase* of 32 trials where the robot performs always the same trajectory for the same button. This phase allows the subject to get used to the robotic motions, and practice the prediction and button pressing. It also allows to validate that the improvement in the subject's prediction is not only due to them learning the robot's motion. Further improvement after that habituation phase will then only be explained by the robot being more legible. Preliminary results indicate that 32 trials are sufficient for habituation.

After habituation, we start the *optimization phase* of 96 trials with the optimization algorithm presented in Section 3.1. The two policies that generate trajectories for the two different buttons are optimized in two independent processes.

Task Representation: Cost Function

The cost function that the robot optimizes during the 96 trials after the habituation phase consists of three components:

$$J = \underbrace{T_{\text{robot}} + T_{\text{subject}}}_{\text{Efficiency}} + \underbrace{\gamma \delta_{\text{buttons}}}_{\text{Robustness}} + \underbrace{\alpha \left| \dddot{\mathbf{q}}_{1\ldots N, 1\ldots T} \right|}_{\text{Energy}} \tag{3.9}$$

Efficiency: The time between the onset of the robot's movement (t_0) and the pushing of the button by the human (T_{subject}) and the robot (T_{robot}).

Robustness: Whether the subject pressed the correct button (δ_{buttons} =0) or not (δ_{buttons} =1). γ is an arbitrary high cost, it was set to 20 in this experiment, expressing that a failure is equivalent to a penalty of 20*s* in terms of efficiency.

Energy: The sum over the jerk, i.e. the third derivative of the joint positions ($\dddot{\mathbf{q}}_{t_i}$). at each time step i in the trajectory. The time step Δ_t used to calculate the derivatives was arbitrary set to 0.2. The scaling factor α is chosen such that the cost of the jerk is about 1/20 of the total cost in the initial trajectories.

The joint task completion time depends mainly on how fast the human is able to predict the intention of the robot (proximate cause). But we use the total time because: 1) the ultimate motivation behind our research is to make human-robot interaction more efficient. 2) our set-up easily allows us to determine the button pressing times, but not the exact time at which the human predicts the robot's intention.

User Study

We perform two experiments with human subjects. In Experiment A, the task consists in pressing two corresponding buttons. First the robot decides on a button to press and then the subject needs to press the corresponding button as soon as possible.

In Experiment B, the task consists in selecting a button based on the location a robot will deposit a bottle. First the robot decides on a target location to put a bottle after grasping it, the subject needs to press the button corresponding to the location of the target location before the robot grasps the bottle.

In both experiments the time taken by the user to press the button is used as the main cost of the collaborative task. In practice, if the motion of the robot is more informative then the subject will be able to predict sooner the target of the robot.

Both experiments follow the protocol described Section 3.3. Subjects for all the experiments in this study are INRIA staff, PhD students in computer science, and under-grad students of cognitive science[3].

[3] A video of both experiments is available on our Vimeo Channel https://vimeo.com/237417895

Experiment A: Joint button pressing

In this task, the robot reaches for and presses one of two buttons. The subject is instructed to press a button of the same color as early as possible, whilst avoiding mistakes, i.e. pressing another button than the robot intended to. We used the set-up presented in Figure 3.2, where the subject sits on the side of the robot. An illustration of the setup is also depicted in Figure 3.3.

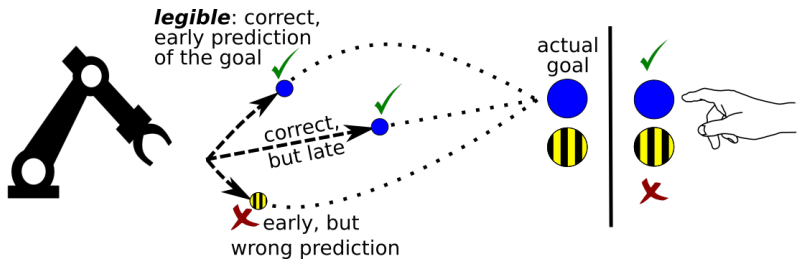

Figure 3.3: Illustration of the button pressing experiment or reaching experiment. The robot chooses a button to press and the subject has to press the button of the same color on his or her side of the box.

The two initial policies, i.e. one for each button, have been recorded through programming by demonstration, as described in Ijspeert, Nakanishi, Hoffmann, et al. (2013). The starting position is the same and the path to each button is straight. Therefore, the trajectories are hard to differentiate in the first part of the movement.

Results For illustration purposes, the top graph in Figure 3.4 shows an example experiment for one subject, visualizing both the values of the time it takes the subject to push the button (T_{subject}) and whether the same buttons are pushed. The transition from the habituation to the optimization phases is depicted as a dashed line.

The main results of the reaching task experiment are summarized in the two lower graphs in Figure 3.4, which highlight statistics at important transitions during learning: the start (trial 1 to 8), the last trial of the habituation phase (25-32), and the first (33-40), intermediate (81-88) and final (121-128) block of trials during the optimization phase. We also measure the trajectory completion at prediction time, i.e the relative amount of trajectory (timewise) observed by the subject when it presses the button. This measure is calculated using the formula $100(1 - \frac{T_{robot} - T_{subject}}{T_{robot}})$. The complete results are shown in Figure 3.5.

The box plots show the average value of T_{subject} over all 9 subjects and over blocks of 8 trials. To allow comparison between subjects without introducing variance due to the natural overall differences in their button pressing time T_{subject}, we normalized the results of each subject by their intrinsic time after habituation, which is computed as the average of the last 8 values of T_{subject} in the habituation phase. Thus, the normalized mean over the last 8 trials of the habituation phase is 100 for each subject by definition.

Finally, the bottom graph in Figure 3.4 shows the number of prediction errors per block of 8, averaged over all subjects.

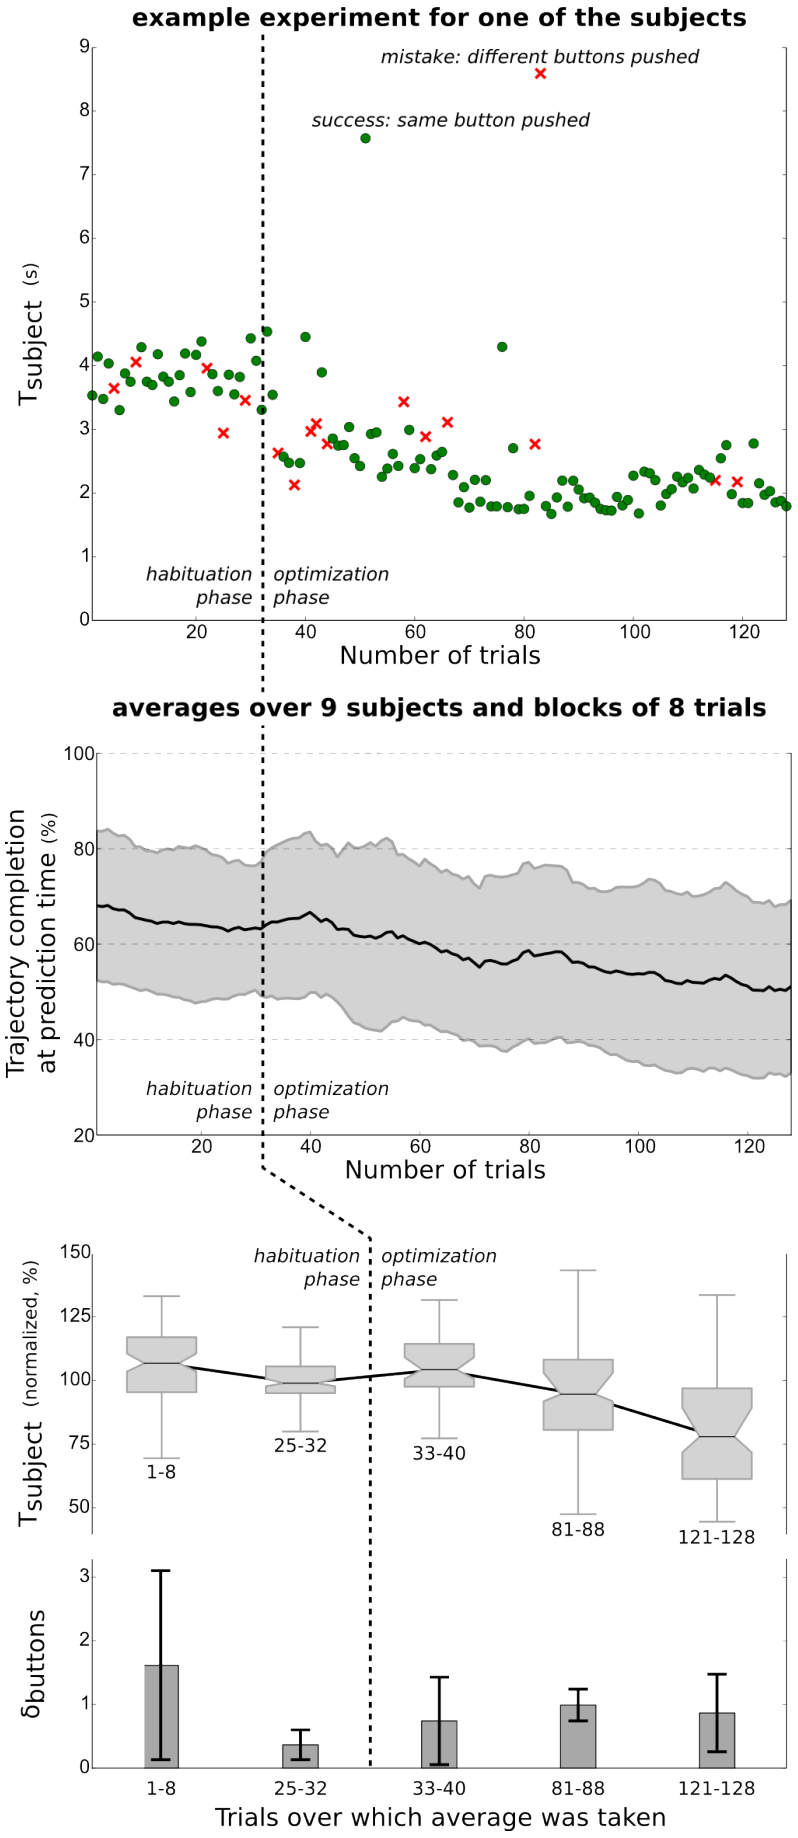

Figure 3.4: Results of the reaching experiment (1/2).
Top) Example experiment for one subject, where T_{subject} is plotted against the number of trials. Successful and failed trials are depicted as circles and crosses respectively.
Middle) Average over all 9 subjects ($\mu \pm \sigma$) of the trajectory completion at prediction time, i.e, the relative amount of trajectory (timewise) observed by the subject when it presses the button. This value is calculated using the formula $100(1 - \frac{T_{robot} - T_{subject}}{T_{robot}})$.
Bottom) Normalized T_{subject} (see main text for normalization method), averaged over all 9 subjects and blocks of 8 trials; average number of failures, i.e. when different buttons were pushed, averaged over all 9 subjects and blocks of 8 trials. The two graphs show the values at certain key frames during learning.

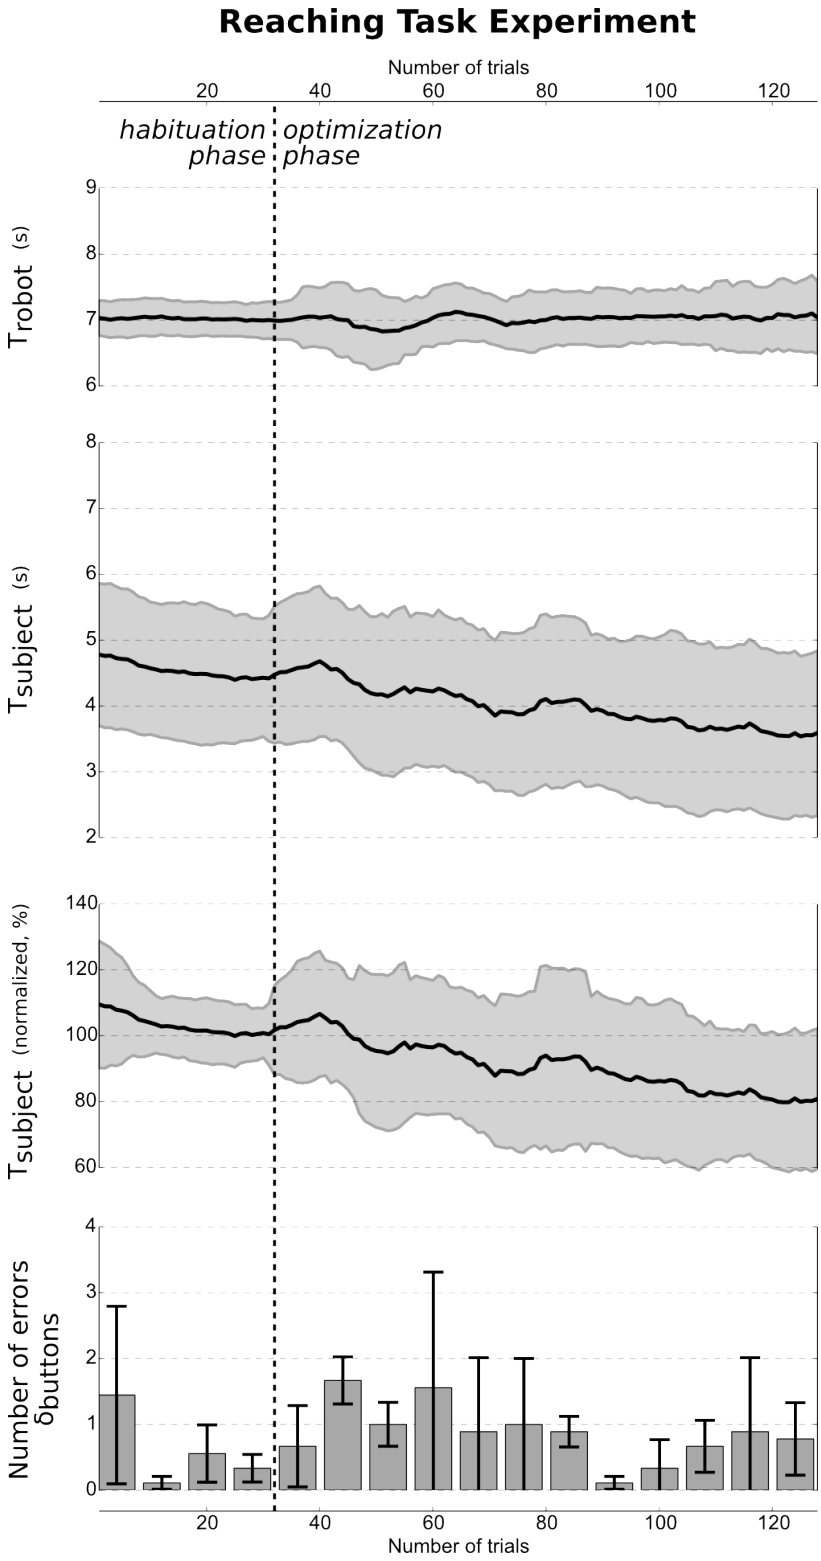

Figure 3.5: Results of the reaching experiment (2/2). The start of the optimization phase is indicated by the vertical dashed line.
Top row) Average ($\mu \pm \sigma$) of the robot button pushing time (T_{robot}). It varies little for the DMP policy (left) and even less for the viapoint policy (right). For the latter this is to be expected, as the duration of pressing the button is not dependent on the parameters of the policy in which exploration and optimization takes place.
Second row) Average ($\mu \pm \sigma$) of the subject button pushing time T_{subject}, over all 9 subjects. Variance is quite high because some subjects push quickly overall, whereas others are more careful.
Third row) Again the average subject button time, but this time normalized with respect to the average value of T_{subject} during the last 8 trials of the habituation for each subject. This reduces the variance caused by the overall differences between subjects.
Bottom row) Number of times the incorrect button was pushed, averaged over blocks of 8 trials and all 9 subjects.

Discussion The main conclusion we derive from Figure 3.4 and Figure 3.5 is that optimizing the robot's motion leads to a substantial (20%) and significant ($p = 5e^{-8}$, Wilcoxon signed-rank test) drop in T_{subject}, i.e. the time it takes for the user to press the button, between the end of the habituation phase (25-32) and the end of the optimization (121-128). As T_{robot} is consistent throughout the experiment, this drop in T_{subject} also induces a drop in the trajectory completion at prediction time (from 70% to 50%). This improved efficiency is not merely due to subjects simply guessing a button, because the number of mistakes does not increase over time ($p = 0.26$, Wilcoxon signed-rank test between end of habituation and end of optimization).

There is also a relatively small but significant ($p = 0.001$) decrease of the prediction time during the habituation phase, which indicates that the differences in the initial trajectories before optimization already enable the subject to predict the robot's intention. The fact that T_{subject} is further improved by 20% during the optimization shows that the optimized trajectories are more easily distinguishable, i.e. legible, than the initial trajectories.

After the habituation phase, subject's performance get lowered (higher prediction time and higher number of mispredictions). This effect arises from the variance of the parameters. As we do not model legibility, the robot can perform deceptive motions (Dragan, Holladay, et al. 2014) while exploring the parameter space of the trajectories. This type of motion, which leads to higher cost under our cost function in 3.9, will slowly disappear after some iterations. Only the most legible trajectories remain, as confirmed by the drop in prediction time and the low misprediction rate.

In summary, the optimization algorithm effectively improves human-robot collaboration by producing motions that are easier to predict by the subject. By penalizing errors and the joint robot/human execution time, the robot learns policies that enable the human to distinguish the robot's intentions earlier without more errors.

Experiment B: Pick-and-place

In the second task, the robot reaches for and grasps a single object. After grasping the object, the robot has to place it inside one of two boxes located on each side (see Figure 3.6). The 7 subjects are instructed to press the button corresponding to the aimed box as early as possible while avoiding mistakes.

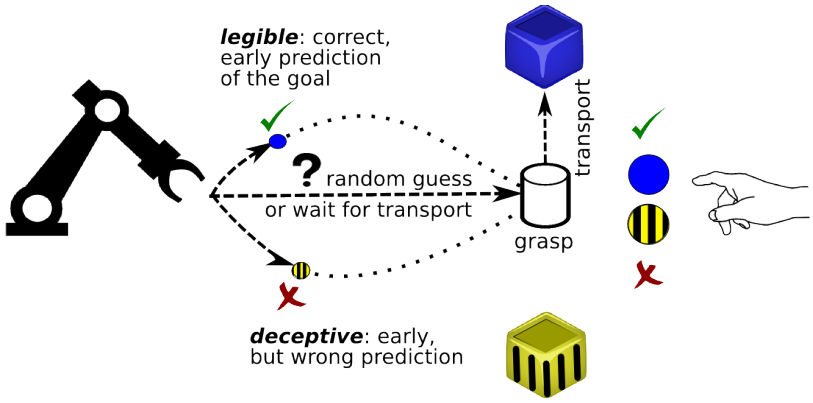

Figure 3.6: Illustration of the pick-and-place experiment. The robot starts with an initial trajectory toward the object that is identical for both placing location. Therefore, users' guesses during the habituation phase is purely random.

To improve the task, the subject must predict which side the object will be moved to and press the corresponding button. This experiment differs from the joint button task because the motions to approach the object are initially identical for both outcomes, guessing before that point results in 50% chances of success. As the robot aims at eliciting an early response from the subject, differentiating the approach motion is necessary to improve joint coordination. Thus, our hypothesis is that at the end of the optimization phase, the subject should be able to predict the robot's intention before it even grasps the object.

Results As in the previous experiment, Figure 3.7 summarizes the results by showing the results for one example subject (top), the average prediction times over all 7 subjects (middle), and the number of errors (bottom).

The results for the example subject show that during the habituation phase, the subject waits for the robot to actually start moving the object (approx. 1.4*s*) towards the box to predict the ultimate goal. Because the initial trajectories for each box are identical during habituation, guessing before that point results in 50% chances of mistakes. However when the optimization starts, the two trajectories start distinguishing themselves and there is co-adaptation between the robot and the human on the intent of each trajectory. After some trials and errors, the prediction time of the human drops to a consistent 1.5*s*, meaning the trajectories for the left or right box differs in their early parts. Despite such early decision, this subject has close to 0% of prediction errors. In comparison to the previous task, the improvements due to the optimization are more pronounced.

The prediction times have a bimodal distribution. Subjects either wait until the robot starts transporting the object, or make a prediction early on during the reaching phase. Rather than averaging over this bimodal distribution data, we compute the ratio of early/late prediction times, averaged over blocks of 8 trials and all 7 subjects. The threshold is the average over all the prediction times for one subject. In the middle graph, we see that the late prediction ratio decreases from [0.8-1.0] during the habituation phase to <0.05 at the end of the optimization phase. Furthermore, this early prediction is not accompanied by an increase in the number of errors, as the bottom graph shows. Thus, the robot learned behaviors that enabled subjects to predict the correct box before the grasp was even performed.

Because the robot's reaching behavior is the same for both boxes during the habituation phase, subjects can expect a 50% error rate when pressing the button when the robot is still reaching. The increase in the ratio of late predictions and the decrease of errors during the habituation phase may indicate that subjects learn that early guesses lead to errors, and should thus be avoided.

Discussion During habituation, subjects must wait for the robot to transport the object before being able to robustly predict its intention. After optimization however, all subjects are able to predict the intention of the robot very early on during the reaching phase, when the robot has not yet grasped the object. Experiment B thus confirms the observations in Experiment A, that earlier intention recognition is achieved, but without an increase in errors.

As the results suggest, our method successfully increases the legibility of the robot motions. By optimizing the weights of the DMPs, it creates variations of the movements that allows a faster prediction of the target.

However, one question that arises is are those legible motions *universal*, i.e recognizable by any individual that observes them once? As the variations in the motions are unconstrained, the number of possible trajectories is infinite which render this hypothesis rather unlikely. Moreover, as the human brain is very efficient at recognizing patterns, it is much more likely that the subjects were actually learning to recognize the specificity of the motion that leads toward one of the goal.

This question has motivated us to extend our work and to continue our quest toward *universal* legibility.

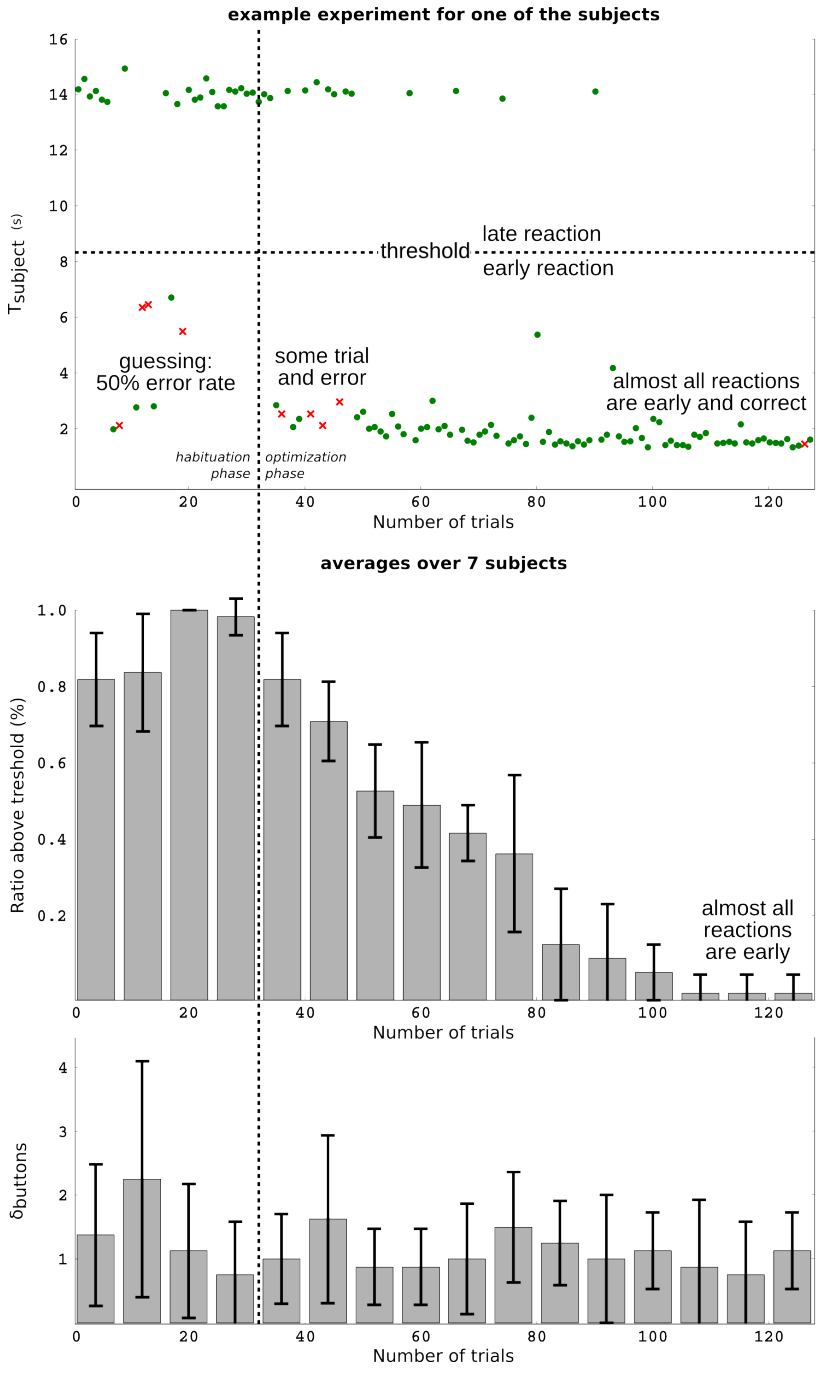

Figure 3.7: Results for Experiment B. Top) Prediction times of an example subject during learning. Middle) Ratio of late prediction times (see threshold in top plot), averaged over the 7 subjects. Bottom) Number times the incorrect button was pushed, averaged over 7 subjects. The start of the optimization phase is indicated by the vertical dashed line.

4
Toward Universal Legibility

Contents

Although the answer to our initial question "Can a robot learn to generate legible motion from user interactions?" is positive, the resulting trajectories were nevertheless different from those observed in (Dragan and Siddhartha S. Srinivasa 2013). As an example, Figure 4.1 plots two views of the robot's trajectory. We clearly see a substantial upward movement at the beginning of the trajectory for button 1.

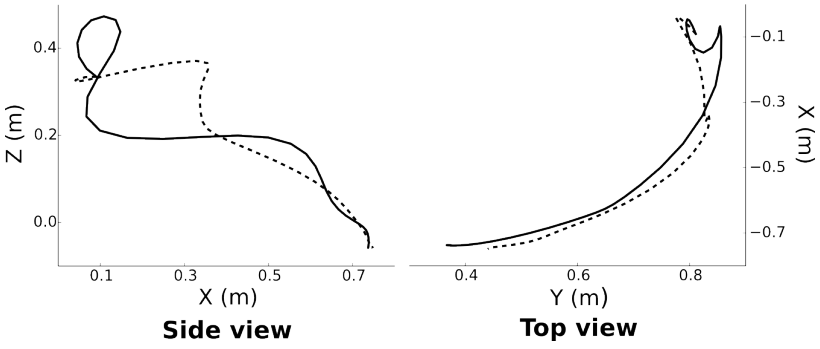

Figure 4.1: Side and top view of generated trajectories after optimization for a single subject. Black/dashed: trajectory for button 1/2 respectively.

