

Predictive and Robust Robot Assistance for Sequential Manipulation Tasks

Theodoros Stouraitis, Michael Gienger

2023

Preprint:

This is an accepted article published in IEEE Research and Automation Letters (RA-L). The final authenticated version is available online at:
<https://doi.org/10.1109> Copyright 2023 IEEE

Predictive and Robust Robot Assistance for Sequential Manipulation

Theodoros Stouraitis¹ and Michael Gienger¹

Abstract—This paper presents a novel concept to support physically impaired humans in daily object manipulation tasks with a robot. Given a user’s manipulation sequence, we propose a predictive model that uniquely casts the user’s sequential behavior as well as a robot support intervention into a hierarchical multi-objective optimization problem. A major contribution is the prediction formulation, which allows to consider several different future paths concurrently. The second contribution is the encoding of a general notion of constancy constraints, which allows to consider dependencies between consecutive or far apart keyframes (in time or space) of a sequential task. We perform numerical studies, simulations and robot experiments to analyse and evaluate the proposed method in several table top tasks where a robot supports impaired users by predicting their posture and proactively re-arranging objects.

Index Terms—Physical Human-Robot Interaction; Human-Aware Motion Planning; Optimization and Optimal Control

I. INTRODUCTION

ASSISTIVE robots can play a crucial role for more than 200 million humans that need assistance with activities of daily living (ADLs) [1], such as getting out of bed, going to the restroom, preparing a meal, *etc.* Towards addressing this societal challenge [2], the areas of physical Human-Robot-Interaction (pHRI) [3] and rehabilitation robotics [4] have developed approaches for ergonomic [5], safe [6], and personalised [7] support. Key challenges studied are capturing human motion [8], anticipating human behaviour [9], modelling robot behaviour influence to the human behaviour [10], and coping with variations [11], [12] of human behavior due to factors like disabilities, partial observability, *etc.*

In this paper, we study these challenges in the context of manipulation. Our focus is on temporally extended ADLs that require sequential manipulation capabilities. As a typical example, consider an impaired human (immobilized left arm) preparing a beverage, as shown in Fig. 1 and Fig. 2. This task requires bimanual operations, and due to the injury the human has rely on the robot support. To provide optimal support, the robot needs to predict the human’s actions (discrete decisions) and their outcome (continuous states) [13], assess the possible human postures [14], and intervene to assist. Key to this scenario is that the prediction needs to consider a set of possible futures, as accurate predictions of the human’s actions might not be possible.

Manuscript received: March 15, 2023; Revised: July 23, 2023; Accepted: August 24, 2023. This paper was recommended for publication by Editor Aniket Bera upon evaluation of the Associate Editor and Reviewers’ comments. (Corresponding author: Theodoros Stouraitis)

¹Theodoros Stouraitis and Michael Gienger are with Honda Research Institute Europe (HRI-EU), 63073 Offenbach am Main, Germany. (e-mails: stoutho@gmail.com, michael.gienger@honda-ri.de)

The authors would like to thank Stephan Hasler for his help with the experimental setup.

Digital Object Identifier (DOI): see top of this page.



Fig. 1. Illustration of the presented concept in a supported pour-bottle-into-glass task. A robot intervenes to move the glass (transparent view) to a location on the table that enables an impaired user to comfortably pour the content of a bottle into it. The supported human (solid view) is enabled to perform the task optimally given an ergonomic metric. In the non-supported case, the (transparent) human needs to get into an uncomfortable posture.

In this work, we assume that the action schedule (task-level coordination [15]) is known, and we focus on interventions that (i) adjust continuous geometric quantities, *e.g.* change the pose of the mug, (ii) parallel actions, *e.g.* while the human picks the bottle the robot fetches the mug. Following the “assist-as-needed” principle¹, we view the robot interventions as preparatory actions [16] that enable or ease subsequent actions and motions of the impaired human. This translates into a method that has as backbone a physics-based prediction mechanism, which allows us to investigate the effects of possible (robot) interventions. The prediction mechanism can be considered as part of combined task and motion planning (TAMP) [17], where the task schedule (linear or tree sequence of modes in the form of constraints) is given. Adapting online the sequences of modes [18], [19] is beyond the scope of this work. Thus, we ask ourselves the question “*how to turn an action schedule into a prediction of the human’s plan while considering: (a) set of possible futures, and (b) the influence of intervention by a robot?*”

We focus on the low-level geometric choices in keyframes, rather than computing trajectories as in Trajectory Optimization methods. Keyframes represent important transitions within the action sequence such as contact transitions or the start of a new task-level trajectory. They describe the subgoal structure of the task. Based on keyframes we developed an efficient constraint resolution method that; (i) can consider both a linear and a tree-like sequence of modes, (ii) holds robot actions and their effects, (iii) scales to large sequences of high dimensional systems, like the DoFs of the human body, and (iv) is reactive in the sense that predictions can be updated within few hundreds of milliseconds. We call this, reactive Tree-

¹“assist-as-needed” indicates that assistance should be minimal to motivate the voluntary contribution of the impaired human (patient).




Fig. 2. Illustration of a sequential manipulation where the robot assists a human to prepare a beverage by pushing the milk and rotating the mug.

