

A Human-Inspired Controller for Fluid Human-Robot Handovers

José R. Medina*

Felix Duvallet*

Murali Karnam

Aude Billard

Abstract—Handovers are seamless events that occur frequently and naturally between people. Although previous works have focused on the design of robot handover controllers synthesizing one of the phases of a handover (approaching, passing, or retracting) in an isolated manner, we take a different approach and treat the handover as a single continuous entity. In this paper we present a novel bi-directional and human-inspired handover controller using insights we learned from human-human experiments. We model the dynamics of the handover in a continuous and time-independent way yielding robust and fluid behavior. We implemented our approach on a robot platform consisting of a 7-DoF robotic arm with a 16-DoF humanoid hand. Our results show that resulting human-robot handovers are smooth and reduce internal forces with the human compared to traditional switching approaches.

I. INTRODUCTION

A handover occurs anytime two parties exchange an object without the use of an external tool (such as a table). These fast and seamless events occur naturally between people, in a wide range of environmental conditions with and without constraints, and for objects ranging in size and mass. To become useful partners in human-robot teams, robots will need to approach this level of capability, and be able to perform fast natural handovers of arbitrary objects as both giver and receiver.

Once a handover begins, it generally consist of three parts: an *approach* phase where one or both parties move to the handover location (where the object is reachable by both), a *passing* phase where the object is physically transferred from the giver to the receiver, and a *retraction* phase where the giver and receiver move away from each other. While much work exists that attempts to understand the dynamics of this handover process in humans, to date robot controllers attempting to replicate these processes have fallen short of providing a complete handover implementation that is fluid and seamless. One explanation for this gap is that prior studies that replicated these different phases generally did so by considering each phase in isolation (e.g., moving the arm in a predictable way, or controlling the grip force during object transfer but without any robot arm motion). This results in a handover that is un-naturally sequential: the handover proceeds from one phase to the next, rigidly waiting for a specific condition before transitioning to the next phase (e.g., a waiting for a force trigger before completely opening the hand).

We take a different approach, treating the handover as a single entity where we design jointly the motions of the

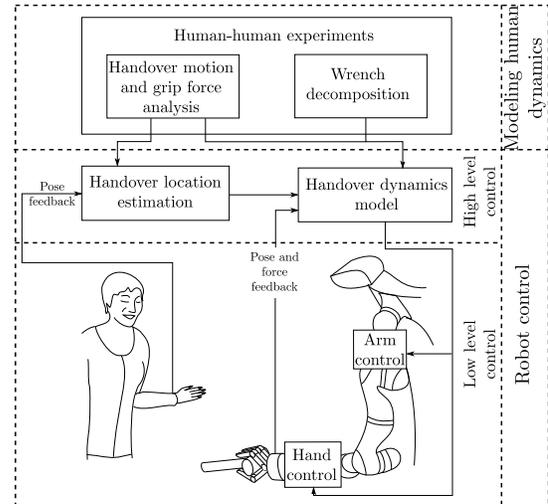


Fig. 1. Our goal is to enable natural and fluid handovers of objects between people and robots. Using studies of human-human handovers, we analyzed the role of motion, grip forces, and interaction wrenches to model the dynamics of fluid handovers. We then developed a controller that estimates the location of the handover, moves the arm towards the handover pose in a natural way, and releases the object in an intuitive manner. All of these occur in a continuous and phase-less way.

complete robot (in our case, arm and hand) to perform a fluid handover. Our controller is human-inspired, using insights acquired from human-human handover experiments. The resulting handover is natural, bi-directional, and adaptable to many object shapes and weights.

The components for our approach, shown in Figure 1, are discussed in the rest of this paper. After reviewing prior work in the area of handovers (Section II), we present the mathematical formulation of the problem in Section III. We then describe our human-human handover experiments and the insights derived from them in Section IV; importantly we show that people control both their hand and arm during all handover phases depending on the proportion of the object load they support, implying that a robot controller should understand, decompose, and robustly react to interaction forces in order to perform a fluid handover. We describe such a human-inspired controller in Section V, and present results from human-robot handover experiments in Section VI.

II. RELATED WORK

Given their frequency and importance in day-to-day life, it is not surprising that object handovers are well studied by researchers. Insights about handovers come from two major areas of literature: observations of human-human handovers that attempt to understand the complex dynamics underpin-

*The first two authors contributed equally to this paper. All authors are with the LASA laboratory of the École Polytechnique Fédérale de Lausanne, Switzerland. {firstname.lastname}@epfl.ch

ning handovers, and user studies (with a robot implementation) that attempt to replicate (and analyze) some of the complex dynamics on a robot.

Across these broad areas of study, the major types of studies by researchers are studies into the *forces* exerted during handovers (e.g., the grip forces exerted on the object by both parties) and the *motions* of the handover participants (e.g., trajectories to the handover location).

A. Studies of handover forces

Studies on forces exerted during handover analyze the relationship between grip force (the amount of force applied on the object by each participants' fingers) and load force (the proportion of the object's weight sustained by each participant). For example, Mason and MacKenzie [1] found that human-human handovers involve a gradual modulation of the grip force while both participants implicitly share (and transfer) the load force. The main coordination mechanism is haptic feedback in order to coordinate the transfer force rates; the grip force increases for the receiver and decreases for the passer.

Another consistent finding in prior studies is that the time required to transfer the object from the passer to the receiver is between 300 and 500 ms [1, 2, 3]. Even Endo et al. [3], who outfitted the receiver with a glove in some of the trials to attenuate the tactile information, found that while the glove did increase the duration of the average contact period slightly (from 324 ms to 334 ms), it did not have a significant effect on the grip force profiles.