Therefore, one question that arises is whether the robot learns to generate universally legible behavior, or rather idiosyncratic behavior that a human learns to interpret. If we consider again the cyclist example, presented in the introduction (see Figure 2.4), a universally legible signal, such as raising the left arm before turning left, would

be immediately understood by people, independently of their origin or cultural background. On the other hand, an idiosyncratic signal would be only understood by the cyclist teammates that have agreed beforehand on the meaning of the said signal.

Similarly, a robot may learn arbitrary but recognizable variations of the movement, which the human may learn to be predictive of moving to the left. This idiosyncratic behavior will have to be re-learned by other humans working with the same robot. In universally legible behavior, the intention is already understood during the first interaction(s).

Further anecdotal evidence is that some subjects reported being able to infer the intention of the robot from differences in the *sound* produced by its motors. Differences in sound arise due to the different velocity profiles of the trajectories for the two buttons. This is clearly a very different type of legibility from that studied in (Dragan, Holladay, et al. 2014; Dragan and Siddhartha S. Srinivasa 2013; Zhao et al. 2016). Although this can be seen as another learned idiosyncrasy, it also suggests that legibility could be obtained by other means than only observing spatial variations of trajectories. This idea is also highlighted in Glasauer's work (Glasauer et al. 2010) where they prove that minimum jerk velocity profiles are more legible than trapezoidal joint velocity one. Combining those elements could lead to even more legible trajectories.

For this reason, we have designed another set of experiments, following the same protocol described in Chapter 3 which aims at avoiding such idiosyncratic behavior, and measuring the effects on learning legibility.

Learning Legible Motion with a Less Expressive Policy

The overall experimental set-up is the same as in Experiments from Chapter 3. Therefore, we only explain the differences, which are the policy representation, and a slightly modified cost function.

Methods

To avoid the idiosyncratic behavior observed with the DMPs, we designed a policy that allows for much less variations. The DMPs were defined in joint space (7 joints) with 3 basis functions that are varied per joint, leading to a policy that has ϑ=21 parameters. To reduce this number, the second policy representation generates trajectories that pass through a viapoint, which itself is parameterized by only two parameters, as visualized in Figure 4.2. Therefore, we refer to this policy as the *viapoint policy*.

We have limited this study to only the button pressing experiment, i.e experiment A of Chapter 3.

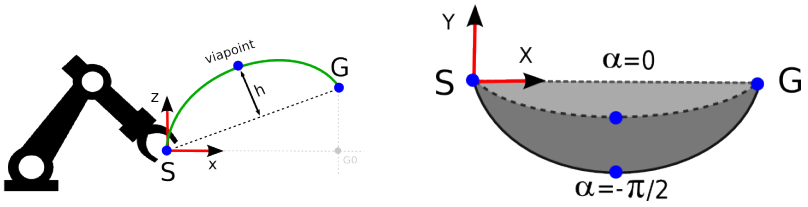

Figure 4.2: Viapoint policy representation. Top: the trajectory is generated from the start S to the goal G (the location of the button), through a viapoint whose distance to the line $S--G$ is determined by the parameter h. The rotation around the x-axis is determined by α.

The trajectories are generated from a start point S (initial robot configuration) to an end point G (such that the button is pushed), which are fixed throughout the experiment. The height of the parabolic path is defined as a parameter h. The rotation around the x-axis, parallel to the ground, is defined as the parameter α. We represent this rotation seen from above. This policy constraints the generated trajectories for more smoothness. We expect them to resemble the ones obtain in Dragan's work (Dragan, K. C. Lee, et al. 2013). However we do not encode explicit informations about their legibility. Thus during the exploration of the parameter space some of the generated trajectories might be really deceptive.

The cost function for the viapoint policy is the same as in Eq. 3.9, except that the penalty on the jerk is now in task space, not joint space. As before, the optimization of this cost function takes place within the space of the policy parameters `, which is now of dimensionality 2 (instead of 21 as with the DMP). We again use 9 subjects. To avoid any habituation effect from the first experiment we have chosen new participants.

Results

The main results of the experiment using the viapoint policy are summarized in Figure 4.5, which has the same format as Figure 3.4. The complete results for this experiment are shown in the right column of Figure 4.7 Figure 4.7 allows for a direct comparison of the two policies.

Discussion

We again observe a drop of the prediction time during optimization. Similarly to the results obtained with the DMP policy, experiment A, this creates a drop in the trajectory completion at prediction time (from 80% to 60%). The number of prediction errors increases during the optimization process before stabilizing at the end. The average number of errors is still sufficiently low, and not significantly different compared to the end of habituation ($p = 0.73$), to prove that the subjects are not simply guessing the next target. The decrease in prediction time during the habituation is significant ($p = 0.005$) as well as the decrease after the optimization ($p = 2.1e^{-5}$).

Qualitatively, these are thus the same results as in the experiment with the DMP policy. As for the DMPs, we represent in Figure 4.3 two views of the trajectories. As expected, this policy produces smoother trajectories to the targets. In this case, the trajectories look

like what we would expect from a legible behavior, i.e an exaggeration on the right for the right target and on the opposite side for the left one.

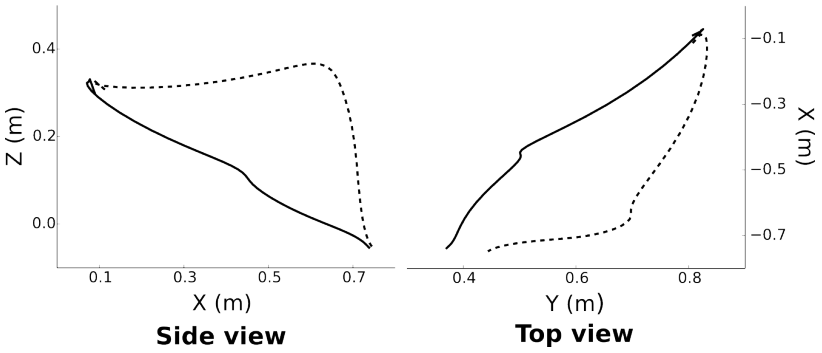

Figure 4.3: Side and top view of generated trajectories after optimization for a single subject. Black/dashed: trajectory for button 1/2 respectively.

The higher variance at the end of the optimization compared to the DMP ones suggests that not all the subjects obtain such legible behaviors. In Figure 4.4 we represent another example of optimized trajectories. The trajectories seen from above (top view) look rather deceptive.

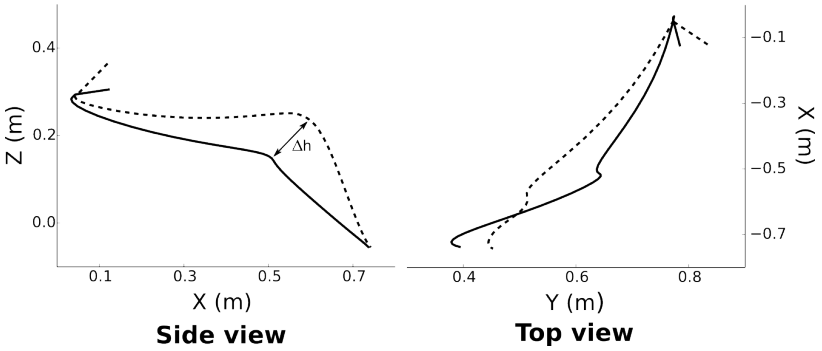

Figure 4.4: Side and top view of generated trajectories after optimization for a single subject. Black/dashed: trajectory for button 1/2 respectively. The generated trajectory seems more deceptive when looking at the top view. Yet trajectories are distinguishable in term of height as represented by Δh in side view.

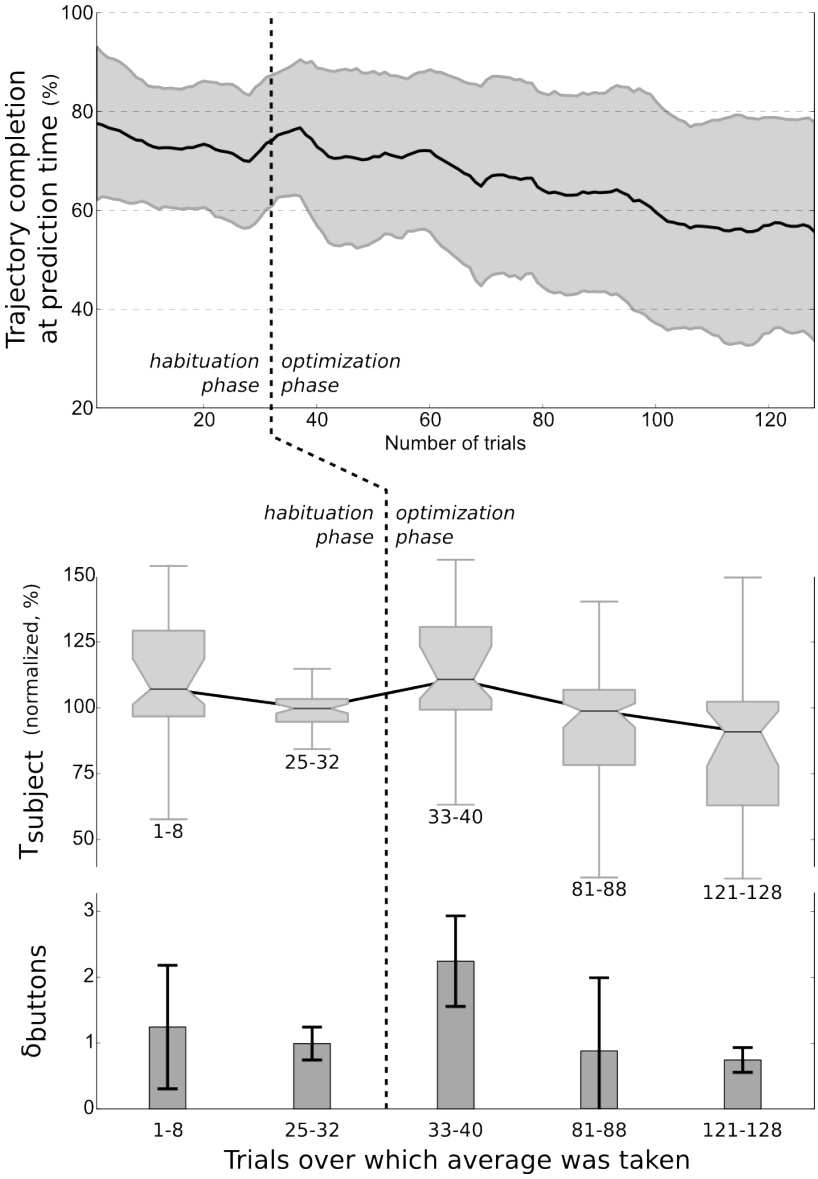

Figure 4.5: Results of the experiment with the viapoint policy (1/2)
Top) Average over all 9 subjects ($\mu \pm \sigma$) of the trajectory completion at prediction time, i.e, the relative amount of trajectory (timewise) observed by the subject when it presses the button. This value is calculated using the formula $100(1 - \frac{T_{robot} - T_{subject}}{T_{robot}})$.
Bottom) Normalized T_{subject}, averaged over all 9 subjects and blocks of 8 trials; average number of failures, i.e. when different buttons were pushed, averaged over all 9 subjects and blocks of 8 trials. The two graphs show the values at certain key frames during learning.

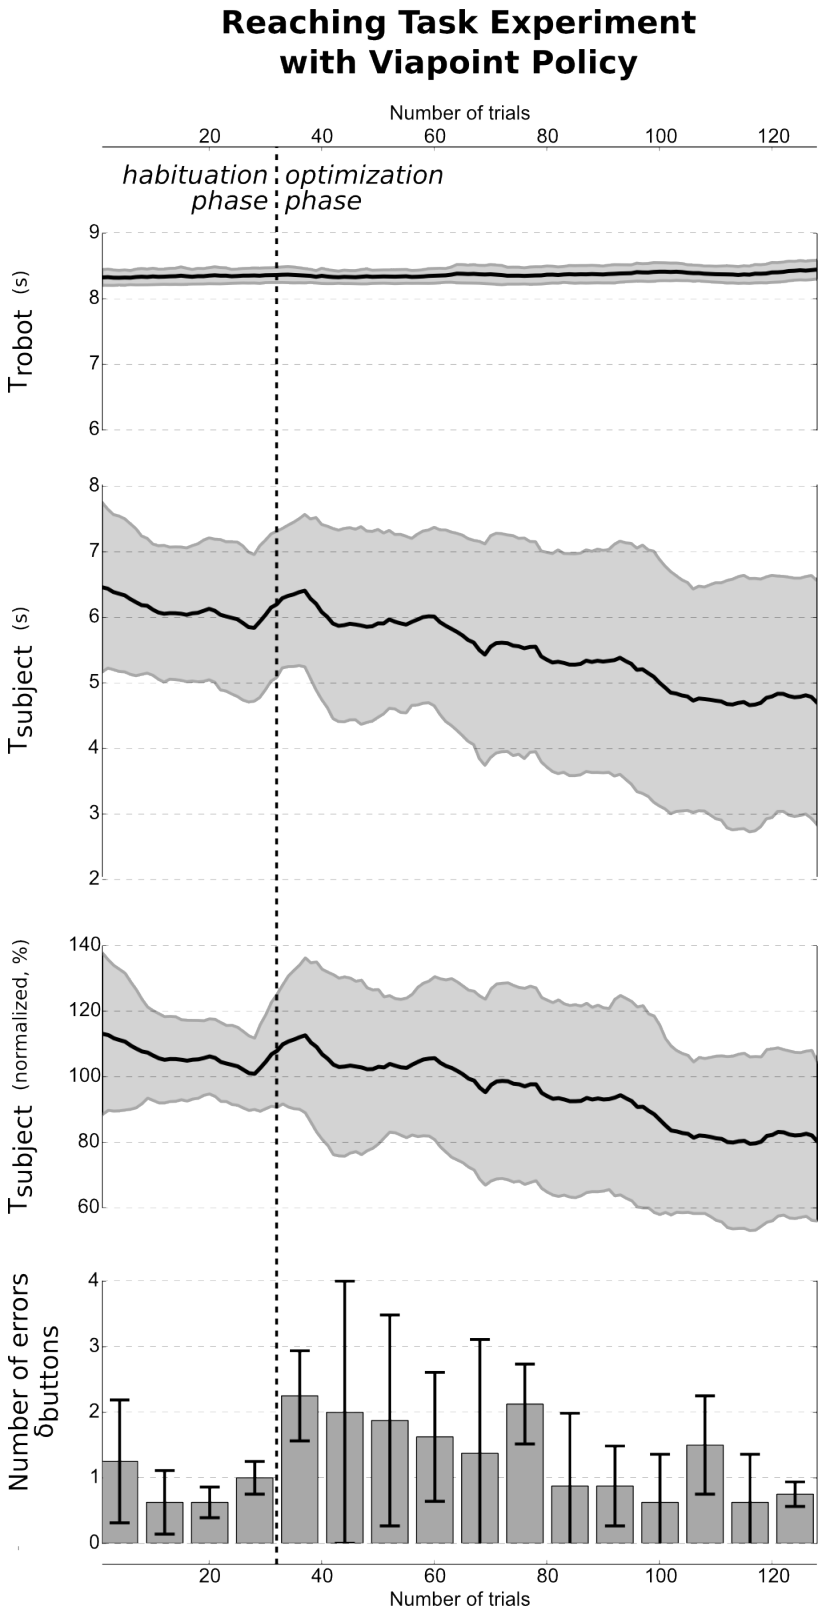

Figure 4.6: Results of the reaching experiment with viapoint policy (2/2). The start of the optimization phase is indicated by the vertical dashed line.
Top row) Average ($\mu \pm \sigma$) of the robot button pushing time (T_{robot}). It varies little for the DMP policy (left) and even less for the viapoint policy (right). For the latter this is to be expected, as the duration of pressing the button is not dependent on the parameters of the policy in which exploration and optimization takes place.
Second row) Average ($\mu \pm \sigma$) of the subject button pushing time T_{subject}, over all 9 subjects. Variance is quite high because some subjects push quickly overall, whereas others are more careful.
Third row) Again the average subject button time, but this time normalized with respect to the average value of T_{subject} during the last 8 trials of the habituation for each subject. This reduces the variance caused by the overall differences between subjects.
Bottom row) Number of times the incorrect button was pushed, averaged over blocks of 8 trials and all 9 subjects.

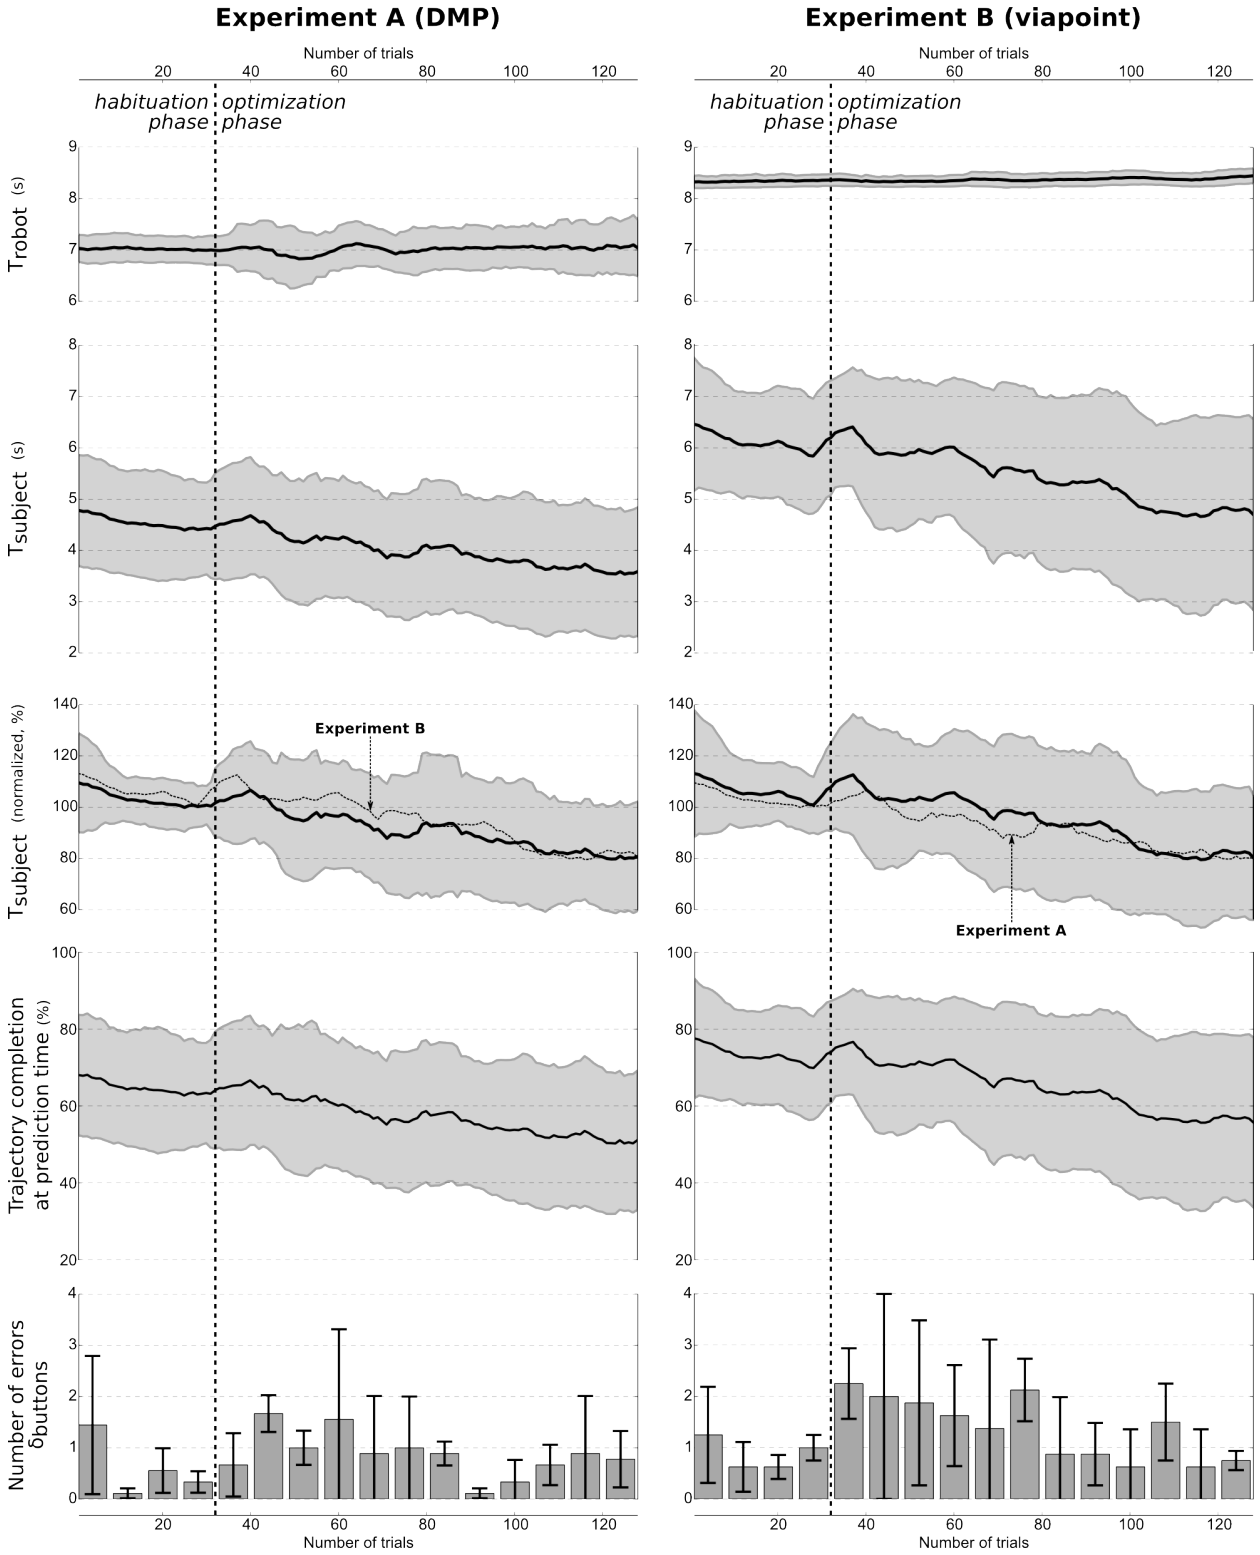

Figure 4.7: Comparison between DMP and viapoint policies. Graph axis are similar to Figure 4.5

One hypothesis is that by constraining the trajectories to resemble legible behavior we increase the number of local minima of the optimization. Consider that the global minima is achieved when the trajectories meet what we expect a legible motion to be. Because of the sampling in the parameter space that solution might not be found during the optimization. Moreover subjects might learn a deceptive or idiosyncratic motion as they do with the DMP policy. Thus most of them decrease their prediction time at the end of the optimization. However the ones with the biggest drop obtain trajectories similar to those represented in Figure 4.3.

The experiment in the next section will investigate how well trajectories generated by the two different optimized policies (DMP and viapoint) transfer to novel users.

Transferability of Legibility

Previous experiments verify that robots are able to improve the legibility of their behavior from interactions with humans. One question that remains is whether the learned legibility do transfer to other observers. We now present two experiments in which we investigate whether the adaptations that have been learned during interactions with one subject also improve the legibility for other subjects. The first experiment (Experiment C) uses the same protocol both previous experiments, but starts with trajectories that have been previously optimized. The second experiment (Experiment D) does not involve optimization, but rather presents several previously optimized trajectories in a random order. Experiment C is aimed at determining whether humans can learn to interpret the idiosyncratic motions of robots, whereas D aims at analyzing which type of trajectories enable humans to immediately recognize intentions, without the need to learn how to interpret them.

Experiment C: Pre-optimized policies

Do subjects learn quicker when starting with policies that have been optimized previously with another subject?

Method We ran the same experimental protocol with the habituation and optimization phases, as described in Chapter 3, with 4 novel subjects each for both policy parameterizations. In contrast to the optimizations described previously, the initial trajectories are now trajectories that have been previously optimized for other subjects. The initial trajectories were not chosen randomly but correspond to the most legible ones for each parameterization, i.e. the ones that lead to the biggest drop in term of prediction time.

Results The results of Experiment C are plotted in Figure 4.8. Whereas previous experiments showed smaller improvements during habituation (7% and 10% for DMP and viapoint respectively) and large improvements during optimization (a further 20% and 20%), we here see the inverse. The improvement during habituation is now 37% and 43% (both $p < 1^{-5}$), whereas during optimization they are small and not significant ($p = 0.47$ and $p = 0.52$). The complete results of Experiment C are shown in the right columns of Figure 4.9 and Figure 4.10.

Discussion The results in Figure 4.8 suggest that subjects can quickly learn to recognize the intentions of the robot from trajectories optimized for another subject, for both the DMP and the viapoint policy. Because the improvement in T_{subject} during habituation is much more pronounced than during the experiments with non pre-optimized policies, we deduce that these trajectories are indeed more legible.

From a comparison between T_{subject} of the two policies and their equivalent in Experiment C we observe some interesting behaviors. First the difference in T_{subject} for the DMP on the first eight trials is significant ($p < 0.03$, Mann-Whitney U test) with T_{subject} being lower for the DMP policy. We also note that the subject's predictions happen at 70% of the trajectory with the DMP policy and 90% with the viapoint policy and that this difference is significant ($p < 0.03$, Mann-Whitney U test). Initial trajectories for the DMP policy are close to straight line to the target (learned by demonstration). According to the definition of legibility this suggests that optimized trajectories might be less legible when shown to novel users without habituation. However humans adapt very quickly and by the end of the habituation the optimal time is reached and does not vary throughout the optimization. Moreover at the end of the habituation the prediction is performed at 50% of the robot trajectory when subject are shown optimized trajectories compared to 60% with the straight lines. We then deduce that optimized trajectories are more legible. This is however a contradiction with the fact that they started as less legible. As stated, the definition of legibility from Dragan et al. (Dragan, K. C. Lee, et al. 2013), cannot handle such contradictions because it does not account for the possibility of habituation. At the end of the optimization phase the difference between the two policies is not significant neither in term of time ($p = 0.42$) nor in term of trajectory completion at prediction time ($p = 0.08$).

