Reconstructing Human Hand Pose and Configuration using a Fixed-Base Exoskeleton

A. Pereira[1], G. Stillfried[1], T. Baker[1], A. Schmidt[1,2], A. Maier[1], B. Pleintinger[1], Z. Chen[1], T. Hulin[1], and N. Y. Lii[1]

***Abstract*— Accurate real-time estimation of the pose and configuration of the human hand attached to a dexterous haptic input device is crucial to improve the interaction possibilities for teleoperation and in virtual and augmented reality. In this paper, we present an approach to reconstruct the pose of the human hand and the joint angles of the fingers when wearing a novel fixed-base (grounded) hand exoskeleton. Using a kinematic model of the human hand built from MRI data, we can reconstruct the hand pose and joint angles without sensors on the human hand, from attachment points on the first three fingers and the palm. We test the accuracy of our approach using motion capture as a ground truth. This reconstruction can be used to determine contact geometry and force-feedback from virtual or remote objects in virtual reality or teleoperation.**

I. Introduction

Hand exoskeletons allow us to interact with virtual or remote environments intuitively. It is necessary to acquire the position and orientation of the human hand and the joint angles of the fingers, for determining contact or penetration of remote or virtual environments, positioning the slave manipulator or virtual human hand, and calculating force feedback. Several dexterous haptic interfaces use encoders on mechanical joints fixed to joints of the fingers [1], or sensors in a data glove [2], [3], which often results in bulky interface designs when force-feedback is also required. We present an approach to accurately estimate pose and configuration of the human hand from only the positions of well-chosen attachment points to the exoskeleton.

We consider a *grounded* device with *omnidirectional force feedback*. Fixed to the ground, high forces can be rendered, e.g. reaction forces from fixed objects in virtual/remote environments, and the user's weight supported, preventing fatigue. Omnidirectional feedback allows one to feel various material properties, e.g. surface friction, or even drag, lift and buoyancy forces from fluids. Examples of such systems are [1], [4], [5] and DLR's novel device Exodex Adam [6]. This (Fig. 1), is a modular exoskeleton attached to the first 3 fingers and the palm, adjustable for a range of human hands.

A. Why do we need hand pose and configuration?

Accurately rendering contact forces in a virtual environment requires knowledge of the contact geometry. Not only the contact points, but also the orientation of the fingers affects force feedback: different parts of the distal phalanx have different frictional [7] and mechanical properties [8], which affect the contact forces to be rendered. In [4], [5], the attachment point of the human to the exoskeleton is located near the fingertip, and the point of contact with remote or virtual objects can be approximated as the attachment point. In this case, no human hand model is used. Fig. 2 shows a limitation of this approach: contact with a virtual object can occur at various places on the finger; assuming a fixed contact point could lead to predicting contact incorrectly.

Furthermore, in teleoperation, mapping grasping or manipulation to a slave end effector of vastly different kinematics to the human is nontrivial: joint-to-joint mapping is not possible. In [9] the transformation on the manipulated object caused by the motion of the human hand's contacts is mapped to a desired movement of the robot hand contact points. Abstracting further, the system can estimate intention of the human grasp/manipulation from the hand pose and configuration, and choose appropriate slave motion, similar to the gesture-based manipulation approach in [10]. If the operator power-grasps, the slave can assume that grasp stability is more important than manipulability and grasp accordingly; in case of a precision grasp, the slave's grasp could prioritise easy manipulation of the object in the directions afforded by the human's grasp.

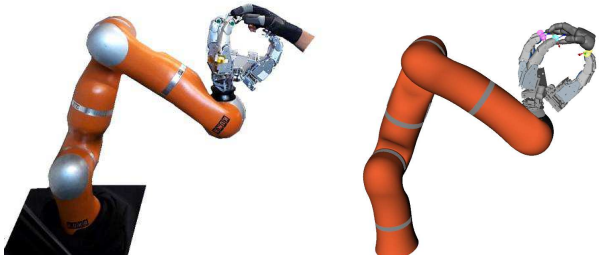

Fig. 1. Left, a user connected to Exodex Adam; right, visualisation of exoskeleton with online prediction of user's hand pose and configuration.

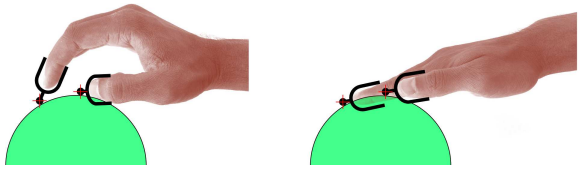

Fig. 2. The thimbles on the fingers that lock into the exoskeleton are black; red crosses are attachment points to the exoskeleton, at identical positions in both cases. Left: 2nd finger does not touch the object but thumb does; right: 2nd finger is inside the object while thumb does not touch it. Hence attachment points alone are insufficient to render contact.

[1] Institute of Robotics and Mechatronics, German Aerospace Centre (DLR), 82234 Weßling, Germany. Contact: aaron.pereira@dlr.de

[2] Department of Mechanical, Maritime and Materials Engineering, Delft University of Technology, Mekelweg 2, Delft, Netherlands

B. Why attach to the palm as well as the fingertips?

Firstly, the palm position is needed to determine the hand pose: the fingertip positions alone are insufficient. [5] measures the distance from device to palm with ultrasonic sensors, but it is not known if this suffices to reconstruct the hand pose accurately. By adding attachment points on the palm, as in the exoskeleton Exodex Adam, its position in space can be found and used in reconstructing the hand pose and configuration (see Fig. 2)

Secondly, most grounded dexterous exoskeletons available today, such as the aforementioned, ignore force feedback on the palm. However, the palm is crucial in-hand manipulation and manual exploration [11], [12]. Applying forces to the palm can significantly improve the range of haptic feedback and immersiveness of the user experience. The user can execute, and have feedback from, both precision grasps with the fingertips and power grasps with the whole hand.

