Overapproximative Human Arm Occupancy Prediction for Collision Avoidance

Aaron Pereira and Matthias Althoff

Abstract—Predicting the occupancy of a human in real time is of great interest in human–robot coexistence for obtaining regions that a robot should avoid in safe motion planning. The human body is composed of joints and links, suiting approximation by a kinematic chain, but the control strategy of the human is completely unknown, meaning the potential occupancy grows very fast and it is difficult to compute tightly in real time. As such, most previous work considers only specific, known, or probable movements, and usually does not account for a range of human dimensions. Focusing on the human arm, we analyze archetypal movements performed by test subjects to create a dynamic model. Motion-capture data of subjects are fitted, for modeling purposes, to two abstractions: a 4-degree of freedom (DOF) model and a 3-DOF model, to obtain dynamic parameters. We validate our approach on movements from a publicly available database. The prediction is shown to be computationally fast, and reachable sets of the abstraction are shown to enclose all possible future occupancies of the arm for different subjects, tightly but overapproximatively. The 3-DOF model has advantages over the 4-DOF in terms of speed, though the 4-DOF model is tighter at smaller time horizons. Such an overapproximative representation is intended for certifiable safety-guaranteed collision avoidance algorithms for robots.

Note to Practitioners—Motivated by the need to keep humans safe when working alongside robots, our earlier work proposes a method of trajectory planning where the robot certifies each movement as safe before it performs it. For this to prove that unsafe collisions cannot occur, an *overapproximative* prediction of the human is needed, meaning that no possible future position of the human is outside the predicted region, or *reachable occupancy*. However, making this prediction both small enough (so that it does not include unreachable regions) and fast enough for real-time use is not straightforward. We find the limits of human motion by asking a range of test subjects to perform movements as fast as possible. We calculate the reachable occupancies based on these limits and show that our predictions are indeed overapproximative, fast, and not wasteful of volume. One can then use the aforementioned approach to guarantee safety; future challenges are reliably sensing the human's pose and implementing our approach on an industrial robot.

Index Terms—Human-centered automation, manufacturing, rescue, safety, security, theoretical foundations.

I. INTRODUCTION

WITH increasing automation in industry, in particular in small and medium enterprises, robots that can be certified to work safely with humans are necessary. Speed, power, and separation monitoring are some measures taken to improve safety [1], [2], but even such cumbersome requirements do not guarantee safety. Collision-free coexistence of humans and robots requires predicting human motion, which is no trivial task since human movements are highly unforeseeable and fast.

We address the problem of predicting and formally bounding any possible future spatial occupancy of human arm movements. Consideration of all possible movements is realized by reachability analysis, which computes the set of possible trajectories of a dynamic system given sets of initial states, inputs, and parameter values [3]. In order to ensure collision avoidance, we model human behavior as differential inclusions, which are uncertain models that allow one to obtain several possible future behaviors [4]. We compute overapproximations of the set of possible future occupancies, since exact algorithms exist only for special problems such as linear systems with purely real or imaginary eigenvalues [5]. The obtained occupancies can be used to provide formal guarantees in motion planning, e.g., as in [6].

As human motion is unpredictable and fast, the large possible set of future human occupancies may unnecessarily impede the movement of surrounding robots. We overcome this often overly-cautious limitation with a two-step approach. First, we generate a long-term plan that is based on the most likely movement of surrounding humans. Of this long-term plan, we execute only the first section and attach a fail-safe maneuver, which brings the robot to a safe state before it reaches positions possibly occupied by humans. In this paper, we consider the robot safe when it is stationary, although other safety criteria could be used, such as limiting the impact force [7]. We refer to the execution of the first section of the long-term plan plus the attached fail-safe maneuver as the *short-term plan*, visualized in Fig. 1. By adapting the duration of the first part of the long-term plan, we adjust the required time horizon of the prediction to limit the expansion of human occupancies [8]. Both the long-term and short-term plans are constantly adjusted based on the pose estimation of surrounding humans. Due to continuous replanning, it should be avoided that the fail-safe maneuver is ever fully executed, and thus the robot behaves mostly as intended, except in
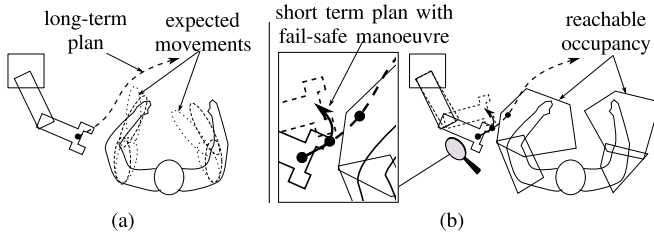
Fig. 1. (a) Long-term motion planning around expected human movement. (b) Short-term fail-safe maneuver to account for unexpected movement, based on *reachable occupancies* enclosing all possible motions. The first section of the long-term trajectory has, at all times, at least one fail-safe maneuver attached at the end. If the next section of the long-term trajectory is not verified safe, the fail-safe maneuver will execute.

case of unexpected human motion. This approach is simulated in [9].

Our work differs from previous work in safe human–robot coexistence (HRCoex) in that we can guarantee collision avoidance given that our uncertain models encompass all possible human behavior, i.e., no human motion falls outside our models' prediction. We show this using reachset conformance checking [10].

Much previous work is concerned with limiting the impact between humans and robots in a reactive manner, i.e., after detecting a collision using, e.g., joint torque [11] or artificial skin [12]. Another approach is to build soft inherently safe robots [13]. However, those concepts can become unsafe when handling sharp objects or when moving fast. Other concepts realize precollision control by minimizing impact energy [14], a danger index [15], or motion scaling based on proximity [16]. However, for many applications, such as with high-inertia robots, it is desirable to guarantee a safety condition, such as noncontact, while the robot is in motion. This is the aim pursued in our work.

Since path planning of robots in dynamic environments (see, e.g., [17]) is much more researched compared with set-based predictions of humans, this paper solely focuses on set-based prediction. We categorize our literature review of human motion prediction into approaches that generate: 1) a single behavior; 2) probability distributions of behaviors; and 3) the set of all possible behaviors. Finally, we review techniques for obtaining human motion models.

A. Single Motion Prediction

Approaches predicting a single future behavior must first determine what natural human movement is. Human movements are often believed to be optimal with respect to some unknown cost function of the dynamics [18]. Morasso [19] observes straight spatial hand trajectories in point-to-point movements and deduces that the human control system operates in Cartesian space rather than joint space. Flash and Hogan [20] present the minimum spatial jerk optimization criterion that fits well to observed motion data. Available muscular effort is also a factor in human movement; in [21], the authors present a muscle-effort minimization criterion that predicts the natural movement of humans holding weights

accurately. The effect of obstacles on natural human behavior is inconclusive [22].

B. Probabilistic Motion Prediction

Since one cannot reliably predict the intention of the human, several researchers use probabilistic models. Koppula and Saxena [23] anticipate high-level movement using a temporal conditional random field. Markov chains are used to model human dynamics in [24], and probabilistic reachability analysis, based on the dynamics obtained, is used in a collision avoidance algorithm. Ding *et al.* [25] present another probabilistic prediction of the human occupancy using a hidden Markov model (HMM). Mainprice and Berenson [26] compute a voxel map of occupancy probability, which is used in a cost function to guide a path-planning algorithm. They infer future movement from the initial arm motion using a Gaussian mixture model trained on a database of previous movements [27].

C. Set-Based Motion Prediction

When considering certifying safety, one obvious question is what happens when humans do not behave as predicted by a single behavior or a probability distribution of behaviors? This may happen during a reflex movement when touching something sharp or grabbing a toppling object. The authors in [25] point out that unforeseen or unusual movements would not be accounted for by their approach, and propose formal techniques, specifically reachability analysis, as a complementary technique to account for such movements. There exists no previous work that can rigorously predict the set of future occupancies given an uncertain dynamic model of human motion. The mapping of uncertain joint angles to the occupancy of a kinematic chain has been considered in [28] in an overapproximative way. A nonformal approach to compute the set of future occupancies is described in [29], which considers the future occupancy of the human given some maximum joint velocities, where these velocities are assumed. In this paper, we use reachability analysis to predict the set of future occupancies. Solutions based on reachability analysis have been proposed to guarantee safety of power systems [3], automated vehicles [30], and recently for robot manipulators in [8]. One challenge of this approach is that the set-based calculations required are time consuming, and for systems with nonlinear dynamics and high dimensionality, it can be very difficult to calculate occupancies online and within real time. Here, we use a Lagrangian technique that propagates the reachable set in consecutive time intervals. For linear continuous dynamics in particular, large state spaces with potentially more than 100 continuous state variables can be efficiently computed [31], [32].

D. Human Modeling

In order to predict human motion, one requires a model of the human. The modeling task can be integrated into the prediction approach, as it is often done in machine learning, e.g., the aforementioned HMMs [25], or one obtains the model

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

PEREIRA AND ALTHOFF: OVERAPPROXIMATIVE HUMAN ARM OCCUPANCY PREDICTION FOR COLLISION AVOIDANCE 3
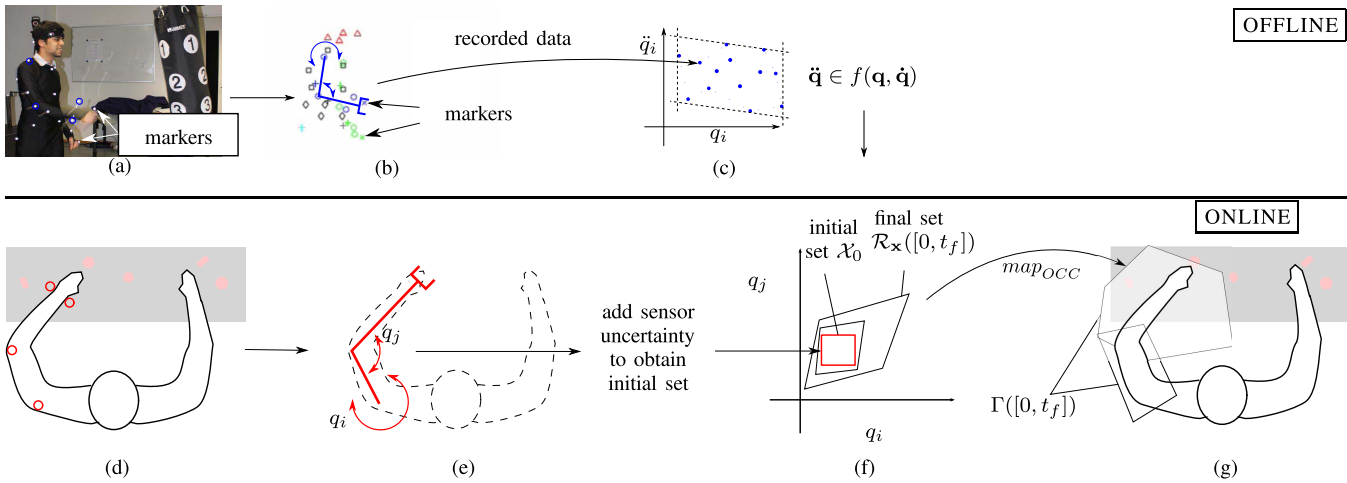