For the viapoint policy the situation is the exact opposite. During the first eight trials the difference in T_{subject} is not significant ($p = 0.23$) neither is the difference in term of trajectory completion ($p = 0.41$). Thus optimized trajectories are at least as legible as straight lines without habituation. However at the end of the optimization there is a significant difference in term of time ($p < 0.03$) and therefore in term of trajectory completion trajectory with a drop of almost 10% (from 50% to 40%). The trajectories selected for Experiment C were

the most legible one, i.e. the one that lead to the greatest drop in the subject's prediction time between habituation and optimization. This observation supports the hypothesis that the optimization of the viapoint policy presents some local minima.

Are the viapoint trajectories more legible than the DMPs? For the DMP based trajectories, when looking at trials 8 to 16, the difference, in term of prediction time, between the straight lines of the habituation phase of the non pre-optimized and the already optimized trajectories of Experiment C are not significant($p < 0.03$). This means that after 8 trials of habituation subjects were able to perform similarly to those who observed straight lines to the target. But by the trials 16 to 24 they perform significantly better. For the viapoint policy it is sufficient to wait for the trials 8 to 16 to see a significant improvement in the prediction time. Thus we can conclude that the viapoint policy requires less habituation trials to perform better than the two other type of trajectories.

Between DMP and viapoint policies we note, at the end of the optimization, a difference in term of trajectory completion (50% with the DMP trajectories versus 40% with the viapoint ones). However this difference can be explained by the fact that T_{robot} is slightly different between the two policies. In fact, in term only of prediction time, both DMP and viapoint policies perform similarly (they both converge to $3.5s$). Therefore, a direct comparison between them in term of prediction time might not be suitable as the subject's prediction time depends also on the speed of the movement of the robot.

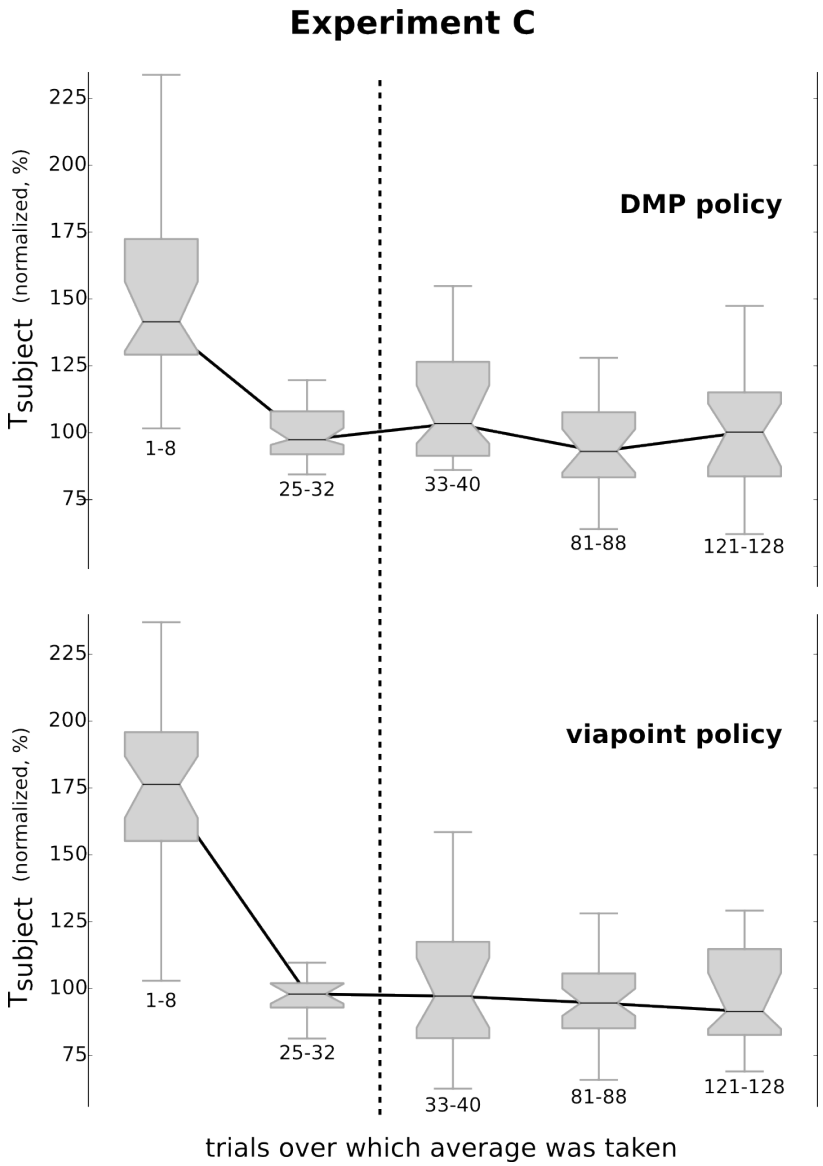

Figure 4.8: Box plots for the normalized prediction times, when starting with previously optimized trajectories from the beginning, averaged over all subjects, and blocks of 8 trials. Top) DMP based trajectories. Bottom) viapoint trajectories.

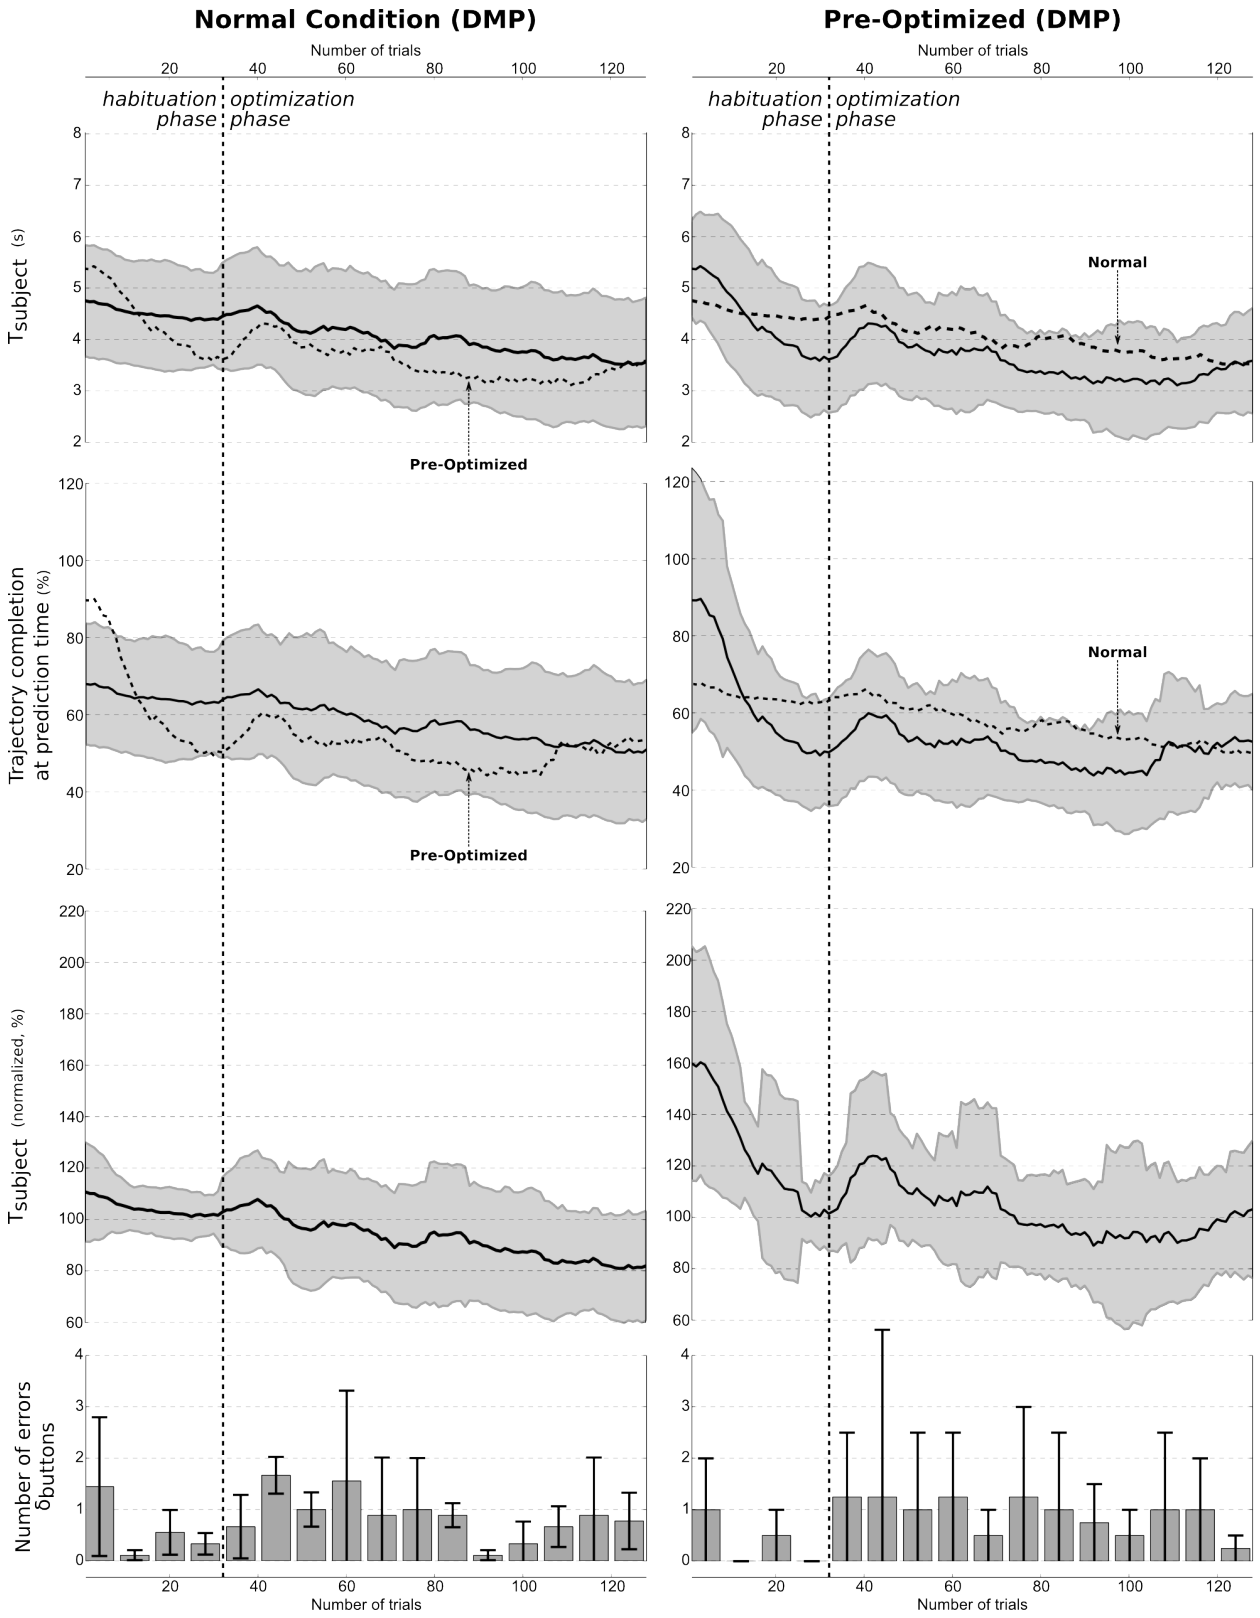

Figure 4.9: Comparison between normal and pre-optimized conditions for the DMP policy.

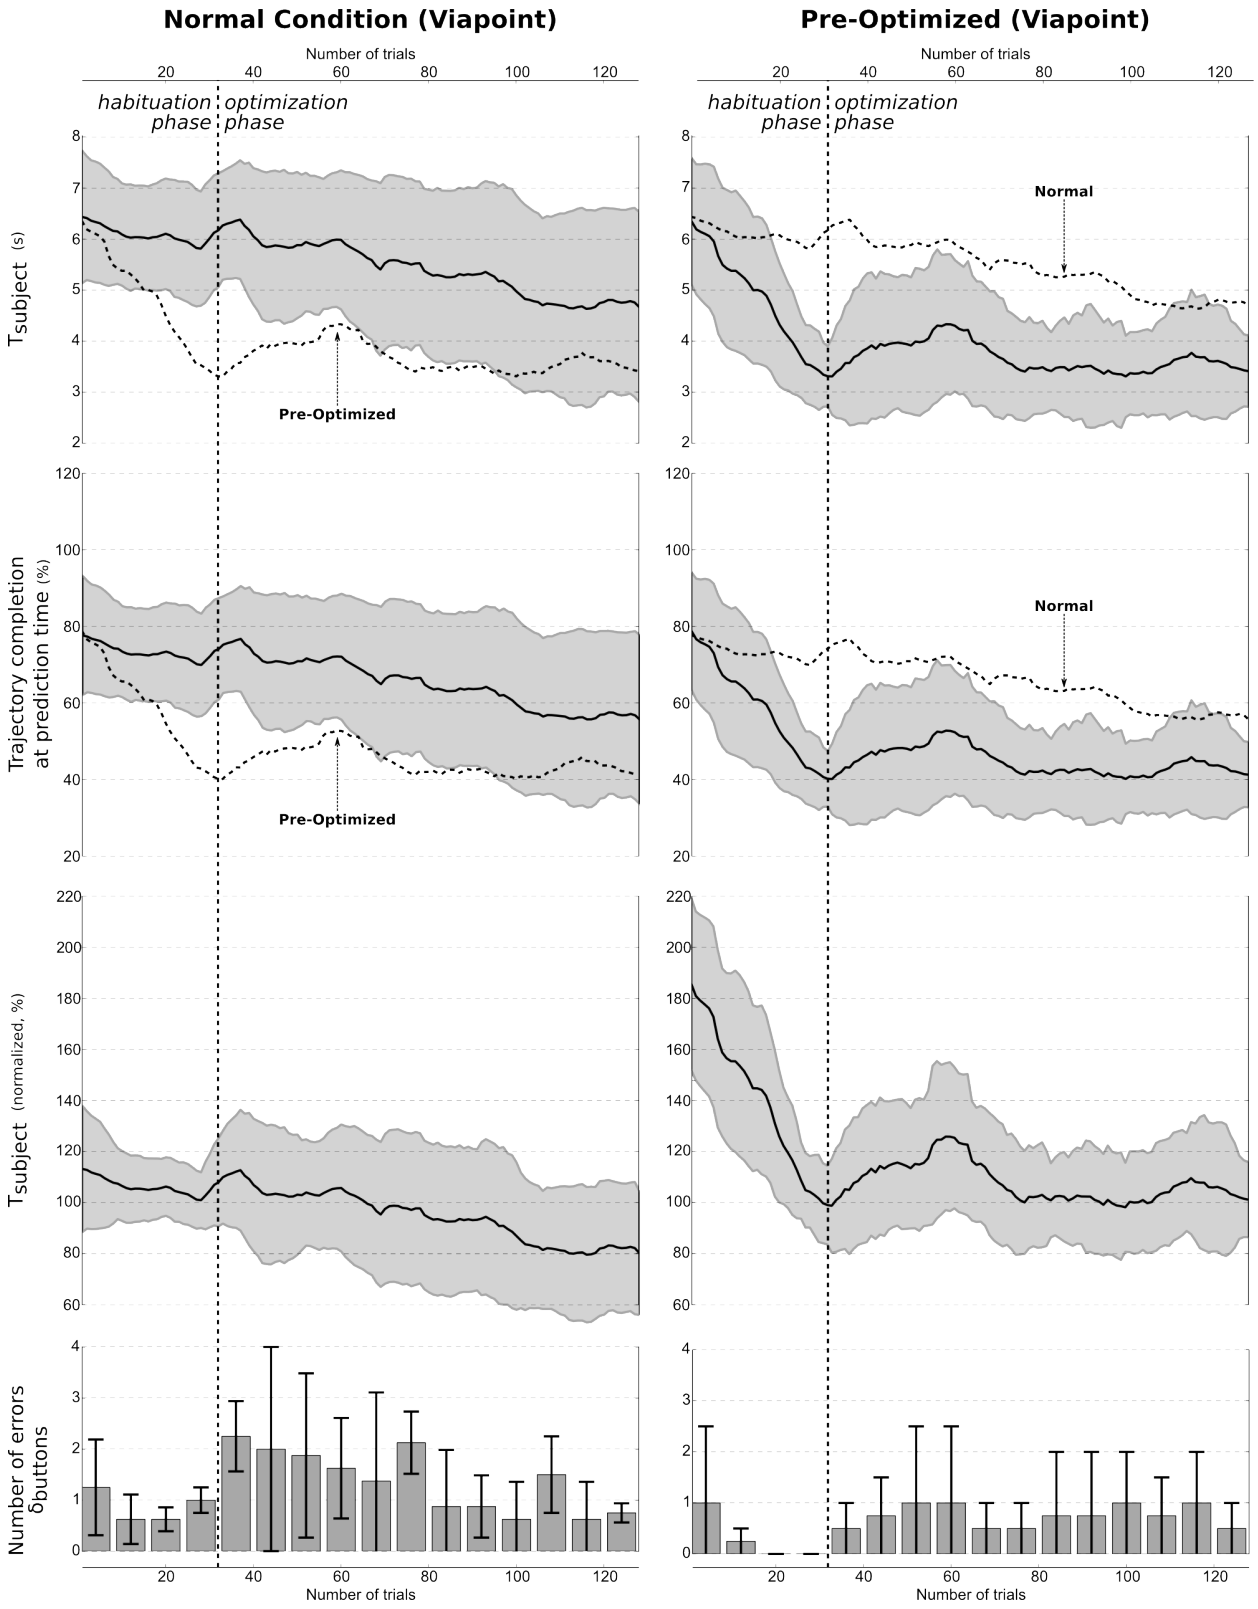

Figure 4.10: Comparison between normal and pre-optimized conditions for the viapoint policy.

Experiment D: Cold start

The aim of this experiment is to determine if subjects can immediately recognize the intention of the robot from trajectories optimized for other subjects. Therefore, we use neither a habituation nor optimization phases for one particular trajectory, but rather present different previously optimized trajectories only a few times. A limited number of presentations is necessary, because the human may learn to interpret potential idiosyncrasies of the movements, which we want to avoid in this experiment.

Method For both buttons, five types of trajectories are presented:

- trajectories generated by two optimized DMP policies that lead to the largest reduction in T_{subject}. We refer to them as DMP_1 and DMP_2
- as above but with two viapoint policies noted VP_1 and VP_2
- straight line minimum-jerk trajectories (S) with end-effector pointing toward the button, as a baseline.

Run	1	2	3	4	5	6	7	8	9	10
Targets	**R**	**R**	**B**	**R**	**B**	**B**	**B**	**R**	**B**	**R**
Types	DMP_2	S	S	DMP_1	VP_1	VP_2	DMP_1	VP_2	DMP_2	VP_1

Table 4.1: Illustration of one random sequence for experiment D. A complete run comprises a repetition of four such random sequences. This makes a total of 40 trials.

The order of the buttons (denoted **R** and **B**) and trajectory types is random within a sequence of 10 trials. The sequence is repeated 4 times which lead to a complete run comprising 40 trials. An example of a random sequence is presented in Table 4.1. Zhao et al. (2016) show that straight line minimum-jerk trajectories, with end-effector pointing toward the target, are the most legible type of motion for a high number of possible targets. By comparing the DMPs and the viapoint based trajectories to this kind of straight lines, we hypothesize that for the two-target case scenario the other types of trajectories convey more informations and thus are more legible. For this experiment, 8 novel subjects were used.

Results The results of Experiment D are summarized in Figure 4.11. The top graph, depicts T_{subject} for all types of trajectories. Each bar represents the average over all users and buttons. Differences between buttons were not significant ($p > 0.33$, Wilcoxon signed-rank test), and thus pooled. Differences between the DMP and the two other type of trajectories are significant ($p < 0.03$, Welch's t-test). However the difference between the viapoint policy and the straight line is not significant ($p = 0.21$). The bottom graph depicts the same results for the number of errors. The differences between the viapoint policy and the two other type of trajectories is significant ($p < 0.03$, Welch's t-test). However there is no significant difference between the DMP and the straight lines ($p = 0.33$).

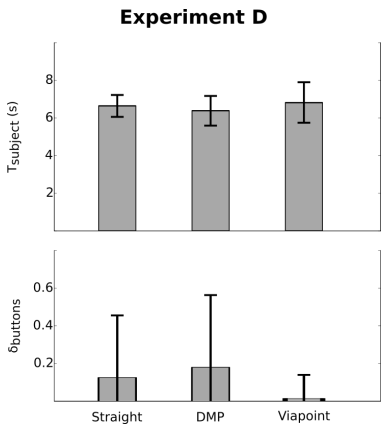

Figure 4.11: Times (top graph) and prediction errors (bottom graph) for the three type of trajectories.

Discussion The results in Figure 4.11 are in accordance with the observations made in Experiment C. In term of prediction time all trajectories perform similarly. We recall that T_{robot} differs between the DMP and the viapoint policies. Thus comparing them only on time might be biased. However there is no ambiguity when looking at errors. The number of errors for the DMP policy is similar to that of the straight trajectory, but the number of errors for the viapoint policy is far lower. This means that subjects are able to recognize the intention of the robot from the viapoint policy much more robustly than from the two other policies. Because subjects are able to do so immediately *without* habituation or previous training, this indicates that the viapoint policy is more legible than the two other policies.

From those results we conclude that reusing optimized trajectory on novel subjects allows for a faster learning of the robot's sense of legibility. Even with DMP based trajectories, where the robot's motion can be considered as idiosyncratic, subject were able to recognize faster the robot's intention. Moreover only the habituation phase is sufficient to reach the performances of the initial subjects for whom trajectories have been optimized. After habituation, no further improvement is achieved. The legibility of previously optimized trajectories could not be further increased by further optimization with another user. Another conclusion is that the viapoint policy is significantly more legible than the two other type of trajectories as it requires less habituation and leads to a lower error rate when presented without habituation. Although we cannot conclude that the viapoint trajectory is indeed *universally* legible, it is a first step toward generating more *universally* legible motion.

This chapter concludes our study on the legibility of robotic motions and the first part of the thesis. By incorporating more legibility in the robot motion we are aiming at reducing the cognitive load of the human coworker during the interaction. When uncertain about the robot's intention, the human might perform incorrect support actions which creates stressing situations and increases the global time of a complete task. Therefore, improving the legibility of the robotic motions impact the productivity of the human-robot team as well as the overall human comfort.

Obviously, there are also other factors that impact the comfort of the human coworker. As the robot interact physically with the human, its motions and behaviors could be harmful at a physical level. This physical aspect of the interaction is the motivation behind the second part of this thesis. In the next chapters, we will introduce and detail our work on the ergonomics of human-robot interaction.

5
Introduction to Physical Ergonomics

As we have seen in the previous chapters of the thesis, the legibility of the robot motions is an important factor to improve for an efficient human-robot team. Improving the legibility of the robot motions reduces the human cognitive load and allows for a faster and smoother interaction.

However, the comfort of the human coworker not just concerns psychological effects but sometimes reduces to more physical constraints. Due to its large number of degrees of freedom, the human body is a wonderful machine that allows us to accomplish a single task in very different manners. Observe two people performing the same task, you will most probably notice differences in the way they accomplish them. Take as an example the two ladies in the front row in Figure 5.1. Although they are playing the same score, their postures are actually different. The main problem with this multitude of possible postures is that not all of them are, in fact, adapted to perform the task. This is truly an issue for autodidact musicians that were never taught the importance of a good posture during the practice of a musical instrument. It might have a strong impact on their capacity to play and the sound they produce. In the worst cases, it can even lead to sever injuries (Brockman et al. 1992; Medoff 1999; Rietveld 2013).

This clearly raises concerns on the self-awareness about one's posture and is not restricted to music players. Physical ergonomics is a key issue as 30% of the European workers suffer from chronic lower back pain, leading to enormous social and economic costs[1]. Most industrial companies have made a lot of effort to improve the ergonomics of the workstations (Rosecrance et al. 2000; Wagenaar et al. 1994). The automotive industry is probably one of the sector that have invested the more. Many automotive manufacturing companies use existing health and safety standards, such as OSHA, to develop their own set of guidelines (Desa et al. 2013). Evaluating the results and benefits of the proposed solutions is, however, a very long and tedious process (Fredriksson et al. 2001).

Figure 5.1: Cello players in an orchestra. Image downloaded from pixabay.com.

[1] According to the European Agency for Safety and Health at Work https://osha.europa.eu/en.

Despite the many efforts “Musculoskeletal disorders" (MSDs) are the single largest category of work-related diseases in many industrial countries (Punnett et al. 2004). Over the years, assessment techniques have been developed to evaluate worker's body posture at their workstation (Takala et al. 2010). Methods for postural assessment can be classified according to the accuracy of data collection and the degree of intrusion caused by the required sensors to collect them (Beek et al. 1998; Wells et al. 1997). Observational methods, based on observations of the worker, are straightforward to use and applicable to a large variety of situations (Hignett et al. 2000; Karhu et al. 1977; Kee et al. 2001). Their accuracy is however limited as they are mainly developed for paper-based observations without the need of specific tools or sensors. With the development of non-intrusive human tracking system, like depth sensors, automatic postural assessment based on observational methods start to emerge (Diego-Mas et al. 2013). The development of such techniques will have a significant impact in industry by drastically reducing the cost of ergonomic solutions. Moreover, it will also benefit fields like human-robot interaction by allowing a more thorough analysis of the human coworker.

In human-robot interaction safety and potential hazards is a central aspect of the interaction (De Santis et al. 2008). Humans and robots are starting to work in shared environments allowing an increasing diversity of tasks they can accomplish together. As we are developing robots to assist workers in industry, we have here a good opportunity to increase the well-being at work, by providing solutions to prevent work related diseases and injuries. Some efforts have already been made to improve acceptability by including social behaviors, like proxemics[2] (Koay et al. 2014; Mitsunaga et al. 2008; Takayama and Pantofaru 2009). Most advanced robotic planners combine multiple notions, including safety and social acceptance, to create human-aware robotic responses (Sisbot, Luis F. Marin-Urias, et al. 2010). In contrast to a human-human interaction, where both parties try to maximize their own comfort, sometimes at the expense of the other, an assistive robot should always maximize the comfort of its coworker.

[2] Proxemics is defined as physical and psychological distancing from others (Edward Twitchell Hall 1966)

Consider, as an example, a situation where the robot has to handover an object. By carefully planning the motion, we ensure that the trajectory is safe, i.e the robot does not enter in contact with the human while moving toward its final position. Yet, is it enough to consider it as a safe interaction? Probably not if you look at the left picture of Figure 5.2.

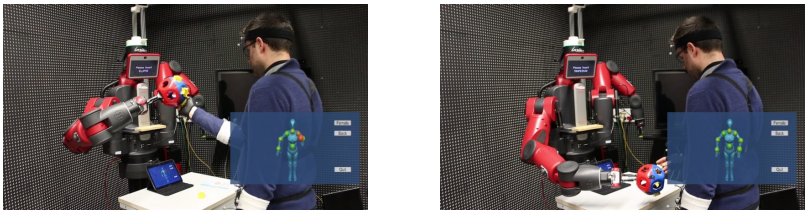

Figure 5.2: Difference between two interactions. An interaction that might hurt the receiver over the long term (left picture) opposed to a safer interaction (right picture)

The scenario represented in this image is clearly uncomfortable for the worker who has to overextend his wrist to insert the piece in the object. If this happens only once, it could be acceptable. Repetition of this situation, however, can be hazardous[3]. The right image of Figure 5.2 shows, on the other hand, an interaction that seems more ergonomic and more comfortable for the worker.