Two additional useful findings by Chan et al. [2] are that the passer has a "post-unloading" phase towards the end of the handover, where the giver applies a positive grip force even though their load force is approximately zero. This, according to the authors, implies that the giver takes responsibility for the *safety* of the object. Additionally, their studies found that the receiver adjusts the load transfer rate depending on the weight of the object, and concluded that the receiver is in charge of the handover *timing*. The recommendations resulting from this study were used to develop a linear controller that varied the robot's grip force according to the sensed load force during a robot-to-human handover [4]. However, the robot controller implementations based on these studies focus solely on the grip force modulation, with no arm motion during the handover.

These studies have important implications for designing robot handover controllers:

- the load force drives the dynamics of the passing phase of handovers,
- handovers must be fast and respond quickly to changes in the load force.

This means our robot must be able to quickly measure the proportion of the object's load it supports, and quickly adapt its grip force on the object during the passing phase of a handover.

B. Studies of handover motions

Studies on handover motions for robots fall into two broad categories: planning-based approaches that optimize

a complete trajectory, and controller-based approach that continuously optimize some controller input.

Planning-based approaches formulate one or more parts of the handover as a search problem. For example, such approaches have been applied to finding configuration of the robot arm the agree with human preferences [5], plans involving handovers in constrained environments [6], or handover locations that optimize the shared effort of the human and robot [7].

Controller-based approaches for robot handovers also exist, and generally attempt to reach a desired goal position [8], often by following a desired trajectory [9].

While the implementations in these studies resulted in natural trajectories *to* the handover location, the robot stops moving before transitioning to the passing phase (once both parties make contact with the object) and in general these approaches did not address the handover release strategy. Our human studies show that people move slightly during the in-contact (passing) phase, and thus it is important for robots to consider the *entire* handover as one fluid motion, without hard transitions between the handover phases.

III. HANDOVER FORMULATION

We consider the problem of handing over an object from one party to another. We assume that each party consists of an arm and a hand, that the object is rigid and its physical properties are known, and that initially one of the parties is grasping the object. In our setting, we do not consider the signaling problem and assume that the intention to perform a handover has already been established. We consider a handover successful when the object load is transferred from the *giver* to the *receiver* without being supported anywhere else (e.g., falling on the ground or placed on the table). In the following, subscripts g, r will refer to the giver and the receiver respectively.

We are interested in modeling the motion of the giver and receiver as they move towards the handover location (a phase we will informally refer to as the approach phase) and transfer of the object's load (similarly, the passing phase). We denote the arm end effector poses of the giver and receiver by $\mathbf{x}_g = [\mathbf{p}_g^T \ \mathbf{o}_g^T]^T$, $\mathbf{x}_r = [\mathbf{p}_r^T \ \mathbf{o}_r^T]^T$, where $\mathbf{p}_g, \mathbf{p}_r \in \mathbb{R}^3$ and $\mathbf{o}_g, \mathbf{o}_r \in \mathbb{R}^4$ are the Euclidean positions and orientations in the axis-angle representation respectively. We also denote the hand configurations by $\mathbf{x}_{h,g}, \mathbf{x}_{h,r} \in \mathbb{R}^h$.

We represent the rigid body dynamics of an object o with mass $m_o \in \mathbb{R}$ by

$$M_o(\mathbf{x}_o)\ddot{\mathbf{x}}_o + C_o\dot{\mathbf{x}}_o + \mathbf{g}_o = \mathbf{u}_o, \quad (1)$$

where $\dot{\mathbf{x}}_o = [\dot{\mathbf{p}}_o^T \ \dot{\boldsymbol{\omega}}_o^T]^T$ with $\dot{\mathbf{p}}_o, \dot{\boldsymbol{\omega}}_o \in \mathbb{R}^3$ are respectively the translational and rotational velocity of its center of mass (expressed in the inertial frame), the mass matrix is $M_o(\mathbf{x}_o) = \text{diag}\{m_o I_3, J_o(\mathbf{x}_o)\}$ where $J_o(\mathbf{x}_o) \in \mathbb{R}^{3 \times 3}$ is the inertia tensor, $\mathbf{u}_o = [\mathbf{f}_o^T \ \boldsymbol{\tau}_o^T]^T \in \mathbb{R}^6$ is the resulting wrench applied on the object with $\mathbf{f}_o, \boldsymbol{\tau}_o \in \mathbb{R}^3$ denoting the force and torque respectively, $\mathbf{g}_o = [m_o \mathbf{g}^T \ 0_3^T]^T$ is the

wrench on the object due to gravity, and

$$C_o = \begin{bmatrix} 0_{3 \times 3} & 0_{3 \times 3} \\ 0_{3 \times 3} & \boldsymbol{\omega}_o \times J_o(\mathbf{x}_o) \end{bmatrix}$$

represents Coriolis and centrifugal terms.

When one of the parties is grasping the object, the applied wrench on the object \mathbf{u}_o corresponds to the sensed wrench at the *wrist* \mathbf{u}_s considering the kinematic constraints of the grasp, i.e.

$$\mathbf{u}_o = G\mathbf{u}_s ,$$

where G usually denotes the *grasp matrix*. It is given by

$$G = \begin{bmatrix} I_3 & 0_{3 \times 3} \\ P & I_3 \end{bmatrix} ,$$

where $P = [\mathbf{p}]_{\times}$ is the cross product matrix for \mathbf{p} , the vector from the object's center of mass to the end-effector position.