Fig. 2. Top: Offline, we capture archetypal movements of subjects and fit to a kinematic model; from the resulting joint positions, velocities, and accelerations, we obtain a dynamic model. Bottom: Online, we fit sensor data (not necessarily motion capture) to the kinematic model to obtain the state in joint space, and calculate the reachable set until a future time t_f with the dynamics obtained offline. We then convert this to a reachable occupancy in Cartesian space. (a) Motion capture experiments. (b) Fitting to kinematic model. (c) Obtaining dynamic model. (d) Sensing of human pose. (e) Fitting to kinematic model. (f) Reachable set computation. (g) Occupancy mapping.

upfront. In our approach, we obtain the human model upfront by collecting motion capture data of archetypal movements and fitting them to a kinematic model. Digital human models (DHMs) where both skeleton geometry and muscle placement are taken into account, such as the model presented by Holzbaur *et al.* [33], are well studied and provide insight into how humans move. Demircan *et al.* fit the former model to human motion-capture data in real time, achieving position errors of a few centimeters [34] and showing how athletes control their muscles to maximize acceleration in a particular direction [35]. In contrast to these approaches, our goal is not to reduce the error between predicted and observed motion, but to *bound* the possible occupancy of the human and to calculate this possible occupancy with real-time constraints. Our model should be as accurate as possible but must be overapproximative, so that the occupancy given by the model for a certain time interval includes *all* possible reachable occupancies of human arms of any shape, size, and mass, to ensure that robots can safely account for *all* humans. As such, DHMs have two drawbacks: 1) they do not account well for variation in size and mass parameters and 2) they are too complex to calculate in real time with set-based arithmetic.

E. Contribution

A previous attempt to rigorously predict the set of possible future occupancies of humans [36] has the drawback that at very short prediction horizons, the reachable occupancy is large. Furthermore, since only a simple model using joint position, velocity, and acceleration limits is considered for prediction, the reachable set includes states that should, in fact, be excluded.

We address these issues by developing another kinematic parameterization; this parameterization introduces less over-approximation at short prediction horizons, yielding a tighter prediction while still accounting for all movement, but raises challenges in accounting for singularity. These challenges are

addressed in Section IV. We also show that joint accelerations are highly dependent on position and velocity and exploit this to give a tighter prediction that can be used over longer prediction horizons, without compromising real-time feasibility. This is explained in Section V-C. This paper expands upon [36] as mentioned above and presents a more in-depth evaluation of the conservativeness and tightness of the results, including for longer time horizons.

F. Organization

In the following section, we state the problem and our approach. We show how the data are collected in Section III. In Sections IV and V, we present the kinematic and dynamic aspects of the models and the case for a simplified abstraction, as well as the conversion from the model state to the human arm occupancy. We verify the overapproximative models in Section VII on unknown movements and conclude in Section VIII.

II. PROBLEM STATEMENT AND APPROACH

In this paper, we consider the scenario of a human sitting or standing at a workstation, for example, on an assembly line (see Fig. 2). In such a scenario, the fastest movement (and the movement to be avoided by robots) is that of the arms. Thus, in this paper, we intentionally focus on human arm movement to demonstrate the proposed approach of set-based prediction.

The pose estimation used to start the prediction can avail of various sensing technologies such as cameras [37] or depth sensors [38], as we account for varying detection accuracies in the error of the set of initial states. In our evaluation, we validate our method with infrared motion capture data from the publicly available CMU Graphics Lab Motion Capture Database.[1]

[1] obtained from `mocap.cs.cmu.edu` on August 11, 2015

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

4   IEEE TRANSACTIONS ON AUTOMATION SCIENCE AND ENGINEERING

The occupancy of the human body is fast changing and hard to predict: bodies have different dimensions, shapes, and a large number of DOFs. To account for all possible uncertainties, we perform an extensive user study, described in Section III, collecting motion capture data of test subjects performing various extreme movements as illustrated in Fig. 2(a). The complete raw data are fitted to a kinematic model [Fig. 2(b)] resulting in possible combinations of joint positions $\mathbf{q}$, velocities $\ddot{\mathbf{q}}$, and accelerations $\ddot{\mathbf{q}}$, as shown in Fig. 2(c).[2] These data parameterize two uncertain models of human arm motion as presented in Section IV: a 3-DOF model and a 4-DOF model, which are critically discussed later.

The uncertain models are used during HRCoex to predict the human occupancy as in the approach of Fig. 1. Initially, we read sensor data of the human [Fig. 2(d)] and fit them to the same kinematic model used during the user study [Fig. 2(e)]. Thus, we obtain initial joint angles and velocities, which we enlarge based on assumed measurement uncertainties, yielding the initial set of states for our prediction. Our reachable set computation then predicts the set of possible joint angles and velocities until the prediction horizon t_f [Fig. 2(f)]; the prediction horizon depends on the length of the short-term plan as previously mentioned. Finally, the reachable set of joint positions and velocities is mapped to the space occupied by a human arm [Fig. 2(g)]. The process in Fig. 2(d)–(g) is constantly repeated so that the prediction is constantly updated. If the occupancy of the robot would never enter the predicted occupancy of the human, the proposed short-term plan is verified and can be executed. To formally state our problem, we introduce reachable sets and two mappings.

Definition 1 (Reachable Set): Given a system with state $\mathbf{x}(t)$, input $\mathbf{u}(t)$, and dynamics $\dot{\mathbf{x}}(t) = f(\mathbf{x}(t), \mathbf{u}(t))$, where t is time. The possible initial states $\mathbf{x}(0)$ and inputs $\mathbf{u}(t)$ are bounded by sets, $\mathbf{x}(0) \in \mathcal{X}_0, \forall t : \mathbf{u}(t) \in \mathcal{U}(t)$. The reachable set of states $\mathbf{x}$ at $t = t_f$ is [see Fig. 2(f)]

$$\mathcal{R}_{\mathbf{x}}(t_f) := \left\{ \mathbf{x}(t_f) = \mathbf{x}(0) + \int_0^{t_f} f(\mathbf{x}(t), \mathbf{u}(t)) \mathrm{d}t \; \middle| \right.$$
$$\left. \mathbf{x}(0) \in \mathcal{X}_0, \forall t : \mathbf{u}(t) \in \mathcal{U}(t) \right\}.$$

Given a sensor reading $\mathbf{s}(t) \in \mathcal{S}$, we find the state $\mathbf{x}(t) \in \mathcal{X}$ of a kinematic model from the inverse kinematic mapping

$$\mathrm{map}_{IK}(s) : \mathcal{S} \to \mathcal{P}(\mathcal{X}),$$

where we use the power set $\mathcal{P}(\mathcal{X})$ to account for nondeterminism due to noise in sensor readings. From the sensor reading, we can obtain the initial set of states $\mathcal{X}_0 = \mathrm{map}_{IK}(\mathbf{s}(0))$ for the reachability analysis as defined in Definition 1. We further define the occupancy mapping, which maps states to a subset of space occupied by the arm:

$$\mathrm{map}_{OCC}(\mathbf{x}) : \mathcal{X} \to \mathcal{P}(\mathbb{R}^3). \qquad (1)$$

After combining the reachable set of joints and joint angles in Definition 1 with the occupancy mapping in (1), we obtain the reachable occupancy.

[2]The vector of joint positions is $\mathbf{q}$; the position of the ith joint is q_i.

Fig. 3. From left to right: positions A, B, C, D, and E.

Definition 2 (Reachable Occupancy): Reachable occupancy at time t is defined as

$$\Gamma(t) \supseteq \{\mathrm{map}_{OCC}(\mathbf{x}) | \mathbf{x} \in \mathcal{R}_{\mathbf{x}}(t)\}.$$

We define $\Gamma(t)$ as a superset, since the map_{OCC} mapping for sets introduces overapproximation (detailed in Section V-E). We further define the reachable set and reachable occupancy on a time interval as the union of the reachable sets/occupancies of all times in the time interval, i.e., $\mathcal{R}_{\mathbf{x}}([t_1, t_2]) := \bigcup_{t \in [t_1, t_2]} \mathcal{R}_{\mathbf{x}}(t)$ and $\Gamma([t_1, t_2]) := \bigcup_{t \in [t_1, t_2]} \Gamma(t)$. The above introduced definitions make it possible to formulate our problem statement.

Problem 1 (Occupancy Prediction): The predicted occupancy $\Gamma([0, t])$ has to fulfill the following requirements.

1) $\Gamma([0, t])$ encloses the actual arm for any movements of any test subject that could occur in a collaborative environment from time 0 to time t.
2) The computation time of $\Gamma([0, t])$ is far less than t, so that it can be used online in a collision avoidance algorithm.
3) $\Gamma([0, t])$ is as tight as possible, to minimize false-positive detected collisions.

We intend $\Gamma([0, t])$ to be forwarded to a safety controller as proposed in [6] and applied by the authors to a formally safe robot controller in [8]. We now show how the arm model is built, starting with the data collection.

III. DATA COLLECTION

Our set-based prediction bases its validity on collecting extreme movement data from subjects of both genders with a wide range of ages and physical activity levels; these data are used offline to obtain motion parameters, specifically, maximum joint positions, velocities and accelerations, and acceleration ranges as a function of joint positions and velocities [see Fig. 2(a)–(c)].

We collect motion capture data from 38 persons, 12 female and 26 male, ranging in age between 18 and 49 with a median age of 24; 50% did three or more hours of sport a week. The archetypal movements captured were designed to optimally cover the entire workspace of the human arm. The subjects were required to perform the following movements as fast as possible (positions shown in Fig. 3), since unexpected movements in an HRCoex scenario, such as reflex movements, are often fast.

1) Punch to the front then recover to position E.
2) Punch to the front, ending in position A.
3) Position A to position B, elbow allowed to bend.
4) Position A to position B, elbow not allowed to bend.
5) Position D to position C, via position A.
6) Position D to position C, via position B.

We capture the marker positions using a six-camera Vicon motion capture system at a rate of 120 Hz and apply inverse kinematics to find joint positions at each point in the time series. A Kalman filter[3] is used to remove noise and to obtain the joint velocity. Acceleration is found by numerical differentiation of velocity.

To the best of the authors' knowledge, no maximum accelerations of a kinematic parameterization of the human arm are published (maximum torques are more common, e.g., in [33]), though studies of individual joints have been made, in particular the elbow. The choice of filter and hence the parameters obtained are validated by the fact that the elbow joint accelerations agree well with those measured by Amis *et al.* [39] using a camera. We obtain maximum elbow accelerations of -524 and 676 rad/s^2, respectively, while Amis *et al.* measured mean maxima of -800 and 780 rad/s^2 around cessation of motion (flexion is in the positive direction). Though they even observed movements of up to 1200 rad/s^2, they mention that the latter figure was painful to the test subject; furthermore, in their test setup, the arm was laid along a flat surface, while in our tests, the arm was unconstrained. In [40], the authors also measure accelerations of the elbow (and of Euler angles at the shoulder) from test movements, but since this is with a view to building an exoskeleton for rehabilitation and not obtaining the limits of human motion, the movements are not extreme movements and accelerations are much lower (around 30 times lower) than those we observed.