Consistent-Constraint (TCC) resolution method. The action schedule defines a set of constraints, and TCC predicts a sequence (linear or tree) of keyframes that satisfy them. The interdependencies between these keyframes are formed with spatio-temporal constraints that create links between bodies in task space, such as objects, end-effectors, *etc.* Interventions are represented with open variables that are optimised to satisfy secondary objectives, such as postural preferences. The contributions of our work are:

- The reactive Tree-Consistent-Constraint (TCC) method that can simultaneously compute a set of possible futures based on temporal and concurrent dependencies.
- A computation formalism based on Damped Least Squares (DLS) to cast cooperative sequential manipulation as a single underactuated system where robot assistance (*i.e.* interventions) lies in the null-space of manipulation sequences.
- Demonstrating the capabilities of TCC to predict and optimize keyframes of uncertain human manipulation sequences including assistive robot interventions.

In Section II, we review the related work, in Section III we formally set the problem and in Section IV, we present our method. In Section V we evaluate the proposed method and demonstrate its capabilities in real world setups.

II. RELATED WORK

A. Human-aware prediction, planning and control

An extensively studied prediction area is ground-level 2D trajectories [20]. In contrast to that, we focus on predicting articulated full body motions in manipulation sequences. Prior model-based approaches demonstrate local adaptation of the robot posture that accounts for the overloading human joint torques [5], [21]. Others learned models to predict human motion and to adapt the robot plans. In [22] human-robot interaction primitives are used to select and coordinate the robot movement with the recognized human actions, while in [6], [23] safe (in terms of collision) concurrent (human and robot) reaching motions are demonstrated. These works studied prediction of articulated bodies, yet they only consider short-term interactions without sequential dependencies, *i.e.* interdependencies in a series of actions.

Most relevant to our work are cooperative TAMP methods [13], [24] that plan decisions and motions of both the human and the robot (articulated bodies) towards finding offline the optimal robot manipulation plans. In our approach predictions can be updated on-the-fly, as in [24], yet ours also enables simultaneous consideration of a set of possible future sequences, rather than a single linear sequence.

B. Spatio-temporal constraints

Prior work has proposed various constraint formulations to efficiently synthesize whole-body motion plans. These motions exhibit complicated dependencies across concurrent tasks (same time-slice) [25], and along the horizon (between keyframes) [18]. This led to the question: "How to model concurrent and sequential dependencies with constraints?"

Explicit geometric constraints (distances, friction cone, *etc.*) have been used in [26] to form and solve a multi-robot co-manipulation problem. Task-space-region-chains [25] were proposed to link together a set of workspace-goal-regions (similar to task-intervals [27]) and form a manifold where RRT-based planners can search for whole-body manipulation paths. Both approaches can consider concurrent dependencies, *e.g.* closed kinematics chains, multiple robots postures, *etc.* Yet, they did not consider sequential representations of several connected postures.

Temporal dependencies—a crucial modelling aspect of sequential motions—were analytically modelled by linking consecutive footsteps given symmetric gait patterns in [28]. The Contact-Consistent Elastic Strips framework [29] combined the concept of task-intervals [27] or task-space-region [25] with contact-consistent constraints that couple the consecutive robot poses to generate multi-contact locomotion plans. More generally temporal dependencies can be encoded as *constancy constraints* [18] either by pair-wise coupling constraints or by sharing of DoFs across time-slices [30]. We use task-intervals [27] and we broaden the notion of *constancy constraints* to capture both temporal [18] and concurrent dependencies. This enables us to optimize manipulation planning trees, rather than linear sequences.

III. PROBLEM STATEMENT

Our goal is to investigate the influence of robot assistive actions in sequential manipulation tasks that include a set of objects and are performed by an impaired human. To do that, we break-down the problem into two focal points.

Following the multi-modal motion planning (MMMP) notation, we denote the configuration space (C-space) of K movable objects² as $\mathcal{O} = O^1 \times \dots \times O^K \subseteq SE(3)^K$, the C-space of the human is $H \subseteq \mathbb{R}^\alpha$. Hence, the configuration space of the prediction problem is $\mathcal{P} = H \times \mathcal{O} \subseteq SE(3)^K \times \mathbb{R}^\alpha$, in which we are searching for paths (to be executed by a human) that have a known initial configuration $\mathbf{p}_0 \in \mathcal{P}$ and comply with constraints of the sequential manipulation task. These consider task-scenario specifications—*e.g.* impairment, connected objects, stable contacts, constant grasps, *etc.*—and temporal transitions—*e.g.* whether an object can move or

²The superscripts on O s, θ s and B s are indexes and not exponents.

not, whether an object can slide in a hand or not, *etc*—such that the path is feasible. The former partitions the configuration space, while the latter shapes how the partitions of the configuration space are linked together. We represent specifications and transitions with constraints as $\mathbf{T}_{\mathcal{P}}(\cdot) \leq 0$, where $\mathbf{T}_{\mathcal{P}}(\cdot) : \mathcal{P}^{\gamma} \rightarrow \mathbb{R}^{\gamma}$ and γ is the total number of constraints (more details are provided in Section IV).

Many possible futures: A sequential manipulation task is conditional on the preferences of the human that performs it. For example, a human might prefer to first pick up the milk and then the mug (see Fig. 2) or vice versa, or a human might prefer to power-grasp the mug, rather than picking it up from the handle using a pinch-grasp. We describe these preferences by the latent variable θ , which is one out of a finite set of many predefined preferences² $\Theta = \{\theta^1, \theta^2, \dots, \theta^v\}$, also referred to as modes in MMMP. The human preferences are not likely to be known in advance, hence a prediction of N future manipulation steps should be able to consider many preferences concurrently.