When two parties g and r are grasping the object, the object-centered dynamics become

$$\mathbf{u}_o = G_g\mathbf{u}_{s,g} + G_r\mathbf{u}_{s,r} .$$

Note that we are considering the grasp of the giver and the receiver as a single wrench input and we are not modeling the wrenches due to each finger/contact point, as typically done in grasping. In this case, we can decompose the wrench input of each party (denoted by i) as

$$\mathbf{u}_{s,i} = G_i^{-1}(A_{\text{mot},i}\mathbf{u}_{\text{mot},o} + A_{\text{load},i}\mathbf{g}_o) + \mathbf{u}_{\text{int},i} , \quad (2)$$

where $\mathbf{u}_{\text{int},i}$ are the *internal wrench* components between the giver and the receiver inducing no object motion nor compensating gravity wrenches, and the matrices $A_{\text{mot},i}, A_{\text{load},i} \in \mathbb{R}^{6 \times 6}$ allocate wrench components inducing object motion $\mathbf{u}_{\text{mot},o} = M_o(\mathbf{x}_o)\ddot{\mathbf{x}}_o + C_o\dot{\mathbf{x}}_o$ and due to gravity \mathbf{g}_o respectively and fulfill

$$\begin{aligned} \mathbf{u}_{\text{mot},o} &= (A_{\text{mot},g} + A_{\text{mot},r})\mathbf{u}_{\text{mot},o} \\ \mathbf{g}_o &= (A_{\text{load},g} + A_{\text{load},r})\mathbf{g}_o . \end{aligned} \quad (3)$$

They are given by

$$A_{\text{mot},i} = \begin{bmatrix} \beta_{f,i}I_3 & 0_{3 \times 3} \\ 0_{3 \times 3} & \beta_{\tau,i}I_3 \end{bmatrix} \quad A_{\text{load},i} = \begin{bmatrix} \alpha_i I_3 & 0_{3 \times 3} \\ 0_{3 \times 3} & 0_{3 \times 3} \end{bmatrix} , \quad (4)$$

where $\alpha_i, \beta_{f,i}, \beta_{\tau,i} \in [0, 1]$ and $\sum_i \alpha_i = 1$, $\sum_i \beta_{f,i} = 1$, $\sum_i \beta_{\tau,i} = 1$. This unusual wrench distribution separating wrenches due to motion and gravity is necessary for handovers, as they are characterized by the transfer of the object's load.

Since our decomposition directly represents the proportion of the load share supported by the giver and receiver in Equation (3), we are able to represent the overall objective of transferring the load from the giver to the receiver as an evolution from $\alpha_g = 1.0$ (the giver supports the object's load) to $\alpha_g = 0.0$ (the receiver now supports the load). Therefore,

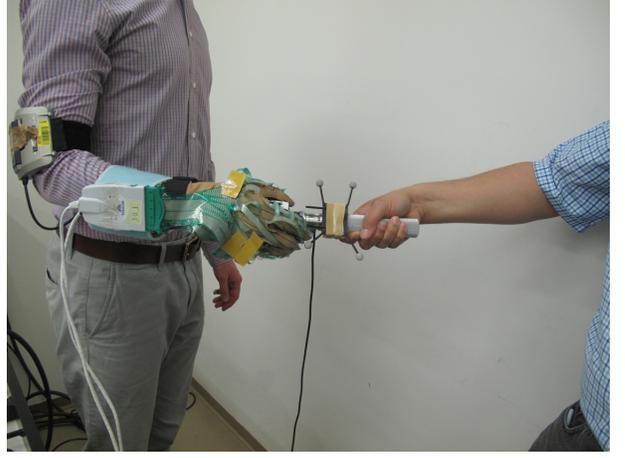


Fig. 2. Experimental setup used in the human-human experiments consisting of a CyberGlove with TekScan patches, OptiTrack markers, and an instrumented object with a force/torque sensor.

the problem considered in this work consists of the design of function

$$\begin{bmatrix} \dot{\mathbf{x}}_g \\ \dot{\mathbf{x}}_r \\ \dot{\mathbf{x}}_{h,g} \\ \dot{\mathbf{x}}_{h,r} \end{bmatrix} = \mathbf{f} \left(\begin{bmatrix} \mathbf{x}_g \\ \mathbf{x}_r \\ \mathbf{x}_{h,g} \\ \mathbf{x}_{h,r} \\ \alpha_g \end{bmatrix} \right) \quad (5)$$

that describes the motion of both parties fulfilling objective $\lim_{t \rightarrow \infty} \alpha_g(t) \rightarrow 0$.

IV. INSIGHTS FROM HUMAN-HUMAN HANDOVER STUDIES

To better understand the dynamics at play during a fluid handover, we collected data from human-human handovers of a sensorized object between two participants such as the one shown in Figure 2. We then analyzed these trials to extract insights about the motion of the giver and receiver during the entire handover process, as well as the interaction forces exerted on the object while it is being transferred. This study serves two key purposes. First, it enables us to infer the important variables and structural constraints that underlie the handover dynamics of Equation (5). Second, they provide a corpus of training data with which we will learn the parameters for the chosen model.

In these studies, two participants exchanged a sensorized object consisting of two handles joined together with ATX Nano 25 force/torque sensor. Each handle has a mass of $m_{\text{handle}} = 0.28$ [kg] and the object has a total mass of $m_o = 2m_{\text{handle}}$ [kg]. We also placed an OptiTrack marker on the object, and recorded its pose during the entire handover. One participant was instrumented with a data glove (the CyberGlove with TekScan patches on the fingers and palm) and an OptiTrack marker attached to their forearm. The second participant had an OptiTrack marker attached to their wrist but was not instrumented otherwise.

Both participants took turns acting as the giver and receiver. For each handover, one of the participants grasped

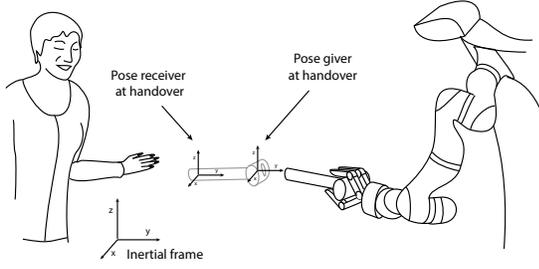


Fig. 3. We represent all poses in the frame of the configuration at the handover time. This enables us to represent handovers in a generalizable frame of reference.

the object with a power (whole-hand) grasp, waited for the “go” signal, and then handed the object over to the other participant. The only grasping constraint was that both parties should grasp the object on opposite sides of the force/torque sensor (we did not constrain users to use precision grasps). When the handover was complete we stopped recording data for that trial and moved to the next handover. We collected a total of 25 trials of handovers in both directions, with three different participants holding the object both horizontally and vertically.