C. Approach Overview

We find the pose and configuration of the human hand only using the 5 positions of the attachment points to the exoskeleton. These are found using forward kinematics of the exoskeleton and robot arm to which it is attached. They are fitted to a kinematic model of the human hand, to derive its joint angles and pose. Since there are more degrees of freedom (DOFs) in the human hand than those constrained at the attachment points to the exoskeleton, there is redundancy; we optimise away from the human hand joint limits in the nullspace to find the most "natural" position.

Relevant to our study is [13], an evaluation of the CyberGlove. A measurement is a "hit" if the CyberGlove measures within a given tolerance of the ground truth. We compare our mean absolute joint error and hit rate to the CyberGlove. Where in [13] ground truth is determined using physical fixtures along which the subjects were instructed to align their fingers, we find ground truth from motion tracking.

II. Hand Kinematic Model

We use a human hand model from [14] developed by some of the authors, derived from magnetic resonance imaging (MRI) data using the methods presented in [15]. The full kinematic model in [14] has 22 DOFs and the pose of the hand in space adds an additional 6 (we modelled the translation and rotation of the hand's base coordinate system by 3 orthogonal prismatic joints starting at the fixed base of the robot arm, followed by 3 revolute joints). We made some simplifications to reduce dimensionality. Since the 4$^{\text{th}}$ and 5 $^{\text{th}}$ fingers are not attached to the exoskeleton, we disregard them. Fig. 3 shows the resulting model. Joints DIP2 and DIP3 are the *distal interphalangeal joints*; joints PIP2 and PIP3 the *proximal interphalangeal joints*; joint IP1 the *interphalangeal joint* of the thumb; and the *metacarpophalangeal joints*, i.e. the 2-DOF joints between metacarpals and proximal phalanges, are designated by MCPX.Y, where X is the number of the finger and Y is 1 for abduction/adduction and 2 for flexion/extension. IMC and CM stand for *intermetacarpal* and *carpometacarpal* respectively.

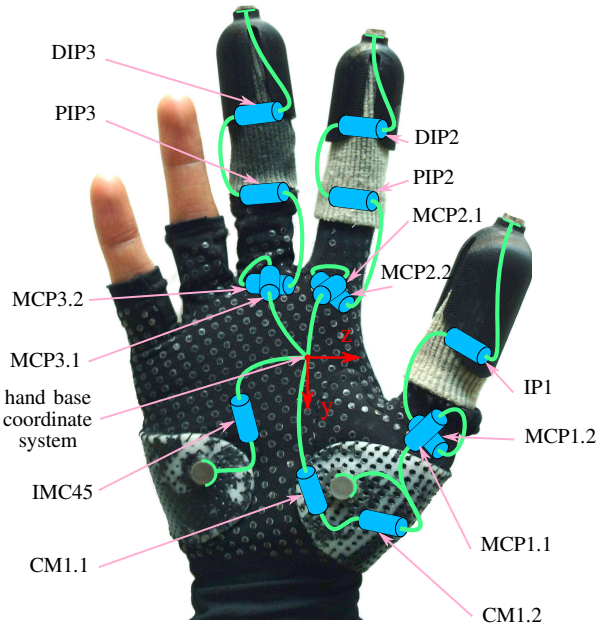

Fig. 3. Schematic of hand model, simplified to exclude 4$^{\text{th}}$ and 5$^{\text{th}}$ fingers. Cylinders are revolute joints, green lines are links. Not shown are the three prismatic and revolute joints linking the world coordinate system with the base coordinate system of the hand (the latter shown here in red). The hand is in a glove to which flexible plastic parts with magnets are attached; index and middle fingers and thumb are in finger sleeves with magnets at the ends.

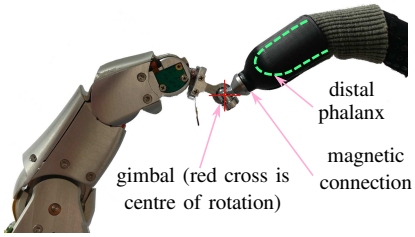

Fig. 4. Attachment to Exodex Adam. Magnets on the subject's glove/finger sleeves attach to gimbals on the exoskeleton, forming a rigid connection from the distal phalanges/the palm to the centre of the gimbals. The positions of the gimbals' centres are found from the exoskeleton forward kinematics.

The distal and proximal interphalangeal joint angles on the index and middle fingers have often been assumed proportional; a ratio of $\frac{2}{3}$ (DIP:PIP) is derived from empirical measurements in [16]; a ratio of $\frac{1}{2}$ is suggested for power grasps in [17]. We assume the DIP joint angle to be the PIP joint angle scaled by a factor k; how to set k is investigated in Sec. III. This brought the DOFs of the hand model to 18.

A. Attachment Points on Exoskeleton

The human hand is connected to the exoskeleton at the tips of the thumb, the index and middle fingers, and at two points in the palm. The exoskeleton consists of five identical robot "fingers" from the Five-Fingered Hand [18], attached to an adjustable, but non actuated base, which in turn is attached to a KUKA LWR 4+. Since this paper focusses on the approach of deriving the the human hand pose, the exoskeleton design is not discussed in further detail; readers are referred to [6].

The attachment to the exoskeleton is through a 3-axis gimbal mounted on a fixed-directional magnetic clutch, see Fig. 4, allowing a wide range of orientations and placed such as to avoid gimbal lock during operation. The human's index and middle fingers and thumb fit snugly into silicone coated sleeves with a hard plastic thimble on the end. They are long enough that the distal phalanx does not slip or rotate with respect to the thimble, but short enough to constrain movement of the distal interphalangeal joint of the finger as little as possible. The human also wears a glove to which two rigid but flexible plastic parts are attached as shown in Fig. 3. Magnets attached to these plastic parts and to the end of the thimbles lock into the gimbals on the exoskeleton. The position of the gimbals' centre of rotation can be determined by forward kinematics from the robot arm's and robot fingers' joint positions. These are referred to as the "attachment points" of the human and the exoskeleton.