The positions of the shoulder, elbow, and wrist were found from markers placed on the arm; details are found in the Appendix. We call the section of the kinematic chain from the shoulder to the elbow the *upper arm* and the section of kinematic chain from the elbow to the hand the *forearm*.

IV. KINEMATIC MODEL OF THE HUMAN ARM

The human body, made of joints and links, is often modeled as a kinematic tree, e.g., [34] and [29]. Here, we present two simple but overapproximative models that account for the occupancy of the human arm.

A. 4-Degree-of-Freedom Model

We review previous work to motivate the choice of our model. Fig. 4(a) shows a 7-DOF arm model similar to [21] and [42]. The movement of the shoulder is determined by several joints of the shoulder complex, but we consider only movement of the spherical glenohumeral (GH) joint at the shoulder: the range of movement of the other joints is limited, and hence can be considered part of torso movement. References [41] and [43] also ignore the shoulder complex to determine interrelation of joints when calculating the workspace of the human arm; these models are identical apart from the orientation of the first joint axis (a similar model is used by [21] with yet another first joint axis orientation).
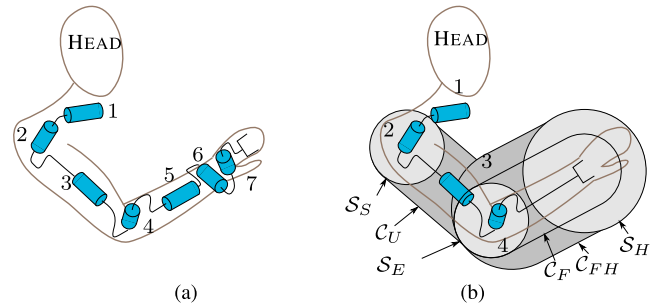

Fig. 4. (a) 7-DOF human arm and (b) 4-DOF model excluding wrist and hand movement, similar to [29], [41], and [21].

Representing the spherical GH joint as a chain of three coincident orthogonal revolute joints with the third joint axis collinear with the upper arm axis allows the first joint axis to be chosen freely; this is discussed in Section IV-C.

Among other studies, [41] disregards hand and wrist movement, and rotation along the forearm axis. As these movements do not greatly affect the occupancy, we also remove these joints to reduce complexity; this leaves us with a 4-DOF model.[4] We account for all movements of the hand caused by joints 5–7 in a static sphere $\mathcal{S}_H$ of radius $r_H = 0.205$ m (the length from the wrist to the middle fingertip of a 95th percentile British male [44]), around the wrist [see Fig. 4(b)]. The upper arm is enclosed in a capsule $\mathcal{C}_U$ constructed by a sphere $\mathcal{S}_S$ of radius r_S swept from the shoulder to the elbow, and the forearm is similarly enclosed in a capsule $\mathcal{C}_F$ constructed by a sphere $\mathcal{S}_E$ of radius r_E swept from the elbow to the wrist. We make a further simplification in order to reduce the number of volumes in the reachable occupancy to collision check (see Section V-E): since $\mathcal{C}_F$ and $\mathcal{S}_H$ are fixed relative to each other, we enclose them both in a capsule $\mathcal{C}_{FH}$ of radius 0.205 m [see Fig. 4(b)]. r_E and r_S are taken as 0.1 m (human arms are actually less broad, but this figure accounts for clothing).

We refer to the position of the ith revolute joint in Fig. 4 as q_i. In this 4-DOF model, the axes of the first and second joints are orthogonal and intersecting, as are the axes of the third and fourth joint. This leads to a coordinate singularity of q_1, described in Section IV-C, and a singularity at the elbow when the elbow is fully extended, where q_3 (rotation around the upper arm, see Fig. 4) is undefined [45]. We describe our solution to the latter singularity for the 4-DOF model in Section V-D; in the following, we introduce a 3-DOF model in which this singularity does not exist.

B. 3-Degree-of-Freedom Model

We develop a kinematic parameterization of arm movement that removes the aforementioned singularity at the elbow by merging the third and fourth joints of the previous model into a single prismatic joint: the forearm is enclosed in a sphere $\mathcal{S}_F$ with its center on the forearm [36]. The shoulder is enclosed

---

[3]State is $[\mathbf{q}^\top, \ddot{\mathbf{q}}^\top]^\top$ where $\mathbf{q}$ is the vector of joint positions. Error covariance is a diagonal matrix of 0.001 rad for rotational DOFs and 0.0005 m for the third joint in the 3-DOF model, which is prismatic. Process covariance is a matrix of zeros for the position part of the state and a diagonal matrix of 0.1 rad/s for the velocity part in rotational DOFs (0.05 m/s for prismatic DOF).

[4]In the following description, the subscripts S, E, H, F, and U refer to shoulder, elbow, hand, forearm, and upper arm, respectively.
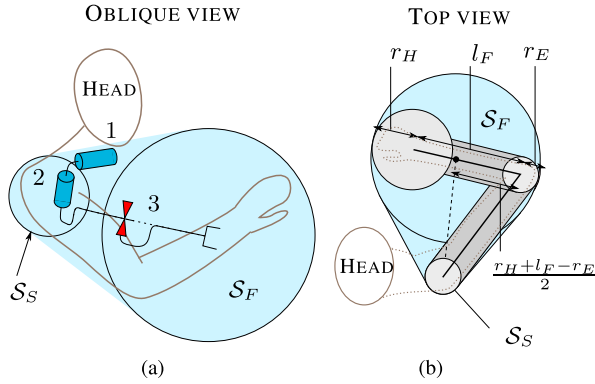
OBLIQUE VIEW   TOP VIEW

Fig. 5. (a) Kinematics of the 3-DOF model. q_1 and q_2 are joint angles of the first two rotational joints. r_3 is the extension of the prismatic joint. (b) Evaluating r_F the radius of $\mathcal{S}_F$.

in another sphere $\mathcal{S}_S$ at the origin, and the arm occupancy is the convex hull of these spheres, as shown in Fig. 5(a).

The following description refers to Fig. 5(a). As the shoulder is in $\mathcal{S}_S$ and elbow is in $\mathcal{S}_F$, the arm is contained in the convex hull by the property of convexity. The center of $\mathcal{S}_F$ is the end effector of the kinematic chain, and thus the prismatic joint axis is along the line from the origin to the center of $\mathcal{S}_F$; the distance from the origin to the end effector is the translational position of the prismatic joint, denoted r_3 to distinguish from rotational joint position q_3 in the 4-DOF model.

The center point of $\mathcal{S}_F$, which we call the *end effector* of the 3-DOF model, lies on the forearm, as shown in Fig. 5(b). However, we must define the exact position. Since we wish to keep the radius r_F of $\mathcal{S}_F$ as small as possible to reduce overapproximation, we choose the end effector to be halfway between the elbow and the tip of the hand (calculated as the point extended past the wrist by the maximum hand length from Section IV-A, 0.205 m; we also extend the forearm beyond the elbow to include the thickness r_E of the forearm). Thus, the forearm movement is reduced to movement of a prismatic joint at the end of a simple kinematic chain.

C. Placement of First Joint Axis

In both 3-DOF and 4-DOF models, the first two joints in the kinematic chain are orthogonal and both revolute. This introduces a coordinate singularity when the elbow (or in the 3-DOF model, the end effector) lies on the first joint axis. At this point q_1 is undefined, and near this point, q_1 varies highly for small variations in task space [see Fig. 6(a)]. As the orientation of the first axis can be freely chosen, this singularity can be avoided. To find an axis through the origin that the elbow (or end effector) does not intersect, we plot all marker data (see Fig. 7). The axis chosen for the 3-DOF model is the long axis of the clavicle,[5] i.e., along the z-axis in Fig. 17 of the Appendix (similar to [43]), rotated around the y-axis by 5° and then around the x-axis by 5°; that of

[5]The reader can verify with a little gymnastics that it is very difficult to align the humerus with the clavicle, as the acromial process blocks it.
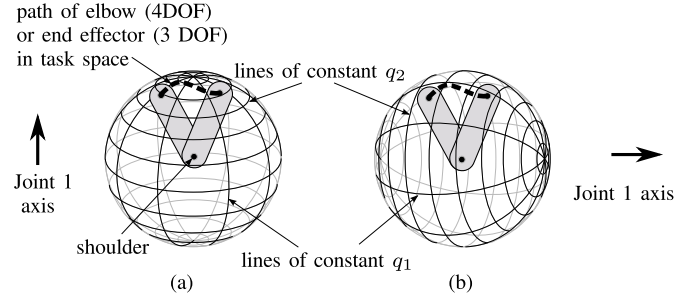

Fig. 6. (a) Choice of joint 1 axis means that the inverse kinematic solution shows rapidly changing q_1; when the end effector is aligned with the joint 1 axis, the inverse kinematics is even undefined. (b) More appropriate choice of axis to minimize $\dot{q}_1$ and $\ddot{q}_1$.
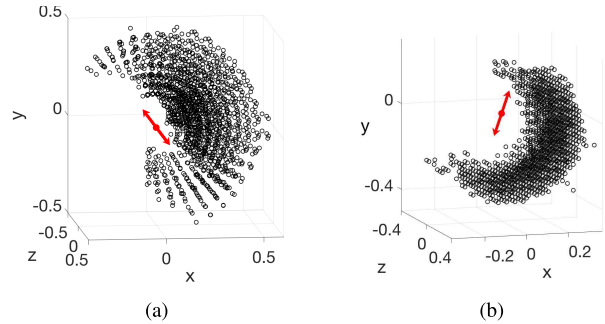

Fig. 7. Sampled workspaces of (a) 3-DOF model end effector and (b) 4-DOF model elbow, with their respective axes in red. The axes do not intersect the workspace. Coordinate system is that of the shoulder (scale in meters).

the 4-DOF model is the clavicle axis rotated around the y-axis by 20°—both were chosen by observation of the sampled workspace.

D. Accounting for Measurement Uncertainty

The measurements of the lengths of the forearm and upper arm are subject to uncertainty; our models account for variation of these lengths differently. The rest of this section refers to Fig. 8. In the 3-DOF model, the forearm is enclosed if $\mathcal{S}_F$ encloses the forearm plus measurement error; the upper arm is automatically enclosed as both elbow and shoulder are inside the convex hull of $\mathcal{S}_F$ and $\mathcal{S}_S$.

For the 4-DOF model, the length of our capsule $\mathcal{C}_U$ is set at the length of the upper arm measured from the movement data plus measurement uncertainty. In this way, the upper arm is certainly enclosed. Setting the length of capsule $\mathcal{C}_{FH}$ to that of the measured forearm length and hand, plus uncertainty, accounts for uncertainty in the length of the forearm (note that, if we had not enclosed $\mathcal{C}_F$ and $\mathcal{S}_H$ in $\mathcal{C}_{FH}$, this would not have been the case). However, the occupancy of the forearm in the 4-DOF model also depends on the length of the upper arm, and is therefore not automatically enclosed (see Fig. 8). We account for measurement uncertainty in the upper arm length by adding a prismatic pseudojoint between the second and third revolute joints, as shown in Fig. 9. The interval of joint values of this pseudojoint is the measured length of the upper arm plus/minus measurement uncertainty.