First focal point: Develop a prediction method that *predicts* a set \mathcal{S} of feasible future paths in \mathcal{P} . This set is

$$\mathcal{S} = \{\mathbf{h}_{1:N}, \mathbf{o}_{1:N}^k : \mathbf{T}_{\mathcal{P}}(\mathbf{p}_0, \mathbf{h}_{1:N}, \mathbf{o}_{1:N}^k; \Theta) \leq 0\}, \quad (1)$$

where $\mathbf{h}_{1:N}$ and $\mathbf{o}_{1:N}^k$ are linear sequences of configurations for the human and the movable objects, respectively.

Influence of interventions: Let us consider that an assistive agent (robot)—with C-space $Q \subseteq \mathbb{R}^{\beta}$ —can also manipulate the same movable objects. Including the robot agent results in an extension of the configuration space to $\mathcal{X} = \mathcal{P} \times Q \subseteq SE(3)^K \times \mathbb{R}^{\alpha+\beta}$, with total number of dimensions represented by $d_{\mathcal{X}} \in \mathbb{R}$. Thus, we extend $\mathbf{T}_{\mathcal{P}}(\cdot)$ to $\mathbf{T}_{\mathcal{X}}(\cdot) \leq 0$ with $\mathbf{T}_{\mathcal{X}}(\cdot) : \mathcal{X}^{\gamma} \rightarrow \mathbb{R}^{\gamma}$, such that it can also encode robot specific constraints. Using the above we can now define a *prediction set* \mathcal{S}' of feasible future paths that are influenced by the robot's interventions, as

$$\mathcal{S}' = \{\mathbf{h}_{1:N}, \mathbf{o}_{1:N}^k : \mathbf{T}_{\mathcal{X}}(\mathbf{x}_0, \mathbf{h}_{1:N}, \mathbf{o}_{1:N}^k, \mathbf{q}_{1:N}; \Theta) \leq 0\}, \quad (2)$$

where $\mathbf{x}_0 \in \mathcal{X}$ is a known initial configuration, and our goal is to optimize the robot's interventions such that \mathcal{S}' has lower cost than \mathcal{S} according to a metric of the human state.

Second focal point: Find a single linear sequence of robot configurations $\mathbf{q}_{1:N}$ that minimizes a metric for all possible future paths of the human. This can be expressed as

$$\min_{\mathbf{q}_{1:N}} \mathbf{c}(\mathbf{h}_{1:N}) \quad (3a)$$

$$\text{s.t. } \forall \mathbf{h}_{1:N} \subset \mathcal{S}' \quad (3b)$$

where \mathcal{S}' contains all the possible future paths of the human which according to (2) depends on the robot plan $\mathbf{q}_{1:N}$, *i.e.* robot interventions. Cost function $\mathbf{c}(\cdot)$ is a metric of the human state, *e.g.* comfort, ergonomics, effort, *etc.* Note that (3) is a Robust Optimization (RO) problem as we seek to find a solution that minimizes a metric for all possible futures. Specifically, this is a *multiforecast model predictive optimization* formulation (see sec. 9.9.3 in [31]).




Fig. 3. Illustration of the constraints used to describe the first three keyframes of Fig. 2 with the robot and its actions being omitted for clarity. *Task constraints* are used only within keyframes and *constancy constraints* are used between keyframes, hence they are illustrated with arrows to link timeframes and diamonds to indicate which dimensions remain constant.

IV. METHOD

To achieve the two focal points stated above, we propose a model-based optimisation approach that has the two following key aspects. First, the constraints $\mathbf{T}_{\mathcal{X}}(\cdot) \leq 0$ are *task constraints* for each keyframe and *constancy constraints* that encapsulate a notion of invariance between keyframes. Fig. 3 shows these types of constraints being used to describe a sequence of manipulation actions. Second, we developed a barrier method that uses a Levenberg–Marquardt update to resolve all the constraints concurrently (even across parallel future paths). Solve an optimization problem with all constraints shown in Fig. 3 results in the corresponding keyframes of Fig. 2 (three first). This allows us to regard all possible futures (via the constraints) as primary objectives and treat the robot's assistance as secondary objectives and hence, find a single linear sequence of robot interventions that improves the human state for all possible futures. Next, we describe these two aspects in detail.

A. Spatio-temporal constraints

To ease the description that follows, we introduce L number of bodies that belong to the kinematic chains of the human or the robot (functions with domain H and Q)³, and we denote their C-space² as $B^1, \dots, B^L \subseteq SE(3)$.