B. Terminology

The elements of $\boldsymbol{q}_{\text{pose}}$ are the x, y and z displacements, followed by roll, pitch and yaw angles, of the hand base coordinate system with respect to the world coordinate system. The *configuration* $\boldsymbol{q}_{\text{conf}} \in \mathbb{R}^{12}$ refers to the angles of the joints in the hand; in order (see Fig. 3), CM1.1, CM1.2, MCP1.1, MCP 1.2, IP1, MCP2.1, MCP2.2, PIP2, MCP3.1, MCP3.2, PIP3, IMC45. Values of DIP2 and DIP3 are the values PIP2 and PIP3 scaled by k, as described in Sec. II. The *state* $\boldsymbol{q} = [\boldsymbol{q}_{\text{pose}}^\top, \boldsymbol{q}_{\text{conf}}^\top]^\top \in \mathbb{R}^{18}$ refers to the combination of both. The positions of the i^{th}attachment point of the exoskeleton $\boldsymbol{x}_i \in \mathbb{R}^3$ is measured at the gimbal centre of rotation. The vector of attachment point positions in the world coordinate frame of the exoskeleton is defined: $\boldsymbol{x} = [\boldsymbol{x}_1^\top, \boldsymbol{x}_2^\top, ...\boldsymbol{x}_5^\top]^\top \in \mathbb{R}^{15}$.

C. Inverse Kinematics

We use the iterative *Jacobian Transpose Method* [19], [20] for inverse kinematics calculation, where the error between the desired attachment points (i.e. those measured on the Exodex Adam exoskeleton) and the actual attachment points (derived from forward kinematics on the hand) is iteratively minimised. However, we employ a two-step approach, first adjusting the hand pose, then the hand configuration.

Explicitly, when $\boldsymbol{q}'$ is hand state at the previous timestep, and $\boldsymbol{x}$ is the vector of attachment point positions at the current timestep, $\boldsymbol{J}_{\text{pose}}$ and $\boldsymbol{J}_{\text{conf}}$ are the partial derivatives matrices of $\boldsymbol{x}$ with respect to $\boldsymbol{q}_{\text{pose}}$ and $\boldsymbol{q}_{\text{conf}}$, and $\boldsymbol{A}_{\text{pose}}$ and $\boldsymbol{A}_{\text{conf}}$ are gain matrices, the hand state $\boldsymbol{q}$ at the current timestep is updated as in Alg. 1. It was found that this two-step approach allowed higher gains $\boldsymbol{A}_{\text{pose}}$ and $\boldsymbol{A}_{\text{conf}}$ without causing instability (and hence faster convergence), than updating all elements of the state at once.

In line 10, the joint values are saturated at the maximum and minimum allowed values $\boldsymbol{q}_{\min}$ and $\boldsymbol{q}_{\max}$ given in [14]. However, the minimum values of PIP2 and PIP3 are set to a small value above zero (0.03 rad), disallowing full extension and hyperextension. Near full extension, the restoring values $\boldsymbol{A}_{\text{conf}}\boldsymbol{J}_{\text{conf}}(\hat{\boldsymbol{q}})\boldsymbol{\delta x}$ when moving back into flexion in line 7 would be small (since this is essentially a singularity) and convergence would be slow; in hyperextension, these would try to extend rather than flex the joints.

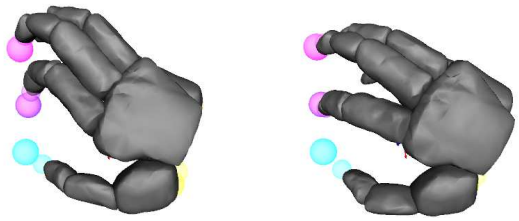

Fig. 5. Left: a hand without nullspace optimisation in an unnatural position; right: a hand with a correction acting on all finger joints to pull away from joint limits, projected into the nullspace.

Algorithm 1 Update hand state

Require: observed attachment point positions $\boldsymbol{x}$, previous hand state $\boldsymbol{q}'$

Ensure: updated hand state $\boldsymbol{q}$

1: $\boldsymbol{x}_{\text{calc}} \leftarrow$ `forwardKinematics`$(\boldsymbol{q}')$ // *calculate positions of attachment points on hand given previous state*
2: $\boldsymbol{\delta x} \leftarrow \boldsymbol{x} - \boldsymbol{x}_{\text{calc}}$ // *deviation measured to calculated*
3: $\boldsymbol{q}_{\text{pose}} \leftarrow \boldsymbol{q}'_{\text{pose}} + \boldsymbol{A}_{\text{pose}}\boldsymbol{J}_{\text{pose}}(\boldsymbol{q}')^\top \boldsymbol{\delta x}$
4: $\hat{\boldsymbol{q}} \leftarrow [\boldsymbol{q}_{\text{pose}}^\top, \boldsymbol{q}'^\top_{\text{conf}}]^\top$
5: $\hat{\boldsymbol{x}}_{\text{calc}} \leftarrow$ `forwardKinematics`$(\hat{\boldsymbol{q}})$ // *calculate positions of attachment points on hand given intermediate state*
6: $\boldsymbol{\delta}\hat{\boldsymbol{x}} \leftarrow \boldsymbol{x} - \hat{\boldsymbol{x}}_{\text{calc}}$ // *deviation measured to calculated*
7: $\boldsymbol{q}_{\text{conf}} \leftarrow \boldsymbol{q}'_{\text{conf}} + \boldsymbol{A}_{\text{conf}}\boldsymbol{J}_{\text{conf}}(\hat{\boldsymbol{q}})^\top \boldsymbol{\delta}\hat{\boldsymbol{x}}$
8: $\boldsymbol{\delta q}_{\text{cor,N}} \leftarrow$ `nullspaceCorrection`$(\boldsymbol{J}(\hat{\boldsymbol{q}}), \hat{\boldsymbol{q}}, \boldsymbol{q}_{\max}, \boldsymbol{q}_{\min})$
9: $\boldsymbol{q}_{\text{unsat}} \leftarrow [\boldsymbol{q}_{\text{pose}}^\top, \boldsymbol{q}_{\text{conf}}^\top]^\top + \boldsymbol{\delta q}_{\text{cor,N}}$
10: $\boldsymbol{q} \leftarrow$ `saturateAtJointLimits`$(\boldsymbol{q}_{\text{unsat}}, \boldsymbol{q}_{\max}, \boldsymbol{q}_{\min})$