A. Data Processing

We processed the pose data to smooth it and compute its derivatives using a Savitzky-Golay filter, and determined the *handover poses* for both participants (the pose at the time when the participants are closest to each other), $\mathbf{x}_{g,ho}$ for the giver and $\mathbf{x}_{r,ho}$ for the receiver, as shown in Figure 3. These handover poses provide a way to represent the handover in a generalizable frame of reference with respect to the object pose at the handover, and for notational simplicity in the remainder of this analysis the poses of both the giver \mathbf{x}_g and the receiver \mathbf{x}_r will be expressed relative to their respective handover pose, such that $\mathbf{p}_{i,ho} = \mathbf{0}$ and therefore $\|\mathbf{p}_i - \mathbf{p}_{i,ho}\| = \|\mathbf{p}_i\|$ represents the distance to the handover pose.

The force/torque sensor enables us to estimate both the internal wrench applied on the object \mathbf{u}_{int} , as well as the load share α_i . The wrench captured by the force/torque sensor $\mathbf{u}_{s,o}$ expressed in the *inertial frame* is

$$\mathbf{u}_{s,o} = \tilde{\mathbf{u}}_{s,r} + \mathbf{u}_{handle} ,$$

or considering the other side of the sensor

$$\mathbf{u}_{s,o} = -(\tilde{\mathbf{u}}_{s,g} + \mathbf{u}_{handle}) ,$$

where \mathbf{u}_{handle} is given by (1) but with mass m_{handle} and $\tilde{\mathbf{u}}_{s,i} = G_i \mathbf{u}_{s,i}$. Given observed object poses \mathbf{x}_o , grasp matrices G_i and the object’s physical properties, the computation of \mathbf{u}_{handle} and $\mathbf{u}_{s,i}$ is straightforward. To estimate the load share, the computation of the wrench distribution matrices from (2) is necessary. The parameters $\beta_{f,i}$ and $\beta_{\tau,i}$ concerning the motion-inducing wrenches are given by the least squares solution of $\|\tilde{\mathbf{u}}_{s,i} - A_{mot,i} \mathbf{u}_{mot,i}\|$ fulfilling con-

straints from (4), yielding

$$\beta_{f,i} = \max \left\{ \min \left\{ \frac{\tilde{\mathbf{f}}_{s,i}^T \mathbf{f}_{mot,o}}{\|\mathbf{f}_{mot,o}\|^2}, 1 \right\}, 0 \right\} ,$$

$$\beta_{\tau,i} = \max \left\{ \min \left\{ \frac{\tilde{\boldsymbol{\tau}}_{s,i}^T \boldsymbol{\tau}_{mot,o}}{\|\boldsymbol{\tau}_{mot,o}\|^2}, 1 \right\}, 0 \right\} . \quad (6)$$

This solution follows from applying Gauss’s principle of least constraint for multibody systems [10]. The *load share* is computed similarly although subtracting a priori the motion-inducing wrench components already considered in the computation of $\beta_{f,i}$, i.e.

$$\alpha_i = \max \left\{ \min \left\{ \frac{(\tilde{\mathbf{f}}_{s,i} - \beta_{f,i} \mathbf{f}_{mot,i})^T \mathbf{g}}{\|\mathbf{g}\|^2}, 1 \right\}, 0 \right\} . \quad (7)$$

Given the results from Equations (6) and (7), the internal wrench $\mathbf{u}_{int,i}$ follows straightforwardly from Equation (2). Note that for computing α_i , the only information necessary about the object’s physical properties is its mass.

To estimate the grip force applied by the instrumented party, we use the Tekscan patches located on the hand to compute the total pressure applied by the participant on the object. For simplicity, in our analysis the estimated grip force \bar{u}_i is given by the sum of the measurements across all Tekscan patches.

The CyberGlove measures the 22-dimensional joint configuration of the hand of the instrumented participant (similarly to the grip force, we can compute the hand configuration for only *one* participant). For simplicity, in our analysis we use this information to compute the derivative of the sum across all measurements of the hand configuration, yielding approximately $\dot{\mathbf{x}}_{h,i}$.

B. Data Analysis and Important Insights

After processing the data, we analyzed it for useful insights which will determine the important signals that will be used by our controller. In order to identify potential dependencies between all variables involved in the handover we first performed a Granger causality test [11]. This is a statistical hypothesis test that evaluates if the addition of a new time series component improves the prediction for another time series, compared to using the past values alone. In our setting, if some variable Granger-causes (G-causes) one of the target variables involved in the handover (such as the desired trajectory or the grip force), then it might be beneficial to include it as an input in the controller. We used the MVGC Matlab Toolbox [12].

1) *Insight: The giver and receiver are coupled:* We evaluated G-causality during the motion phases (approach and retract), between the distance of the giver’s and receiver’s end-effectors to the handover position ($\|\mathbf{p}_g\|$ and $\|\mathbf{p}_r\|$ respectively), and the change in the hand configuration \mathbf{x}_h . Our results, shown in Figure 4, indicate a dependency from the giver’s motion to that of the receiver’s. This suggests that the giver leads the motion during the approach phase, and the receiver’s motion is coupled to it.