When the robot passes, or presents, an object to the worker the final object pose influences the way he or she grasps it (Kölsch et al. 2003; Rosenbaum et al. 1996), and by extension, the comfort of the posture at receiving time. Improving this posture requires analyzing how the human receiver grasps a specific object, and orient it to improve the receiving posture (Aleotti et al. 2012). Another way is to define a cost function that captures all possible hand-over configurations which allows to find the optimal comfortable solution. Common cost functions comprise safety, visibility, and postural comfort (Sisbot, Luis F. Marin-Urias, et al. 2010) or naturalness & appropriateness (Cakmak, Siddhartha S Srinivasa, et al. 2011).

[3] Evidence suggests that there may be interdependence between force and repetition with respect to MSD risk. Repetition seems to result in modest increases in risk for low-force tasks but rapid increases in risk for high-force tasks. This interaction may be representative of a fatigue failure process in affected tissues (Gallagher et al. 2013)

Reasoning in terms of human capacities to solve a task or plan safe human motions requires the use of a realistic human model. Most researches make use of a generic human model such as the one included in OpenRAVE[4] (Mainprice et al. 2010; Sisbot and Alami 2012). The main drawback of this approach is that it fully ignores biological differences between people that might limit the proposed solutions. Figure 5.3 shows how the result of a simple handover can impact the posture at receiving time if planned for a non adapted human model.

Only a few recent studies propose to consider biomechanical models adapted to the current user (Bestick et al. 2015; Suay et al. 2015). Creation of those model has been greatly simplified with the development of efficient skeleton tracking solutions based either on depth sensor data, such as a Kinect (Papadopoulos et al. 2014), or more recently, directly from a simple RGB flux (Cao et al. 2017).

[4] OpenRAVE provides an environment for testing, developing, and deploying motion planning algorithms in real-world robotics applications (Diankov et al. 2008).

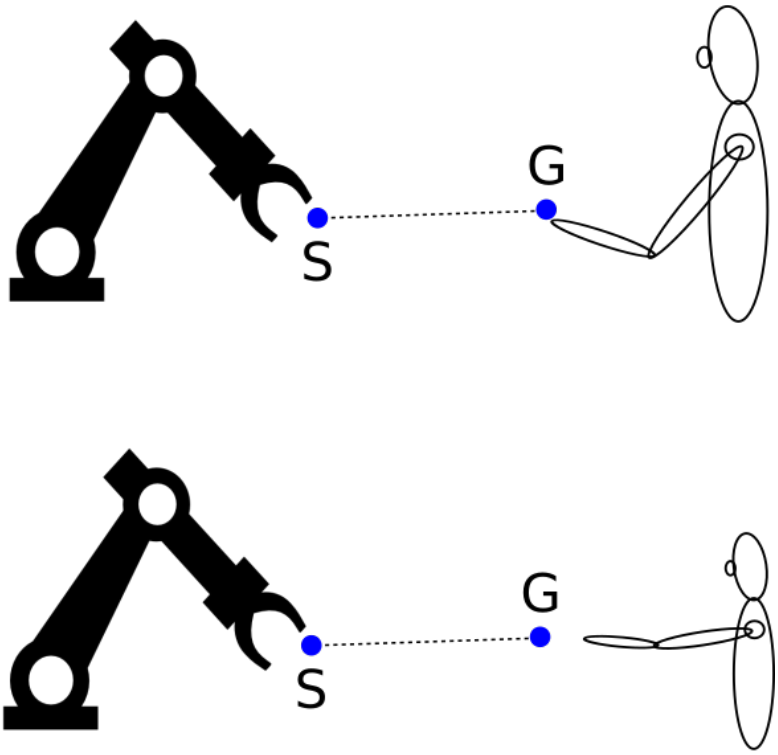

Figure 5.3: Representation of how the difference between people might impact the planned solution for handing over an object.
Top drawing) A solution to hand over an object where the robot starts from a position *S* and deliver the object at position *G*.
Bottom drawing) The same solution applied to a smaller model. Even by fully extending its arm the receiver is unable to grasps the object.

One question is, *how can we include physical ergonomics in the interaction?*

There are multiple levels where the ergonomics could leverage the interaction. One aspect, for example, is to plan the robot motions for maximum ergonomics on the human side as it has been already addressed in the literature. Most researches optimize the robot's placement and behaviors with respect to specific cost functions that acknowledge human needs (Ikuta et al. 2003; Sisbot, Luis F Marin-Urias, et al. 2007; Sisbot, Luis F. Marin-Urias, et al. 2010). However, many do not make use of the extensive amount of available cost functions provided by the ergonomics field. Most of them also reason on a generic human model, with the issue illustrated in Figure 5.3. For this reason, we have decided to incorporate simple cost functions from the ergonomics field to plan the robot motion and actions while making use of a personalized human model for each of our users.

In Chapter 6 we present an optimization algorithm to plan the robot motion for maximum physical ergonomics on the human side, while considering constraints linked to the task at hand[5]. This approach also tries to maximize the acceptance of the robotic system by introducing notions derived from proxemics (Mitsunaga et al. 2008), subject laterality, and visibility (Pandey et al. 2014). Other measures such as legibility could also be considered, as described in the previous Chapters, but here we focus more on physical comfort.

[5] This work is published in the International Conference on Intelligent Robots and Systems (IROS 2017) (Busch, Maeda, Mollard, et al. 2017)

Our idea is to find an optimal human body posture to solve a task considering three important aspects of the interaction that are *safety*, *acceptability* and *task constraints* as illustrated in Figure 5.4. Those three factors were chosen according to their importance in the existing literature and are very similar to previous researches in the human-robot interaction field (Cakmak, Siddhartha S Srinivasa, et al. 2011; Sisbot and Alami 2012). By first extracting a personalized human model from the Kinect skeleton of the user, we optimize the human posture to fulfill a task under those defined constraints.

The key component of our solution is the physical ergonomic cost function. We label this cost under the *safety* category as it drives the human body away from postures that present a risk of MSDs. The ergonomics literature introduces a large variety of assessment techniques. We needed a score function that provides a quantitative measure from joint angles of the human body. Therefore, after a careful review of available choices we have selected the *Rapid Entire Body Assessment (REBA)* method[6] (Hignett et al. 2000). As we will see in Chapter 6 we still had to adapt a bit the REBA cost which was originally developed for pen and paper assessment of a worker's posture.

Once the optimal human body posture is selected, we derive a robot motion that would lead the human coworker toward this optimal posture. Our hypothesis, verified in this work, is that the user will intuitively behave similarly to the predictions made by our algorithm. In the case of handing over objects or joint manipulation, the robot can induce the human location and posture by planning the positioning of the object within the shared workspace.

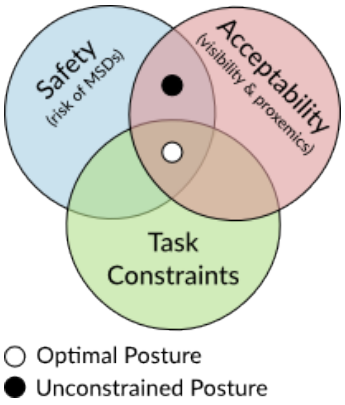

Figure 5.4: Illustration of the trade off between *safety*, *acceptability* and *task constraints*. The set of optimal postures lies at the intersection between the three. Each constraint can be seen as an independent cost function. Optimizing for the three of them is performed by optimizing the weighted sum of the cost functions.

[6] Although the REBA method is a good fit, other techniques could have been also selected. In our solution, the ergonomics component is sufficiently modular to be replaced by any other technique to fit specific needs.

One could wonder why did we choose such a *model-based* approach after considering a *model-free* method to optimize for legibility? Using *model-free* trials and errors, similar to the way we have handled legibility, would have consisted of relying on the subject's conscious feedback, e.g. asking him or her where the robot should hand-over an object or place it on a table (Cakmak, Siddhartha S Srinivasa, et al. 2011). Although this approach is interesting to understand why do people prefer to interact with the robot in some specific ways, we believe it is not sufficient to ensure that the interaction will remain safe.

As we have stated before, MSDs are often due to a repetition of a wrong posture over a long period of time. When you ask for a conscious subject's feedback people often tend to provide an instantaneous response that represent their feelings on the moment. The problem is that a given posture could let the person feel perfectly safe but lead to MSDs if repeated. As an example, think about the way you sit in front of your computer. You know you are probably doing it wrong, but the lack of pain or discomfort do not forces you to reconsider it. In worst cases there might even be a conflict between people preferences and ergonomic recommendations as illustrated in the *automobile seat designer's dilemma*[7] (Kolich 2003)

[7] Reed et al. (1991) described the *automobile seat designer's dilemma* as the need for a balance between prescribing a physiologically appropriate seated posture and accommodating a driver in a preferred posture. They reasoned that prescribed postures sometimes compromise long-term comfort.

Because of this, we believe a *model-based* approach to be more appropriate in this case. Moreover, as the only differences between different users appear in their morphologies, a parameterized model can overcome this effect.

Is only planning the motions enough to achieve an ergonomic human-robot interaction?

We believe the choices of actions and the sequences in which they are performed also impact the ergonomics. Take as an example the situation where both the robot and the human can lift an heavy box. This is a burdensome task that does not require any specific skills. Such a task should automatically be assigned to the robot, relieving the human from an excessive physical overload.

Task allocation and coordination between humans and robots is an extensively studied topic in HRI (Johannsmeier et al. 2017; Toussaint and Lopes 2016; Tsarouchi et al. 2016). A strong emphasis is made on the communication with the human partner and negotiating the planned sequence of action (Clodic et al. 2009; Foster et al. 2008) Johannsmeier et al. (2017) proposes a framework to assign a full sequence of atomic tasks to a human-robot team based on the capabilities of each agents. The geometric part of the tasks is, however, known in advance and supposed to be immutable, i.e objects arrive on conveyors without variability in their locations. On the other hand, Toussaint and Lopes (2016) introduces a Monte Carlo Tree Search (MCTS)[8] approach to solve Task and Motion Planning (TAMP) problems simultaneously(Lozano-Pérez et al. 2014; Plaku et al. 2010; Srivastava et al. 2014). This approach allows to define cost functions based on the geometry of the scene, such as velocity and contact constraints, to adapt to changing environment when planning the task allocation. For example, if the next task requires to use the screwdriver but this one is too far for the human to directly grasp it, the robot should go pick up the screwdriver and hand it over to the human.

In Chapter 7, we present our work on ergonomic task allocation[9]. We have based our work on the planning solution proposed by Toussaint and Lopes (2016). With that approach we can integrate the cost functions introduced in Chapter 6 in the optimization process. Doing so allows us to plan the actions, not only considering capabilities of both the robot and human agents and geometric feasibility of the task, but also the ergonomics of the interaction.

[8] Monte Carlo Tree Search (MCTS) is a method for making optimal decisions in artificial intelligence (AI) problems, typically move planning in combinatorial games. MCTS combines the generality of random simulation with the precision of tree search (Browne et al. 2012).

[9] At the time of the thesis writing this work was submitted to the International Conference on Robotics and Automation (ICRA 2018) and still under review (Busch, Toussaint, et al. 2018)

6
Postural Optimization for Ergonomics

Contents

We took the approach that ergonomic interaction can be achieved by optimizing the robot motion with cost functions that capture the needs of the human coworker. To capture the ergonomic cost, we rely on a well-accepted measure in the industry called the Rapid Entire Body Assessment technique (Hignett et al. 2000). In Section 6.2 we give a detailed explanation on how to assess a body posture using the REBA method. While there exist many other assessment technique that could have been potential candidates for our method, we have chosen to use the REBA metric due to its generality and the provision of a quantitative assessment.

As we have seen in the introduction, personalized human model is also an important feature to consider. Therefore, the first step of our solution is to generate a personalized model as described in Section 6.1. In Section 6.3 and 6.4 we detail the different cost function we

have introduced in the optimization and the implementation details we had to consider respectively.

Finally, we conclude this work with a user study on an insertion task with 40 participants where we prove the benefits of including ergonomics in the interaction. Both qualitative results, from questions asked to the participants, and quantitative results, from the assessment of their recorded posture during the interaction, show significant improvements.

Personalized Human Model

Reasoning in terms of human kinematics requires the use of a personalized human model. In order to evaluate the risk associated with a given posture, this model should provide joint values from pose of body segments. To this end, we create a personalized human model as a URDF[1] with 32 degrees of freedom (DOF) by calculating the length of each body segments on the observed subject's skeleton from a Kinect.

[1] The Universal Robotic Description Format (URDF) is an XML file format used in ROS to describe all elements of a robot (Garage 2017).

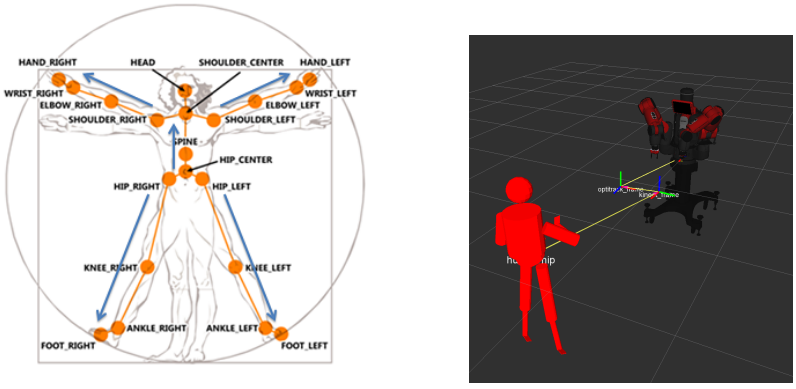

(a) Kinect skeleton

(b) Animated model

Figure 6.1: Images of the personalized human model.
Left) The Kinect skeleton on which our model is based. Image taken from microsoft website.
Right) A view of our model animated in Rviz.

This model can be used ofline to plan the robot motion as detailed in Section 6.3 or online to track the subject and automatically assess his or her current posture.

Rapid Entire Body Assessment (REBA)

Automatic postural assessment, based on observational methods, has started growing in interest with the development of cheap and easy to use human tracking systems (Li et al. 1999). Although the REBA technique has been initially developed for pen-paper observations (Hignett et al. 2000), the fact that it handles static as well as dynamic postures and that it relies on quantitative values makes it suitable for an automatic assessment. It describes ergonomic preferences with a table as follows. First, a *part table* assigns a score to each parts of the body (trunk, neck, upper arms, ...), based on their inclinations at the time of the assessment. Then, an overall score is calculated from *correspondence tables* that take into account the importance of the body segment. Indeed, a score of 3 for the trunk represents a greater risk than the same score for the upper arm. Details of the

calculation are given in the Appendix A. This assessment technique is the main component of our algorithm for postural optimization detailed in Section 6.3.

Postural Optimization

To respect the three main aspects of interaction, that are *safety*, *acceptability*, and *task constraints*, we define several cost functions and constraints to find an optimal body posture during the interaction. In the following sections we detail, for each aspect, the associated cost function or constraint.

Safety and Comfort

Human safety is a central piece of the interaction. In term of body posture, safety corresponds to minimizing the risk of MSDs. To achieve this we find the body posture that minimizes the score returned by the REBA assessment technique as detailed in Section 6.2.

With our work we want to adopt such meaningful industrial standards, but make them applicable in our optimization algorithm. This requires us to propose a differentiable version of the REBA cost function, as described in the following.

Differentiable REBA (dREBA) In its original formulation, the REBA score is a stepwise linear function. We propose to fit a differentiable model to the REBA score, dREBA, and use this as cost function in the LGP formulation. We model dREBA as a sum of weighted polynomial functions,

$$\text{dREBA}(q,t) = \delta_{payload} + \sum_{i=1}^{n} w_i Q_i(q_i,t) \ , \tag{6.1}$$

where n is the number of joints considered in the REBA techniques and $Q_i(q_i,t)$ is a 2nd order polynomial of the joint i as a function of the joint value q_i at time t. First, the coefficients of each polynomial $Q_i(q_i,t)$ are calculated to minimize the squared error to the the *part tables* of the original REBA score table, for each joint separately. Then, the weights w_i are learned from the total REBA score, including the *corresponding tables*. For this, the squared error of our surrogate and the original REBA score is minimized on a set of sample body configuration: We define a set of random poses, ensuring each class of the original REBA cost (from 1 to 10) are evenly represented.

Computing the gradients of the fitted dREBA cost function (6.1) is straightforward,

$$\nabla_{q_i} \text{dREBA}(q,t) = w_i \nabla_{q_i} Q_i(q_i,t) \ . \tag{6.2}$$

Table 6.1 shows the root mean-square error (RMSE, bottom row) of the dREBA approximation for each of the 10 REBA classes (upper row). As we can see, the RMSE is relatively high for classes of high REBA cost but acceptable for those of low REBA cost.

1	2	3	4	5	6	7	8	9	10
0.72	0.91	0.86	1.08	1.02	1.72	1.86	2.18	2.18	1.50

Table 6.1: RMSE for each REBA classes

In our experiments, we use dREBA to enable optimization of the REBA score. However, as the original REBA score is the accepted standard we will report all the scores with respect to the original REBA score.

Payload cost calculation The term $\delta_{payload}$ in (6.1) corresponds to the cost for carrying an object. Although convenient for pen and paper assessment, as it is simpler to work with integer values, the function of (A.1) from the original REBA calculation means that carrying objects lighter than $5kg$ is costless and transitions are sharp. As we are working in the continuous domain, we find it more suitable to also consider a continuous function to represent the payload coast. Therefore, we propose the following linear cost term for load,

$$\delta_{payload} = \frac{load}{\omega_{load}}\,, \tag{6.3}$$

where $\omega_{load} = 5kg$. The function defined in (6.3) is strictly equivalent to the original calculation in (A.1) at the boundaries fixed by the REBA method but has the advantages to set a cost even for light weight objects.

Task constraints

For some tasks, the pose of a specific body part, like the hand, is constrained to a specific location. For example when we use a tool our hand is tied to it, which itself is linked to the object we are acting upon. This part of the algorithm is the only specific part that needs to be specified for each task. Most common task constraints would be to set parts of the worker's body to a certain pose.

Considering the forward kinematic of our human model FK, the pose in Cartesian space of body part i is represented as FK_i. Constraining a body part to a certain location in space is equivalent to minimizing the distance between the forward kinematic and the desired pose FK_i^{des}. This distance is expressed as two subcomponents, one for the position as Euclidean distance and one for the orientation as distance between two quaternions. The distance between two quaternions is defined as $(Q_1, Q_2) = 1 - \langle Q_1, Q_2 \rangle^2$. Constraining multiple body parts is performed by summation of the distances,

$$\begin{aligned} C_{task}(FK, FK^{des}) = \sum_i (&\|FK_{i,pos} - FK_{i,pos}^{des}\| \\ &+ k(1 - \langle FK_{i,rot}, FK_{i,rot}^{des} \rangle^2)). \end{aligned} \tag{6.4}$$

Because position and orientation distances do not share the same unit, it requires the addition of a scaling factor $k \in \mathbb{R}$. In practice, we have set $k \approx \pi$. As the quaternion distance varies between 0 and 1, this then corresponds to the distance between the two angles traveled on the unit circle. We also note that FK^{des} might be expressed with respect to the object frame. This is particularly interesting when the object is free to move in space and its pose is part of the optimization.

Acceptability

The concept of *acceptability* is not easy to define. It is usually divided into social acceptability: *How does the society perceive robotic systems?*, and practical acceptability: *How do people perceive the robot when interacting with it?* (Nielsen 1994). In this study we focus on the latter. We incorporate visibility, proxemics and user's laterality (right or left-handed) in our algorithm in order for the worker to feel at ease during the interaction.

Visibility It is an important notion to consider. When the robot moves an object, it should always do so in the worker's field of view. Otherwise, it might surprise the user. We have put the visibility notion under *acceptability* although it does not uniquely belong to this category. For example, a non trained worker would most probably look at the object he is acting upon in order to fulfill his task (*task constraints*). Such as novices in computers look at their keyboard when typing. Another one might get hurt if he uses some tools and cannot visually verify the results of his actions (*safety*).

To optimize for visibility, the model's head pose is directly linked to the pose of the object X_{obj} acted upon. Pose of the object can be considered as fixed, attached to the model's hand by a predefined transformation or set as a variable of the optimization. We denote the sight vector coming from the head as $\vec{H}_x$, H corresponding to the forward kinematic of the head FK_{head}. We consider that this vector originates from the center of the head and is parallel to the ground when the head is straight. The visibility cost corresponds to the angle between the vector from the center of the head to the object $\vec{X_{obj}H}$ and $\vec{H}_x$. It is calculated using the dot product between them as $\arccos(\vec{H}_x \cdot \vec{X_{obj}H})$. This way we ensure that the object is approximately at the center of the vision cone when this cost is close to 0.

For some objects, the orientation also matters. For example, the handle of a tool should be directed toward the user to ease the grasping. To represent this we consider a feature vector $\vec{v}_{obj}$ that need to be aligned with the directional vector of the head $\vec{H}_x$. At the moment, this vector is predefined for each object used in the experiments. The corresponding cost is also defined by the dot product $\arccos(\vec{H}_x \cdot \vec{v}_{obj})$. The final cost for visibility corresponds to the sum of the two dot products,

$$C_{visibility}(FK_{Head}, X_{obj}) = \arccos(\vec{H}_x \cdot \vec{X_{obj}H}) + \arccos(\vec{H}_x \cdot \vec{v_{obj}}). \quad (6.5)$$

Proxemics People have special concerns about their personal space either in terms of intimacy and safety. We consider proxemics to address such concerns. Close interaction like handing-over an object would be socially unacceptable and create uncomfortable situations if it happens too close to the worker, especially for non-trained ones. In our algorithm we have added the proxemics notion as a constraint to ensure that the optimized posture will not violate the worker's intimate space. The proxemics constraint corresponds to the minimal distance from the human waist to the object. Only planar distance (x and y coordinates) are here constrained. This limit correspond to the boundary of the subject's intimate space, i.e 0.45m (Edward T. Hall 1963),

$$X_{Obj} \in proxemics\_constraints. \quad (6.6)$$

Laterality Every worker is unique. Some are right-handed, other left-handed. Therefore, the model should also take into account this specificity. We leave the possibility to select the worker's laterality. This impacts which hand is constrained to act upon an object.

Algorithm for postural optimization

Finally, we group all the previous cost functions and constraints (6.1-6.5) to define a multi-criterion optimization, whose pseudocode is detailed in algorithm 1, for optimizing the human body posture while solving a task. Our formalism is very generic and considers many tasks beyond the ones presented in this study. The task constraint function and its parameters are the parts that have to be redefined depending on the specific task to solve.

The final cost function is a weighted sum of the different elements, some of them with different units for which we need different weights $w_i \in \mathbb{R}$. This weights will most probably be task dependent. For example, in some industrial tasks, visibility might be a critical aspect. In this case increasing the weight that account for visibility will lead to the desired behavior. The outputs of the algorithm q_{opt} and ${}^{H}p_{opt}$ correspond to the optimal human posture in joint space and the optimal object pose with respect to the human reference frame respectively.

Algorithm 1: Postural Optimization Algorithm

procedure OPTIMAL_BODY_POSTURE($laterality$, w, FK^{des}, obj_dof, ${}^{H}X_{obj}$)
  $model\_laterality \leftarrow laterality$
  $q \leftarrow$ random *model* joint values
  $p \leftarrow []$
  if obj_dof is “fixed" **then**
    $p \leftarrow {}^{H} X_{obj}$
  else if obj_dof is “position" **then**
    $p \leftarrow$ random position
  else if obj_dof is “orientation" **then**
    $p \leftarrow$ random orientation
  else if obj_dof is “full" **then**
    $p \leftarrow$ random pose
  $f(q,p) = w_1\, C_{posture}(q) + w_2\, C_{task}(FK(q), FK^{des}) + w_3\, C_{visibility}(FK_{Head}(q), p)$
  $q_{opt}, {}^{H} p_{opt} \leftarrow$ **min** $f(q,p)$
  subject to $p \in$ *proxemics constraints*
  return $q_{opt}, {}^{H} p_{opt}$

Implementation Details

In this section we define the implementation choices we have made to track the human body to automatically assess the worker's posture and to optimize the body posture using the optimization defined in Section 6.3.

Human tracking and automatic assessment

While during the calibration of the model the environment can be set up such that the human can stand directly in front the camera, during the execution, we found that the presence of the robot, constant occlusions, and space constraints make it difficult to reliably assess the REBA score using the same camera.

Therefore, at runtime we opted to track the human using optical markers. Each subject has been equipped with a suit covered by OptiTrack markers to track the main body segments (head, hands, elbows, torso and waist). The model's frames are recorded during the whole interaction. To obtain the model joint values we compute its inverse kinematic on the recorded data. As the legs are not tracked, we set them in rest position, i.e. with joint values leading to the minimal REBA score.

The waist frame is the most important frame of the OptiTrack suit. It corresponds to the model base frame. We also use it to express the position and orientation of the subject in the robot workspace.

Figure 6.2: Optitrack suit for human tracking. The suit tracks the Cartesian location of the waist, trunk, head and both shoulder, elbows and wrists. To obtain the q vector of the model we perform an inverse kinematic.

Finally, we have adapted the protocol for calculating the REBA score, detailed in Section 6.2, for an automatic calculation. This allow us to record the subjects' postures and to provide quantitative results, based on their REBA score, during the user study.

Optimization of the cost functions

To minimize the costs defined in Algorithm 1 we use a simple gradient descent algorithm from *scipy* library. This is sufficient as we consider only static postures and not trajectories[2].

[2] This approach was the one considered in (Busch, Maeda, Mollard, et al. 2017). In Chapter 7 we will see how we have integrated the cost functions detailed in Section 6.3 in a more sophisticated motion planner in order to also account for trajectories.

From human to robot space

The solution provided by Algorithm 1 corresponds to the optimal body posture in joint space, q_{opt} and the optimal object pose ${}^{H}p_{opt}$ in the human space. We must now generate a robot trajectory that positions and orients the object in its optimal location ${}^{H}p_{opt}$ such that the human is in the optimal posture q_{opt}. The relation between q_{opt} and ${}^{H}p_{opt}$ is defined as the forward kinematic FK. The robot trajectory can be designed by expressing the pose of the object w.r.t the robot's reference frame, ${}^{R}p_{opt}$. To this end, at runtime we locate the current position of the human w.r.t the robot's frame using, for example, optical markers or depth cameras to compute a transformation ${}^{R}T_{H}$. This process is illustrated in the left part of Figure 6.3.