The intermediate values $\hat{\boldsymbol{q}}$ were used for the nullspace optimisation in line 8, in order that $\boldsymbol{J}_{\text{conf}}$ need not be recalculated with the new values of $\boldsymbol{q}$ – this is the most computationally expensive step, apart from the calculation of the nullspace correction. The latter is explained next.

D. Nullspace Optimisation

With $3 \times 5 = 15$ constraints for an 18-DOF system, the inverse kinematics problem is under-defined; i.e. there are several possible inverse kinematics solutions for a given set of exoskeleton positions. See Fig. 5: sometimes the inverse kinematics returns unnatural solutions, typically with one or more human hand joints at or near position limits. For this reason, nullspace optimisation was run in parallel to move the fingers away from the joint limits (line 8, Alg. 1). A correction acting on all finger joints is defined as follows:

$$\delta q_{i,\text{cor}} = \alpha_i \frac{q_{i,\max} + q_{i,\min} - 2q_i}{q_{i,\max} - q_{i,\min}},$$

where α_i is a "stiffness" for the i^{th} joint and q_i, $q_{i,\max}$, $q_{i,\min}$, and $\delta q_{i,\text{cor}}$ are the i^{th} elements of the vectors $\boldsymbol{q}$, $\boldsymbol{q}_{\max}$, $\boldsymbol{q}_{\min}$, and $\boldsymbol{\delta q}_{\text{cor}}$, respectively. This is then projected into the nullspace: $\boldsymbol{\delta q}_{\text{cor,N}} = \boldsymbol{N}\boldsymbol{\delta q}_{\text{cor}}$, and added to the finger joint and base angles (line 9 of Alg. 1); the effect is that the

finger joints move away from their limits while the positions of the attachment points stay the same. The static nullspace projector $\boldsymbol{N}$ was calculated as in [21]:

$$\boldsymbol{N} = \boldsymbol{I} - \boldsymbol{J}^{\top}(\boldsymbol{J}^{\boldsymbol{W}+})^{\top}, \tag{1}$$

$$\boldsymbol{J}^{\boldsymbol{W}+} = \boldsymbol{W}^{-1}\boldsymbol{J}^{\top}(\boldsymbol{J}\boldsymbol{W}^{-1}\boldsymbol{J}^{\top})^{-1}, \tag{2}$$

where $\boldsymbol{J}^{\boldsymbol{W}+}$ is the weighted pseudoinverse of the Jacobian (the dependency on $\boldsymbol{q}$ is dropped for brevity) with weighting matrix $\boldsymbol{W}$. Since the dimension of the nullspace is 3 (18 DOF - 15 constraints), different projections are possible. A weighted inverse can yield a projection where $\boldsymbol{\delta q}_{cor,N}$ corrects the joint angles in the human hand away from their limits, more than correcting its pose.

The *unweighted* pseudoinverse (i.e., $\boldsymbol{W}$ is the identity matrix $\boldsymbol{I}$) can be calculated in polynomial time using single-value decomposition (SVD) [22]. We define:

$$\boldsymbol{W}^{-1} = \boldsymbol{\Omega}\boldsymbol{\Omega}^{\top}, \quad \boldsymbol{H} = \boldsymbol{J}\boldsymbol{\Omega},$$

and (2) reduces to an unweighted inverse:

$$\boldsymbol{J}^{\boldsymbol{W}+} = \boldsymbol{\Omega}\boldsymbol{H}^{\top}(\boldsymbol{H}\boldsymbol{H}^{\top})^{-1} = \boldsymbol{\Omega}\boldsymbol{H}^{\boldsymbol{I}+} \tag{3}$$

If $\boldsymbol{W}$ is chosen as a diagonal matrix with entries $[w_1, w_2, ...]$ along the diagonal, then $\boldsymbol{\Omega}$ is also a diagonal matrix with entries $[w_1^{-0.5}, w_2^{-0.5}, ...]$. We chose $w_i^{-0.5}$ for $i \in \{1, ..., 6\}$ (i.e. the elements relating to pose) and for IMC45 as 100, for PIP2 as 10, and for the rest as 1. The stiffness α_i was tuned to 0.1 for CM1.1 and 1.2, and MCP1.2; 0.04 for MCP2.1 and 2.2, 0.02 for MCP3.1 and 3.2, and 0.01 for all other joints.

Despite the high dimensionality, with this rearrangement we were able to calculate the nullspace projector in parallel with the robot control (the computation took too long to run in the same model as the inverse kinematics and robot control in real time). Note: by using the SVD to find the pseudoinverse of $\boldsymbol{J}$ and by approximating it in Sec. II-C by its transpose, we avoid calculating inverses of large matrices.