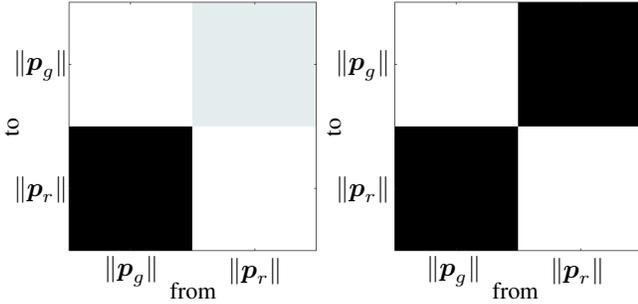


Fig. 4. Granger causality (left: score, right: statistical significance) for the distance of both end-effectors to the handover position during the approach

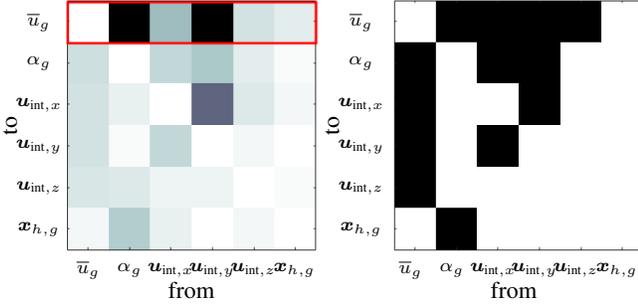


Fig. 5. Granger causality between the grip force of the giver \bar{u}_g , load share α_g , the 3-dimensional internal wrench \mathbf{u}_{int} , and the hand configuration of the giver $\mathbf{x}_{h,g}$ during the handover. A darker square indicates higher causality score. The red rectangle highlights the scores of all variables G-causing the grip force \bar{u}_g .

2) *Insight: The load share drives the evolution of the passing phase:* Since the grip force of the giver \bar{u}_g will be the key control input regulating the load transfer during the handover, we are interested in studying its relationship to the load share α_g , the internal wrenches sensed by the giver \mathbf{u}_{int} (we hereafter drop the g index for simplicity) which are not part of the load force, and the hand motion of the giver $\mathbf{x}_{h,g}$. The G-causality results for the passing phase, shown in Figure 5, indicate that the two highest scores are for the internal wrench in the y direction $\mathbf{u}_{\text{int},y}$ and the load share α_g . Since in our experiments $\mathbf{u}_{\text{int},y}$ is aligned with the direction of motion, it thus corresponds to pushing and pulling on the object by the participants in the direction of the handover.

To investigate this relationship further, we analyzed grip force as a function of the load share. We find that the grip force increases almost linearly as a function of the load share, as shown in Figure 6. This result suggests that humans modulate grip force as a function of the amount of perceived load share, and aligns with prior studies modeling the relationship between the two variables [1, 4].

3) *Insight: Both parties move during the passing phase:* In addition, it is also interesting to study arm motion during the passing phase and its potential role in the load transfer. More precisely, motions in the z -axis are especially relevant as they directly influence the load force. Figure 7 shows the difference between the giver and the receiver positions

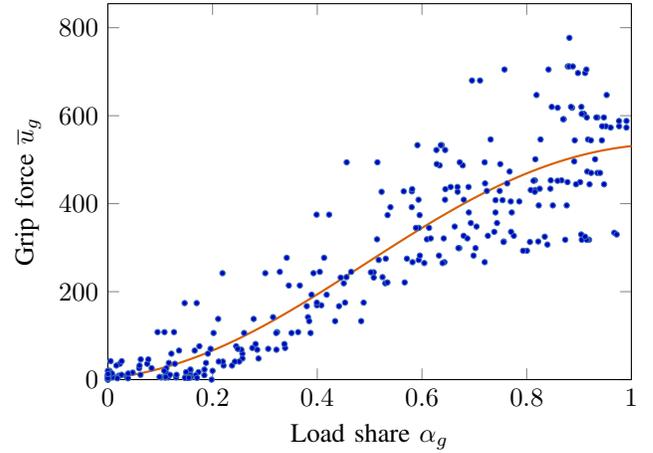


Fig. 6. Relationship between the giver's grip force \bar{u}_g as a function of the load share α_g during handovers. We show both the raw points (plus) and a fit using Gaussian Process Regression, indicating the grip force decreases almost linearly as the load share decreases.

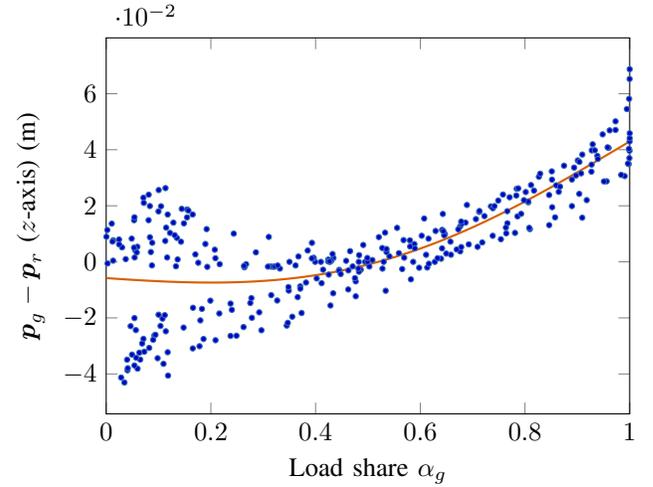


Fig. 7. Difference between the giver and the receiver position in the z axis during the passing phase. As we can see the giver and receiver move relative to each other as they transfer the object.

in the z -axis as a function of the load share, fitted with a Gaussian Process (GP) regression for visualization purposes. Interestingly, the load transfer does not happen from a constant position: at the beginning of the passing phase ($\alpha_g \approx 1$) the giver's position is higher, while at the end ($\alpha_g \approx 0$) the receiver's position is higher. These results show that there is motion during the passing phase of handovers: the giver lowers their hand and the receiver raises their hand, indicating that humans rely in part on arm motion in order to perform handovers.

V. A FLUID HUMAN-INSPIRED HANDOVER CONTROLLER

Exploiting the insights and data from the study from the previous section, in this section we detail the motion design of our proposed handover controller, the estimation of the handover location, and the low-level control necessary to perform fast bi-directional handovers.

A. Desired trajectories

In our analysis we observed that during the approach phase both the giver and the receiver exhibit a goal-oriented behavior towards the handover pose. In addition, the giver's position leads the motion while the receiver's motion (and other variables such as the grip force) exhibit a subordinate behavior. Concerning the passing phase, our analysis indicates that it is driven by the perceived load share of each participant, which highlights the importance of arm motion during this phase. Coupled Dynamical System (CDS) [13] are well-suited to capture these constraints a priori, where a slave DS evolves depending on the state of the master DS.