Transformed into the robot's space, the optimized solution represents the optimal posture for the human to interact with the robot. For simplicity, take the case where ${}^{R}p_{opt}$ needs to be satisfied only at the end of the interaction (e.g. in handover) and the robot starts from a pre-defined home position ${}^{R}p_{init}$. Depending on the location of the human w.r.t. the robot, and the degree of redundancy of the robot, two situations are possible. If the final object position is within the workspace of the robot, many possible robot trajectories that lead to the satisfaction of ${}^{R}p_{opt}$ can be designed. Conversely, if the final object position is outside the workspace of the robot, no trajectory can be generated.

To account for both cases and quickly generate a trajectory, we solve a motion planning problem where the initial and final end-effector positions of the robot are connected by a straight line (in Cartesian space), as an initial guess. The orientation of the end-effector is interpolated by slerp between the initial and final poses. If an IK solution on this straight trajectory is not found, we use stochastic optimization to slightly perturb the final desired position of the end-effector and also to reshape the straight line to search for a valid robot's IK trajectory [3]. Perturbing the final desired position of the end-effector will make the exact reproduction of ${}^{R}p_{opt}$ and force the human model to deviate from its optimal posture q_{opt}. Thus, we penalize deviations from the original desired location to find a compromise between the optimal posture of the human and a kinematically feasible robot trajectory. This process is illustrated in the right part of Figure 6.3.

[3] details of this method can be found in (Maeda et al. 2016)

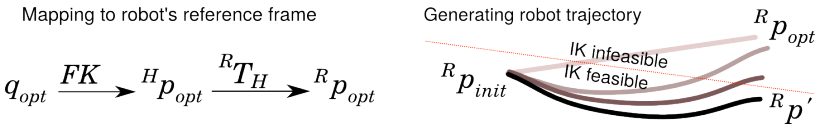

Figure 6.3: Illustration of the trajectory generation.
Left) finding the desired object pose in Cartesian space w.r.t. the robot's reference frame ${}^{R}p_{opt}$ given the solution of the ergonomic optimization q_{opt} and ${}^{H}p_{opt}$.
Right) a straight trajectory is used as an initial guess to bring the object from the robot's home position to the handover position. If a solution is not possible, stochastic optimization is used to optimize the trajectory.

User Study

We apply the postural optimization on an experiment with real users interacting with our Baxter torso-humanoid robot. Two points are tested: *Are the postures generated with our solution perceived better than other types of postures?* and *Does the optimal robot's behavior leads the real user to a safer posture?*

With the object illustrated in Figure 6.4 we mimic an industrial scenario where the robot helps the worker by positioning and orienting an object in which the worker has to insert specific shapes. In total five shapes were considered, located on the blue side of the ball. Name of the shape to insert was written on the robot's screen. Between each insertion the robot was sent back to a neutral pose to let the user prepare the next shape.

We have conducted an extensive user study with 39 participants (18 females and 21 males, 2 left-handed, aged 35 ± 11.26). Before the experiment, subjects were asked to rate their experience with robotic systems on a scale from 1 (no experience) to 5 (daily usage). Result is 2.18 ± 1.27. Three robot's behaviors have been compared. We refer to them as *fixed*, *relative* and *optimized* conditions. Each participant has been presented a sequence of two behaviors from the set of all permutations. Each run of a behavior is composed by the five shapes to insert in a random sequence. Behaviors are detailed in the following paragraphs:

Fixed In this behavior the robot does not take into account the worker position and deliver the object to a fixed pose. We have used this behavior as a baseline.

Relative Research from ergonomics tells us that handing-over an object should happen at torso height at two-third of the receiver's arm maximum elongation (Granjean 1969; Kölsch et al. 2003). For orientation, we constrained the insertion hole to face the user. At execution, the user's torso frame is tracked and the pose of the object on the reference frame of the robot is computed using the required homogeneous transformations.

Optimized Using the optimization from Sec. 6.3 we plan the user's body posture to receive the object and calculate where to deliver it with reference frame transformations. Computing the optimization of the receiving pose is achieved within $\sim 5s$ on a i7-3720QM 2.60GHz computer.

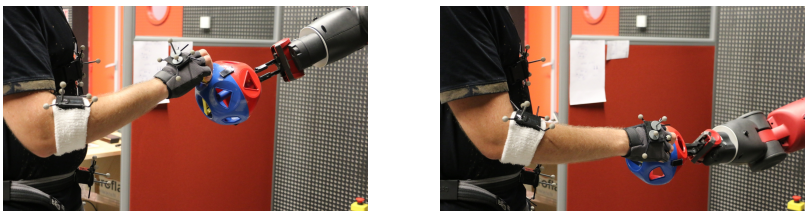

(a) Relative pose (b) Optimized pose

Figure 6.4: Difference between the pose obtained with the *relative* (left picture) and the *optimized* conditions (right picture).

After observing the two behaviors, the subjects were asked to order them according to their preferences and to complete a survey composed of a 12 Likert scale items for each observed behavior. Affirmations were ranked from 1 (I strongly disagree) to 5 (I totally agree). They are based on the System Usability Scale methodology (Brooke 1996) with one affirmation over two written with the negative form. The survey was divided in three subcategories representing the considered aspect of the interaction, i.e *safety*, *acceptability* and *task constraints*. This separation was not made visible to the subjects for whom the 12 affirmations were shown in a random order. After answering the 12 items, subjects had the possibility to leave comments on the observed behaviors. Complete results and survey questions are detailed in Figure6.5. For clarity, the affirmations in the negative form are drawn with negative values from -5 to -1. Subjects' body poses were recorded, using the OptiTrack suit, to evaluate their posture during accomplishment of the task.

The data of the experiment are available on a Zenodo repository (Busch et al. 2017b) alongside a GitHub repository that contains all the necessary materials to read them (Busch et al. 2017a). A video of the experiment is available on our Vimeo channel[4] (Busch et al. 2017c).

[4] From our Github repository `https://github.com/3rdHand-project/PosturalOptimizationDataReader/` you can access both the video and the dataset.

Qualitative Results

From the survey results Figure6.5 and the recorded data we want to validate the three following hypothesis:

- The *Fixed* condition should have the lower rank for all the considered aspects as it does not respect the *task constraints* and do not account for user's *safety* (H_0).
- In terms of task constraints and acceptability the *relative* and *optimized* conditions should produce similar results. Main differences between them should arise from *safety* related affirmations (H_1).
- On the recorded data the *optimized* condition should have the lowest REBA value (H_2).

Almost all subjects who have observed the *optimized* condition have put it first in their preference ordering. Only one person, who got the *relative* and the *optimized* conditions, has chosen the *relative* one first. Second preferred method is the *relative* condition. The *Fixed* condition has never been ranked first. For all the affirmations in Figure6.5 there is a tendency for the *optimized* condition to be preferred

over the two other ones. Results have been compared using Mann-Whitney U test and show a significant preference for the *optimized* condition when summing up all the affirmations ($p - values < 0.05$ for the three combinations). The order of the conditions does not influence the results ($p - values \approx 0.5$ for the three method). Due to the small number of left-handed people we cannot check for any significant differences between laterality.

Considering the results for the three subcategories, there is no significant differences between the *optimized* and the *relative* conditions for the questions on *task constraints* and *acceptability* ($p - values = 0.19$ and 0.33 respectively). All the other combinations are significantly different.

Figure 6.4 shows the difference between the pose obtained with the *relative* and the *optimized* conditions. In the *relative* condition, we observe that the interaction happens at a higher level and forces the subject to bend his wrist.

Quantitative Results

During all the process the subjects' body motions have been recorded and their posture evaluated using the REBA method. We only consider the moment of insertion and average over the few time-steps it has taken the subject to do this action. The moment of insertion was automatically detected from the Cartesian distance between the subject's dominant hand (the one that have been used to insert the shape) and the center of the ball. We arbitrarily set that distance to be less than $20cm$ for the time-step to be accepted as an insertion. Figure6.6 shows the averaged REBA score over all the subjects for the three behaviors for some parts of the body. Differences between the conditions have been tested using a paired t-test on paired data from subjects, merging conditions in opposite order. For example, statistical differences between the *fixed* and the *optimized* conditions have been obtained from people that have observed the *fixed* then the *optimized* conditions and people that have observed the *optimized* then the *fixed* conditions.

The data in Figure6.6 show a significantly lower score on *shoulder*, *wrist* and *total* values for the *optimized* condition. The *Relative* condition presents the highest wrist score. In their comments some subjects have written that their wrist was not in a correct pose when they had to insert the shapes. Especially one subject that had a wrist tendinitis said he felt some discomfort.

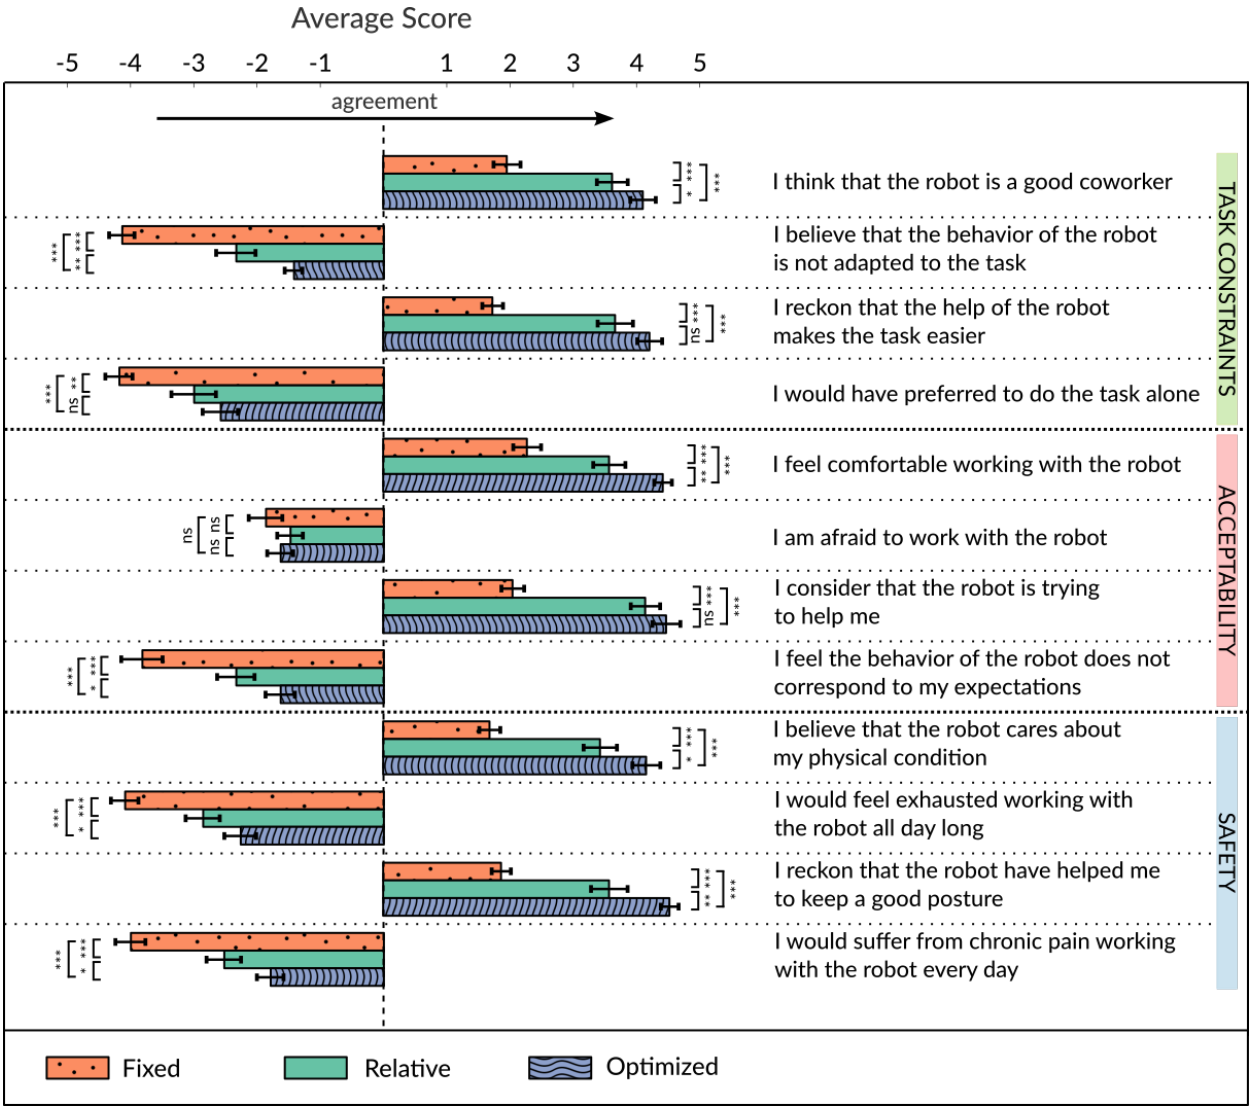

Figure 6.5: Average results and standard error of the mean for the responses to the survey. Each subject have observed two of the three behaviors in a random order. Affirmations are Likert scale items, based on the SUS methodology, ranked from 1 to 5 (Brooke 1996). The negative values for the affirmations in negative form only appear in the figure for clarity matters. Significance have been tested using Mann-Whitney U test and noted according to the standard defined by the APA(American Psychological Association). The *optimized* condition is rated significantly better in almost all of the items.

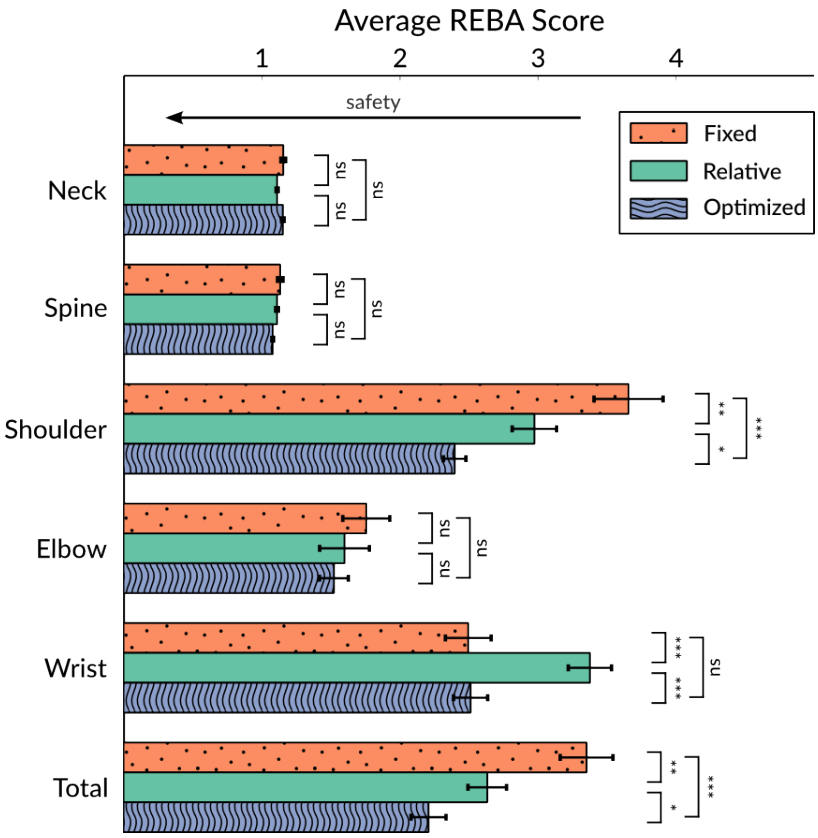

Figure 6.6: Average REBA score and standard error of the mean of the recorded body posture. Only dominant arm have been included in the figure for clarity, but data for the opposite arm are also available. Significance have been tested using Wilcoxon test on paired data. The recorded postures for the *optimized* condition present a lower REBA score. The *relative* condition leads to an higher wrist score that can create some discomfort.

Analysis

From the results of the user study we validate the three hypothesis enunciated previously. The *fixed* conditions is the worst possible behavior both in terms of subject's perception and REBA score (H_0). Although *relative* condition seems to be an acceptable solution it produces a significantly higher body posture score, especially on the wrist. In terms of *acceptability* and *task constraints* it is perceived similarly to the *optimized* solution returned by the optimization. Yet in terms of user's *safety* it performs significantly worst (H_1).

From the analysis of the recorded data, we observe that the *optimized* condition significantly reduces the body posture cost. There is 1.1 point of reduction (3.35 to 2.21 in average) compared to the *fixed* condition and 0.4 point (2.63 to 2.21) compared to the *relative* condition. Although this is not a huge improvement, we have to note that only the dominant arm is used to solve this task. We can expect better results for some task that put a strain on all the body parts. Nevertheless, the repetition of a suboptimal posture can increase the risk of MSDs. Interestingly, the higher wrist score for the *relative* condition can be linked to the comments made by the subjects. During the insertion their wrist was put in extension due to the orientation of the insertion hole. This have been correctly detected by the assessment technique and it is the main weakness of the *relative* condition.

Obviously the *relative* condition results could be improved by changing the orientation that causes the wrist discomfort. But then this would require to empirically find a “correct" solution for each task considered. For some of them this might be not trivial. Therefore, the genericity of our solution allows to automatically find better postures in any task.

The results of the user study suggest that our algorithm is successfully able to generate more ergonomic interactions. By optimizing the robot motion for ergonomics, not only are the obtained posture perceived as better solution compared to naive approaches, the real subject’s posture are also rated better according to the REBA method.

However, there are a few limitations to this approach. First, the robot motion are calculated independently from the optimization. This means that, when the robot cannot reach the optimal trajectory, we select a sub-optimal posture as detailed in Section 6.4.3. Second, the optimization is performed for a static posture and not for a whole trajectory.

In the next Chapter we address those limitations by including the studied cost functions in a more advanced motion planner. This planning solutions allows us to plan the robot and human motions simultaneously, and also to reason in terms of sequences of actions instead of a single task.

7
Ergonomic Task Allocation

Contents

In the previous chapter, we have considered optimizing robotic motions for maximum ergonomics on the human side. However, not only the motions are important but also the choices of tasks and the order they are performed. In this Chapter, we reuse the dREBA surrogate introduced in Section 6.3.1 of Chapter 6 and include it in the Logic-Geometric Program (LGP) formalism introduced by Toussaint (2015) and solved using a Multi-Bound Tree Search (MBTS) approach (Toussaint and Lopes 2016).

In Section 7.1 we present the LGP formalism, used to solve TAMP problems simultaneously. We have applied this formalism in a simulation experiment, detailed in Section 7.2, and a user study, which results are analyzed in Section 7.3.

Ergonomic Task Allocation and Path Optimization

A LGP is an optimization problem over both, a symbolic sequence of actions and a (piece-wise) smooth motion of a system (Toussaint 2015). The logic (e.g., STRIPS-like rules) defines which sequence of actions is feasible; and a sequence of actions defines which geometric constraints the motion has to fulfill. Both are optimized jointly

with respect to a cost function, typically control costs of the resulting motion. Formally, an LGP is of the form

$$\begin{aligned}
&\min_{x,s_{1:K}} \int_0^T c(x(t),\dot{x}(t),\ddot{x}(t))\,dt + f_{\text{goal}}(x(T)) \\
&\text{s.t.} \quad \forall_{t\in[0,T]}\; h_{\text{path}}(x(t),\dot{x}(t) \mid s_{k(t)}) = 0 \\
&\qquad \forall_{t\in[0,T]}\; g_{\text{path}}(x(t),\dot{x}(t) \mid s_{k(t)}) \leq 0 \\
&\qquad \forall_{k=1}^{K}\; h_{\text{switch}}(x(t_k) \mid s_k, s_{k\text{-}1}) = 0 \\
&\qquad \forall_{k=1}^{K}\; g_{\text{switch}}(x(t_k) \mid s_k, s_{k\text{-}1}) \leq 0 \\
&\qquad \forall_{k=1:K}\; s_k \in \text{succ}(s_{k\text{-}1}) \\
&\qquad s_K \models \mathfrak{g}_{\text{goal}}\ ,
\end{aligned}$$

where $x(t)$ is the path, $s_{1:K}$ the symbolic state sequence, $t_{1:k}$ the time points of symbolic state transitions, h_*, g_* the constraints on the motion given $s_{1:K}$, and $\mathfrak{g}_{\text{goal}}, f_{\text{goal}}$ define goal constraints and objectives.

Solving an LGP is hard; most existing combined task and motion planning solvers employ sampling-based methods (Lozano-Pérez et al. 2014; Plaku et al. 2010; Srivastava et al. 2014). As we aim for the optimization with respect to an ergonomic measure we adopt the optimization-based LGP formulation. In (Toussaint and Lopes 2016) a Multi-Bound Tree Search (MBTS) method is proposed to approximate LGP solutions. MBTS uses several bounds of the full LGP that are themselves optimization problems but with less constraints or only over sub-paths and therefore much faster to evaluate. These bounds are used to prune subtrees when they are found infeasible and prioritize search for symbolic sequences.

To apply the LGP framework to our problem we need to formulate

1. the joint symbolic decision space for the specific human-robot cooperation task,
2. the constraints on the human or robot pose depending on such decisions,
3. the cost function for both, the human and the robot.

Concerning the symbolic decisions, we use STRIPS-like operators to define the action space depending on the specific task. For example, for the first task of our experiment section (see Section 7.2) we will define two possible actions, `grasp` and `place`:

- `grasp(t,e,o)`
  - *description*: at time `t`, the end-effector `e` (robot or agent) grasps the object `o`
  - *precondition*: `free(e) on_table(o)`
  - *effect*: `grasped(e,o) ¬free(e) ¬on_table(o)`

- `place(t,e,o,p)`
  - *description*: at time `t`, the end-effector `e` (robot or agent) places the object `o` on table `p`
  - *precondition*: `grasped(e,o)`
  - *effect*: `placed(o,p)` ¬`grasped(e,o)` `free(e)` `on_table(o)`

These operators define feasible symbolic successor states $\text{succ}(s_{k\text{-}1})$, where s_k is a first-order logic state initialized as `placed(screwdriver, table_left)`. Note that the variable `e` refers to possible end-effectors[1].

[1] In this way these operators define the decision space for all "agents" (end-effectors of human and robot)

The decisions imply geometric constraints: A `grasp` implies a kinematic switch of attaching the object `o` to the end-effector `e`, with a respective equality constraint $h_{\text{switch}}(...)$ ensuring that the object does not jump. A `place` action implies a kinematic switch of attaching the object `o` to the table `p`, detaching from the end-effector `e`, both constrained geometric by not having a jump in the object pose. The symbolic goal $\mathfrak{g}_{\text{goal}}$ is `grasped(human_right_hand, screwdriver)`.

Concerning the cost function, we optimize for maximal ergonomics in the side of the human, using the dREBA approximation introduced in Section 6.3.1 of Chapter 6. Additionally we enforce smooth motions by optimizing for the sum-of-squared accelerations of both, the robot and human motion.

Simulation Experiment

We apply the method detailed in Section 7.1 in a simulation experiment where the robot has to place the screwdriver on a table for the human to grasp as illustrated in Figure 7.1. There are three tables of random heights ranging from 0.6 to 1.6 meters located in front of the human. Our hypothesis is that depending whether the ergonomic cost function introduced in Section 6.3 is turned on or off, the choice of table for placing the screwdriver will differ and the overall ergonomic cost will be reduced even if it increases locally. Moreover, the choice of table should impact the human posture cost.

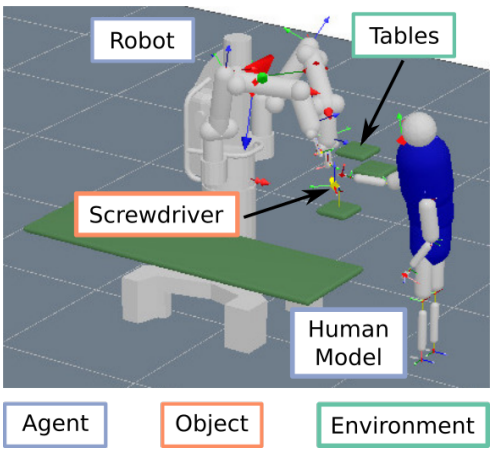

Figure 7.1: Simulated view of the experiment. The robot places the screwdriver on one of the three tables for the human to grasp it. Tables heights are randomly set at each run of the experiment.

The actions available to the human and robot are `pick` and `place` as described in Sec. 7.1. The MBTS solver return a sequences of actions, optimal in terms of the provided cost functions as in the example given in Figure 7.2.

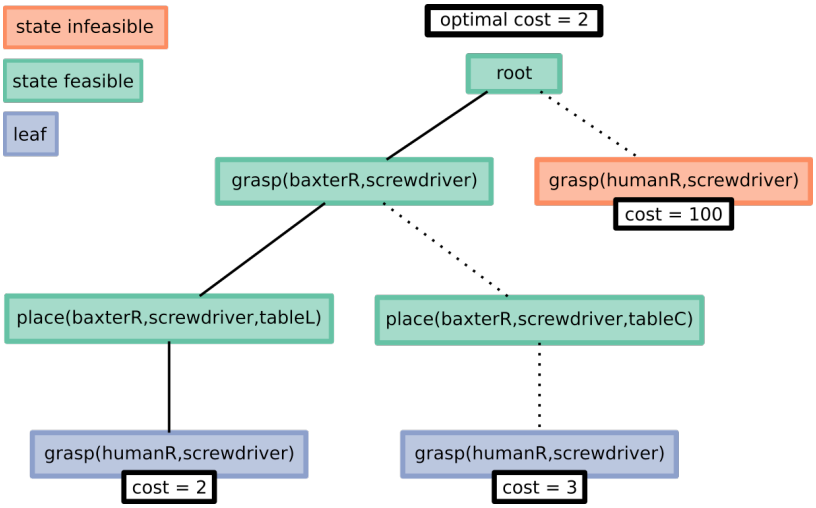

Figure 7.2: Example of a tree search made by the MBTS solver to find the optimal sequence of actions under the given constraints. The path in plain lines corresponds to the optimal sequence. A state is declared as infeasible a constraint is broken. In the example case given, the screwdriver is too far away from the human agent which renders it impossible to grasp. The written cost are given for illustrative purposes and do not correspond to real values returned by the solver.

Experimental setup

To validate the benefits of including the ergonomic cost function in the LGP framework, we have considered three experimental conditions,

1. the *non-ergonomic* condition with ergonomic costs turned off,
2. the *ergonomic* condition with ergonomic costs on,
3. and *re-optimized* condition where the path is optimized with ergonomic cost turned on, but the symbolic action sequence $s_{1:K}$ is fixed to the choice found with non-ergonomic optimizing.

ur main hypothesis is that by optimizing for the dREBA surrogate, the choice of actions and motions will lead the human model to postures that are more ergonomic in terms of REBA score. Therefore, we use the REBA score of the model posture as a measure to validate this effect. Another hypothesis is that optimizing the path along the full sequences of actions will lead to more ergonomic postures compared to stepwise optimization. The objective of the re-optimized condition is then to serve as another baseline that results from methods that have a separation between the high-level task planner and the low-level motion planning. By improving over this baseline we show that it is worth to pay the cost of using the more computationally expensive simultaneous optimization of task and motion planning that is proposed in this work. For example, in this first experiment $s_{1:K}$ refers to the categorical choice of table, which may be chosen sub-optimally with the non-ergonomic optimizer.