III. Evaluation

The accuracy of the inverse kinematics prediction is tested experimentally. It is challenging to obtain a ground truth measurement of the human hand when attached to such an exoskeleton. Vision and depth-based systems, such as LeapMotion[1] perform poorly, since the equipment attached to the hand makes it difficult for the software to identify the fingers. Instead, we used a 5-camera infrared motion-tracking system from Vicon Ltd[2] along with retroreflective markers arranged into rigid patterns, affixed to rigid bodies on the hand, as detailed in [23] and shown in Fig. 6.

A. Robot Control

The Exodex Adam exoskeleton is torque controlled and is compliant to human force cues: when the human moves their arm or fingers, the exoskeleton, including the robot arm, moves with them with minimal resistance. The system employs gravity compensation, meaning the weight of the device is not carried by the user, and the measured torques are fed forward with a gain to the commanded torques, meaning the perceived inertia of the device is lower than its actual inertia as in [24]. The fingers additionally have friction compensation as described in [25]. A full discussion of the control is outside the scope of this paper.

[1] `www.leapmotion.com` retrieved 20.7.18

[2] `www.vicon.com`, retrieved 20.7.18

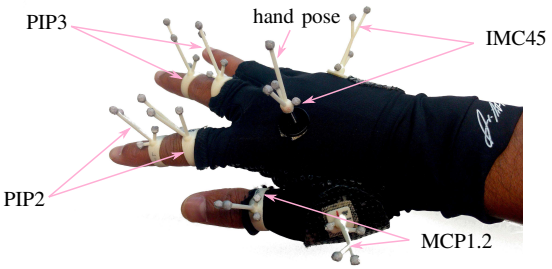

Fig. 6. Markers placed on the human hand and the joints they measure.

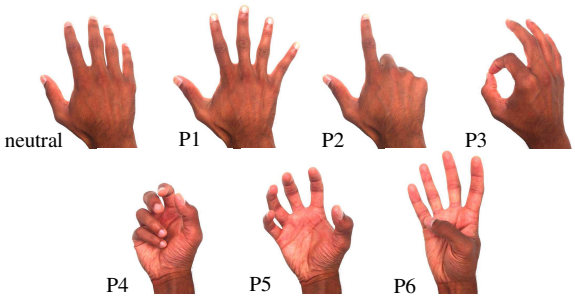

Fig. 7. Hand positions. Top, from left to right: neutral position, fingers outstretched (P1), pointing with index finger (P2), diver's "OK" (P3). Bottom, from left to right: power-grasp of small object (P4), power-grasp of larger object (P5), touch base of pinky with tip of thumb (P6).

B. Experiment Methodology

A standard set of hand positions and gestures which were judged relevant to telemanipulation and interaction with a virtual environment were chosen:

P1	Fingers outstretched
P2	Pointing with index finger
P3	Diver's "OK" finger
P4	Power-grasp of small (e.g. apricot-sized) object
P5	Power-grasp of larger (e.g. grapefruit-sized) object
P6	(Try to) touch base of pinky with tip of thumb

These are all shown in Fig. 7 and demonstrated in the video. Subjects start in a neutral position, move to the first position (or as close as possible), back into the neutral position, then to the next position, etc. The parts of the human hand state measured by the motion-tracking were:

- hand base coordinate system ($\boldsymbol{q}_{\text{pose}}$)
- metacarpophalangeal flexion joint, thumb (MCP1.2)
- proximal interphalangeal joint, index finger (PIP2)
- proximal interphalangeal joint, middle finger (PIP3)
- cupping joint in the palm (IMC45)

These were chosen since they led to the most robust tracking in pre-trials, have rigid bodies to which the markers can be

attached without much skin movement, and have only one degree of freedom (e.g. the metacarpophalangeal joints of the fingers, which have both flexion and abduction, were excluded) – an exception is the metacarpophalangeal joint of the thumb, since abduction/adduction is limited in motion. While theoretically possible to track these 5 simultaneously, the rate of misidentification of markers due to occlusion made this impractical. Instead, each was tracked individually.

Tracking was performed at 100Hz. We recorded the ground truth from the motion capture as well as the positions of the exoskeleton attachment points. Offline, we ran the inverse kinematics algorithms using different values of k, and with and without nullspace optimisation, with the same input values (i.e. the time series of the attachment point positions). This was performed as if it were running in real time, i.e. at 833Hz and with the values from the nullspace optimisation delayed by one cycle period as described in Sec. II-D.

Subjects were 4 males and 1 female, median age 25, none of whom were the original subject used to create the model in [14]. Lengths in the hand model were scaled by the ratio of the subject's hand length (distance from wrist to tip of middle finger) to that of the original subject.

C. Normalising Measurements

Since there were variations in the exact positioning of the markers on the hand of each subject, the offset between the zero positions of the measured angles and the calculated angles was taken as the average offset during the neutral positions prior to each of the 6 positions. There are offsets between the frame of the marker measuring hand pose and of the hand base coordinate system, and between the motion tracking frame and the exoskeleton world frame. A measured transformation matrix of the marker position $\boldsymbol{T}_m$ and a calculated position $\boldsymbol{T}_c$ are related by $\boldsymbol{T}_c = \boldsymbol{A}\boldsymbol{T}_m\boldsymbol{B}$, where $\boldsymbol{A}$ represents the transformation from the world frame of the exoskeleton to the motion-tracking frame. $\boldsymbol{B}$ represents the pose offset of the marker placed on the human hand. The matrices $\boldsymbol{A}$ and $\boldsymbol{B}$ are estimated by regression on $\boldsymbol{T}_c$ and $\boldsymbol{T}_m$ at the neutral positions, for each subject.