Concerning the giver's arm position, we assume the following structure

$$\dot{\mathbf{p}}_g = \mathbf{f}_g(\mathbf{p}_g - \mathbf{g}_g(\alpha_g)) . \quad (8)$$

where \mathbf{f}_g is a continuous and continuously differentiable function with a single attractor given by $\mathbf{g}_g(\alpha_g)$. In addition, we consider a priori that $\mathbf{g}_g(\alpha_g = 1) = \mathbf{p}_{g,ho}$. This way, during the approach phase, the giver follows the attractor given by the handover location, i.e. $\dot{\mathbf{p}}_g = \mathbf{f}_g(\mathbf{p}_g - \mathbf{p}_{g,ho})$, as an autonomous system. After contact with the other party, α_g shifts the attractor as observed in Section IV-B3. This formulation inherently assumes that the load share α_g is the master DS driving the advancement of the task, as highlighted in Section IV-B2 and $\mathbf{g}(\alpha_g)$ represents the *coupling function* between the motion of the giver and the load share.

From the insights of Section IV-B1, the receiver's arm motion is given by a slave DS that depends on the giver's motion and the load share

$$\dot{\mathbf{p}}_r = \mathbf{f}_r(\mathbf{p}_r - \mathbf{g}_r(z)) , \quad (9)$$

where

$$z = \begin{cases} \|\mathbf{p}_g - \mathbf{p}_{g,ho}\|, & \text{if } \alpha_g = 1 \quad (\text{approach}) \\ \alpha_g - 1, & \text{if } 0 \leq \alpha_g \leq 1 \quad (\text{passing}) \end{cases}$$

Again, \mathbf{f}_r is a continuous and continuously differentiable function with a single attractor given by $\mathbf{g}_r(z)$. Note that the coupling function $\mathbf{g}_r(z)$ is expressed in term of the auxiliary variable z , which encompasses both the advancement of the giver's arm during the approach phase and the transfer of the load during the passing phase. The end-effector angular motion of both parties is coupled to z in a similar manner.

To synthesize desired hand motions, instead of modeling hand configurations $\mathbf{x}_{h,g}$, $\mathbf{x}_{h,r}$ directly, we modulate the grip force during the handover with a slave DS that depends on the auxiliary variable z

$$\dot{\bar{u}}_i = f_{h,i}(\bar{u}_{h,i} - g_{h,i}(z)). \quad (10)$$

We characterize the handover dynamics in Equations (8) to (10) using data gathered in our human-human handover study. We learn the giver's arm position dynamics (Equation (8)) by estimating the joint distributions $P(\dot{\mathbf{p}}_g, \mathbf{p}_g)$ and $P(\alpha_g, \mathbf{p}_g)$ with two Gaussian Mixture Models inferred

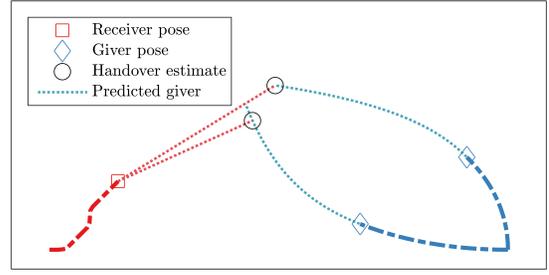


Fig. 8. Depiction of the receiver's estimation of the handover pose for two different trajectories of the giver (in blue). We model the giver's motion as a linear dynamical system, and predict the handover pose as the one closest to the receiver along the predicted trajectory.

from our study data. Functions \mathbf{g}_g and \mathbf{f}_g are given by the expected values of the respective conditional probabilities:

$$\begin{aligned} \mathbf{g}_g(\alpha_g) &= \mathbb{E}[P(\mathbf{p}_g | \alpha_g)] \\ \mathbf{f}_g(\mathbf{p}_g - \mathbf{g}_g(\alpha_g)) &= \mathbb{E}[P(\dot{\mathbf{p}}_g | (\mathbf{p}_g - \mathbf{g}_g(\alpha_g)))] . \end{aligned}$$

Estimating the remaining dynamics (Equations (9) and (10)) is done in a similar fashion. For a more detailed explanation we refer the reader to [13].

B. Low-level control

To reproduce the desired trajectories of the arm on a torque-controlled manipulator while at the same time ensuring a compliant behavior during the passing phase, we employ an impedance-based control scheme

$$\boldsymbol{\tau} = \mathbf{J}^\top (K_d(\dot{\mathbf{x}}_{d,i} - \dot{\mathbf{x}}_i) + \alpha_i \mathbf{g}) + \boldsymbol{\tau}_{\text{ext}} ,$$

where $\boldsymbol{\tau}$ is the torque input, $\boldsymbol{\tau}_{\text{ext}}$ are the external torques, \mathbf{J} is the Jacobian, K_d is a velocity-tracking gain and $\dot{\mathbf{x}}_{d,i}$ with $i \in \{g, r\}$ is the desired twist of the i -th party given by the learned translational velocities (8)(9) and angular velocities. The term $\alpha_i \mathbf{g}$ compensates for the proportion of the object's weight supported by the robot.

To simplify the control of a humanoid hand, we leverage recent work in grasp synergies [14] to project the high-dimensional control problem into a lower-dimensional subspace. We modulate the desired grip force \bar{u}_i by jointly modeling the hand-object around the contact point as a spring in synergy space [4]. A simple mapping from desired grip force to a synergy weighting (for instance, scaling the first principal component) results in a grasp that scales according to the desired grip force. While in principle this mapping is object-dependent, in our studies we only consider objects of similar size. Furthermore, for simplicity we ensure this mapping provides a conservative estimate of the force necessary to successfully grasp any object we hand over. Object-specific mappings (as well as mappings that vary over different object grasp points) remain as future work.