Task constraints: We use *task constraints* to analytically model each manipulation step (time-slice). These can be written as $\mathbf{g}(\mathbf{x}_i) \leq 0$, where $\mathbf{g}(\cdot) : \mathcal{X} \rightarrow \mathbb{R}^{\gamma_g}$ with $\gamma_g \in \mathbb{R}$ and i indicates the current keyframe in a sequence. The explicit form, in terms of input arguments, of $\mathbf{g}(\cdot)$ is

$$\mathbf{g}(\{\mathbf{o}_i^k\}, \mathbf{h}_i, \mathbf{q}_i, \{\mathbf{b}_i^l\}) \leq 0, \quad (4)$$

where $k \in [1, \dots, K]$ is the index to an object, $l \in [1, \dots, L]$ is the index to a body, and $\mathbf{o}_i^k \in O^k$, $\mathbf{h}_i \in H$, $\mathbf{q}_i \in Q$, and $\mathbf{b}_i^l \in B^l$ are configuration vectors. With the explicit form we formulate *task constraints* that describe manipulation actions, *e.g.* grasping and pushing actions— $\mathbf{g}(\mathbf{o}_i^k, \mathbf{b}_i^l)$, and relations in task space between objects and bodies, *e.g.* collision between objects— $\mathbf{g}(\mathbf{o}_i^k, \mathbf{o}_i^{k'})$. With *task constraints* we also enforce kinematic and dynamic functions of multi-bodied systems, *e.g.* forward/inverse kinematics of the human $\mathbf{g}(\mathbf{h}_i, \{\mathbf{b}_i^l\})$ and the robot $\mathbf{g}(\mathbf{q}_i, \{\mathbf{b}_i^l\})$. *Task constraints* are also used to represent human impairments, *e.g.* use of cane for support, limited joint ranges at the knee, *etc.* Each keyframe can have many *task*

³Note, that B s are redundant, they are not represented explicitly in our method and are only used here to ease the description of constraints.

Algorithm 1: From task to constancy constraints

```

Input:  $\mathcal{T}$ ; /* Set of task constr. */
1 foreach  $i \in [1, \dots, N]$  do
2    $j \leftarrow \text{par}(i)$ ; /* Get parent index */
3   foreach  $\omega \in [o^1, \dots, o^K] \vee [b^1, \dots, b^L]$  do
4     // Check if parent & child have same task
5     // constr.
6     if  $\mathbf{g}(\omega_i) == \mathbf{g}(\omega_j)$  then
7       remove  $\mathbf{g}(\omega_j)$  from  $\mathcal{T}$ ;
8       add  $\bar{\mathbf{f}}(\omega_i, \omega_j)$  in  $\mathcal{T}$ ;
9     end
10    // For any object or any body
11    foreach  $\beta \in [o^1, \dots, o^K] \vee [b^1, \dots, b^L]$  do
12      // Check if relative task constr. between
13      // objects and/or bodies is constant
14      if  $\mathbf{g}(\omega_i, \beta_i) == \mathbf{g}(\omega_j, \beta_j)$  then
15        remove  $\mathbf{g}(\omega_j, \beta_j)$  from  $\mathcal{T}$ ;
16        add  $\bar{\mathbf{f}}(\omega_i, \beta_i, \omega_j, \beta_j)$  in  $\mathcal{T}$ ;
17      end
18    end
19  end
20 return  $\mathcal{T}$ ; /* Updated constr. set */

```

constraints that together specify critical points of the sequence (see Section I) or mode switches (*a.k.a.* transition manifolds) in MMMP or TAMP terminology [32] (*a.k.a.* jumps in hybrid dynamical systems).

Constancy constraints: We use *constancy constraints* to capture dependencies between keyframes, such as fixed pose between a hand and a bottle in three keyframes (i) grasping a bottle, (ii) holding it to pour into a mug and (iii) place it on the table, *i.e.* *constancy constraints* describe which motion dimensions between different keyframes remain constant.

Generally, a system’s evolution within a mode (*a.k.a.* system dynamics or flow in hybrid dynamical systems) can be described using *running constraints* [18]. Their discrete form is $f(\mathbf{x}_i, \mathbf{x}_{i+1}, \mathbf{u}_i) \leq 0$, with \mathbf{u}_i representing the control input. In case certain dimensions of a linear sequence remain constant (*e.g.* stable grasps, contacts, *etc.*) or bounded (*e.g.* in-hand motions; an object sliding in the hand), these can be simplified to *constancy constraints* $\bar{f}(\mathbf{x}_i, \mathbf{x}_{i+1}) \leq 0$ that encode temporal indifferences between consecutive in time keyframes [18]. Yet, when considering tree-like sequences (many possible concurrent futures), there is also the need to capture dependencies that occur between different keyframes of the same time-slice. Thus, we define *constancy constraints* as

$$\bar{\mathbf{f}}(\mathbf{x}_i, \mathbf{x}_j) \leq 0, \quad (5)$$

where $\bar{\mathbf{f}}(\cdot) : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}^{\gamma_{\bar{\mathbf{f}}}}$ with $\gamma_{\bar{\mathbf{f}}} \in \mathbb{R}$ specify a pairwise coupling between i and j indexes of any two keyframes ($i \neq j$) even of the same time-slice.

Sequential invariance: Consider a keyframe where a human holds a bottle with the right hand, this keyframe can be followed by a keyframe where the human places the bottle on the table while still holding it with the right hand, but not by a keyframe where the human holds the same right hand a glass. In such a case, some dimensions of \mathcal{X} (see Section III) between adjacent⁴ keyframes remain the same—invariant—*e.g.* the distance between the right hand and the bottle (holding). Hence, such invariant dimensions can be enforced with *constancy constraints* based on Algorithm 1.

⁴Adjacent keyframes are specified via parent-child relationship.