D. Experiment Results

Tab. I shows mean absolute error and hit rate (defined in Sec. I-C) averaged over the 6 held positions for each subject, both with and without nullspace correction, and for different k values. We took the tolerance as $\pm 15°$, as this was used in [13] for joints MCP1.2, PIP2 and PIP3. For the base coordinate system, we find mean absolute position error and angular error from the measured to the calculated (corrected with $\boldsymbol{A}$ and $\boldsymbol{B}$ as described above). These are shown in Tab. II, broken down by movement. In all cases, values are averaged over the 1 second during which the positions are held. Fig. 8 shows the tracking of joint PIP2 of subject 3 and Fig. 9 shows the base coordinate system position, roll pitch and yaw angles, positional error and angular error.

IV. Discussion

The finger angles had a mean absolute joint error of between $7 - 12°$. For the thumb, index and middle fingers,

TABLE I

Comparison of mean absolute error (deg, above) and hit rate (%, below) with (N.S) and without (X) nullspace, for differing values of k. Highest hit rates in bold

	MCP1.2		PIP2		PIP3		IMC45	
k	X	N.S.	X	N.S.	X	N.S.	X	N.S.
0	16.3	8.1	7.1	7.7	10.3	11.5	9.3	8.2
0.167	15.7	8.4	7.1	7.5	10.9	11.6	9.0	8.1
0.333	15.1	8.3	7.2	7.3	11.5	11.5	8.8	8.0
0.5	14.6	8.2	7.3	7.4	12.1	11.2	9.2	8.0
0.667	14.6	8.3	7.4	7.4	12.0	11.1	9.0	7.9
0.833	14.6	9.0	7.7	7.7	12.0	11.0	8.9	7.8
1	14.6	9.0	8.0	7.8	11.9	10.7	8.7	7.8

	MCP1.2		PIP2		PIP3		IMC45	
k	X	N.S.	X	N.S.	X	N.S.	X	N.S.
0	50.0	83.3	83.3	90.0	73.3	63.3	76.7	83.3
0.167	53.3	**86.7**	83.3	90.0	70.0	70.0	80.0	83.3
0.333	60.0	**86.7**	83.3	**93.3**	63.3	73.3	80.0	83.3
0.5	60.0	**86.7**	83.3	**93.3**	60.0	73.3	80.0	**86.7**
0.667	63.3	**86.7**	83.3	**93.3**	60.0	**76.7**	80.0	83.3
0.833	63.3	**86.7**	76.7	90.0	60.0	73.3	80.0	83.3
1	66.7	**86.7**	76.7	86.7	60.0	73.3	80.0	83.3

TABLE II

Mean absolute error in position (mm, above) and orientation (deg, below), for differing values of k, and (end column) without nullspace optimisation

k	P1	P2	P3	P4	P5	P6	Average	w/o N.S.
0	12.5	18.7	17.3	19.1	11.1	27.7	17.8	17.6
0.167	12.0	17.8	16.9	18.8	10.4	27.7	17.3	17.5
0.333	11.3	17.2	16.5	18.2	9.7	28.2	16.9	17.5
0.5	11.1	17.0	16.3	17.8	9.8	28.4	16.7	17.5
0.667	10.9	17.0	16.4	17.6	10.0	28.4	16.7	17.6
0.833	10.9	17.1	16.6	17.7	10.1	28.7	16.9	17.6
1	11.1	17.3	16.7	17.8	10.3	28.9	17.0	17.8

k	P1	P2	P3	P4	P5	P6	Average	w/o N.S.
0	19.8	26.4	17.2	25.0	14.2	22.9	20.9	20.0
0.167	18.5	25.3	16.5	23.9	13.2	21.4	19.8	19.9
0.333	17.1	24.3	15.9	22.6	12.4	20.3	18.7	19.7
0.5	16.5	23.6	15.5	21.4	11.7	19.7	18.1	19.4
0.667	15.5	22.8	14.8	20.6	11.4	19.8	17.5	19.4
0.833	15.5	22.7	15.4	20.4	11.4	19.9	17.5	19.3
1	15.6	22.7	16.1	20.8	11.4	20.1	17.8	19.3

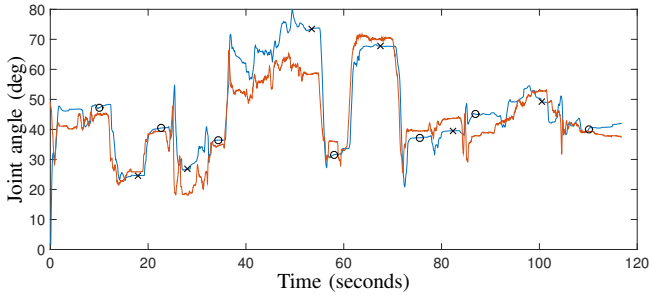

Fig. 8. Tracking joint PIP2 on subject 3, $k = \frac{2}{3}$. Ground truth is orange, calculated value is blue. ○ are neutral positions, × are positions P1–P6 in order.

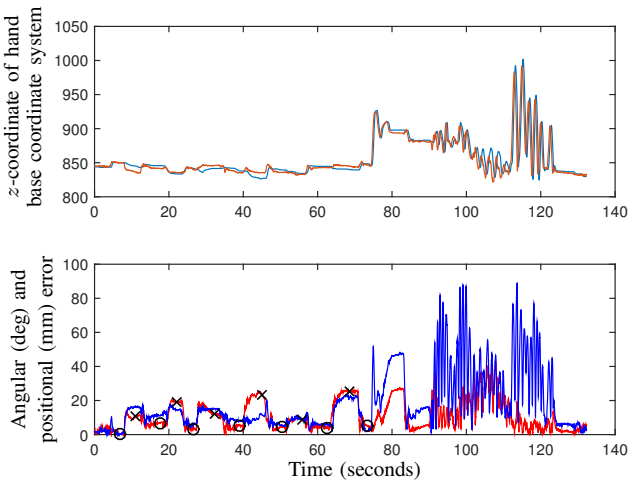

Fig. 9. Tracking of the base coordinate frame, subject 4, $k = \frac{2}{3}$. Above: z-coordinate tracked (orange) and calculated (blue). Below, position error (blue) and angular error (red). ○ are neutral positions, × are the positions P1–P6 in order. At 90s, the subject moves their hand around quickly; the lag between the calculated and actual positions results in an error.

this is a good result; for the cupping angle IMC45, the range of joint motion in the model is only 29°, so we cannot confidently state that the value of 8° error is significantly better than random. Though error does not improve in PIP2 and PIP3 with nullspace optimisation, hit rate does.