C. Handover pose estimation

Our handover controller formulation requires an estimate of the handover pose (the pose where both parties meet) in advance. We model the motion of the human as a dynamical system and use it to predict the handover pose continuously.



Fig. 9. We evaluated our controller on a KUKA LWR robot equipped with an Allegro hand and the ATI Gamma force/torque sensor, with an Optitrack marker on the human’s wrist. Our approach works with any graspable object with a known mass.

Using the pose data of the human for a pre-defined time window, we estimate a linear dynamical system $\dot{x}_i = Ax_i$ using least squares approximation. Using the model, the human hand motion is predicted and the position closest to the robot is estimated as the current handover location. The estimate evolves over time, and converges to the human pose, assuming that it is in the workspace of the robot. A schematic of finding the handover estimate is shown in Figure 8.

VI. HUMAN-ROBOT HANDOVER EXPERIMENTS

To validate our proposed controller we implemented it on a robotic platform and compared it to a baseline controller based on thresholds.

1) *Experimental setup*: We implemented our human-inspired controller for fluid handovers on a 7-DoF KUKA LWR 4+ robot, equipped with the SimLabs Allegro Hand (a 16-DoF humanoid hand). The robot is also equipped with a ATI Gamma force/torque sensor mounted at the wrist, which measures the interaction force on the object as it is passed from the giver to the receiver. We use an OptiTrack system to measure the 6-DoF pose of the participant’s wrist.

2) *Experimental procedure*: We evaluate the controller in handovers where the robot is the giver. The robot begins with the object in its hand. For the evaluation we used the plastic water bottle shown in Figure 9 weighting $m_o = 0.556$ [kg]. The specific grasp is not important; in our experiments we have the robot close its hand on the object, resulting in a power grasp of the object. We perform a calibration step with the robot is a static configuration to compute the baseline forces/torques from the object; this accounts for the mass of the hand and object. We then begin the handover with a signal to the human subject to move towards the robot to receive the object. We consider the handover complete when the robot no longer has the object in its hand.

3) *Experimental conditions*: We evaluate two different handover controllers on the robot:

- Our **fluid human-inspired** controller (described in Section V) which continuously estimates the handover pose and moves the robot end-effector towards it. The controller simultaneously monitors the load share estimate,

using it to control the robot hand’s grip force. Note that thanks to our decoupled formulation, the robot can still be moving when it begins to release the grasp, for example if an unknown third party should take the object as the robot is moving.

- A **threshold-based** baseline controller which switches between discontinuous phases. When the robot is the giver, it first applies a constant grip force (the same as in our controller) to grasp the object. Then the robot approaches the handover pose and stops at a distance of 0.1 [m] from it. The robot waits until the sensed load share is lower than a hand-tuned threshold of $\alpha_g = 0.2$ and then completely opening the hand in an open-loop fashion. The switching-based controller for giving the object follows the same idea. Note that since the load share is computed directly from the wrist force/torque sensor, in the absence of forces due to acceleration, the load share threshold approach is equivalent to simply detecting changes in the raw force readings along the vertical axis.

For each of these two controllers, we recorded 4 handovers performed with a single healthy subject.

4) *Measures*: To evaluate the performance of the handovers we consider two measures:

- The *internal wrench* norm between the robot and the human u_{int} computed as in (2). Internal wrenches represent counteracting wrench components which are unnecessary to accomplish the task.
- The duration of the *passing phase*, defined as the time during which both parties are supporting the object. We compute this duration as in [1, 2, 3]: the passing phase is the time during the handover where the load share is shared between the giver and receiver (i.e. $\epsilon_\alpha < \alpha_g < 1 - \epsilon_\alpha$ for a threshold ϵ_α).

Since internal wrenches represent forces on the object that serve no purpose towards the completion of the handover, they can be seen as nonessential, and thus a lower internal wrench norm is desirable to enable natural handovers. Similarly, a faster handover is desirable since both parties spend less time in contact with the object negotiating the load transfer.

A. Results

Our handover experiments indicate that our controller significantly reduces internal forces between the robot and the human, compared to the threshold-based controller. One instance of the handover dynamics for two handovers is shown in Figure 10. As we can see, waiting for the load share to reach a threshold before opening results in high internal forces on the object, whereas actively controlling the grip force and the arm using the sensed load share results in a more natural handover that remains in motion during the passing phase.

The average internal force norm and duration of the passing phase across trials are shown in Figure 11. Our proposed controller reduces internal forces w.r.t the threshold-based controller both in terms of the average of means and