Minimal description: We ensure that *task* and *constancy constraints* are described using a minimal set of coordinates, *e.g.* an object placed upright on the table requires a 1-dof height and a 1-dof inclination constraint, and not a 6-dof pose constraint. Further, constraints are modelled as inequalities to allow task coordinates to vary within some given bounds, *e.g.* the placement of an object on a table may be anywhere within the extents of the table surface.

By exploiting the two above properties, we attain the maximal null space for each keyframe and we link the keyframes’ null spaces along the sequence (linear or tree) to consider favorably interdependencies between keyframes and holistically optimize objectives across keyframes. This enables our method to reach better minima for the whole sequence as shown in Section V-A.

B. Tree-Consistent-Constraint Problem

By assembling a sequence of *task* and *constancy constraints* as described above, we can specify the schedule of manipulation tasks (*a.k.a.* skeleton [13]) with which we can predict how the human might execute the sequential task. We treat the prediction of the sequential task as *primary objective*, *e.g.* will the human pick the bottle or the mug, and the assistance as *secondary objective* (inspired by the “assist-as-needed” principle), *e.g.* where should the robot push the bottle. Based on this strict hierarchy of objectives, we realize the RO problem (predict and assist) defined in (3) as

$$\text{lexmin}_{\mathbf{x}_{1:N}} \{ \mathbf{c}_1(\mathbf{x}_{1:N}), \mathbf{c}_2(\mathbf{x}_{1:N}) \} \quad (6a)$$

$$\text{s.t. } \bar{\mathbf{f}}(\mathbf{x}_{1:N}, \mathbf{x}'_{1:N}) + \mathbf{c}_1(\mathbf{x}_{1:N}) = 0, \quad (6b)$$

$$\mathbf{g}(\mathbf{x}_{1:N}) + \mathbf{c}_1(\mathbf{x}_{1:N}) = 0, \quad (6c)$$

which is a lexicographic multi-objective problem [33], *i.e.* a problem with strict priority of objectives. $\bar{\mathbf{f}}(\cdot)$ and $\mathbf{g}(\cdot)$ can be stacked together as $\mathbf{T}_{\mathcal{X}}(\cdot)$ (see Section III with $\gamma = d_{\bar{\mathbf{f}}} + d_{\mathbf{g}}$). The cost term $\mathbf{c}_1(\cdot)$ is the *primary objective* that is minimized first and it enacts as the barrier function to replace the inequalities (5) and (4) with equalities (6b) and (6c). Here, it is also used as an innate metric of motion (*e.g.* velocity minimization) that is often used in motion generation [34], *e.g.* a human posture is minimally changed between a standing keyframe and the following keyframe where the bottle is grasped. The cost term $\mathbf{c}_2(\cdot)$ includes *secondary objectives* that are minimized only after the $\mathbf{c}_1(\cdot)$ and is an expectation over all possible sequences (see eq. 9.7 in [31]). Its metric functions depend on comfort, ergonomics, *etc.*, which is implicitly influenced by the robot interventions. By solving (6) we can obtain a sequence (linear or tree) of N configurations $\mathbf{x}_{1:N} \in \mathbb{R}^{N \times d_{\mathbf{x}}}$. This folds within; predictions of multiple futures of the human and objects states (see (1)), and simultaneously a single sequence of optimal robot interventions for all possible futures (see (3)).

Solver: To solve (6), we treat it as a non-linear least-squares problem⁵ with hierarchy and we construct its feasible region using a barrier function based on displacement intervals [27].

⁵Non-linear least-squares is a commonly used unconstrained optimization formulation of IK and estimation (*e.g.* SLAM [35]) problems in robotics.

By linearizing the non-linear constraint functions using a Taylor expansion, we can obtain the Jacobian⁶ of the constraints \mathbf{J} , and we impose the strict hierarchy by projecting secondary objectives in the null space of the primary ones [36]. Following this strategy the problem in (6) can be efficiently solved with a series of Levenberg–Marquardt updates, based on

$$\Delta \mathbf{x}_{1:N} = \underbrace{\mathbf{J}^\# (\Delta \mathbf{T}_{\mathcal{X}} + \nabla \mathbf{c}_1^T \mathbf{J}^\#)}_{\text{prediction term}} + \underbrace{\mathbf{A} \mathbf{P} \nabla \mathbf{c}_2^T}_{\text{assistance term}}, \quad (7)$$

where

$$\mathbf{J}^\# = (\mathbf{J}^T \mathbf{W}_{\mathbf{x}_{1:N}} \mathbf{J} + \lambda \mathbf{I})^{-1} \mathbf{J}^T \mathbf{W}_{\mathbf{x}_{1:N}}, \quad (8)$$

is the pseudoinverse Jacobian of all constraints of all keyframes, $\mathbf{W}_{\mathbf{x}_{1:N}} \in \mathbb{R}^{d_{\mathbf{x}} \times d_{\mathbf{x}}}$ is a weight matrix, $\mathbf{P} = (\mathbf{I} - \mathbf{J}^\# \mathbf{J}) \mathbf{W}_{\mathbf{x}_{1:N}}^{-1}$ is a projector on the null space of \mathbf{J} , and \mathbf{A} is a selection matrix that filters per dimension the interventions resulting from $\mathbf{c}_2(\cdot)$ (secondary objective in (6)). $\nabla \mathbf{c}_1(\mathbf{x}_{1:N}) \in \mathbb{R}^\gamma$ acts as the barrier function and is the gradient of $\mathbf{c}_1(\cdot)$ (primary objective in (6)) with respect to $\mathbf{x}_{1:N}$. Scalar λ is a damping parameter that allows to switch between Gauss-Newton (small λ) and gradient descent updates (large λ).