In Tables I and II, it appears that no single value of k minimises the error. However, hit rates were maximal using nullspace optimisation at $k = \frac{2}{3}$, supporting its use in the literature. In the experiments, the joint limits of the exoskeleton's robotic fingers were often reached, particularly in positions P1 and P6, and a virtual spring is implemented in the control so that the joints never reach their limits. The human therefore perceives a force at the attachment points, which was seen to deform the distal joints (PIP2, PIP3, IP1). A potential solution would be to increase the length of the finger thimbles such that the distal joints are constricted, but this would hinder free movement.

Compared to CyberGlove, our approach performs well. On the CyberGlove, hit rates between $86-100\%$ were observed on the PIP2 and PIP3 joints and between $33-96\%$ for the MCP1.2 (the value of 33% was at 90° flexion). Our sensorless approach has hit rates of 93.3% and 76.7% on PIP2 and PIP3 and 86.7% on MCP1.2. The accuracy on the thumb metacarpophalangeal joint is particularly promising, given the thumb's central role in manipulation [26]. Nullspace optimisation appears to improve hit rate for all joints.

For the pose of the hand, the worst tracking was observed in P6 in the position and P2 in the angle. P6 is a difficult position to recreate on the exoskeleton and it may be that the human hand had to contort into unnatural positions, which were not found by the algorithm.

1) Unreliable Ground Truth: During high flexion angles in the 2nd and 3nd fingers, the markers often tended to slip such that the angle was undermeasured, contributing to poor tracking. Some markers could also not be placed directly on the skin, but only on the glove, so slippage occurred between glove and skin. Finally, to measure the thumb angle, one marker was on the plastic magnet attachment on the palm, which often did not flex sufficiently when the thumb flexed.

2) Evaluation and Possibilities for Improvement: Judging from our previous work in gesture-commanded teleoperation [10] and preliminary trials, this accuracy is sufficient for gesture and intention recognition, see video attachment. With regard to teleoperation and virtual reality (VR), since the attachment points to the exoskeleton are close the human fingertips, and these can be determined from the joint encoders very precisely, the joint angle error leads to a small position error in the attachment with virtual/remote objects.

This accuracy was achieved with no sensor attached to the hand. It could be further improved by adding angular sensors to the gimbal at the exoskeleton attachments, so that the orientation of the distal phalanges and the points on the palm could be measured. Sensor fusion with e.g. a data glove is expected to improve accuracy, in particular near singularities (e.g. full extension of the human fingers), in hyperextension of the fingers, and in measuring the joint IMC45.

The approach is iterative, so the hand state estimate is updated incrementally. This leads to a lag between measured and computed state, visible in Fig. 9. Increasing computing power to compute more than one iteration per control cycle, and tuning gains (or adding an integral gain term) could improve tracking performance. Additionally, the flexibility of the individual's fingers played a large role. While subject 2 could not hyperextend finger joints, in subject 3 hyperextension was often observed. This could be addressed by 1) considering the DIP2 and DIP3 joints as independent degrees of freedom and using sensors as mentioned above to measure them, or 2) restricting hyperextension with a mechanical stop in the glove–however, this restricts free human hand movement, which is a key feature of this exoskeleton.

Finally, Exodex Adam, the exoskeleton we use, has the capacity for an extra robot finger. This could be attached to the 5th finger of the human, allowing this finger's joints to be estimated. Inter-finger relationships in [17] could be exploited to estimate joint angles in the 4th finger.

V. Conclusion and Further Work

We present and evaluate an approach for real-time estimation of human hand pose and configuration on a grounded exoskeleton connected to the fingers and the palm. Inverse kinematics on a redundant system is augmented by nullspace optimisation, which improves accuracy. Possibilities for further improving accuracy and extending the approach are identified. The next step is using these predictions to provide realistic force feedback in a variety of virtual and remote environments.

Acknowledgment

The authors would like to thank our colleagues at DLR: Hannes Höppner for advice on marker tracking, Maxime Chalon for help with hardware, and Alexander Dietrich for advice regarding nullspace projections.

REFERENCES

[1] A. Frisoli, F. Rocchi, S. Marcheschi, A. Dettori, F. Salsedo, and M. Bergamasco, "A new force-feedback arm exoskeleton for haptic interaction in virtual environments," in *World Haptics Conference*, 2005, pp. 195–201.

[2] M. Borghetti, E. Sardini, and M. Serpelloni, "Sensorized glove for measuring hand finger flexion for rehabilitation purposes," *IEEE Transactions on Instrumentation and Measurement*, vol. 62, no. 12, pp. 3308–3314, 2013.

[3] T. L. Baldi, S. Scheggi, L. Meli, M. Mohammadi, and D. Prattichizzo, "Gesto: A glove for enhanced sensing and touching based on inertial and magnetic sensors for hand tracking and cutaneous feedback," *IEEE Trans. Human-Machine Systems*, vol. 47, no. 6, pp. 1066–1076, 2017.