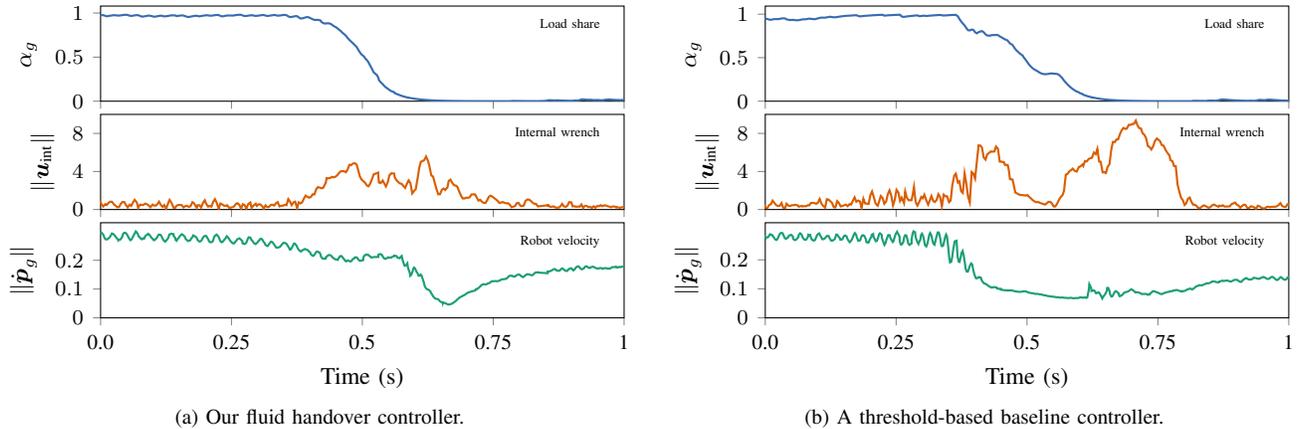


Fig. 10. Dynamics of robot-to-human handovers with our strategy (10a) and a threshold-based strategy (10b). As we can see, in both cases the giver’s load share (α_g) decreases during the handover, whereas the magnitude of the internal wrench force ($\|\mathbf{u}_{\text{int}}\|$) is much lower for our controller since we actively control the grip force (leading to a smoother change in load share). Note that since our controller formulation does not depend on a discrete phase switch, the robot is always in motion ($\|\dot{\mathbf{p}}_g\| > 0$), even during the object transfer.

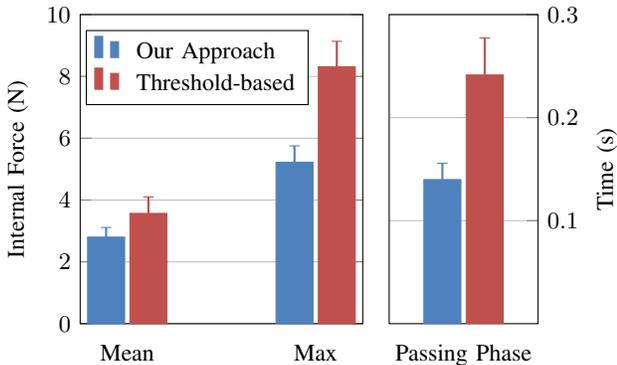


Fig. 11. Comparison of internal forces and duration of the passing phase between our fluid controller and a threshold-based controller. As we can see, our controller yields faster handovers with lower internal force, resulting in a more natural handover.

the average of maximum values. This results in handovers which require less unnecessary work. The duration of the passing phase is also significantly reduced. Continuously controlling the arm and the hand during the handover yields a more responsive interaction and a faster passing of the object as shown in the video available at <https://youtu.be/Ac4kgipC7A0>.

VII. CONCLUSIONS

In this work we presented a novel human-inspired bidirectional human-robot handover controller. Its design is supported by insights from existing studies as well as novel insights observed in our human-human experiments: our results indicate the existence of motion during the passing phase as well as a coupling between the motions of the giver and the receiver. Our controller is based on a phase-less handover dynamics model that produces smooth and fast handovers. Experiments show that our controller is smooth, fast, and reduces internal forces on the object compared to traditional switching-based approaches.

ACKNOWLEDGMENTS

We gratefully acknowledge the funding provided by the European Commission (Horizon 2020 Framework Programme, H2020-ICT-643950) for the SecondHands project. Thanks to Klas Kronander and Guillaume deChambrier for their help with the robot experiments.

REFERENCES

- [1] A. H. Mason and C. L. MacKenzie, “Grip forces when passing an object to a partner,” *Experimental Brain Research*, 2005.
- [2] W. P. Chan, C. A. Parker, H. M. Van der Loos, and E. A. Croft, “Grip forces and load forces in handovers: Implications for Designing Human-Robot Handover Controllers,” in *Human-Robot Interaction*, 2012.
- [3] S. Endo, G. Pegman, M. Burgin, T. Toumi, and A. M. Wing, “Haptics in between-person object transfer,” in *Haptics: Perception, Devices, Mobility, and Communication*, 2012.
- [4] W. P. Chan, C. A. Parker, H. M. Van der Loos, and E. A. Croft, “A human-inspired object handover controller,” *The International Journal of Robotics Research*, 2013.
- [5] M. Cakmak, S. S. Srinivasa, M. K. Lee, J. Forlizzi, and S. Kiesler, “Human Preferences for Robot-Human Hand-over Configurations,” in *International Conference on Intelligent Robots and Systems*, 2011.
- [6] J. Waldhart, M. Gharbi, and R. Alami, “Planning handovers involving humans and robots in constrained environment,” in *International Conference on Intelligent Robots and Systems*, 2015.
- [7] J. Mainprice, M. Gharbi, T. Simeon, and R. Alami, “Sharing effort in planning human-robot handover tasks,” in *Proceedings of RO-MAN*, 2012.
- [8] V. Micelli, K. Strabala, and S. S. Srinivasa, “Perception and control challenges for effective humanrobot handoffs,” in *RSS 2011 RGB-D Workshop*, 2011.
- [9] M. Prada, A. Remazeilles, A. Koene, and S. Endo, “Dynamic Movement Primitives for Human-Robot interaction: Comparison with human behavioral observation,” in *International Conference on Intelligent Robots and Systems*, 2013.
- [10] S. Erhart and S. Hirche, “Internal force analysis and load distribution for cooperative multi-robot manipulation,” *IEEE Trans. Robot.*, vol. 31, no. 5, pp. 1238–1243, 2015.
- [11] C. Granger, “Some recent development in a concept of causality,” *Journal of Econometrics*, 1988.
- [12] L. Barnett and A. Seth, “The MVGC multivariate granger causality toolbox: A new approach to granger-causal inference,” *Journal of Neuroscience Methods*, vol. 223, pp. 50–68, 2014.
- [13] A. Shukla and A. Billard, “Coupled dynamical system based arm-hand grasping model for learning fast adaptation strategies,” *Robotics and Autonomous Systems*, vol. 60, no. 3, pp. 424 – 440, 2012.
- [14] M. Santello, M. Flanders, and J. F. Soechting, “Postural hand synergies for tool use,” *The Journal of Neuroscience*, vol. 18, no. 23, pp. 10 105–10 115, 1998.