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

PEREIRA AND ALTHOFF: OVERAPPROXIMATIVE HUMAN ARM OCCUPANCY PREDICTION FOR COLLISION AVOIDANCE 7
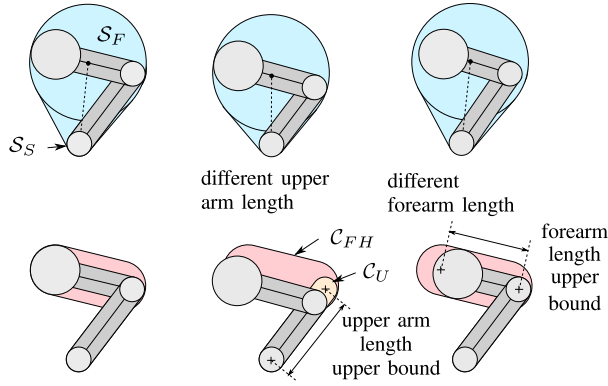

Fig. 8. 3-DOF (top) and 4-DOF (bottom) models of the arm coping with measurement uncertainty in arm length. The 4-DOF model cannot account for uncertainty in upper arm length, as the variable upper arm length means the forearm may not be enclosed by $\mathcal{C}_{\text{FH}}$.
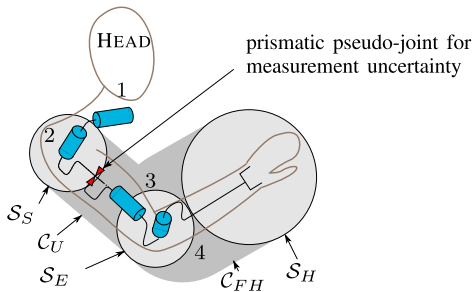

Fig. 9. In the 4-DOF model, measurement uncertainty in the upper arm is accounted for by the addition of a prismatic pseudojoint.
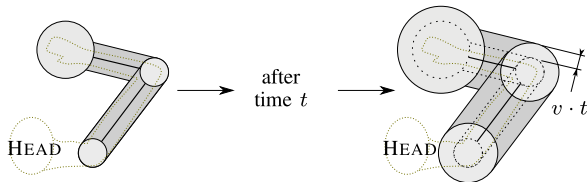

Fig. 10. Prediction $\Gamma_{\text{ISO}}(t)$ based on hand speed constant from [1].

E. ISO-Compliant Model

To compare our overapproximative prediction, we introduce an ISO-based reachable occupancy $\Gamma_{\text{ISO}}([0, t))$. The hand speed constant v = 1.6 m/s is given in [1] as the maximum speed of the upper body of a nonwalking human. We take the 4-DOF model from Section IV-A and enlarge it by $v \cdot t$ as shown in Fig. 10 to account for movement under the assumed maximum velocity.

V. DYNAMIC MODEL

Having abstracted the human arm to a kinematic chain, we define the *state* $\mathbf{x}$ of a kinematic chain as $[\mathbf{q}^\top, \dot{\mathbf{q}}^\top]^\top$, where $\mathbf{q}$ is the vector of joint positions and $\dot{\mathbf{q}}$ is the vector of joint velocities. We discuss two models for determining the dynamics of the kinematic chain in Section IV from the motion data gathered as described in Section III—one model is torque based and the other is acceleration based. Finally, we motivate the use of the acceleration-based model.

A. Torque Model

Holzbaur *et al.* [33] collate the maximum muscle forces from a variety of studies for use in their DHM, and thus it would be intuitive to calculate the range of torques of each joint as the sum of the range of contributions from each muscle, if the moment arms were known. Alternatively, Otis *et al.* [46] measure the maximum flexion, abduction, and rotation torques for a range of shoulder joint positions and angular velocities, and Amis *et al.* [39] study the forces on the arm during maximal elbow movements. The dynamics of a torque-based model of the human arm is identical to that of a kinematic chain and can be described as [47]

$$\tau = H(\mathbf{q})\ddot{\mathbf{q}} + C(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} + F(\dot{\mathbf{q}}) + g(\mathbf{q}, \boldsymbol{\alpha}) + \tau_{\text{ext}}$$

where $\mathbf{q}$ is the vector of joint positions, H is the inertia matrix, C is the matrix of Coriolis and centrifugal forces, F is the friction vector, and g is the gravitational vector as a function of the joint positions and orientation $\boldsymbol{\alpha}$ of the shoulder coordinate system with respect to the world. τ is the vector of the total torque on the joints and τ_{ext} is the vector of external torques, including those from the movement of the base relative to the world. Calculating the reachable set of such a system is challenging as: 1) mass, inertia, and friction parameters for calculating H, C, and F are uncertain, which introduces large overapproximations; 2) external torques τ_{ext} are unknown, as is the direction of the gravity vector $g(\mathbf{q}, \boldsymbol{\alpha})$, which varies with shoulder orientation; and 3) the system is highly nonlinear and computation times are not real-time-capable for this application with current methods [3], [48].

We therefore opt for a model that is based on acceleration, but which accounts for joint position and velocity limits. This is described in the following sections.

B. Accounting for Joint and Velocity Limits

It is quicker to compute the intersection of reachable sets from four simple models than the reachable set from one complex model. This is done as described in [49, Prop. 1], where the reachable sets $\mathcal{R}^{(1)}, \dots \mathcal{R}^{(n)}$ from n models $\mathcal{M}_1, \dots \mathcal{M}_n$ are all overapproximative and enclose the exact reachable set $\mathcal{R}^e$, then $\mathcal{R}^e \subseteq \bigcap_{i=1}^n \mathcal{R}^{(i)}$. Here, $\mathcal{Q}(0)$ and $\dot{\mathcal{Q}}(0)$ are the sets of initial positions and velocities in joint space, respectively, $\oplus$ is the Minkowski sum ($A \oplus B = \{a + b | a \in A, b \in B\}$), and we find the maximum and minimum positions, velocities, and accelerations, $\mathbf{q}_{\text{inf}}$, $\mathbf{q}_{\text{sup}}$, $\dot{\mathbf{q}}_{\text{inf}}$, $\dot{\mathbf{q}}_{\text{sup}}$, and $\ddot{\mathbf{q}}_{\text{inf}}$, $\ddot{\mathbf{q}}_{\text{sup}}$ from the archetypal movements. The reachable sets we predict with our models are as follows:

1) a zeroth-order model of maximum joint position: $\mathcal{R}_\mathbf{q}^{(1)}([0, t]) = [\mathbf{q}_{\text{inf}}, \mathbf{q}_{\text{sup}}]$;
2) a first-order model of maximum joint velocity: $\mathcal{R}_\mathbf{q}^{(2)}([0, t]) = \mathcal{Q}(0) \oplus [\dot{\mathbf{q}}_{\text{inf}}, \dot{\mathbf{q}}_{\text{sup}}]t$;
3) a second-order model of maximum joint accelerations, $\mathcal{R}_\mathbf{q}^{(3)}([0, t]) = \text{CH}(\mathcal{Q}(0)\mathcal{Q}(0) \oplus \dot{\mathcal{Q}}(0)t \oplus [\ddot{\mathbf{q}}_{\text{inf}}(t^2/2), \ddot{\mathbf{q}}_{\text{sup}}(t^2/2)])$;
4) another second-order model $\mathcal{R}_\mathbf{q}^{(4)}([0, t])$ based on acceleration as a function of state, described in the next section.

In order to obtain the reachable set of the time interval $[0, t]$ in the third model, we must compute the convex hull, indicated by the operator CH, of the initial set $\mathcal{Q}(0)$ and the reachable set at time t found using a second-order approximation of motion (this encloses the reachable set of the interval, as it is a special case of [30, Prop. 3.3]). We enclose this convex hull in a Cartesian product of intervals. $\mathcal{R}_{\mathbf{q}}^{(4)}([0, t])$ is obtained as a zonotope and also enclosed in a Cartesian product of intervals. We use Cartesian products of intervals, as these can be quickly and easily intersected: the intersection $\mathcal{R}_{\mathbf{q}}([0, t]) = \bigcap_{i=1}^{4} \mathcal{R}_{\mathbf{q}}^{(i)}([0, t])$ gives us a tighter overapproximative interval of joint positions than any one model alone, where the joint position, velocity, and acceleration limits are not exceeded.

C. Acceleration Model

Though ample work exists on maximum arm torques and muscle forces, to the best of the authors' knowledge, no complete study into the maximum accelerations of all human arm joints exists. We present a model where acceleration is a linear function of the state $\mathbf{x} = [\mathbf{q}^{\top}, \dot{\mathbf{q}}^{\top}]^{\top}$. Let k be the number of joints, $\mathbf{C}$ be a matrix of size $k \times 2k$ of coefficients to be determined, $\mathbf{I}_k$ be an identity matrix of dimension k, $\mathbf{0}_{k \times 1}$ be a vector of zeros of length k, $\mathbf{0}_{k \times k}$ be a matrix of zeros of size $k \times k$, and $\mathbf{u}_{\min}$ and $\mathbf{u}_{\max}$ be a vector of uncertain inputs to the model. The acceleration model can be described by a set of linear differential inclusions of the form

$$\begin{bmatrix} \dot{\mathbf{q}} \\ \ddot{\mathbf{q}} \end{bmatrix} = \dot{\mathbf{x}} \in \begin{bmatrix} \mathbf{0}_{k \times k} & \mathbf{I}_k \\ & \mathbf{C} \end{bmatrix} \mathbf{x} \oplus \begin{bmatrix} \mathbf{0}_{k \times 1} \\ [\mathbf{u}_{\min}, \mathbf{u}_{\max}] \end{bmatrix}. \quad (2)$$

Methods for determining reachable sets of a linear first-order system of differential inclusions such as (2), for any time interval when given an initial set, are fast and are detailed in [30]. The reachable set is obtained as a zonotope to avoid the wrapping effect; we then enclose it in a Cartesian product of intervals. This enclosure in a product of intervals is necessary for efficient intersection with reachable sets from other models, as mentioned in Section V-B, and for overapproximative conversion to a reachable occupancy, as shown in Section V-E.

If $\mathbf{C}$ is a matrix of zeros, in order that (2) accounts for all of our observed data, $\mathbf{u}_{\min}$ and $\mathbf{u}_{\max}$ would need to be the minimum and maximum accelerations observed over all our data. This is a wide range, and our reachable set will grow too quickly. An acceleration model that exploits the correlation between state $\mathbf{x}$ and range of joint accelerations would be preferable, as this can combine some of the accuracy and insight offered by the torque model with the speed of calculation of the linear model. The dependence of joint torque on joint angle is well known and studied for the single and multijoint cases. The dependency in a single joint is due to muscle physiology and several models exist [50] whereas multijoint interdependencies could be due to muscle-on-muscle impingement and intermuscular force transmission [51]. From observation of our data, some relationships between joint positions, velocities, and accelerations can be observed for certain joints (see Table I). Illustrated in Fig. 11 is the dependency of $\ddot{q}_4$ on q_4 during one movement (this is intuitive, as at the positive joint limit, only negative acceleration is possible, and *vice versa*).