In this way, we can guarantee that the prediction of all possible futures is attained first, which corresponds to the first focal point in Section III. The assistance term (second focal point in Section III) lies in the null space of the (multiple) predictions, hence, the interventions are optimal for all possible futures concurrently (expectation over futures), which satisfy the RO problem stated in (3).


V. EVALUATIONS AND EXPERIMENTS

In this section, we first study the importance of the proposed constraint formulation (see Section IV-A) in terms of interventions by performing an ablation study in two scenarios. These scenarios also demonstrate the ability of the method to consider both linear and tree sequences. Further, we show the our method’s capability to personalize the predicted human action based on various factors such as impairments, injuries, strength, *etc.* Finally we realize the assistive capabilities of the method in a table-top scenario with a bimanual robot that provides personalized support to improve the ergonomic posture of a human while preparing a beverage. See the attached video for simulations and experiments of the above described capabilities.

Implementation setup: Our method generate keyframes sequences for a beverage pouring task, omitting continuous trajectories. All evaluations are conducted on a 64-bit Intel Quad-Core i5 3.80 GHz computer with 16GB RAM. The concepts are implemented in the `Rcs C++` library⁷.

A. Comparative studies in simulation

Here, we compare three methods against ours to investigate how different formulations of constraints $\mathbf{T}_{\mathcal{X}}(\cdot) \leq 0$ influence optimality when using the solver in (7):



(a) Linear sequence in simulation. From left to right, keyframes; (2) pick a bottle, (3) pour into glass, (4) place bottle on table, and (5) pick glass, while using the right hand only, as the left holds the cane for support. Intervention: relocate glass before keyframe (3).

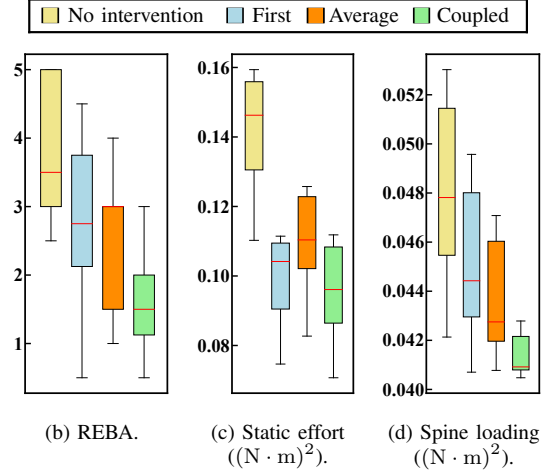


Fig. 4. Average metrics on keyframes 3-5 (post-intervention) for the four methods on the task shown in (a).

- (i) *No intervention*: Only prediction of the human’s future actions is realized without any intervention. Essentially, we solve the problem defined in (1).
- *Heuristic-based*: These method optimizes the configuration of each keyframe separately without considering transitions from one keyframe to another. Then, for all free dimensions that are common between linked keyframes, strategy (ii) *First* applies the configuration of the first keyframe to all consecutive ones. Strategy (iii) *Average* computes the average configuration of the keyframes and applies them to all linked keyframes.
- (iv) *Coupled*: Our approach that solves (3) and (6) to concurrently predict the human’s future actions and find the optimal intervention across linked keyframes.


For all methods we use the following optimization criteria (across keyframes) that are:

- *Human posture*: A sum-of-square (SOS) minimization of all human joints to a preferred configuration (posture). This term penalizes joint limit violations and the preferred posture has REBA score equal to 1.
- *Static effort*: A SOS minimization of all human static joint torques for a given load at the right hand of the human. This incentivizes to avoid stretching the arm.
- *Spine loading*: A SOS minimization of the spine joint torques with respect to a given load at a head of the human. This incentivizes to avoid bending.

The scenarios are; first, a linear sequential manipulation task shown in Fig. 4(a), where a human uses a cane for support. The glass location in keyframes 3 (pouring) and 5 (picking) are the same, hence linked with *constancy constraint*. Yet, due to different orientation and height of the pouring and the picking actions, the optimal human postures and glass poses in

⁶The *constancy constraints* $\bar{\mathbf{f}}(\cdot)$ make the Jacobian non-banded but sparse.

⁷<https://github.com/HRI-EU/Rcs>



(a) Tree sequence in simulation. Pick the mug with right or left hand using power-grasp and hand it over to the free hand to pinch-grasp the handle. Intervention: relocate glass to a common pose for both keyframes 2 and 4.

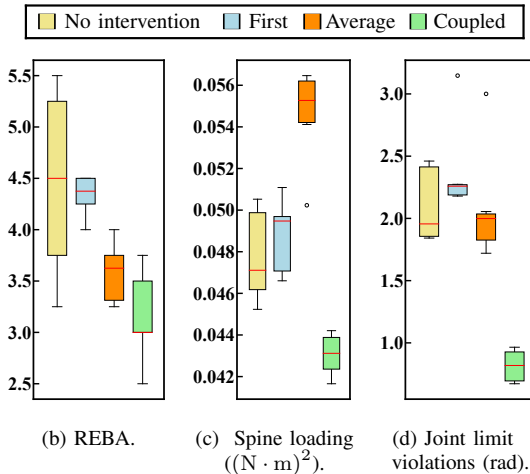


Fig. 5. Average metrics on keyframes 2-5 (post-intervention) for the four methods on the task shown in (a).

keyframes 3 and 5 differ. The robot’s intervention takes place between keyframe 2 and 3. The location of the glass (green) is modified so that the above criteria are optimized.