[4] T. Endo, H. Kawasaki, T. Mouri, Y. Ishigure, H. Shimomura, M. Matsumura, and K. Koketsu, "Five-fingered haptic interface robot: Hiro iii," *IEEE Transactions on Haptics*, vol. 4, no. 1, pp. 14–27, 2011.

[5] T. Endo and H. Kawasaki, "Collision avoidance and its experimental investigation for a side-faced-type multi-fingered haptic interface," in *IEEE Int. Conf. Systems, Man, and Cybernetics*, 2014, pp. 3984–3989.

[6] N. Y. Lii, G. Stillfried, Z. Chen, M. Chalon, B. Pleitinger, and A. Maier, "Handexoskelett sowie Roboterarm mit solchem Handexoskelett," Patent DE102 017 220 996.8, 23-11-2017, pending.

[7] H.-Y. Han, A. Shimada, and S. Kawamura, "Analysis of friction on human fingers and design of artificial fingers," in *Proc. Int. Conference on Robotics and Automation*, vol. 4, 1996, pp. 3061–3066.

[8] F. Martinot, A. Houzefa, M. Biet, and C. Chaillou, "Mechanical responses of the fingerpad and distal phalanx to friction of a grooved surface: Effect of the contact angle," in *14th Symp. Haptic Interfaces for Virtual Environment and Teleoperator Systems*, 2006, pp. 297–300.

[9] G. Salvietti, L. Meli, G. Gioioso, M. Malvezzi, and D. Prattichizzo, "Multicontact bilateral telemanipulation with kinematic asymmetries," *IEEE/ASME Trans. Mechatronics*, vol. 22, no. 1, pp. 445–456, 2017.

[10] N. Y. Lii, Z. Chen, M. A. Roa, A. Maier, B. Pleintinger, and C. Borst, "Toward a task space framework for gesture commanded telemanipulation," in *IEEE Int. Symp. Robot and Human Interactive Communication (RO-MAN)*, 2012, pp. 925–932.

[11] A. Chabrier, F. Gonzalez, F. Gosselin, and W. Bachta, "Analysis of the directions in which forces are applied on the hand during manual manipulation and exploration," in *Proc. IEEE World Haptics Conf.*, 2015, pp. 280–285.

[12] K. Hirota and K. Tagawa, "Interaction with virtual object using deformable hand," in *Proc. IEEE Virtual Reality*, 2016, pp. 49–56.

[13] G. D. Kessler, L. F. Hodges, and N. Walker, "Evaluation of the cyberglove as a whole-hand input device," *ACM Trans. Comput.-Hum. Interact.*, vol. 2, no. 4, pp. 263–283, 1995.

[14] R. G. Camarero, T. Hulin, and B. Vodermayer, "The stamas simulator: A kinematics and dynamics simulator for an astronaut's leg and hand exoskeleton," in *STAMAS Workshop - Smart technology for artificial muscle applications in space*, 2015.

[15] G. Stillfried, U. Hillenbrand, M. Settles, and P. van der Smagt, "Mri-based skeletal hand movement model," in *The Human Hand as an Inspiration for Robot Hand Development*, R. Balasubramanian and V. J. Santos, Eds. Springer, 2014, pp. 49–75.

[16] H. Rijpkema and M. Girard, "Computer animation of knowledge-based human grasping," *SIGGRAPH Comput. Graph.*, vol. 25, no. 4, pp. 339–348, 1991.

[17] S. Cobos, M. Ferre, M. A. S. Uran, J. Ortego, and C. Pena, "Efficient human hand kinematics for manipulation tasks," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots and Systems*, 2008, pp. 2246–2251.

[18] H. Liu, K. Wu, P. Meusel, N. Seitz, G. Hirzinger, M. H. Jin, Y. W. Liu, S. W. Fan, T. Lan, and Z. P. Chen, "Multisensory five-finger dexterous hand: The DLR/HIT hand II," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots and Systems*, 2008, pp. 3692–3697.

[19] A. Balestrino, G. D. Maria, and L. Sciavicco, "Robust control of robotic manipulators," *IFAC World Congress*, vol. 17, no. 2, pp. 2435 – 2440, 1984.

[20] W. A. Wolovich and H. Elliott, "A computational technique for inverse kinematics," in *IEEE Conf. Decision and Control*, 1984, pp. 1359–1363.

[21] A. Dietrich, C. Ott, and A. Albu-Schäffer, "An overview of null space projections for redundant, torque-controlled robots," *The International Journal of Robotics Research*, vol. 34, no. 11, pp. 1385–1400, 2015.

[22] G. H. Golub and C. F. van Loan, *Matrix Computations*, 4th ed. Johns Hopkins Univ. Press, 2013.

[23] D. Gierlach, A. Gustus, and P. van der Smagt, "Generating marker stars for 6d optical tracking," in *IEEE RAS EMBS Int. Conf. Biomedical Robotics and Biomechatronics (BioRob)*, 2012, pp. 147–152.

[24] K. Hertkorn, T. Hulin, P. Kremer, C. Preusche, and G. Hirzinger, "Time domain passivity control for multi-degree of freedom haptic devices with time delay," in *Proc. IEEE Int. Conf. Robotics and Automation*, 2010, pp. 1313–1319.

[25] L. Le Tien, A. Albu-Schäffer, A. D. Luca, and G. Hirzinger, "Friction observer and compensation for control of robots with joint torque measurement," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots and Systems*, 2008, pp. 3789–3795.

[26] G. Cotugno, K. Althoefer, and T. Nanayakkara, "The role of the thumb: Study of finger motion in grasping and reachability space in human and robotic hands," *IEEE Trans. Systems, Man and Cybernetics; Systems*, vol. 47, no. 7, pp. 1061–1070, 2017.