TABLE I

CORRELATIONS BETWEEN JOINT POSITIONS AND VELOCITIES, AND ACCELERATION, OVER ENTIRE DATA SET (HIGHEST VALUES FOR EACH JOINT ACCELERATION IN BOLD). (a) 4-DOF SYSTEM. (b) 3-DOF SYSTEM

(a)

r^2	q_1	q_2	q_3	q_4	$\dot{q}_1$	$\dot{q}_2$	$\dot{q}_3$	$\dot{q}_4$
$\ddot{q}_1$	**-0.36**	-0.12	-0.10	**0.20**	0.05	0.14	-0.07	**0.18**
$\ddot{q}_2$	**-0.11**	**-0.32**	0.02	0.08	-0.07	0.05	-0.05	**0.15**
$\ddot{q}_3$	**-0.09**	0.04	**-0.17**	-0.02	0.03	0.01	0.04	**-0.08**
$\ddot{q}_4$	0.07	0.01	-0.02	**-0.25**	**-0.16**	**-0.15**	0.06	0.07

(b)

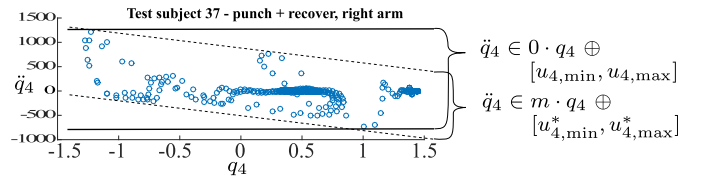| r^2	q_1	q_2	r_3	$\dot{q}_1$	$\dot{q}_2$	$\dot{r}_3$
$\ddot{q}_1$	**-0.39**	-0.06	**-0.15**	0.04	0.12	**-0.22**
$\ddot{q}_2$	-0.06	**-0.33**	-0.03	**-0.06**	0.05	**-0.19**
$\ddot{r}_3$	-0.04	0.03	**-0.23**	**0.15**	**0.16**	0.07

Fig. 11. By considering dependency of $\ddot{q}_4$ on q_4 (dashed lines), the set of uncertain inputs $[u_{\min}^*, u_{\max}^*]$ in the bottom model can be chosen smaller than $[u_{\min}, u_{\max}]$ in the top model.

Problem 2 (Lowest Possible Input Range): Given the dynamical system in (2), we wish to find the matrix $\mathbf{C}$ such that the range of $[\mathbf{u}_{\min}, \mathbf{u}_{\max}]$ is minimized and all the training data are consistent with this system.

We approach the problem one joint acceleration at a time. We rearrange (2) to obtain

$$\ddot{q}_i(t) \in \sum_{j=1}^{j=2k} C_{i,j} x_j(t) \oplus [u_{\min,i}, u_{\max,i}]. \quad (3)$$

We consider only the three dimensions most highly correlated with $\ddot{q}_i$. Letting corr_i be the indices of these dimensions, $\mathbf{x}_{\text{corr}_i}$ is a projection of $\mathbf{x}$ onto the dimensions corr_i. For example, Table I(b) shows that $\ddot{q}_1$ in the 3-DOF model correlates most highly with dimensions 1, 3, and 6, so $\text{corr}_1 = [1, 3, 6]$ and we have $\mathbf{x}_{\text{corr}_1} = [q_1, r_3, \dot{r}_3]^{\top}$. Note: just because acceleration is not a function of all elements of the state that does not mean the model is not overapproximative; for example, as seen in Fig. 11, the differential inclusion that is not dependent on q_4 is still overapproximative (it still accounts for all data points), but it requires a larger range of $[\mathbf{u}_{\min}, \mathbf{u}_{\max}]$. Letting C_{i,corr_i} be the entries in row i of $\mathbf{C}$ in the columns of corr_i, (3) becomes

$$\ddot{q}_i(t) \in C_{i,\text{corr}_i} \mathbf{x}_{\text{corr}_i}(t) \oplus [u_{\min,i}, u_{\max,i}]. \quad (4)$$

We construct k data sets $\mathcal{S}_i, i \in \{1, \ldots k\}$, where the data points are values of $[\mathbf{x}_{\text{corr}_i}^{\top}, \ddot{q}_i]^{\top}$ at every point in time over the entire motion capture data. From these data sets, our goal is to find C_{i,corr_i} such that $u_{i,\max} - u_{i,\min}$ is minimized, for
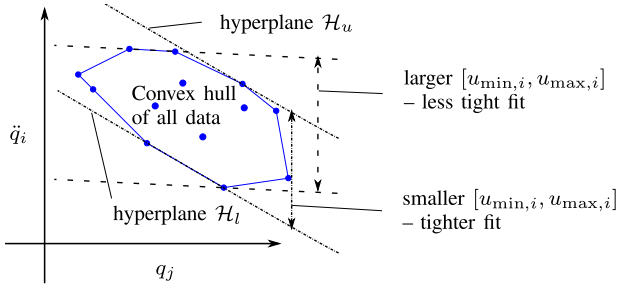
This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

PEREIRA AND ALTHOFF: OVERAPPROXIMATIVE HUMAN ARM OCCUPANCY PREDICTION FOR COLLISION AVOIDANCE 9

Fig. 12. Two pairs of hyperplanes. The vertical distance between them is to be minimized, as this is the range of $[u_{\min,i}, u_{\max,i}]$, thus the range of possible values of $\ddot{q}_i$.
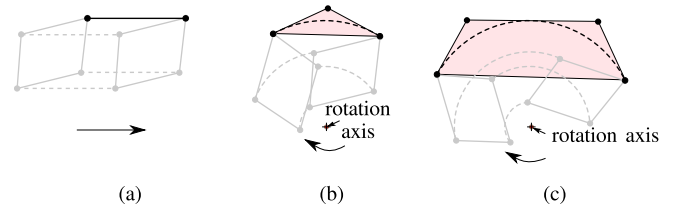

Fig. 13. (a) Point on the SSV under the action of a prismatic joint is enclosed in a line segment. (b) Occupancy of a point on the SSV under the action of a revolute joint (dashed line) is enclosed by three points in the new SSV (solid line) or (c) if the angle subtended is larger, we use more points. The previous SSV had four points. The new one would have (a) 8 points, (b) 12 points, or (c) 16 points (not all shown, for clarity of illustration).

each joint i. Since (4) must hold for all elements $\mathbf{s}_i \in S_i$, we rearrange it to show

$$u_{\max,i} \geq \{[-C_{i,\text{corr}_i}, 1]\mathbf{s}_i | \mathbf{s}_i \in \mathcal{S}_i\}$$
$$u_{\min,i} \leq \{[-C_{i,\text{corr}_i}, 1]\mathbf{s}_i | \mathbf{s}_i \in \mathcal{S}_i\}. \tag{5}$$

We can find the minimal $u_{i,\max} - u_{i,\min}$ by formulating the constraints in (5) as a linear programming problem [52, Ch. 2.17].

D. Interdependence of Third and Fourth DOFs

When the elbow is fully extended [i.e., $q_4 \approx -(\pi/2)$], the 4-DOF arm has a singularity, meaning q_3 (rotation around the upper arm, see Fig. 4) is undefined [45]. The inverse kinematics near this singularity is affected: miniscule changes in Cartesian position lead to large changes in q_3, and thus extremely high values of $\ddot{q}_3$ are observed. We account for this by ignoring these erroneously high values when building the dynamic model, and assume that the accelerations and velocities of q_3 near the singularity are the same as those away from the singularity; they simply cannot be measured. We define the boundary for cutoff at $q_4 = -(\pi/2) + \epsilon$.[6] Values of $\ddot{q}_3$ and $\dot{q}_3$ where $q_4 < -(\pi/2) + \epsilon$ are not used in determining the dynamic model.

We are now able to determine the reachable set of the kinematic chain given the initial set of states, found online by inverse kinematics, and the dynamics, determined in advance.

E. Representation in Space

We now consider the mapping map_{OCC} from Section II, shown in Fig. 2(g). To the best of our knowledge, there exists only one approach that computes overapproximative occupancies of kinematic chains efficiently, which is that of Täubig *et al.* [28]. Here, the occupancy of each link on a kinematic chain is represented as a sphere-swept volume (SSV), defined as the Minkowski sum of a convex hull of points with a zero-centered sphere. The set of points of the SSV is generated sequentially from a given initial SSV enclosing the ith link in its own (ith) coordinate system. For each joint down to the base of the kinematic chain, the algorithm from [28] sequentially generates a new SSV, which encloses the link in the $i-1$th coordinate system under all positions of the ith

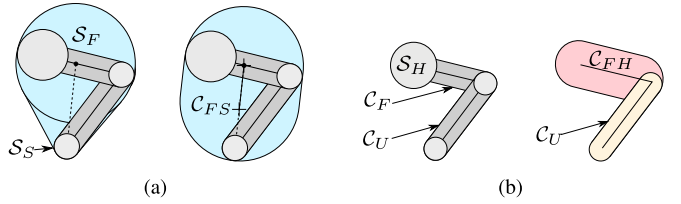[6] in our evaluation, we choose $\epsilon = 0.225$ rad

Fig. 14. Top view of Figs. 4 and 5, only arm model and not human body is shown. (a) 3-DOF model enclosed in capsule $\mathcal{C}_{FS}$ so as to use method in [28]. (b) Hand and forearm of 4-DOF model enclosed in $\mathcal{C}_{\text{FH}}$ to save on collision checking.

joint. This SSV in its turn is then used to generate a new SSV enclosing the link in the $i-2$th coordinate system under all positions of the $i-1$th joint, and so on until the (fixed) base coordinate system. See Fig. 13: for a prismatic joint, each point on the previous SSV corresponds to two points on the new SSV, and the new SSV encloses the previous SSV under the action of the joint exactly. For a revolute joint, the number of points in the new SSV increases by a factor of at least three, and the enclosure is not exact but is conservative; using more points reduces the overapproximation, but of course increases complexity. We use a factor of three when the range of angles $q_{i,\max} - q_{i,\min}$ is less than $(\pi/2)$, four when less than π, five when less than $(3\pi/2)$, and six otherwise. We found that a finer resolution did not yield significant improvements in terms of volume. Hence, the SSV enclosing the end effector has at least $m \cdot 2^p 3^r$ points, where m points define the end-effector occupancy in its own coordinate system, and the kinematic chain has p prismatic and r rotational joints. We refer the reader to [28] for a detailed description of the algorithm.