Second, a tree sequential manipulation task is shown in Fig. 5(a), where a human is on a wheelchair and either power-grasps the mug with the right hand to handed it over to the left for a pinch-grasp from the handle, or power-grasps the mug with the left to pinch-grasp it from the handle with the right. The task is motivated from scenarios where; (i) the grasping actions is ambiguous (power-grasp with left or right hand) and (ii) the power-grasp orientation of the first determines the feasibility of the pinch-grasp (second action). The robot’s intervention takes place between keyframe 1 and both 2 and 3. The location of the mug (gray) is modified so that the objectives of both futures (right and left hand branches) are optimized, *i.e.* both branches can power-grasp the mug with orientations that allow pinch-grasp from the handle in the next action.

The metrics uses to evaluate the methods are averaged across keyframes, (i) REBA score where we subtract 1 as it is the minimum, (ii) *Static effort*, (iii) *Spine loading*, and (iv) joint limit violations. We sample 10 different poses (position and orientation) of the human (with the cane) and of the wheelchair. We report mean and standard deviations with the bar plots shown in Fig. 4 and Fig. 5 for each scenario.

As it can be observed, ours (*Coupled*) outperforms the *no intervention* method, which demonstrates that the interventions of our method improve the human state. The *Coupled*




Fig. 6. Divergent tree sequence, where the right and the left hand branched are independent. Intervention: relocate glass for the right and the left branch separately. See keyframes 2 and 4.

method also performs preferably to the two heuristic methods, which illustrates that having the constraint structure, described in Section IV-A, forms a null-space that can be exploited to reach better optima by our solver (see Section IV-B). The optimality of two heuristic methods depends on the scenario and the metric, which motivates the development of a consistent method, like the *Coupled* (ours).

B. Simulation based case studies

1) Utility of constancy constraints within a time-slice:

Fig. 6 shows a tree sequence with two possible futures along with the graph connectivity of the linked keyframes. This indicates that the two independent branches of the graph (right and left hand) have different optimal poses of the mug. If the wrong future is assumed, then the intervention is not optimal. In contrast, the coupled method (see Fig. 5(a)) finds a common optimal pose for the mug for both futures (right or left hand branches). This is enabled by the *constancy constraint* (see the diamond) that links keyframes 2 and 4.


2) Personalized predictions:

Fig. 7(a) shows a nominal prediction where the robot’s intervention takes place between keyframe 2 and 3. Fig. 7(b) shows the same sequence, however with a different human impairment; an immobilized (broken) right elbow. This impairment leads to a strong bending motion of the upper body to reach the bottle and the glass to compensate for the immobilized elbow. Fig. 7(c) is a result of an analytically modelled impairment (*e.g.* weak left leg) that makes the human lean strongly onto the cane.


In both cases, the impairment consideration leads to a different support intervention which positions the glass in keyframe 3 at a different location to the one in Fig. 7(a). These cases demonstrate the capability of our method to perform personalised predictions and interventions.

C. Robot experiments


Fig. 8 illustrates the system operating interactively in a dynamically sensed environment. The location of the table, bottle, and glass are detected using fiducial markers. The posture of the human is determined by a 3d posture tracking system (Azure Kinect DK) with 5Hz. This enables us to compute the REBA score of the sensed human. The sensed human posture is mapped onto keyframe 1 and it is assumed to be the preferred pose. The predicted human postures (rest keyframes) are updated on-the-fly by solving the TCC problem. One iteration of the optimization takes about 50–70 msec, leading



(a) Nominal prediction with a subject requiring a cane with the left hand for stabilization.



(b) Impaired subject same as (a) and with immobilized (broken) right elbow.



(c) Impaired subject same as (a) and requiring a large stabilizing force on the cane with the left hand.

Fig. 7. Simulation of predicted postures based on human impairments.

to convergence in 1-2 seconds. In Fig. 8(a), the human is in an upright posture, while in Fig. 8(b) the human posture reflects a back injury (leaning forward). As in the simulations above, the predicted keyframes show personalised predictions and the effect of the robot's intervention; to relocate the glass between keyframe 2 and 3.

Further, we validate the utility of the proposed approach to indeed improve the posture of the human. As shown in Fig. 9, based on the sensor-based predictions and intervention, a robot assists a human during the execution of a pour-bottle-into-glass task by pushing the glass to the preferable pose. In contrast to the assisted case, in Fig. 9(a) the human needs to bend forward extensively to complete the task. Further robot support intervention cases can be observed in the video.

VI. SUMMARY AND DISCUSSION


We have proposed a novel concept to support impaired users in daily physical object manipulation tasks with a robot. Our concept is based on a prediction model that uniquely casts the user's sequential behavior as well as a robot support intervention into a hierarchical multi-objective optimization problem. This allows to compute an optimal robot support intervention that rearranges objects in the scenario so that the impaired human is optimally supported. We have shown in simulation studies that our concept is effective, both in terms of the human's comfort metric and in computational efficiency to enable on-the-fly updates of the predictions and

the interventions. Real-world experiments highlight the ability to provide support in dynamically changing situations.