In the *non-ergonomic* condition, the planned trajectory of the agent to grasp the screwdriver might not be natural and lead to non-ergonomic postures like an over bending of the spine. Clearly, this will lead to non-favorable REBA scores. In the *re-optimized* condition, re-optimizing such a path with ergonomic costs, but fixing $s_{1:K}$, allows us an easier comparison with the full *ergonomic* condition, which highlights the effects of the choice of the symbolic sequence $s_{1:K}$.

We fix a random height for the three tables and test the three conditions keeping the same table height. The process is repeated 100 times and for each run we collect the table chosen for placing the screwdriver and the human posture at grasping time. We also force the human to be right handed to remove the effects of changing hands to reduce the costs. The initial body posture is set to be at rest according to the REBA assessment (posture of minimal REBA score).

Results

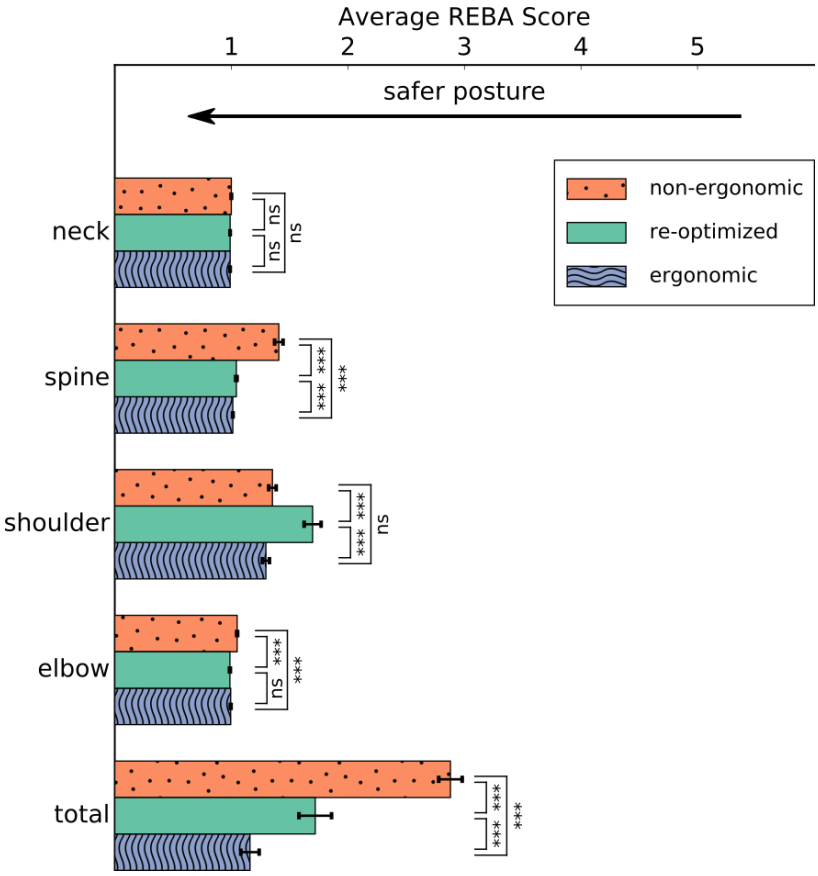

Figure 7.3: Average REBA score and standard error of the mean for the three conditions. The lower the score, the safer is the posture. Significance are verified using Anova statistical test and noted according to the standard defined by the APA(American Psychological Association).

As we observe in the results of Figure 7.3 and 7.6, re-optimizing the path for ergonomics leads, as expected, to a smaller posture score. Still, the best posture scores are obtained in the *ergonomic* condition. Statistical differences are noted directly on the figure and verified with Anova statistical test.

In terms of table heights, the average height of the chosen table in the *non-ergonomic* condition is 1.08 ± 0.22. It is slightly below the average height of the three tables ($1.1 \pm 0.3m$). The average height of the chosen table in the *ergonomic* condition is $0.95 \pm 0.15m$.

By looking at the optimal choices, summarized in table 7.1, we observe that in the *ergonomic* condition, the most chosen table is the centered one. This makes sense from an ergonomic point of view as deviating the arm on the side is more costly. Nevertheless, this choice also depends on the table height as suggested by the number of times the left of right table were chosen. On the other side, in the *non-ergonomic* condition, the most chosen table is the one located on the agent's left side. This choice seems to be almost independent from the table height.

	non-ergonomic	*ergonomic*
left	89	26
center	7	40
right	4	34

Table 7.1: Table choice ratio (%)

This effect arises from the base cost function of the LGP formalism which minimize by default the sum-of-square accelerations of the joints for both the robot and the human. As the time to perform an action is fixed, this leads to select the shortest trajectory for both agents. By extension, placing the screwdriver on the left table often corresponds to the shortest path.

Minimizing the joints acceleration also means that the initial human posture must impact the optimization. To study this effect, we have performed a new experiment, keeping the same experimental setup and reusing the previously generated table heights. The only difference comes from the initial human body posture which is set randomly.

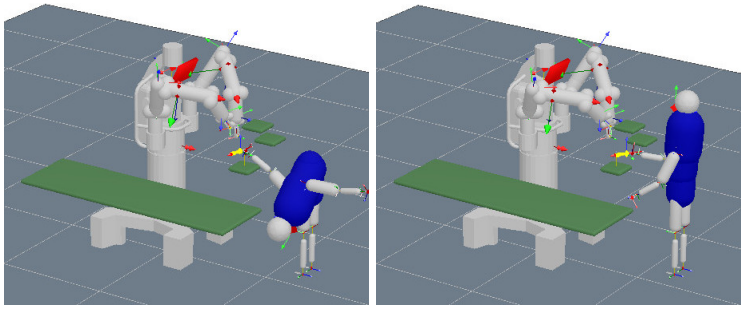

Figure 7.4: Illustration of the impact of the initial body posture.
Left) An example of a grasping posture impacted by the initial body posture in the *non-ergonomic* condition. Without posture correction the shortest path leads to this wrong final posture.
Right) The same grasping but in the *re-optimized* condition which counteracts the effects of the initial body posture.

As we observe in Figure 7.5 this strongly impacts the posture score of the *non-ergonomic* condition. As the human body possesses a large number of degrees of freedom, some joints are not necessary to move to fulfill the task. Therefore the shortest path is to keep them at their initial value. The left image of Figure 7.4 shows an example of a grasping posture impacted by the initial body configuration. In the *ergonomic* and *re-optimized* conditions however, the first human motion is to move back to an ergonomic posture as illustrated in the

right image of Figure 7.4 and especially in the cost profile of Figure 7.6. This action limits the impact of the initial body configuration. As table 7.2 suggests, choices of tables is not impacted in the *ergonomic* condition. For the *non-ergonomic* condition, the impact is significant. As the random posture is often bended, such as the one presented in the left image of Figure 7.5, the shortest path might be a table of lower height. This point is confirmed by the average height of the chosen tables, $1.02 \pm 0.2m$ which is $6cm$ less than in the normal configuration.

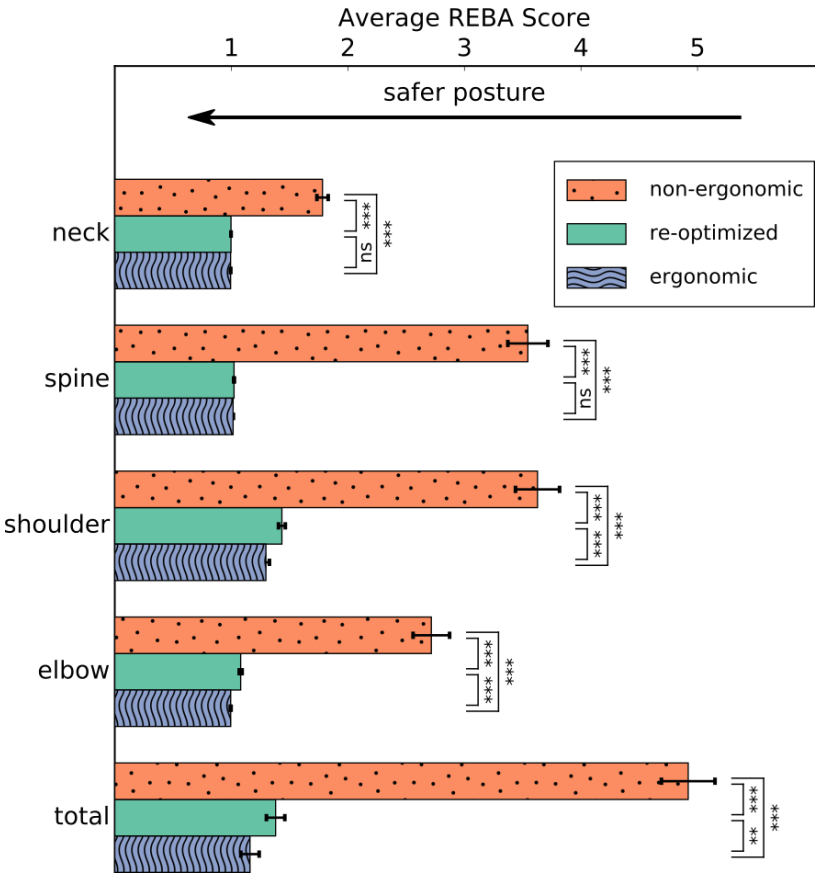

Figure 7.5: Average REBA score and standard error of the mean for the three conditions with random initial posture. As in Figure 7.3, the lower the score, the safer is the posture.

By choosing tables of lower heights, the posture score of the *non-ergonomic* condition slightly gets better. Still, the best REBA scores are obtained in the *ergonomic* condition.

	non-ergonomic	*ergonomic*
left	66	26
center	23	40
right	11	34

Table 7.2: Table choice ratio in the random body configuration (%)

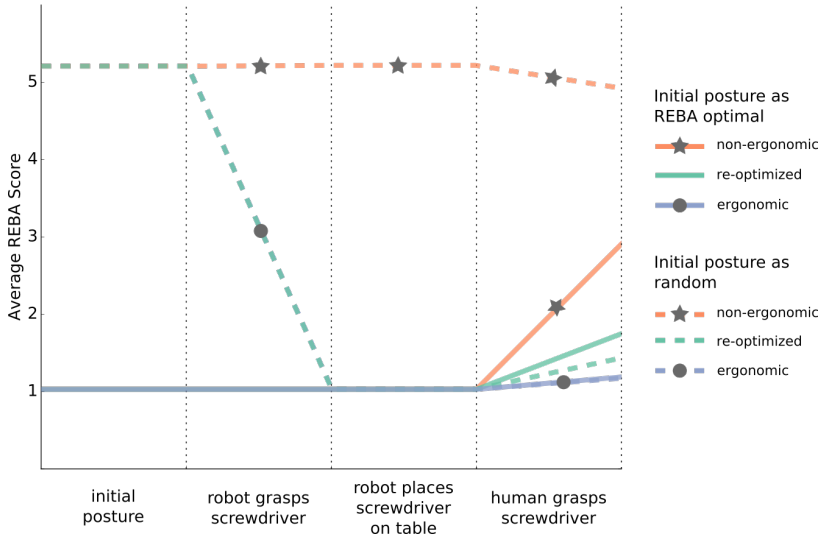

Figure 7.6: Evolution of the averaged REBA score during all the phases of the interaction. Plain lines) Initial posture optimal in terms of REBA score. The human starts moving only to grasp the screwdriver which explain the flat score until the last phase. Dashed lines) Initial posutre chosen as random. In the *ergonomic* and *re-optimized* the human moves after the initial phase to get back to an ergonomic posture.

Discussion

From the results of this simulated experiment we conclude that optimizing the sequence of actions for ergonomics leads to safer body posture than only optimizing actions independently. We can also conclude that optimizing the task allocation and motion planning for dREBA induces lower REBA score as expected.

Another interesting result is that considering an ergonomic cost for the body posture allows to remove the effects of the initial body configuration chosen prior to the optimization.

In next section, we introduce a more complex scenario with a longer sequence of actions. As a proof of concept we also implement the generated path and actions on a Baxter humanoid robot.

Box Assembly Experiment

The second experiment we consider is the assembly of a toolbox. We limit the assembly to screwing the handle (`/toolbox/handle`) and one side of the toolbox (`/toolbox/side_left`) together. We extend the set of possible actions defined in Section 7.1 with the following decisions:

- `handover(t,e1,o,e2)`
  - *description*: : at time `t`, the end-effector `e1` (robot) hands over object `o` to the end-effector `e2` (agent)
  - *precondition*: `grasped(e1,o) free(t,e2)`
  - *effect*: `¬grasped(e1,o) free(e1) grasped(e2,o) ¬free(e2)`

- `hold(t,e,o)`
  - *description*: : at time `t`, the end-effector `e` (robot or agent) holds object `o`
  - *precondition*: `on_table(o) free(e)`
  - *effect*: `held(o) ¬free(e)`
- `screw(t,e,o1,o2)`
  - *description*: : at time `t`, the end-effector `e` (agent) screws objects `o1` and `o2` together
  - *precondition*: `on_table(o1) on_table(o2) free(e) held(o1)`
  - *effect*: `screwed(o1,o2) ¬held(o1)`

Initially the `place` action can be accomplished by both the robot and the agent. However, in our setup the robot is equipped with a vacuum gripper on its left end-effector. This type of gripper is convenient to grasp the toolbox parts but cannot be used to precisely place them vertically on the table. Therefore, we limit the `place` decision to only the agent.

To each of those logic decisions we associate a geometric equivalent that defines constraints between objects as kinematic switches. We set the desired goal to `screwed(/toolbox/handle, /toolbox/side_left)` and perform the path optimization. The following is an example of a decision sequence found by the MBTS solver that leads toward the set goal.

```
grasp(1,baxterL,/toolbox/handle)

handover(2,baxterL,/toolbox/handle,handR)

grasp(3,baxterL,/toolbox/side_left)

place(4,handR,/toolbox/handle,tableC)

handover(5,baxterL,/toolbox/side_left,handR)

place(6,handR,/toolbox/side_left,tableC)

hold(7,baxterR,/toolbox/handle)

screw(8,handR,/toolbox/side_left,/toolbox/handle)
```

The solver also returns a path for both the robot and the agent in terms of joint trajectories and a simulated view of the scenario.

Real robot application

After planning the sequence of actions for the toolbox assembly we convert it into actions to be performed on our Baxter torso humanoid robot. The system also uses a set of predicates to recognize when an action is finished prior to start the next one. Predicates and actions were reused from previous work (Toussaint, Munzer, et al. 2016).

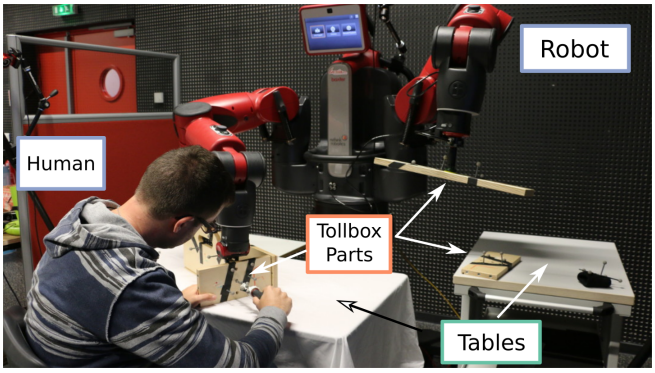

Figure 7.7: Picture of the experimental setup of the toolbox assembly annotated using the same convention as Figure 7.1.

At the moment, the system supposes that the agent is do not have the capacity to move freely in the environement, i.e. he or she cannot walk during the execution to avoid re-planning. As in the simulation experiment, the solver planify the motion of both the robot and human agents. However we cannot predict exactly how the real human coworker will act. Therefore, some actions like grasping or handing over toolbox parts are converted directly from the geometric part of the simulation, i.e trajectory of the robot arm to hand over an object are calculated by the solver. For actions following an agent intervention, such as holding a part placed on the table by the agent, we cannot rely on this approach. As the agent might decide to put the object on the table at a different location than the one planned, the trajectory to hold the object might differ. Therefore we track the objects parts using Optitrack motion capture system and generate a new trajectory of the robot arm to the current location of the object. With this approach we do not need to re-plan the full sequence of actions every time the geometry of the scene slightly varies. This greedy approach might, however, be sub-optimal in some cases. Indeed variation of the geometry could lead to changes in the sequence of actions when a full re-planning is performed.

As in the simulation experiment, we calculate the plan in the *ergonomic* and *non-ergonomic* conditions. In this case, we do not consider the *re-optimized* condition as it produces the same sequence of actions as the *non-ergonomic* one. Between the two tested conditions, not only the trajectory of the handing over is affected but also the planned sequence of actions. Changes appear on the first four actions. In the *non-ergonomic* condition the agent is asked to place the handle on the table only after the robot has grasped the next part:

```
grasp(1,baxterL,/toolbox/handle)

handover(2,baxterL,/toolbox/handle,handR)

grasp(3,baxterL,/toolbox/side_left)

place(4,handR,/toolbox/handle,tableC)
```

This creates an uncomfortable situation where the agent holds the handle and is unnecessary waiting for his or her next action. Most likely he or she will place it on the table without being told to do so which might create some confusion. In the *ergonomic* condition the sequence is smoother as a hand over is immediately followed by a place action:

```
grasp(1,baxterL,/toolbox/handle)

handover(2,baxterL,/toolbox/handle,handR)

place(3,handR,/toolbox/handle,tableC)

grasp(4,baxterL,/toolbox/side_left)
```

This change of the sequence of action is due to the weight of the object being comprised in the calculation of the cost function (6.1). When holding an object, the cost is slightly higher. Therefore, as the solver minimizes costs over time, placing the object earlier on the table leads to a smaller cost. The rest of the sequence of actions is similar in both conditions[2].

[2] A video of the experiment is available on our Vimeo channel https://vimeo.com/232348427.

As the planned sequence of actions differs between the two conditions, we need a proper way to inform the user about the actions he or she is required to perform. To this extent, we have generated a graphical interface from the results of the of the LGP solver.

Visualization of Task Allocation

In human-robot interaction, communication between the robot and the agent is crucial. For the setup we have considered, the communication is limited to inform the agent about the steps to perform the whole task. If one step is awaiting for a worker's action, it should also detail how this action is performed.

To visualize the planned sequence of actions we generate a webpage with a human readable description of each actions and a simulated view of the last frame of the action path as illustrated in Figure 7.8. This webpage is based on *reveal.js* presentation framework[3]. By clicking on the image, the agent launches a video of the full action path. We use the *reveal.js* api to automatically switch to the current action when the previous action is finished. With this web based approach, we can display this interface on mobile devices (tablets or smartphones) or computer screens.

[3] reveal.js is an HTML based presentation framework that exposes a javascript api to interact with the slides *Reveal.js The HTML Presentation Framework* 2017.

The webpage is accessible offline for the agent to navigate through the sequence of actions. The offline version can be seen as a training solution to understand the steps required for a given task. It can also be displayed online to show the current action to reduce the human cognitive load of understanding the action performed by the robot and understanding the action that he or her is expected to perform.

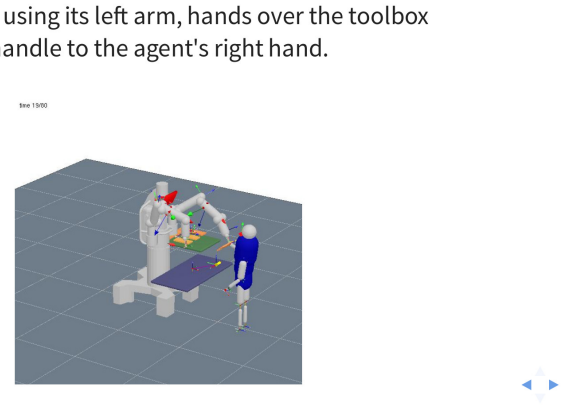

Figure 7.8: A representation of an action generated by the solver and displayed on a webpage to simplify the visualization. The agent can click on the image to start a video of the action. Arrows on the right corner are also clickable to navigate between the previous and the next actions.

Thanks to this display interface we are now able to evaluate our approach by means of a user study with subjects interacting with our Baxter robot.

User study

We apply the calculated plans in a user study with 10 participants. Each subject has performed the assembly in the two conditions and we have recorded his or her posture as in (Busch, Maeda, Mollard, et al. 2017). Postures score are then calculated and averaged over all subjects and over all the timesteps of the interaction.

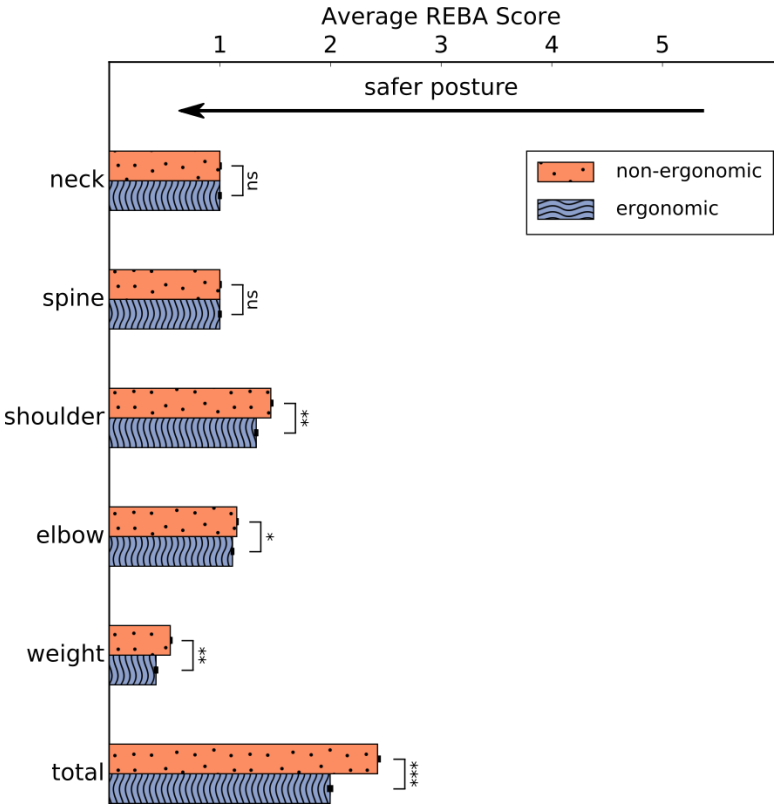

Figure 7.9: Average REBA score and standard error of the mean for the two conditions recorded from subjects of the user study.

Results are presented in Figure 7.9 and show that the *ergonomic* condition induces a lower REBA score on average over all the participants. This reduced score is partially explained by the lower *weight* score as objects were held by the subjects for a shorter amount of time. Nevertheless, the impact on the *shoulder* and *elbow* is also significant which indicates that the objects were handed over at a better position.

Discussion

The results of this experiment confirm the two points made with the simulation results. First, optimizing for dREBA correctly leads to lower REBA score, even in a real robot application. Second, in order to generate ergonomic human-robot interaction, considering single actions optimized with ergonomic measures, e.g. ergonomically planning a hand over, is not sufficient. Reasoning in term of sequence of actions leads to more ergonomic situations and improves the overall human comfort.

Including ergonomics in the LGP formalism allows us to reason both in terms of motions and task allocation. As the results suggests, optimizing for task allocation leads to better ergonomics than motion only. This approach also resolves the two limitations raised in Chapter 6. By using the LGP formalism, we optimize simultaneously the motions for both the human and the robot and we also optimize over a full trajectory rather than a static pose.

This Chapter concludes our work on the ergonomics of the interaction and our contributions for this thesis. In the next Chapters, we make a final review of our contributions and share some ideas on the next steps that should be taken for a more ergonomic human-robot interaction.

8
Conclusion

Throughout this thesis we have analyzed the impact of the robot's behavior and motions on the human coworker and improved the ergonomics of the interaction. As this impact can take a large number of aspects, we have limited our research to two specific scenarios. In a first part spanning Chapter 2 to 4 we have studied the legibility of the robot motions and provided solutions to learn the concept of legibility through interaction with real user. The second part, from Chapter 5 to 7, was focusing more on physical human-robot interaction and how to include physical ergonomics in the choices of actions and planned motions.

This chapter concludes our thesis on leveraging ergonomics in human-robot interaction. By studying two concrete problems that are legibility of the robot motions and physical ergonomics of the interaction we believe we have made a step towards a more human centered and ergonomic interaction. In accordance with the philosophy of the 3rdHand project we have proposed methods that are generic, do not require extra programming to be adapted to new scenarios and could be deployed on any robot, independently of its design. This philosophy is also depicted in our choice of methods and the way we have implemented them.

Our research on legibility shows that it can be obtained in a model-free approach. As any particular task will require different properties of motion, we wanted to achieve such results without any task-specific model of legibility. To such end we took an approach where we have defined a task-independent cost function that rewards efficiency (joint execution time), robustness (task errors), and energy (jerk). These measures can be readily defined for any task. To optimize such cost function through experiment we have relied on a model-free optimization algorithm, PI^{BB}, to efficiently optimize this cost function through trial-and-error interaction of the robot with the human.

In several human-robot experiments, we have shown that indeed, for different types of motions, robots are able to improve their behavior allowing humans to better read the robot's intentions early and robustly. Our results show that people, even after being habituated to robotic motions, can still substantially improve their prediction times if the robot optimizes its motions.

A second conclusion is that, when optimizing with policies that have a high-dimensional parameter vector (which leads to a lot of variance in the types of motions it can generate, such as with the DMP), it is most likely that idiosyncratic behavior arises. Novel subjects can infer the intention of the robot from its behavior, but this requires an extended phase of interaction with the robot. These interactions are necessary for the novel subject to get to know the specific idiosyncrasies the robot has learned with the previous subject.

Furthermore, the robot is still able to learn legible behavior, even if we actively suppress idiosyncratic behavior by allowing only stereotypical curved minimum jerk movements. Already during first interactions, novel subjects are able to read such behavior more efficiently and robustly than when using the DMP policy. This indicates that this behavior is immediately and more generally legible.

Considering a model-free approach to increase the legibility of the robot motions ensure that the robot can adapt to the specificities of any workers or task without having to program new models. This is a very important feature that could help the deployment of robotic structures in small and middle sized assembly lines, which require adaptations to task and production changes.

Our work on physical ergonomics have highlighted how we can optimize the motions and task allocation for maximum postural safety on the human side. By incorporating a simple evaluation function, such as the REBA method, into a motion planner, we are able to generate robotic behaviors that are safer in the long-term and better correspond to the users' expectations.

By the means of several user studies we have proven that our solution leads to safer working postures and is more appreciated by the end users. We also show that, by using the LGP formalism, we are able to dynamically assign tasks to the human or the robot taking into account the ergonomics of the situation.