The 3-DOF model is the convex hull of spheres of two different radii, which is not an SSV and cannot use the previously described method. We therefore enclose the 3-DOF model in a capsule, as shown in Fig. 14(a). While this increases the overapproximation by including extra volume that the human does not occupy, this volume is mostly near the shoulder—an area the robot is likely to avoid anyway due to proximity of the head and torso. In an HRCoex scenario, therefore, this is less restrictive than including extra volume near the extremities. Furthermore, when the elbow is in flexion, the sphere $\mathcal{S}_S$ may intersect or even be enclosed in $\mathcal{S}_F$, in which case the extra volume is minimal to nonexistent. The representation of the 4-DOF model was two capsules and a sphere, however,

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

10   IEEE TRANSACTIONS ON AUTOMATION SCIENCE AND ENGINEERING

reducing this to two capsules, as shown in Fig. 14(b) and described in Section IV-A, reduces the number of collision checks necessary when collision checking against the robot occupancy.

VI. NOTE ON COMPLEXITY

The entire algorithm of generating the reachable occupancy as illustrated in Fig. 2 consists of: A) obtaining the state of the human; B) reachability analysis; and C) conversion to a reachable occupancy. Although the biomechanics of humans does not change, we present the complexity with the consideration that one might refine the model and increase the DOFs k.

A. Obtaining State of the Human

We calculate the inverse kinematics analytically, and then filter it using a Kalman filter to reduce noise and obtain joint velocities. The inverse kinematics is constant in complexity and the Kalman filter involves matrix–matrix multiplication, where complexity is cubic in the size of the matrices. The state of the filter is the joints and their velocities and thus of size $2k$, meaning the largest matrices to multiply are size $2k$, and the filtering has complexity $O(k^3)$.

B. Reachability Analysis

Of the models presented in Section V-B, the first three have trivial $O(k)$ complexity (vector addition). The fourth model involves reachability analysis on linear differential inclusions over a time interval, which again has complexity $O(k^3)$ [30].

C. Conversion to Reachable Occupancy

The algorithm from [28] creates an SSV defined by a set of points, which is generated sequentially from an initial set as described in Section V-E. At each DOF, the set of points is multiplied by two (prismatic DOF) or at least three (rotational DOF) different transformation matrices to generate a new set of cardinality twice or thrice the previous set. Since the transformation matrices are constant size, the number of multiplication operations is proportional to the number of points in the new set obtained. As mentioned in Section V-E, the number of points in the reachable occupancy is at least $m \cdot 2^p 3^r$, where the link is enclosed in an SSV of m points, and p and r are the number of prismatic and rotational DOFs. The complexity of the algorithm is therefore exponential in $k = p + r$. Recall that the 4-DOF model also requires a prismatic pseudojoint to account for uncertainty in arm length (Section IV-D).

For collision checking the reachable occupancies with the occupancy of the robot, the well-known GJK algorithm may be used, where computation time is approximately proportional to the number of points of the SSVs to collision check [53]; we do not need to explicitly compute the convex hull of the points for the SSV for this algorithm, and hence collision checking is also exponential in k. The 4-DOF model has two volumes to collision check, whereas the 3-DOF model has only one. In a manufacturing scenario, a robot must check its occupancy against all humans in the collaborative space, so computation increases linearly with the number of humans in the vicinity.

The approach is therefore overall of exponential complexity, which further motivates our simplified low-DOF models.

VII. VALIDATION AND DISCUSSION

We validate our approach on data from the Carnegie Mellon University Graphics Lab Motion Capture Database,[7] performing conformance checking [10] to confirm that our approach accounts for the required range of motion conservatively, computing the volumes, and testing computation time.

A. Conformance Checking

The validation data are the same as used in [36], grouped into the following categories:
1) *Everyday motions*, e.g., construction work, machining work, manipulating objects, and stumbling (96 motions);
2) *Sports-related motions*, e.g., boxing, throwing balls, and batting balls (67 motions);
3) *Dance-related motions*, e.g., swing dance, Indian dance, and modern dance (58 motions);
4) *Acrobatic motions*—any motion where the feet are in the air simultaneously, including cartwheels, jumps, swings from a trapeze, and backflips (68 motions).

We choose the *Everyday* movements to be similar to motion in an HRCoex scenario when humans are comfortable working alongside robots, but we add *Sport*, *Dance*, and *Acrobatic* movements to test the limits of the prediction. For the reachability analysis, the toolbox CORA [54] was used, and INTLAB [55] for interval computations. We use the Robotics Toolbox [56] to generate the robot kinematics and dynamics and the MPT toolbox [57] for visualization and volume computation of polytopes. Measurement uncertainties are estimated at ± 0.04 rad for revolute joints and ± 0.001 m for both prismatic joints and the length uncertainty of forearm and upper arm; uncertainty in joint velocity measurements is assumed to be zero.

We wish to show that the predictions are a tight overap-proximation: 1) they enclose all human movement relevant to an HRCoex scenario and 2) that smaller occupancies do not. For this reason, we run the prediction for all these movements under the following conditions:
1) 3-DOF model, whose occupancy is labeled Γ_{3DOF};
2) 4-DOF model (Γ_{4DOF});
3) ISO-based prediction as in Section IV-E (Γ_{ISO});
4) 3-DOF model where the range of velocities and accelerations is reduced by half ($\Gamma_{\text{3DOF,50\%}}$);
5) 4-DOF model where the range of velocities and accelerations is reduced by half ($\Gamma_{\text{4DOF,50\%}}$).

We include the latter two models to test whether the parameters we obtained from extreme data in Section III are indeed necessary to include all movement, or whether the range of motion considered is too large. For each timestep of every

---

[7]Publicly available at `mocap.cs.cmu.edu`, accessed August 11, 2015. The authors would be happy to provide details of the specific data files used on request.

TABLE II
NUMBER OF MOVEMENTS WHERE NO MARKERS ARE OUTSIDE OCCUPANCY AT ANY TIME DURING MOVEMENT (ALL MOVEMENTS IN CATEGORY IN BOLD)

Movements	Everyday	Sport	Dance	Acrobatics
$\Gamma_{3DOF}([0, 16.7]ms)$	**96**	55	66	50
$\Gamma_{4DOF}([0, 16.7]ms)$	**96**	57	66	57
$\Gamma_{ISO}([0, 16.7]ms)$	96	43	62	23
$\Gamma_{3DOF,50\%}([0, 16.7]ms)$	95	55	66	48
$\Gamma_{4DOF,50\%}([0, 16.7]ms)$	95	56	66	53
$\Gamma_{3DOF}([0, 33.3]ms)$	**96**	53	**67**	49
$\Gamma_{4DOF}([0, 33.3]ms)$	**96**	57	**67**	57
$\Gamma_{ISO}([0, 33.3]ms)$	90	15	40	13
$\Gamma_{3DOF,50\%}([0, 33.3]ms)$	95	52	63	40
$\Gamma_{4DOF,50\%}([0, 33.3]ms)$	93	50	**67**	46
$\Gamma_{3DOF}([0, 50.0]ms)$	**96**	53	**67**	49
$\Gamma_{4DOF}([0, 50.0]ms)$	**96**	57	**67**	59
$\Gamma_{ISO}([0, 50.0]ms)$	81	7	23	13
$\Gamma_{3DOF,50\%}([0, 50.0]ms)$	95	51	63	37
$\Gamma_{4DOF,50\%}([0, 50.0]ms)$	93	49	65	43

movement, we filter the motion capture data with the same Kalman filter from Section III and evaluate the reachable occupancy during a time interval $[0, t_f]$. We check whether the arm at time t_f (in this case, all the markers on the arm) is entirely contained in $\Gamma([0, t_f])$. We test $t_f = 16.7$, 33.3, and 50.0 (times in milliseconds).[8] Using simple distance checks for markers on rigid bodies, we determine where markers are incorrectly tracked; incorrectly tracked markers are not checked for inclusion in the reachable volume. We do this for both arms.

As can be seen in Table II, Γ_{3DOF} and Γ_{4DOF} correctly account for all movement in the *Everyday* data set, though in the other data sets, a few movements fall outside the prediction. The arm dynamics in *Dance*, *Sport*, and *Acrobatics* movements are faster than in the data used to parameterize the model in Section III; hence, our assumptions on the arm dynamics do not hold. Reducing the range of velocities and accelerations by half in $\Gamma_{3DOF,50\%}$ and $\Gamma_{4DOF,50\%}$ means that some *Everyday* movements, which should be accounted for, are not; the predictions also perform worse on the other data sets. Fig. 15 shows movements approaching the borders of the reachable occupancies Γ_{3DOF} and Γ_{4DOF} (but still accounted for), which are not captured by $\Gamma_{3DOF,50\%}$ and $\Gamma_{4DOF,50\%}$. Our reachable occupancies are therefore tight in the sense that the parametrization from the archetypal movements is sufficient to enclose relevant movement, but not excessive, since attenuating the parameters causes movements to fall outside the borders of the reachable occupancy. Finally, the ISO-based prediction, while simple and based on the currently accepted standards, proves insufficient to conservatively estimate human motion in all categories.

[8]The test data were sampled at 60 and 120 Hz, so t_f corresponded to a whole number of timesteps in all data.
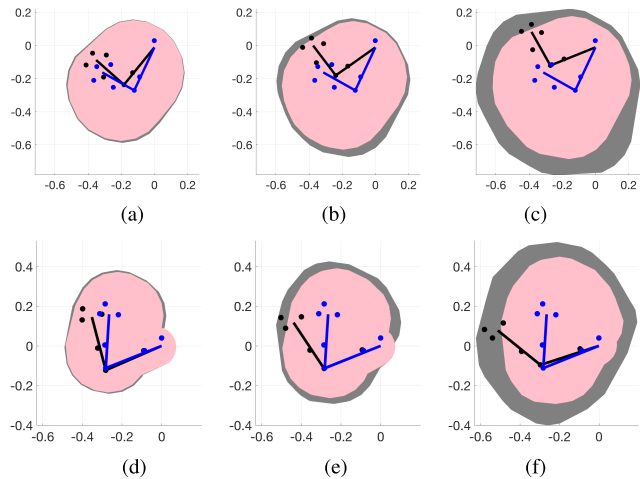

Fig. 15. (a)–(c) 3-DOF model and (d)–(f) 4-DOF model, during *Everyday* movements. The predictions are for the time interval $[0, t_f]$ where (a) and (d) $t_f = 16.7$ ms, (b) and (e) $t_f = 33.3$ ms, and (c) and (f) $t_f = 50.0$ ms, from left to right. Dots are markers on the arm and lines are the forearm and upper arm. Blue is at $t = 0$ and black is at $t = t_f$. All markers lie in the gray occupancies (calculated with parameters from Section III), but not always in the pink occupancies (calculated with parameters from Section III reduced by 50%). This illustrates that the parameters we use are necessary to capture all relevant movement (scale in meters and origin at shoulder).