Although in our work on physical ergonomics we have analyzed the usage of a single evaluation function, the REBA method, the latter could be replaced by other type of ergonomics evaluation functions and achieve similar results. Another aspect of our solution for physical ergonomics lies in the fact that, by optimizing the human motions for ergonomics, we also achieve realistic human posture, without the need of a complex model. This was not the initial focus of our approach, yet it is an interesting contribution as generating human-like motions is sometimes not trivial.

We believe integrating ergonomics and human factors in the human-robot interaction is a major step toward improving the working conditions in industry. As the cobotic systems will start entering workstations, we have here a wonderful opportunity to be proactive and provide solutions in favor of the workers' health and security. This would be profitable also for the industries for which work related injuries are very costly[1].

Even outside of the cobot scenario, our approach is also suitable to be used to optimize existing workstations for maximum ergonomics. It might even be considered to adapt workstations to disable individuals that require specific needs to continue working in good conditions. A line of research or a company could specialize on those questions to provide adaptive workstations on demand.

Finally, the solution we propose might have opportunities to serve not only the robotic community but also ergonomic researches. An interesting foreseen application could be, for example, to see robotic setups as an experimental platform to validate ergonomic functions, by the mean of controlled and validated interaction scenarios.

Although the solutions we have deployed have successfully leaded to a more ergonomic human-robot interaction, there are still a few issues that need to be addressed. In Chapter 9 we detail the main limitations of our approach, the relevant points that need to be clarified and elements that, we believe, represent interesting lines for future research.

[1] In 2012–13, work-related injury and disease cost the Australian economy 61.8 billion dollars, representing 4.1% of GDP (*The Cost of Work-related Injury and Illness for Australian Employers, Workers and the Community: 2012–13* 2015).

9
Limitations and Perspectives

Contents

In this Chapter, we provide the main limitations of our solutions and in the way we have implemented them. We also provide additional details on lines for future research that, we believe, could be of particular interest.

In a first section, we will focus on limitations linked to the legibility part of the thesis. Then we will do the same for the elements on physical ergonomics. Finally, we detail our ideas on combining the two approaches in a unified solution.

Legibility of Robotic Motions

Although we have successfully improved the legibility of the robot motion, there are still some questions that remain unanswered. In the next sections, we detail, point by point, the different elements that require a more thorough analysis and could lead to potential research opportunities.

Universally legible motions

While we have aimed at approaching universal legibility can we confidently conclude that we have have such a goal? In other words, are the trajectories generated via the viapoint policy *universally* legible, i.e. across different robots or human cultures? From our results, it seems that they do perform better, compared to straight lines to the target or DMP based trajectories, in terms of prediction time and errors. However, this is not sufficient to conclude on their potential *universality*. Although prediction time is a good indicator of legibility, there might be other factors that explain its variation. In general, we expect that the transition from idiosyncratic to universally legible behavior may not always be that well defined.

As en example, in our studies, we have limited the number of possible targets. With only two targets, learning idiosyncratic motions is facilitated as you effectively need to focus on learning the trajectory which goes to one of the two targets. Any motion that does not resemble the one you have learned is, obviously, going to the second target. Increasing the number of targets should reduce that effect and probably forces the emergence of more *universally* legible motions. It could also happen that, by increasing the number of targets, direct straight lines to the target would become the most legible motions. This point is the conclusion proposed by Zhao et al. (2016) in their study of legibility with multiple targets, although they only compare model based legible trajectories (Dragan and Siddhartha S. Srinivasa 2013) with straight lines. There might be other type of more legible motions that could only be observed by optimizing for legibility in a model free manner. However, one of the main drawback of our model-free approach is that, in its current definition, it does not scale well with the number of targets. Adding more targets would require a longer optimization phase as the motions to each targets are optimized independently.

Psychological effects

Although our method successfully lead to an increased legibility it is not necessarily linked to an increased comfort of the human coworker. Here we have made the assumption that increasing the robot legibility would reduce the potential prediction errors and, therefore, avoid stressful situations. However, one question that have not been validated by our studies is *how do the subjects perceive those legible motions?* Do they feel safer when interacting with a legible robot, or conversely, more afraid to interact with it? Analyzing this question would require answering surveys, similarly to those we have presented in our work on ergonomics. We could also rely on physiological factors such as cardiac rhythm. An interesting point could be to also analyze the differences between the two policies with respect to this question.

Moreover, when working with real humans we also have to consider that some psychological effects can interfere. For example, at the beginning of the task some subjects might wait for more confidence instead of trying to guess and potentially make mistakes. This fear of making mistakes, which relates to the notion of perfectionism (Frost et al. 1990) could be a major impediment to our method. It is probable that the observers already know the robot's intentions way before they actually press the corresponding button. To counteract this effect, we would need to rely on other factors than their voluntary feedback. One possible option would be too look into brain signals or other physiological factors. Pupil dilation, for example, is a known factor to be correlated with the decision making process (Einhäuser et al. 2010; Preuschoff et al. 2011). We could imagine a setup that records the pupil dilation and uses it as a marker for the subjects' prediction time.

Including social cues

In our research, we have improved the robot legibility by focusing only on the motion to specific targets. This is a very limited subset of all the available social cues[1] that can be used to express intentions. Gaze and proxemics are potential candidates (Fiore et al. 2013), so is execution speed of the motion (Zhou et al. 2017). All of those modalities are signals that could be learned to be triggered in a model-free manner, similarly to an infant that learns through trials and errors how to use those signals.

[1] A social cue can either be a verbal or non-verbal hint, which can be positive or negative. These cues guide conversation and other social interactions. Common social cues in robotics are gaze, i.e using fake eyes to look at the person or proxemics.

The learning of those social cues would, however, require a lot more of interactions to be fully efficient. One point is that, most probably, some of those signals are not task dependent and span between different tasks. Therefore, this type of research would perfectly fit in a life-long learning paradigm.

Inferring social intentions

In the context of our experiments we have improved the legibility of the robot motion to get a better understanding of which target it is aiming at. One question that is not yet answered is *why is it doing this action?*. This amounts to saying *can we infer the robot social intention from its motion?*.

As an example, imagine that the robot is aiming for a bottle of water on the table. Only by observing its motions would you be able to understand that it wants to tidy up the table or that it is going to poor you a glass of water[2]?

[2] This question is still a controversy–even in human-human interaction–and concerns the role of the motor system in understanding social intentions (Beer et al. 2006; Jacob et al. 2005).

As we have seen briefly in the Chapter 2, the *mirror neurons* allow us to recognize another person intention by mirroring the motion onto ourself. The question of deriving the social intention from the same neural mechanism raises more skepticism which led Jacob et al. (2005) to propose the following thought experiment:

> Consider Dr. Jekyll and Mr. Hyde. The former is a renowned surgeon who performs appendectomies on his anesthetized patients. The latter is a dangerous sadist who performs exactly the same hand movements on his non anesthetized victims. As it turns out, Mr. Hyde is Dr. Jekyll. Suppose that Dr. Watson witnesses both Dr. Jekyll's and Mr. Hyde's actions. Upon perceiving Dr. Jekyll, alias Mr. Hyde, execute the same motor sequence twice, whereby he grasps his scalpel and applies it to the same bodily part of two different persons, presumably the very same mirror neurons produce the same discharge in Dr. Watson's brain. Dr. Jekyll's motor intention is the same as Mr. Hyde's. However, Dr. Jekyll's social intention clearly differs from Mr. Hyde's: whereas Dr. Jekyll intends to improve his patient's medical condition, Mr. Hyde intends to derive pleasure from his victim's agony.

Their conclusion is that simulating the agent's movements, by the mean of the mirror cells, allows an observer to understand the agent's motors intention but not necessarily his or her social intention. However, that conclusion relies on the strong assumption that *the way you achieve actions is independent of the social context*, i.e. whatever your goal is to pour me some water or to move the bottle away, you will reach for it in the exact same way.

Becchio et al. (2008) have conducted a user study that seems to discard this hypothesis, by recording the motions of people grasping a bottle to either pass it to a partner or put it in a box. Both partner's hand for the "social condition" and box for the other one were located at the exact same place to avoid differences in motions due to the end-point being different. The recorded motions when grasping the bottle were showing different patterns between the two conditions. This suggests that, unconsciously, we modify our motions to express the sociality of our actions, in similar ways that we would modify our motions to express more legibility. The relatively small number of participants in the study ($n = 13$) forces us to take the results with caution, but this might be a good starting point for a similar study within a robotic setup. We could imagine a scenario very similar to the experiments we have led where we improve the expression of social intentions from the robot motion.

Physically Ergonomic Human-Robot Interaction

In this domain there is also a lot of potentials research opportunities and we will detail in the following sections the points that need improvement or could be the most interesting to focus on.

Human tracking

Our solution for tracking the human worker is based on OptiTrack markers assembled to form a suit. Highly intrusive, this solution is not adapted for an integration in an industrial setup. Moreover, it requires to buy an entire OptiTrack setup which comes at a high cost for small industrial companies. There are, however, other solutions that start emerging.

Obviously, Kinect skeleton tracking is the first solution that comes in mind as it is easy to set up and uses a very affordable sensor. Before moving to an Optitrack based solution, we had based our tracking on the Kinect. The main drawback of the sensors is that it requires the subject to be facing the camera at all time and to be located between 1 to 3 meters in front of the sensor. In practice, this is rather limiting as there might be situations where the worker would not face the camera. One solution could be to use an array of Kinect camera, spanning the whole workspace, and to fuse the skeleton data coming from all the sensors (Yeung et al. 2013). However, it is worth noting that on the 25th of October 2017, Microsoft have issued a statement indicating the end of production of their Kinect sensor[3]. Although the technology, by itself, will probably continues to live on, it is unclear how this will affect the ecosystem. As the source code for the skeleton tracking system is proprietary it will probably be difficult to replace existing solutions based on the Kinect sensor with other similar depth sensors.

[3] The Telegraph has written an article on the subject http://www.telegraph.co.uk/gaming/news/microsoft-shutting-kinect-production/

Therefore, we believe it is necessary to find other solutions to track the subjects reliably, and if possible open-source ones. The openpose library, based on simple RGB flux could be a very interesting alternative and fit all the criterion. Developed by the CMU-Perceptual-Computing-Lab, it is a fully open-source skeleton tracking system based on deep learning and trained on the Coco dataset (Cao et al. 2017). Even if it requires to set up an array of calibrated cameras for 3D reconstruction, it is still a lot cheaper than an OptiTrack installation and do not need the usage of intrusive markers. A code for 3D reconstruction based on several camera is available on their GitHub repository.

Another option, that we thought implementing at some point, would be to create a suit based on inertial sensors[4]. Although still intrusive, inertial sensors could be integrated in the workers' uniforms as it does not require to be visible at all time. As this is not a camera based solution it is also not liable to occlusions and could be used to track workers over a large workspace such as big warehouse.

[4] Inertial sensors measure linear and angular motions usually with a triad of gyroscopes and triad of accelerometers.

Dynamic model

In our studies, we have considered only static poses which do not account for the dynamics of the human body. MSDs are also highly correlated with the force exercised during an action. With greater forces comes a higher risk (Gallagher et al. 2013). In order to study this effect and to consider dynamics in the interaction, we would need to change for a much more complex humanoid model. The Opensim simulator could be a great tool to achieve such a complex model[5]. Simulating the direct dynamic model is a complex task that might not be trivial to insert in the LGP formalism as it is not, per se, a simulation environment[6].

[5] Opensim is an open-source simulator for modeling and analyzing the dynamics of human motions (Delp et al. 2007).

[6] Gravity and external forces are not yet considered. A workaround could be to reason in term of objects velocities to extrapolate the forces.

Estimating inverse dynamic model, however, can be achieved from observations of the workers (Plantard et al. 2017). This is an interesting approach to provide feedback about the risk associated with a particular motion and help preventing musculoskeletal injuries (Pontonnier et al. 2014; Rasmussen et al. 2003).

Ergonomic role switching

An interesting aspect of our approach is the possibility to introduce a role switching based on the ergonomics of a situation. For example if the task requires to lift an heavy box, the robot should perform this action as this is a very burdensome task for the human coworker. Generally speaking, all burdensome tasks should be devoted to the robot, leaving to the human the tasks with great added values, e.g. tasks that require a certain dexterity.

The choice of experiments we have considered do not really highlight this capacity. For now, most of the actions we have designed are agent dependent, i.e. both the robot and human have a set of action they can accomplish and there is not much overlap between them. This is mainly due to the geometric definition of the agents that are quite different. For example, the `grasp` action would introduce different geometric constraints depending if it is performed by the robot or the human. As the LGP solver, in its current implementation, do not allow for a switching between two possible geometric implementation, a workaround is to duplicate the high-level actions with a different implementation for both agents.

One drawback of this ergonomic role switching arises when both agents can accomplish all the necessary actions. As the human coworker presents an extra cost based on ergonomics, the optimal solution would be to leave all the actions to the robot. Obviously, this is not a desired behavior. One way to overcome this issue could be to also optimize for time execution. Introducing a cost based on the time it takes to perform the assembly, under a given sequence of actions, would force the human to actually perform some actions as two agents are faster than a single one. Therefore, there is a trade-off to consider between execution time and ergonomics.

Automatic selection of the optimization weights

Our solutions to maximize physical ergonomics comprise a multi-criteria optimization. To optimize this kind of problem, we consider weighting the different cost functions with respect to their importance to solve a task. For example, posture score need to cost less than task constraints otherwise the human model would simply stay in a resting position throughout the assembly. For non-trained workers, visibility could also be a crucial issue and the associated cost should be increased compared to a similar situation with experienced people.

The problem is how to select appropriately the weightings factors to account for those situations. At the moment, those weighting factors are empirically selected and might be different for each experiments. An automatic weighting selection is required and should address the following issues. It should balance the costs of each components based on their importance at certain moment of the interaction. It should also adapt to the specificity of the tasks, e.g. increasing the safety cost for tasks of high risk. Finally it should be adapted to the current coworker and to its different needs, such as particular disabilities or expertise with the given task.

Online replanning

At the moment, the LGP solver, on which we have based our work, takes a non-negligible amount of time to provide a solution. The solver do not use a greedy approach, i.e. the optimal solution is calculated over the full sequence of actions. One problem arises when the human coworker decides not to follow the calculated plan.

This situation can take two aspects. First, he or she can decide to accomplish the correct high-level action, but with a change in the geometry, for example, by moving to another location in the workspace. To overcome this, we have implemented a greedy approach where we replan only the trajectories with respect to the new locations. This greedy approach, however, presents some drawbacks. In some cases, the new geometry could render the next actions infeasible or leads to suboptimal sequences of actions. As an example, consider a situation where the human coworker is standing too far from the screwdriver. The robot actions would be to pick it up and hand it over. If the human moves closer to the screwdriver, the robot actions are not necessary anymore and, if it performs them anyway, it would create an odd situation.

This situation also highlights the second aspect of the necessity for replanning. By moving closer to the screwdriver, the human would probably decide to pick it up by him or herself. Therefore, this introduces a change in the high-level sequences of actions. Our current implementation do not allow for this possibility. The only solution to solve this issue would be to re-solve the full TAMP problem starting from the new geometric configuration and the current logic state. As this calculation is rather time consuming, the robot would be ineffective while replanning. Therefore, there is a strong need here for an online replanning solution.

Human preferences versus ergonomic recommendations

In Chapter 5 we have briefly talked about the *automobile seat designer's dilemma* (Reed et al. 1991) showing that people might have strong preferences which could enter in conflict with ergonomic recommendations. This can be an important issue as it is linked with a decrease in long-term comfort.

There is here a crucial question which is *should we enforce behaviors that go against people preferences if it is for their own safety?* In our studies, this dilemma did not appear as the preferred behavior was the one in accordance with ergonomic principles. However, this problem arises in industry during ergonomic auditing. An industrial worker might decide not to follow the guidelines simply because it requires a change in the way he or she is performing actions. Those changes might also impact the efficiency, as well, as consciously changing the way you act impacts your execution speed (MacKay 1982).

To overcome this effect, it might be interesting to provide the workers an incentive to follow the recommended guidelines and to help them during the transition process from their old habits. For example, we could rely on gamification by using scores as en encouragement for improvement. Such a process as been successfully used in cars to reduce drivers' distraction (Xie et al. 2016).

Combining Legibility and Physical Ergonomics

In our research we have considered optimizing the robot motions for maximum legibility and maximum physical ergonomics on the human side. However, both aspects are currently segregated and not optimized simultaneously. The main problem to include both of them, in a general solution, is that the approaches we have used are fundamentally different. Our solution for optimizing legible motions is *model-free* when our algorithm for ergonomics is *model-based*. Including the cost function for maximizing legibility in the LGP formalism is currently not a feasible solution as it would require to either simulate an observer or including external rewards, from real user interaction, in the solver. In the first case we would move from a *model-free* to a *model-based* approach with all the drawbacks associated. In the second case, the number of trials required to solve legibility and TAMP problems at the same time would be far too extensive to make it a viable solution

From a logical point of view, however, there are few overlaps between legible and physically ergonomic motions. Legibility of the motion is mainly located at the trajectory level from a starting configuration to a desired goal, while the physical ergonomics is mainly impacted by the final configuration of the robot. In this case, there should not be any trade-off to make between legibility and physical ergonomics. Therefore, a possible solution is simply to optimize for the two separately. For example, we could use the trajectories generated by the LGP solver as initial trajectories to be optimized for legibility, when interacting with real users. Because our *model-free* solution to optimize for legibility do not make any modifications on the robot end-goal, this ensure that physical ergonomics is conserved during the optimization process.

Bibliography

Akella, Prasad et al. (1999). "Cobots for the Automobile Assembly Line". In: *Robotics and Automation, 1999. Proceedings. 1999 IEEE International Conference on*. Vol. 1. IEEE, pp. 728–733.

Alami, Rachid et al. (2006). "Toward Human-Aware Robot Task Planning". In: *AAAI Spring Symposium: To Boldly Go Where No Human-Robot Team Has Gone Before*, pp. 39–46.

Aleotti, Jacopo, Vincenzo Micelli, and Stefano Caselli (2012). "Comfortable Robot to Human Object Hand-Over". In: *Proc. IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN'12))*, pp. 771–776.

Ambrose, Stanley H. (2001). "Paleolithic technology and human evolution". In: *Science* 291.5509, pp. 1748–1753.

Baum, L. Frank (1900). *The Wonderful Wizard of Oz*. George M. Hill Company.

Becchio, Cristina et al. (2008). "The case of Dr. Jekyll and Mr. Hyde: a kinematic study on social intention". In: *Consciousness and Cognition* 17.3, pp. 557–564.

Beek, Allard J. van der and MH Frings-Dresen (1998). "Assessment of Mechanical Exposure in Ergonomic Epidemiology". In: *Occupational and environmental medicine* 55.5, pp. 291–299.

Beer, Jennifer S. and Kevin N. Ochsner (2006). "Social Cognition: a Multi Level Analysis". In: *Brain research* 1079.1, pp. 98–105.

Bestick, Aaron M et al. (2015). "Personalized Kinematics for Human-Robot Collaborative Manipulation". In: *Proc. IEEE International Conference on Intelligent Robots and Systems (IROS'15)*, pp. 1037–1044.

Bishop, Christopher M. (2006). *Pattern Recognition and Machine Learning*. springer.

Bjorklund, David F. and Katherine K. Harnishfeger (1995). "The evolution of inhibition mechanisms and their role in human cognition and behavior." In:

Bley, H. et al. (2004). "Appropriate Human Involvement in Assembly and Disassembly". In: *CIRP Annals-Manufacturing Technology* 53.2, pp. 487–509.

Bowles, Samuel and Herbert Gintis (2011). *A Cooperative Species: Human Reciprocity and its Evolution*. Princeton University Press.

Breazeal, Cynthia L. (2004). *Designing Sociable Robots*. MIT press.

Brockman, Roberta, Raoul Tubiana, and Philippe Chamagne (1992). "Anatomic and Kinesiologic Considerations of Posture for Instrumental Musicians". In: *Journal of Hand Therapy* 5.2, pp. 61–64.

Brooke, John (1996). "SUS-A Quick and Dirty Usability Scale". In: *Usability Evaluation in Industry* 189.194, pp. 4–7.

Browne, Cameron B. et al. (2012). "A Survey of Monte Carlo Tree Search Methods". In: *IEEE Transactions on Computational Intelligence and AI in games* 4.1, pp. 1–43.

Busch, Baptiste, Jonathan Grizou, et al. (2017). "Learning Legible Motion from Human–Robot Interactions". In: *International Journal of Social Robotics*, pp. 1–15.

Busch, Baptiste, Guilherme G. Maeda, Yoan Mollard, et al. (2017). "Postural Optimization for an Ergonomic Human-Robot Interaction". In: *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'17)*.

Busch, Baptiste, Marc Toussaint, and Manuel Lopes (2018). "Planning Ergonomic Sequences of Actions in Human-Robot Interaction". In: *Proc. IEEE International Conference on Robotics and Automation (ICRA'18)*.

Busch, Baptiste et al. (2017a). *Postural Optimization Data Reader*. URL: https://github.com/3rdHand-project/PosturalOptimizationDataReader (visited on 02/22/2017).

– (2017b). *Postural Optimization for a Safe and Comfortable Human-Robot Interaction: Experiment Dataset*. DOI: 10.5281/zenodo.321599. URL: https://doi.org/10.5281/zenodo.321599.

– (2017c). *Postural Optimization for a Safe and Comfortable Human-Robot Interaction: Video*. URL: https://vimeo.com/163699896 (visited on 02/20/2017).

Cakmak, Maya, Siddhartha S Srinivasa, et al. (2011). "Human Preferences for Robot-Human Hand-Over Configurations". In: *Proc. IEEE International Conference on Intelligent Robots and Systems (IROS'11)*. San Francisco, California, USA, pp. 1986–1993.

Cakmak, Maya, Siddhartha S. Srinivasa, et al. (2011). "Using Spatial and Temporal Contrast for Fluent Robot-Human Hand-Overs". In: *Proc. IEEE/ACM International Conference on Human-Robot Interaction (HRI'11)*, pp. 489–496.

Cao, Zhe et al. (2017). "Realtime Multi-Person 2D Pose Estimation using Part Affinity Fields". In: *CVPR*.

Clodic, Aurélie et al. (2009). "Shary: a supervision system adapted to human-robot interaction". In: *Experimental Robotics*. Springer, pp. 229–238.

Colgate, J. Edward and Michael A. Peshkin (1999). *Cobots*. US Patent 5,952,796. URL: https://www.google.com/patents/US5952796.

Colgate, J. Edward, Michael A. Peshkin, and Witaya Wannasuphoprasit (1996). "Cobots: Robots for Collaboration with Human Operators". In:

Craig, John J. (2005). *Introduction to Robotics: Mechanics and Control*. Upper Saddle River, New Jersey, USA: Prentice Hall.

Cross, Emily S. et al. (2012). "Robotic Movement Preferentially Engages the Action Observation Network". In: *Human Brain Mapping* 33.9, pp. 2238–2254.

Darwin, Charles (1888). *The Descent of Man and Selection in Relation to Sex*. Vol. 1. Murray.

Das, Biman and Arijit K. Sengupta (1996). "Industrial Workstation Design: a Systematic Ergonomics Approach". In: *Applied Ergonomics* 27.3, pp. 157–163.

Dayan, Peter and Laurence F. Abbott (2001). *Theoretical Neuroscience*. Vol. 806. Cambridge, MA: MIT Press.

De Graaf, Maartje, Somaya Ben Allouch, and Jan Van Dijk (2017). "Why Do They Refuse to Use My Robot?: Reasons for Non-Use Derived from a Long-Term Home Study". In: *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*, pp. 224–233.

De Santis, Agostino et al. (2008). "An Atlas of Physical Human-Robot Interaction". In: *Mechanism and Machine Theory* 43.3, pp. 253–270.

Decety, Jean and Julie Grèzes (2006). "The Power of Simulation: Imagining One's Own and Other's Behavior". In: *Brain research* 1079.1, pp. 4–14.

Dehais, Frédéric et al. (2011). "Physiological and Subjective Evaluation of a Human-Robot Object Hand-Over Task". In: *Applied ergonomics* 42.6, pp. 785–791.

Delp, Scott L. et al. (2007). "OpenSim: Open-Source Software to Create and Analyze Dynamic Simulations of Movement". In: *IEEE transactions on biomedical engineering* 54.11, pp. 1940–1950.

Desa, Auni Fatin Nadia Chiek et al. (2013). "OHSAS 18001 Performances Efforts and OSHA Performance in Malaysian Automotive Industry". In: *Journal of Environmental Science* 2.2, pp. 217–227.

Diankov, Rosen and James Kuffner (2008). "Openrave: A Planning Architecture for Autonomous Robotics". In: *Robotics Institute, Pittsburgh, PA, Tech. Rep. CMU-RI-TR-08-34* 79.

Diego-Mas, Jose Antonio and Jorge Alcaide-Marzal (2013). "Using Kinect Sensor in Observational Methods for Assessing Postures at Work". In: *Applied Ergonomics* 45.4, pp. 976–985.

Dragan, Anca D., Rachel M. Holladay, and Siddhartha S. Srinivasa (2014). "An Analysis of Deceptive Robot Motion". In: *Proc. Robotics: Science and Systems (R:SS'14)*, pp. 10–19.

Dragan, Anca D., Kenton C.T. Lee, and Siddhartha S. Srinivasa (2013). "Legibility and Predictability of Robot Motion". In: *Proc. IEEE/ACM International Conference on Human-Robot Interaction (HRI'13)*, pp. 301–308.

Dragan, Anca D. and Siddhartha S. Srinivasa (2013). "Generating Legible Motion". In:

Duke, Charles R. (1974). "Nonverbal Behavior and the Communication Process". In: *College Composition and Communication* 25.5, pp. 397–404.

Einhäuser, Wolfgang, Christof Koch, and Olivia L. Carter (2010). "Pupil Dilation Betrays the Timing of Decisions". In: *Frontiers in human neuroscience* 4.

Epley, Nicholas, Adam Waytz, and John T. Cacioppo (2007). "On Seeing Human: a Three-Factor Theory of Anthropomorphism." In: *Psychological Review* 114.4, p. 864.

Fiore, Stephen M. et al. (2013). "Toward Understanding Social Cues and Signals in Human-Robot Interaction: Effects of Robot Gaze and Proxemic Behavior". In: *Frontiers in psychology* 4.

Flash, Tamar and Neville Hogan (1985). "The Coordination of Arm Movements: an Experimentally Confirmed Mathematical Model". In: *Journal of Neuroscience* 5.7, pp. 1688–1703.

Foster, Mary Ellen and Colin Matheson (2008). "Following Assembly Plans in Cooperative, Task-Based Human-Robot Dialogue". In: *Proceedings of Londial* 2008.

Fredriksson, K. et al. (2001). "The impact on Musculoskeletal Disorders of Changing Physical and Psychosocial Work Environment Conditions in the Automobile Industry". In: *International journal of industrial ergonomics* 28.1, pp. 31–45.

Frost, Randy O. et al. (1990). "The Dimensions of Perfectionism". In: *Cognitive therapy and research* 14.5, pp. 449–468.

Gallagher, Sean and John R. Heberger (2013). "Examining the Interaction of Force and Repetition on Musculoskeletal Disorder Risk: a Systematic Literature Review". In: *Human factors* 55.1, pp. 108–124.

Gallese, Vittorio (2001). "The Shared Manifold Hypothesis. From Mirror Neurons to Empathy". In: *Journal of Consciousness Studies* 8.5-6, pp. 33–50.

Gallese, Vittorio and Alvin Goldman (1998). "Mirror Neurons and the Simulation Theory of Mind-Reading". In: *Trends in Cognitive Sciences* 2.12, pp. 493–501.

Garage, Willow (2017). *Unified Robot Description Format (URDF)*. URL: http://wiki.ros.org/urdf (visited on 10/11/2017).

Glasauer, Stefan et al. (2010). "Interacting in Time and Space: Investigating Human-Human and Human-Robot Joint Action". In: *Proc. IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN'10)*, pp. 252–257.

Gombolay, Matthew C. et al. (2015). "Decision-Making Authority, Team Efficiency and Human Worker Satisfaction in Mixed Human-Robot Teams". In: *Autonomous Robots* 39.3, pp. 293–312.

Granjean, Etienne (1969). *Fitting the Task to the Man: an Ergonomic Approach*. London, England: Taylor Francis ltd.

Hall, Edward T. (1963). "A System for the Notation of Proxemic Behavior". In: *American Anthropologist* 65.5, pp. 1003–1026.

Hall, Edward Twitchell (1966). *The Hidden Dimension*. Garden City, NY, USA: Doubleday.

Harrison, Randall (1965). "Nonverbal Communication: Exploration into Time, Space, Action, and Object". In: *Dimensions in communication*, pp. 158–174.

Heider, Fritz and Marianne Simmel (1944). "An Experimental Study of Apparent Behavior". In: *The American Journal of Psychology* 57.2, pp. 243–259.

Hignett, Sue and Lynn McAtamney (2000). "Rapid Entire Body Assessment (REBA)". In: *Applied Ergonomics* 31.2, pp. 201–205.

Hohwy, Jakob (2013). *The Predictive Mind*. Oxford University Press.

Ijspeert, Auke Jan, Jun Nakanishi, Heiko Hoffmann, et al. (2013). "Dynamical Movement Primitives: Learning Attractor Models for Motor Behaviors". In: *Neural Computation* 25.2, pp. 328–373.

Ijspeert, Auke Jan, Jun Nakanishi, and Stefan Schaal (2002). "Movement Imitation with Nonlinear Dynamical Systems in Humanoid robots". In: *Proc. IEEE International Conference on Robotics and Automation (ICRA'02)*.

Ikuta, Koji, Hideki Ishii, and Makoto Nokata (2003). "Safety Evaluation Method of Design and Control for Human-Care Robots". In: *The International Journal of Robotics Research* 22.5, pp. 281–297.

Robots and Robotic Devices – Safety Requirements for Industrial Robots – Part 1: Robots (2011). Standard. Geneva, CH: International Organization for Standardization.

Jacob, Pierre and Marc Jeannerod (2005). "The Motor Theory of Social Cognition: a Critique". In: *Trends in Cognitive Sciences* 9.1, pp. 21–25.

Johannsmeier, Lars and Sami Haddadin (2017). "A Hierarchical Human-Robot Interaction-Planning Framework for Task Allocation in Collaborative Industrial Assembly Processes". In: *IEEE Robotics and Automation Letters (RAL)* 2.1, pp. 41–48.

Kajita, Shuuji et al. (2009). *Introduction à la commande des robots humanoïdes: De la modélisation à la génération du mouvement*. Springer Science & Business Media.

Kaplan, Frédéric (2004). "Who is Afraid of the Humanoid? Investigating Cultural Differences in the Acceptance of Robots". In: *International Journal of Humanoid Robotics* 1.03, pp. 465–480.

Karhu, Osmo, Pekka Kansi, and Iikka Kuorinka (1977). "Correcting Working Postures in Industry: A Practical Method for Analysis". In: *Applied Ergonomics* 8.4, pp. 199–201.

Kee, Dohyung and Waldemar Karwowski (2001). "LUBA: An Assessment Technique for Postural Loading on the Upper Body Based on Joint Motion Discomfort and Maximum Holding Time". In: *Applied Ergonomics* 32.4, pp. 357–366.

Knepper, Ross A. et al. (2017). "Implicit Communication in a Joint Action". In: *Proc. IEEE/ACM International Conference on Human-Robot Interaction (HRI'17)*.

Koay, Kheng Lee et al. (2014). "Social Roles and Baseline Proxemic Preferences for a Domestic Service Robot". In: *International Journal of Social Robotics* 6.4, pp. 469–488.

Kober, Jens and Jan Peters (2011). "Policy search for motor primitives in robotics". In: *Machine Learning* 84.1, pp. 171–203.

Kolich, Mike (2003). "Automobile seat comfort: occupant preferences vs. anthropometric accommodation". In: *Applied ergonomics* 34.2, pp. 177–184.

Kölsch, Mathias, Andrew C Beall, and Matthew Turk (2003). "The Postural Comfort Zone for Reaching Gestures". In: *Proc. Human Factors and Ergonomics Society Annual Meeting*. Vol. 47. 4, pp. 787–791.

Krüger, Jörg, Terje K Lien, and Alexander Verl (2009). "Cooperation of Human and Machines in Assembly Lines". In: *CIRP Annals-Manufacturing Technology* 58.2, pp. 628–646.

Kubrick, Stanley (1968). *2001: a space odyssey*.

Lee, Min Kyung et al. (2011). "Predictability or Adaptivity?: Designing Robot Handoffs Modeled From Trained Dogs and People". In: *Proc. IEEE/ACM International Conference on Human-Robot Interaction (HRI'11)*, pp. 179–180.

Li, Guangyan and Peter Buckle (1999). "Current Techniques for Assessing Physical Exposure to Work-Related Musculoskeletal Risks, with Emphasis on Posture-Based Methods". In: *Ergonomics* 42.5, pp. 674–695.

Lichtenthäler, Christina, Tamara Lorenzy, and Alexandra Kirsch (2012). "Influence of Legibility on Perceived Safety in a Virtual Human-Robot Path Crossing Task". In: *Proc. IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN'12)*, pp. 676–681.

Lozano-Pérez, Tomás and Leslie Pack Kaelbling (2014). "A Constraint-Based Method for Solving Sequential Manipulation Planning Problems". In: *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'14)*, pp. 3684–3691.

MacKay, Donald G. (1982). "The problems of flexibility, fluency, and speed–accuracy trade-off in skilled behavior." In: *Psychological Review* 89.5, p. 483.

Maeda, Guilherme G. et al. (2016). "Acquiring and Generalizing the Embodiment Mapping From Human Observations to Robot Skills". In: *IEEE Robotics and Automation Letters* 1.2, pp. 784–791.

Mainprice, Jim et al. (2010). "Planning Safe and Legible Hand-Over Motions for Human-Robot Interaction". In: *IARP Workshop on Technical Challenges for Dependable Robots in Human Environments (DRHE'10)*, p. 7.

Medoff, Lynn E. (1999). "The Importance of Movement Education in the Training of Young". In: *Medical Problems of Performing Artists* 14.4, pp. 210–219.

Meerbeek, Bernt et al. (2008). "The Influence of Robot Personality on Perceived and Preferred Level of User Control". In: *Interaction Studies* 9.2, pp. 204–229.

Mitsunaga, Noriaki et al. (2008). "Adapting Robot Behavior for Human-Robot Interaction". In: *IEEE Transactions on Robotics* 24.4, pp. 911–916.

Mixamo. *Mixamo*. URL: https://www.mixamo.com/.

Mosar, E. et al. (1969). "Socété d'Ergonomie de Langue Française (SELF)". In: *Le Travail Humain* 32.3-4.

Nielsen, Jakob (1994). *Usability Engineering*. Amsterdam, Netherlands: Elsevier.

Ogorodnikova, Olessia (2008). "Human Weaknesses and Strengths in Collaboration with Robots". In: *Periodica Polytechnica. Engineering. Mechanical Engineering* 52.1, p. 25.

Pagello, Enrico et al. (1999). "Cooperative Behaviors in Multi-Robot Systems Through Implicit Communication". In: *Robotics and Autonomous Systems* 29.1, pp. 65–77.

Pandey, Amit Kumar and Rachid Alami (2014). "Towards Human-Level Semantics Understanding of Human-Centered Object Manipulation Tasks for HRI: Reasoning About Effect, Ability, Effort and Perspective Taking". In: *International Journal of Social Robotics* 6.4, pp. 593–620.

Papadopoulos, Georgios T., Apostolos Axenopoulos, and Petros Daras (2014). "Real-Time Skeleton-Tracking-Based Human Action Recognition Using Kinect Data". In: *MMM (1)*, pp. 473–483.

Plaku, Erion and Gregory D. Hager (2010). "Sampling-Based Motion and Symbolic Action Planning with Geometric and Differential Constraints". In: *Proc. IEEE International Conference on Robotics and Automation (ICRA'10)*.

Plantard, Pierre et al. (2017). "Inverse dynamics based on occlusion-resistant Kinect data: Is it usable for ergonomics?" In: *International Journal of Industrial Ergonomics* 61, pp. 71–80.

Pontonnier, Charles et al. (2014). "Strengths and limitations of a musculoskeletal model for an analysis of simulated meat cutting tasks". In: *Applied ergonomics* 45.3, pp. 592–600.

Premack, David and Guy Woodruff (1978). "Does the Chimpanzee Have a Theory of Mind?" In: *Behavioral and brain sciences* 1.4, pp. 515–526.

Preuschoff, Kerstin, Bernard Marius't Hart, and Wolfgang Einhäuser (2011). "Pupil Dilation Signals Surprise: Evidence for Noradrenaline's Role in Decision Making". In: *Frontiers in neuroscience* 5.

Punnett, Laura and David H. Wegman (2004). "Work-Related Musculoskeletal Disorders: The Epidemiologic Evidence and the Debate". In: *Journal of Electromyography and Kinesiology* 14.1, pp. 13–23.

Putz-Anderson, Vern et al. (1997). "Musculoskeletal Disorders and Workplace Factors". In: *National Institute for Occupational Safety and Health (NIOSH)* 104.

Rasmussen, John et al. (2003). "Anybody-a software system for ergonomic optimization". In: *Fifth World Congress on Structural and Multidisciplinary Optimization*. Vol. 4.

Reed, Matthew P. et al. (1991). *An investigation of driver discomfort and related seat design factors in extended-duration driving*. Tech. rep. SAE Technical paper.

Reeves, Byron and Clifford Nass (1996). "How People Treat Computers, Television, and New Media Like Real People and Places". In: *CSLI Publications and Cambridge*.

Rettberg, Jill Walker (2014). "Seeing ourselves through technology: How we use selfies, blogs and wearable devices to see and shape ourselves". In:

Reveal.js The HTML Presentation Framework (2017). URL: http://lab.hakim.se/reveal-js/#/ (visited on 08/28/2017).

Rietveld, ABM Boni (2013). "Dancers' and Musicians' Injuries". In: *Clinical Rheumatology* 32.4, pp. 425–434.

Rinzler, J.W. and B. Burtt (2010). *The Sounds of Star Wars*. Chronicle Books.

Robotics, Rethink (2017). *Collaborative Robot Safety and Compliance*. URL: http://www.rethinkrobotics.com/safety-compliance/ (visited on 04/19/2017).

Rosecrance, John C. and Thomas M. Cook (2000). "The use of Participatory Action Research and Ergonomics in the Prevention of Work-Related Musculoskeletal Disorders in the Newspaper Industry". In: *Applied Occupational and Environmental Hygiene* 15.3, pp. 255–262.

Rosenbaum, David A., Caroline M. van Heugten, and Graham E. Caldwell (1996). "From Cognition to Biomechanics and Back: The End-State Comfort Effect and the Middle-Is-Faster Effect". In: *Acta Psychologica* 94.1, pp. 59–85.

Sartori, Luisa, Cristina Becchio, and Umberto Castiello (2011). "Cues to Intention: the Role of Movement Information". In: *Cognition* 119.2, pp. 242–252.

Schaal, Stefan, Peyman Mohajerian, and Auke Ijspeert (2007). "Dynamics Systems vs. Optimal Control–a Unifying View". In: *Progress in brain research* 165, pp. 425–445.

Sebanz, Natalie, Harold Bekkering, and Günther Knoblich (2006). "Joint Action: Bodies and Minds Moving Together". In: *Trends in Cognitive Sciences* 10.2, pp. 70–76.

Sisbot, Emrah Akin and Rachid Alami (2012). "A Human-Aware Manipulation Planner". In: *IEEE Transactions on Robotics* 28.5, pp. 1045–1057.

Sisbot, Emrah Akin, Luis F Marin-Urias, et al. (2007). "A Human Aware Mobile Robot Motion Planner". In: *IEEE Transactions on Robotics* 23.5, pp. 874–883.

Sisbot, Emrah Akin, Luis F. Marin-Urias, et al. (2010). "Synthesizing Robot Motions Adapted to Human Presence: A Planning and Control Framework for Safe and Socially Acceptable Robot Motions". In: *International Journal of Social Robotics* 2.3, pp. 329–343.

Srivastava, Siddharth et al. (2014). "Combined Task and Motion Planning through an Extensible Planner-Independent Interface Layer". In: *Proc. IEEE International Conference on Robotics and Automation (ICRA'14)*, pp. 639–646.

Strabala, Kyle W. et al. (2013). "Towards Seamless Human-Robot Handovers". In: *Journal of Human-Robot Interaction* 2.1, pp. 112–132.

Strabala, Kyle et al. (2012). "Learning the Communication of Intent Prior to Physical Collaboration". In: *Proc. IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN'12)*, pp. 968–973.

Stulp, Freek, Jonathan Grizou, et al. (2015). "Facilitating Intention Prediction for Humans by Optimizing Robot Motions". In: *Proc. IEEE International Conference on Intelligent Robots and Systems (IROS'15)*, pp. 1249–1255.

Stulp, Freek, Laura Herlant, et al. (2014). "Simultaneous Online Discovery and Improvement of Robotic Skill Options". In: *Proc. IEEE International Conference on Intelligent Robots and Systems (IROS'14)*, pp. 1408–1413.

Stulp, Freek, Michael Isik, and Michael Beetz (2006). "Implicit Coordination in Robotic Teams Using Learned Prediction Models". In: *Proc. IEEE International Conference on Robotics and Automation (ICRA'06)*, pp. 1330–1335.

Stulp, Freek and Olivier Sigaud (2012). "Policy Improvement Methods: Between Black-Box Optimization and Episodic Reinforcement Learning". URL: https://hal.archives-ouvertes.fr/hal-00738463.

Suay, Halit Bener and Emrah Akin Sisbot (2015). "A Position Generation Algorithm Utilizing a Biomechanical Model for Robot-Human Object Handover". In: *Proc. IEEE International Conference on Robotics and Automation (ICRA'15)*.

Takala, Esa-Pekka et al. (2010). "Systematic Evaluation of Observational Methods Assessing Biomechanical Exposures at Work". In: *Scandinavian Journal of Work, Environment and Health* 36.1, pp. 3–24.

Takayama, Leila, Doug Dooley, and Wendy Ju (2011). "Expressing Thought: Improving Robot Readability with Animation Principles". In: *Proc. IEEE/ACM International Conference on Human-Robot Interaction (HRI'11)*, pp. 69–76.

Takayama, Leila and Caroline Pantofaru (2009). "Influences on Proxemic Behaviours in Human-Robot Interaction". In: *Proc. IEEE International Conference on Intelligent Robots and Systems(IROS'09)*. St. Louis, MO, USA, pp. 494–499.

The Cost of Work-related Injury and Illness for Australian Employers, Workers and the Community: 2012–13 (2015). Tech. rep. Safe Work Australia.

Timmermans, Bert, Sukhvinder Obhi, and Wilfrid Laurier (2012). "Grasping Intentions: From Thought Experiments to Empirical Evidence". In: *Frontiers in Human Neuroscience*, pp. 170–175.

Toussaint, Marc (2015). "Logic-Geometric Programming: An Optimization-Based Approach to Combined Task and Motion Planning". In: *Proc. International Joint Conference on Artificial Intelligence (IJCAI'15)*.

Toussaint, Marc and Manuel Lopes (2016). "Multi-Bound Tree Search for Logic-Geometric Programming in Cooperative Manipulation Domains". In: *Proc. IEEE International Conference on Robotics and Automation (ICRA'16)*.

Toussaint, Marc, Thibaut Munzer, et al. (2016). "Relational Activity Processes for Modeling Concurrent Cooperation". In: *Proc. IEEE International Conference on Robotics and Automation (ICRA'16)*.

Tsarouchi, Panagiota, Sotiris Makris, and George Chryssolouris (2016). "Human-Robot Interaction Review and Challenges on Task Planning and Programming". In: *International Journal of Computer Integrated Manufacturing* 29.8, pp. 916–931.

Unity. *Unity-Game Engine*. URL: https://unity3d.com/.

Vesper, Cordula et al. (2011). "Making Oneself Predictable: Reduced Temporal Variability Facilitates Joint Action Coordination". In: *Experimental Brain Research* 211.3-4, pp. 517–530.

Wagenaar, Willem Albert et al. (1994). "Promoting Safety in the Oil Industry. The Ergonomics Society Lecture Presented at the Ergonomics Society Annual Conference, Edinburgh, 13-16 April 1993". In: *Ergonomics* 37.12, pp. 1999–2013.

Wells, R. et al. (1997). "Assessment of Physical Work Load in Epidemiologic Studies: Common Measurement Metrics for Exposure Assessment". In: *Ergonomics* 40.1, pp. 51–61.

Xie, Jeanne Y., Huei-Yen Winnie Chen, and Birsen Donmez (2016). "Gaming to Safety: Exploring Feedback Gamification for Mitigating Driver Distraction". In: *Proc. Human Factors and Ergonomics Society Annual Meeting*. Vol. 60. 1. SAGE Publications Sage CA: Los Angeles, CA, pp. 1884–1888.

Yeung, Kwok-Yun, Tsz-Ho Kwok, and Charlie CL. Wang (2013). "Improved skeleton tracking by duplex kinects: a practical approach for real-time applications". In: *Journal of Computing and Information Science in Engineering* 13.4, p. 041007.

Zhao, Min et al. (2016). "An Experimental Study for Identifying Features of Legible Manipulator Paths". In: *International Symposium on Experimental Robotics*, pp. 639–653.

Zhou, Allan et al. (2017). "Expressive Robot Motion Timing". In: *Proc. IEEE/ACM International Conference on Human-Robot Interaction (HRI'17)*. ACM, pp. 22–31.

Appendix A
REBA Calculation

The REBA method (Hignett et al. 2000), introduced in Chapter 6 provides a posture score that represents the risk of MSDs from joint angle values. There are two groups, A and B, for the body segments. Calculations of the individual joints composing each groups are illustrated in Figure A.1 and A.2 respectively. *Correspondence tables* A.1 and A.2 provide the scores for the group A (Score A) and B (Score B) respectively. To Score A, obtained from Table A.1, we need to add the payload score $\delta_{payload}$:

$$\begin{aligned} if\ load < 5kg: &\ \delta_{payload} = 0 \\ if\ load\ between\ 5\ to\ 10kg: &\ \delta_{payload} = 1\ , \\ if\ load > 10kg: &\ \delta_{payload} = 2 \end{aligned} \tag{A.1}$$

where *load* corresponds to the weight of the carried object. Finally, the REBA score is obtained from Table A.3.

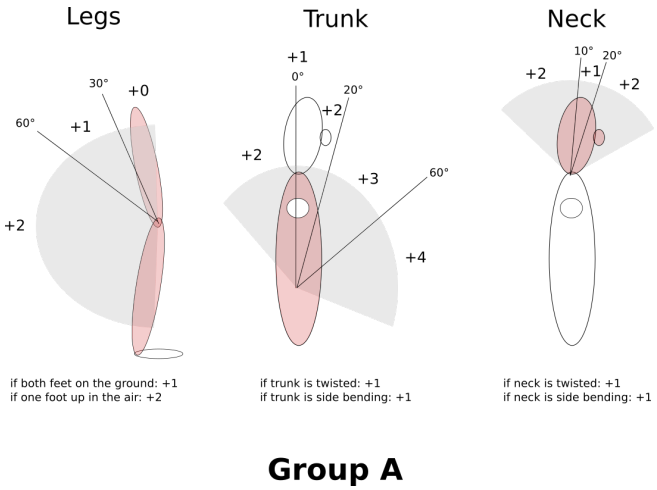

Figure A.1: Illustration of the REBA calculation for individual joints of Group A composed by legs, trunk, and neck. Score A is calculated from Table A.1

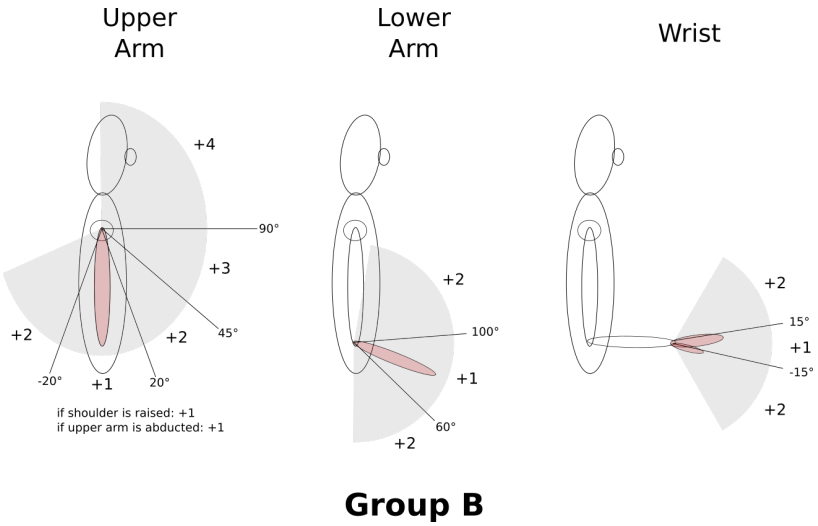

Figure A.2: Illustration of the REBA calculation for individual joints of Group B composed by upper arm, lower arm and wrist. Score B is calculated from Table A.2

Table A		Neck											
		1				2				3			
	Legs												
		1	2	3	4	1	2	3	4	1	2	3	4
Trunk	1	1	2	3	4	1	2	3	4	3	3	5	6
	2	2	3	4	5	3	4	5	6	4	5	6	7
	3	2	4	5	6	4	5	6	7	5	6	7	8
	4	3	5	6	7	5	6	7	8	6	7	8	9
	5	4	6	7	8	6	7	8	9	7	8	9	9

Table A.1: Correspondence table for the group A

Table B		Lower Arm					
		1			2		
	Wrist						
		1	2	3	1	2	3
Upper Arm	1	1	2	2	1	2	3
	2	1	2	3	2	3	4
	3	3	4	5	4	5	5
	4	4	5	5	5	6	7
	5	6	7	8	7	8	8
	6	7	8	8	8	9	9

Table A.2: Correspondence table for the group B

Score A	Table C											
	Score B											
	1	2	3	4	5	6	7	8	9	10	11	12
1	1	1	1	2	3	3	4	5	6	7	7	7
2	1	2	2	3	4	4	5	6	6	7	7	8
3	2	3	3	3	4	5	6	7	7	8	8	8
4	3	4	4	4	5	6	7	8	8	9	9	9
5	4	4	4	5	6	7	8	8	9	9	9	9
6	6	6	6	7	8	8	9	9	10	10	10	10
7	7	7	7	8	9	9	9	10	10	11	11	11
8	8	8	8	9	10	10	10	10	10	11	11	11
9	9	9	9	10	10	10	11	11	11	12	12	12
10	10	10	10	11	11	11	11	12	12	12	12	12
11	11	11	11	11	12	12	12	12	12	12	12	12
12	12	12	12	12	12	12	12	12	12	12	12	12

Table A.3: Correspondence table for the REBA score

Appendix B
Postural Assessment and Feedback

The question on the worker's self-awareness of the risks associated with a given posture still remains. There are multiple projects and companies that develop wearable devices or fabrics that help users monitor their shape or keep track of their posture (Rettberg 2014) Most of them are based on intrusive sensors or clothings that are not really suitable for industrial workers. The growing number of human tracking solutions based on external sensors allows for a better integration of automatic postural assessment techniques in industry.

Most of the time, workers are unaware of the risk induced by a wrong posture. Dangerous postures might not have a direct effect on the body but it is the repetition that increases the risk of MSDs. Therefore, a direct feedback is an interesting feature that could help workers maintaining a good posture at their workstation. To that extent, we have developed a graphical interface to animate a model with the user's observed body posture. The projection on the model is almost real-time.

The graphical interface has been developed in Unity[1] using humanoid models made available by Mixamo[2]. Unity technology allows us to display the interface on any screen, including smartphones or tablets. Figure B.1 shows multiple views of the graphical interface, and also present the two models made available. Using the buttons on the right side, the user can choose to display a male or a female version of the model. He or she can also switch between a front or a back view.

Our solution is decoupled in two modules. First, we record the worker current using inverse kinematic on the tracked frames (see Section 6.4.1). This provides a joint vector that we send to our graphical interface in order to animate the model. Second, we evaluate the risk of the posture using the REBA assessment technique. The REBA values for each individual joints are then also sent to the graphical interface and displayed with spheres located at the evaluated body joints. The color of each sphere represents the risk at its corresponding location. It ranges from "green" (no risk) to "red" (high risk) according to the REBA method.

[1] Unity is a cross-platform game engine, developed by Unity Technologies, which is primarily used to develop video games and simulations for computers, consoles and mobile devices (Unity n.d.).

[2] Mixamo is a 3D computer graphics technology company. Based in San Francisco, the company develops and sells web-based services for 3D character animation (Mixamo n.d.).

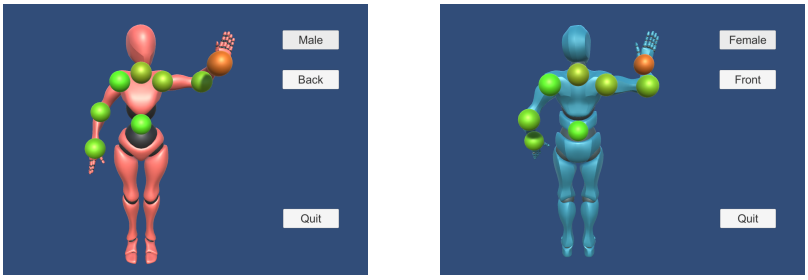

(a) female model, front view

(b) male model, back view

Figure B.1: Screenshots of the graphical interface in two possible configurations. The same body posture is used in both pictures. In front view, by contrast with the back view, the body configuration mirrors the real posture.

After some trials with test subjects, it has been decided that the front view should mirror the user's body configuration, i.e rising the right arm translates to a rising of the model left arm. In back view, however, the body configuration is normally displayed.