TABLE III
VOLUME COMPARISON (m^3): MEAN μ, STANDARD DEVIATION σ; $n = 20$

		3-DOF	3-DOF (from [36])	4-DOF	4-DOF (from [36])	ISO
$\Gamma([0, 16.7]ms)$	μ	0.199	0.201	0.107	0.110	0.086
	σ	0.024	0.024	0.009	0.009	0.003
$\Gamma([0, 33.3]ms)$	μ	0.273	0.278	0.223	0.239	0.116
	σ	0.041	0.041	0.026	0.028	0.004
$\Gamma([0, 50.0]ms)$	μ	0.480	0.488	0.537	0.555	0.151
	σ	0.086	0.087	0.075	0.080	0.006

By including *Acrobatic* motions in the archetypal motions used to determine the model parameters, our dynamic models might also cover such movements. The reachable occupancy would then be larger and the verifier would be extra cautious at the expense of the robot's performance. As workshop behavior guidelines typically prohibit acrobatic-like movements, one may argue that this range of motion can be safely ignored. The exact movements used to obtain the dynamic parameters should be a matter for legislators and standards committees.

B. Volume Calculation

On a sample of the data, we evaluate the average volumes of $\Gamma([0, t_f])$ and compare it with the ISO-based approach (Section IV-E) and the approach from [36], where we do not account for acceleration range as a function of state. The results are in Table III. Fig. 16 shows the growth of the reachable set with t_f.

The set Γ_{ISO} is much smaller than either Γ_{3DOF} or Γ_{4DOF}, however, not conservative. The 4-DOF model is tighter at shorter prediction horizons than the 3-DOF model. That the
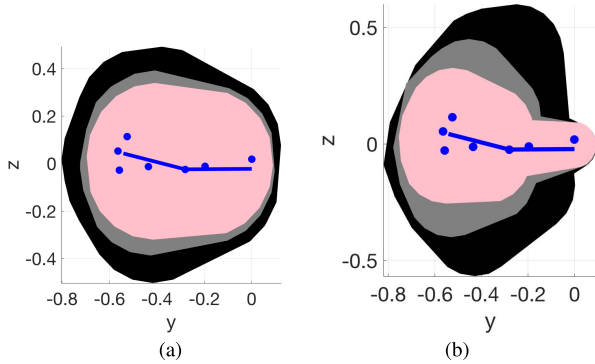
This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

12   IEEE TRANSACTIONS ON AUTOMATION SCIENCE AND ENGINEERING

Fig. 16. Expansion of reachable occupancies in (a) 3-DOF and (b) 4-DOF model. The time intervals for the occupancies are [0, 16.7] ms (pink), [0, 33.3] ms (gray), and [0, 50.0] ms (black). Circles are markers on the arm and lines are forearm and upper arm (scale in meters and origin at shoulder).

TABLE IV

COMPUTATION TIMES (ms)

Model	obtain $\mathcal{X}_0$	obtain $\mathcal{R}_X(t_f)$	obtain $\Gamma(t_f)$	Total
$\Gamma_{3DOF}([0, 16.7]ms)$	0.27	0.76	0.37	1.41
$\Gamma_{3DOF}([0, 33.3]ms)$	0.28	0.77	0.38	1.42
$\Gamma_{3DOF}([0, 50.0]ms)$	0.27	0.77	0.38	1.42
$\Gamma_{4DOF}([0, 16.7]ms)$	0.33	0.81	1.68	2.82
$\Gamma_{4DOF}([0, 33.3]ms)$	0.33	0.81	1.93	3.07
$\Gamma_{4DOF}([0, 50.0]ms)$	0.33	0.81	2.93	4.08

reachable occupancy of the 4-DOF model grows fast at larger prediction horizons is due to the fact that the enclosure of the motion of revolute joints is an overapproximation. The 4-DOF model has four revolute joints whereas the 3-DOF model has only two; hence, the overapproximation in the 4-DOF model is greater. A tighter approximation using a more complex SSV as suggested in [28, Eq. 26] may help, but increases complexity and computation time for generating the reachable occupancy and collision checking.

We observe an improvement in tightness over [36] with the current method: in the previous work, we intersect reachable sets from only three of the four models presented in Section V-B; the reachable set in this paper is strictly smaller than that from [36], meaning the reachable occupancy is smaller. The fourth model in Section V-B is more time consuming than the others. One could consider an "anytime" algorithm, whereby all dynamic models from Section V-B are calculated in parallel and the intersections of only those which are calculated in time are used to compute the reachable set: if the reachability analysis takes too long, only the first three models are used; if the reachability is computed in time, the reachable set from the fourth model is also used to refine the reachable set.

C. Computation Time

Computations were performed on MATLAB R2016b running on a 2.8 GHz Intel i7 processor with 16-GB RAM. Since there is no iterative step in the calculations, hard real time can be implemented. In Table IV, we see the computation time over a sample of data for all steps in the online part of the occupancy prediction, which stays well below the prediction horizon for $t_f \geq 16.7$ ms. Computation time does not vary with prediction horizon, apart from in the conversion from reachable set to reachable occupancy. This is because, as the reachable set expands and the joint intervals become larger, we choose to use more points to define the reachable occupancy, as explained in Section V-E and shown in Fig. 13; this increases computation time.

Overall, therefore, while both models are overapproximative, the 3-DOF model offers advantages in speed, at the expense of tightness at short time horizons. The 3-DOF model also deals better with the measurement uncertainty in arm lengths, as shown in Section IV-D, whereas the 4-DOF model requires an extra DOF to account for this.

A final consideration is that even when standing or sitting, movement of the spine and of the joints in the shoulder complex contributes to overall motion of the arm. Although such motion is small in range and slower compared with arm joints, it should be accounted for. Further human motion studies on torso movement are necessary to determine whether the complex but limited movement may be approximated by the enlargement of uncertainties in the simpler model, for example, as applied in [58], or whether additional DOFs may be necessary, as in [29].

VIII. CONCLUSION

We address the problem of guaranteeing safety in areas where humans and robots work together. From analysis of extreme movement data, we parameterize a 3-DOF and 4-DOF arm model that is then used to formally predict, online and overapproximatively, a volume that tightly bounds the arm occupancy during a certain time horizon.

Tests on publicly available data showed that our occupancies enclose human movement for HRCoex scenarios and also work well in scenarios unlikely to be found in an HRCoex environment. Mean computation times of less than 5 ms demonstrate the suitability of this approach to online dynamic path planning. Future work should focus on reliable state estimation techniques from sensors that can be employed in HRCoex scenarios. In contrast to probabilistic methods, this technique may be used to formally certify safety in shared human–robot workspaces and pave the way for coexistence without cages—a goal long hindered by a lack of guaranteed safety.

APPENDIX
MARKER POSITIONS ON ARM

Marker positions are illustrated in Fig. 17. The fitting of the marker data to the kinematic model was performed so: the orientation of the base coordinate system is that of the clavicle in the ISB recommendations [59], defined by markers T10 and C7 on the midspine and upper spine, and STRN and CLAV
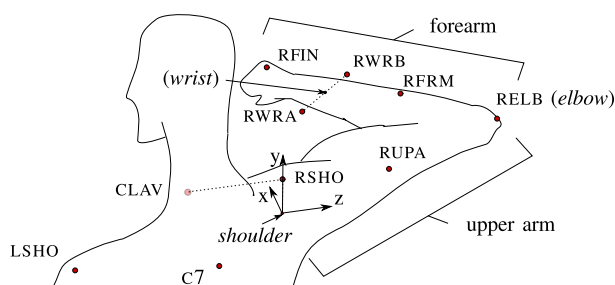
Fig. 17. Markers and local base coordinate system. The marker CLAV is at the top of the sternum on the anterior of the torso; STRN (not shown) is at the base of the sternum and T10 (not shown) is on the back on the midspine. The base coordinate system is 40 mm below RSHO (in the negative y-direction) and its orientation defined by T10, STRN, CLAV, and C7. The z-axis is parallel to the line through RSHO and CLAV.

on the lower and upper ends of the breastbone. The origin is at the center of the (spherical) GH Joint, estimated at 40mm under (in negative y-direction from) the marker RSHO/LSHO (the prefix L/R refers to left/right, respectively), and referred to as the *shoulder*. The joint angles, however, are defined differently from [59], motivated in Section IV-C. The *elbow* position was taken at RELB/LELB. The *wrist* position was taken as the midpoint between the two wrist markers RWRA and RWRB/LWRA and LWRB.

ACKNOWLEDGMENT

REFERENCES

[1] *Safety of Machinery—Positioning of Safeguards With Respect to the Approach Speeds of Parts of the Human Body*, Int. Org. Standardization, Geneva, Switzerland, ISO Standard 13855:2010, 2010.
[2] *Robots and Robotic Devices—Safety Requirements for Industrial Robots—Part 1: Robots*, Int. Org. Standardization, Geneva, Switzerland, ISO Standard 10218-1:2011, 2011.
[3] M. Althoff and B. H. Krogh, "Reachability analysis of nonlinear differential-algebraic systems," *IEEE Trans. Autom. Control*, vol. 59, no. 2, pp. 371–383, Feb. 2014.
[4] G. V. Smirnov, *Introduction to the Theory of Differential Inclusions*. Providence, RI, USA: AMS, 2002.
[5] G. Lafferriere, G. J. Pappas, and S. Yovine, "A new class of decidable hybrid systems," in *Hybrid Systems: Computation and Control*. Berlin, Germany: Springer, 1999, pp. 137–151.
[6] S. Petti and T. Fraichard, "Safe motion planning in dynamic environments," in *Proc. IEEE-RSJ Int. Conf. Intell. Robots Syst.*, Aug. 2005, pp. 2210–2215.
[7] *Robots and Robotic Devices—Collaborative Robots*, Int. Org. Standardization, Geneva, Switzerland, ISO/TS Standard 15066:2016(E), 2016.
[8] A. Pereira and M. Althoff, "Safety control of robots under Computed Torque control using reachable sets," in *Proc. IEEE Int. Conf. Robot. Autom.*, May 2015, pp. 331–338.
[9] M. J. Zeestraten, A. Pereira, M. Althoff, and S. Calinon, "Online motion synthesis with minimal intervention control and formal safety guarantees," in *Proc. IEEE Int. Conf. Syst., Man Cybern.*, Oct. 2016, pp. 2116–2121.
[10] H. Roehm, J. Oehlerking, M. Woehrle, and M. Althoff, "Reachset conformance testing of hybrid automata," in *Proc. Hybrid Syst., Comput. Control*, 2016, pp. 277–286.
[11] S. Haddadin, *Towards Safe Robots: Approaching Asimov's 1st Law*. Heidelberg, Germany: Springer, 2014.
[12] P. Mittendorfer, E. Yoshida, and G. Cheng, "Realizing whole-body tactile interactions with a self-organizing, multi-modal artificial skin on a humanoid robot," *Adv. Robot.*, vol. 29, no. 1, pp. 51–67, 2015.
[13] A. Verl, A. Albu-Schäffer, O. Brock, and A. Raatz, Eds., *Soft Robotics: Transferring Theory to Application*. Berlin, Germany: Springer, 2015.
[14] R. Rossi, M. P. Polverini, A. M. Zanchettin, and P. Rocco, "A pre-collision control strategy for human-robot interaction based on dissipated energy in potential inelastic impacts," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, Sep. 2015, pp. 26–31.
[15] D. Kulić and E. Croft, "Real-time safety for human—Robot interaction," in *Proc. Int. Conf. Adv. Robot.*, 2005, pp. 719–724.
[16] A. M. Zanchettin and P. Rocco, "Path-consistent safety in mixed human-robot collaborative manufacturing environments," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, Nov. 2013, pp. 1131–1136.
[17] S. M. LaValle, *Planning Algorithms*. Cambridge, U.K.: Cambridge Univ. Press, 2006. [Online]. Available: http://planning.cs.uiuc.edu/
[18] S. Albrecht, M. Leibold, and M. Ulbrich, "A bilevel optimization approach to obtain optimal cost functions for human arm-movements," *Numer. Algebra, Control Optim.*, vol. 2, no. 1, pp. 105–127, 2012.
[19] P. Morasso, "Three dimensional arm trajectories," *Biological*, vol. 48, no. 3, pp. 187–194, 1983.
[20] T. Flash and N. Hogan, "The coordination of arm movements: An experimentally confirmed mathematical model," *J. Neurosci.*, vol. 5, no. 7, pp. 1688–1703, Jul. 1985.
[21] V. De Sapio, J. Warren, and O. Khatib, "Predicting reaching postures using a kinematically constrained shoulder model," in *Advances in Robot Kinematics*, J. Lenarčič and B. Roth, Eds. Dordrecht, The Netherlands: Springer, 2006, pp. 209–218.
[22] D. Bortot, H. Ding, A. Antonopolous, and K. Bengler, "Human motion behavior while interacting with an industrial robot," *Work*, vol. 41, pp. 1699–1707, Jan. 2012.
[23] H. S. Koppula and A. Saxena, "Anticipating human activities using object affordances for reactive robotic response," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 38, no. 1, pp. 14–29, Jan. 2015.
[24] F. Rohrmüller, M. Althoff, D. Wollherr, and M. Buss, "Probabilistic mapping of dynamic obstacles using Markov chains for replanning in dynamic environments," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, Sep. 2008, pp. 2504–2510.
[25] H. Ding, K. Wijaya, G. Reißig, and O. Stursberg, "Online computation of safety-relevant regions for human robot interaction," in *Proc. 43rd Intl. Symp. Robot. (ISR)*, 2012, pp. 29–31.
[26] J. Mainprice and D. Berenson, "Human-robot collaborative manipulation planning using early prediction of human motion," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, Nov. 2013, pp. 299–306.
[27] R. Luo and D. Berenson, "A framework for unsupervised online human reaching motion recognition and early prediction," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, Sep. 2015, pp. 2426–2433.
[28] H. Täubig, B. Bäuml, and U. Frese, "Real-time swept volume and distance computation for self collision detection," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, pp. 1585–1592, Sep. 2011.
[29] M. Ragaglia, A. M. Zanchettin, and P. Rocco, "Safety-aware trajectory scaling for human-robot collaboration with prediction of human occupancy," in *Proc. Int. Conf. Adv. Robot.*, Jul. 2015, pp. 85–90.
[30] M. Althoff, "Reachability analysis and its application to the safety assessment of autonomous cars," Ph.D. dissertation, Dept. Fakultät für Elektrotechnik, Tech. Univ. Munich, München, Germany, 2010.
[31] G. Frehse, "Verification tools for linear hybrid automata," in *Handbook of Hybrid Systems Control*, J. Lunze and F. Lamnabhi-Lagarrigue, Eds. Cambridge, U.K.: Cambridge Univ. Press, 2009, pp. 285–296.
[32] E. Asarin, T. Dang, G. Frehse, A. Girard, C. Le Guernic, and O. Maler, "Recent progress in continuous and hybrid reachability analysis," in *Proc. IEEE Conf. Comput. Aided Control Syst. Design*, Oct. 2006, pp. 1582–1587.
[33] K. Holzbaur, W. Murray, and S. Delp, "A model of the upper extremity for simulating musculoskeletal surgery and analyzing neuromuscular control," *Ann. Biomed. Eng.*, vol. 33, no. 6, pp. 829–840, 2005.
[34] E. Demircan, T. Besier, S. Menon, and O. Khatib, "Human motion reconstruction and synthesis of human skills," in *Advances in Robot Kinematics: Motion in Man and Machine*, J. Lenarčič and M. M. Stanišić, Eds. Dordrecht, The Netherlands: Springer, 2010, pp. 283–292.
[35] E. Demircan, T. F. Besier, and O. Khatib, "Muscle force transmission to operational space accelerations during elite golf swings," in *Proc. Int. Conf. Robot. Autom.*, 2012, pp. 1464–1469.
[36] A. Pereira and M. Althoff, "Overapproximative arm occupancy prediction for human-robot co-existence built from archetypal movements," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, Oct. 2016, pp. 1394–1401.

[37] R. Poppe, "Vision-based human motion analysis: An overview," *Comput. Vis. Image Understand.*, vol. 108, nos. 1–2, pp. 4–18, Oct. 2007.

[38] J. Shotton *et al.*, "Real-time human pose recognition in parts from single depth images," in *Machine Learning for Computer Vision*, R. Cipolla, S. Battiato, and M. G. Farinella, Eds., Berlin, Germany: Springer, 2013, pp. 119–135.

[39] A. A. Amis, D. Dowson, and V. Wright, "Analysis of elbow forces due to high-speed forearm movements," *J. Biomech.*, vol. 13, no. 10, pp. 825–831, 1980.

[40] J. Rosen, J. C. Perry, N. Manning, S. Burns, and B. Hannaford, "The human arm kinematics and dynamics during daily activities—Toward a 7 DOF upper limb powered exoskeleton," in *Proc. Int Conf Adv. Robot.*, 2005, pp. 532–539.

[41] N. Klopčar, M. Tomšič, and J. C. Lenarčič, "A kinematic model of the shoulder complex to evaluate the arm-reachable workspace," *J. Biomech.*, vol. 40, no. 1, pp. 86–91, 2001.

[42] A. M. Zanchettin, L. Bascetta, and P. Rocco, "Achieving humanlike motion: Resolving redundancy for anthropomorphic industrial manipulators," *IEEE Robot. Autom. Mag.*, vol. 20, no. 4, pp. 131–138, Dec. 2013.

[43] K. Abdel-Malek, J. Yang, R. Brand, and E. Tanbour, "Towards understanding the workspace of human limbs," *Ergonomics*, vol. 47, no. 13, pp. 1386–1405, 2004.

[44] S. Pheasant and C. M. Haslegrave, "Hands and handles," in *Bodyspace: Anthropometry, Ergonomics and the Design of Work*. Boca Raton, FL, USA: CRC Press, 2006, pp. 143–160

[45] T. Koritnik, T. Bajd, and M. Munih, "A simple kinematic model of a human body for virtual environments," in *Advances in Robot Kinematics: Motion in Man and Machine: Motion in Man and Machine*, J. Lenarčič and M. M. Stanišić, Eds. Dordrecht, The Netherlands: Springer, 2010, pp. 401–408.

[46] J. C. Otis, C.-C. Jiang, T. L. Wickiewicz, M. G. Peterson, R. F. Warren, and T. J. Santner, "Changes in the moment arms of the rotator cuff and deltoid muscles with abduction and rotation," *J. Bone Joint Surgery*, vol. 76, no. 5, pp. 667–676, 1994.

[47] R. Featherstone and D. E. Orin, "Dynamics," in *Handbook of Robotics*, B. Siciliano and O. Khatib, Eds. Springer, 2008, pp. 35–65.

[48] M. Althoff, O. Stursberg, and M. Buss, "Reachability analysis of nonlinear systems with uncertain parameters using conservative linearization," in *Proc. 47th IEEE Conf. Decision Control*, Dec. 2008, pp. 4042–4048.

[49] M. Althoff, D. Heß, and F. Gambert, "Road occupancy prediction of traffic participants," in *Proc. 16th Int. IEEE Conf. Intell. Transp. Syst.*, Oct. 2013, pp. 99–105.

[50] B. M. Nigg and W. Herzog, *Biomechanics of the Musculoskeletal System*. Hoboken, NJ, USA: Wiley, 1999.

[51] D. Hahn, M. Olvermann, J. Richtberg, W. Seiberl, and A. Schwirtz, "Knee and ankle joint torque–angle relationships of multi-joint leg extension," *J. Biomech.*, vol. 44, no. 11, pp. 2059–2065, 2011.

[52] T. S. Arthanari and Y. Dodge, *Mathematical Programming in Statistics*. Hoboken, NJ, USA: Wiley, 1981.

[53] E. G. Gilbert, D. W. Johnson, and S. S. Keerthi, "A fast procedure for computing the distance between complex objects in three-dimensional space," *IEEE J. Robot. Autom.*, vol. 4, no. 2, pp. 193–203, Apr. 1988.

[54] M. Althoff, "An introduction to CORA 2015," in *Proc. Workshop Appl. Verification Continuous Hybrid Syst.*, 2015, pp. 120–151.

[55] S. Rump, "INTLAB - INTerval LABoratory," in *Developments in Reliable Computing*, T. Csendes, Ed. Dordrecht, The Netherlands: Kluwer, 1999, pp. 77–104. [Online]. Available: http://www.ti3.tuhh.de/rump/

[56] P. I. Corke, *Robotics, Vision and Control: Fundamental Algorithms in MATLAB*. Berlin, Germany: Springer, 2011.

[57] M. Herceg, M. Kvasnica, C. Jones, and M. Morari, "Multi-parametric toolbox 3.0," in *Proc. Eur. Control Conf.*, Zürich, Switzerland, 2013, pp. 502–510. [Online]. Available: http://control.ee.ethz.ch/~mpt

[58] M. Althoff and J. M. Dolan, "Reachability computation of low-order models for the safety verification of high-order road vehicle models," in *Proc. Amer. Control Conf.*, 2012, pp. 3559–3566.

[59] G. Wu *et al.*, "ISB recommendation on definitions of joint coordinate systems of various joints for the reporting of human joint motion—Part II: Shoulder, elbow, wrist and hand," *J. Biomech.*, vol. 38, no. 5, pp. 981–992, 2005.

Aaron Pereira received the M.Eng. degree in mechanical engineering from Imperial College London, London, U.K., in 2013. He is currently pursuing the Ph.D. degree in robotics at Technische Universität München, Garching, Germany.

His research interests include formal methods, human robot interaction, trajectory planning, and robot ethics.

Matthias Althoff received the Dipl.-Ing. degree in mechanical engineering and the Ph.D. degree in electrical engineering from Technische Universität München, Munich, Germany, in 2005 and 2010, respectively.

From 2010 to 2012, he was a Post-Doctoral Researcher at Carnegie Mellon University, Pittsburgh, PA, USA, and from 2012 to 2013, he was an Assistant Professor at Technische Universität Ilmenau, Ilmenau, Germany. He is currently an Assistant Professor of Computer Science at Technische Universität München. His research interests include formal verification of continuous and hybrid systems, reachability analysis, planning algorithms, nonlinear control, automated vehicles, robotics, and power systems.