Future work will focus on relaxing the assumptions related to the sequence of the task, on data-driven models of human impairments and on utilization of Augmented Reality to enhance the interpretability of the assistive robot [37].

REFERENCES

- [1] World Health Organization, "World report on disability 2011," 2011.
- [2] B. Sendhoff and H. Wersing, "Cooperative intelligence—a humane perspective," in *IEEE International Conference on Human-Machine Systems (ICHMS)*, 2020, pp. 1–6.
- [3] A. Ajoudani, A. M. Zanchettin, S. Ivaldi, A. Albu-Schäffer, K. Kosuge, and O. Khatib, "Progress and prospects of the human–robot collaboration," *Autonomous Robots*, pp. 43/5/957–975, 2018.
- [4] T. Yan, M. Cempini, C. M. Oddo, and N. Vitiello, "Review of assistive strategies in powered lower-limb orthoses and exoskeletons," *Robotics and Autonomous Systems*, pp. 64/120–136, 2015.
- [5] W. Kim, J. Lee, L. Peternel, N. Tsagarakis, and A. Ajoudani, "Anticipatory robot assistance for the prevention of human static joint overloading in human–robot collaboration," *IEEE robotics and automation letters*, pp. 3/1/68–75, 2017.
- [6] J. Mainprice and D. Berenson, "Human-robot collaborative manipulation planning using early prediction of human motion," in *International Conference on Intelligent Robots and Systems (IROS)*, Nov 2013, pp. 299–306.
- [7] D. F. N. Gordon, C. McCreavy, A. Christou, and S. Vijayakumar, "Human-in-the-Loop Optimization of Exoskeleton Assistance Via Online Simulation of Metabolic Cost," *IEEE Transactions on Robotics*, pp. 38/3/1410–1429, 2022.
- [8] B. Fang, F. Sun, H. Liu, C. Liu, and D. Guo, *Learning from Wearable-Based Teleoperation Demonstration*. Singapore: Springer, 2020.
- [9] H. S. Koppula, A. Jain, and A. Saxena, "Anticipatory planning for human-robot teams," in *Experimental Robotics*. Springer, 2016, pp. 453–470.
- [10] P. Trautman and A. Krause, "Unfreezing the robot: Navigation in dense, interacting crowds," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2010, pp. 797–803.
- [11] Z. Erickson, H. M. Clever, G. Turk, C. K. Liu, and C. C. Kemp, "Deep haptic model predictive control for robot-assisted dressing," in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 4437–4444.
- [12] S. Li, T. Stouraitis, M. Gienger, S. Vijayakumar, and J. A. Shah, "Set-based state estimation with probabilistic consistency guarantee under epistemic uncertainty," *IEEE Robotics and Automation Letters*, pp. 7/3/5958–5965, 2022.
- [13] B. Busch, M. Toussaint, and M. Lopes, "Planning ergonomic sequences of actions in human-robot interaction," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 1916–1923.
- [14] L. van der Spaa, M. Gienger, T. Bates, and J. Kober, "Predicting and optimizing ergonomics in physical human-robot cooperation tasks," in *International Conference on Robotics and Automation (ICRA)*, 2020, pp. 1799–1805.
- [15] M. C. Gombolay, R. A. Gutierrez, S. G. Clarke, G. F. Sturla, and J. A. Shah, "Decision-making authority, team efficiency and human worker satisfaction in mixed human–robot teams," *Autonomous Robots*, pp. 39/3/293–312, 2015.
- [16] L. Y. Chang, G. J. Zeglin, and N. S. Pollard, "Preparatory object rotation as a human-inspired grasping strategy," in *8th IEEE-RAS International Conference on Humanoid Robots*, 2008, pp. 527–534.
- [17] C. R. Garrett, R. Chitnis, R. Holladay, B. Kim, T. Silver, L. P. Kaelbling, and T. Lozano-Pérez, "Integrated task and motion planning," *Annual review of control, robotics, and autonomous systems*, pp. 4/265–293, 2021.
- [18] M. Toussaint, J. Harris, J.-S. Ha, D. Driess, and W. Hönl, "Sequence-of-constraints mpc: Reactive timing-optimal control of sequential manipulation," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 13 753–13 760.
- [19] T. Stouraitis, I. Chatzinikolaïdis, M. Gienger, and S. Vijayakumar, "Online hybrid motion planning for dyadic collaborative manipulation via bilevel optimization," *IEEE Transactions on Robotics*, vol. 36, no. 5, pp. 1452–1471, 2020.
- [20] A. Rudenko, L. Palmieri, M. Herman, K. M. Kitani, D. M. Gavrila, and K. O. Arras, "Human motion trajectory prediction: A survey," *The International Journal of Robotics Research*, pp. 39/8/895–935, 2020.



(a) Upright posture with extended arm.




(b) Back injury.

Fig. 8. Interactive system that utilizes 5 Hz sensor updates to update prediction and interventions on-the-fly.



(a) Unassisted. REBA score during pouring and picking the glass varies between 4 and 5.



(b) Assisted. REBA score during pouring and picking the glass varies between 2 and 3.

Fig. 9. Keyframes of the robot-assisted pour-bottle-into-glass task. By comparing keyframes 3 and 4 for the two cases, we can observe the effect that the robot interventions have in the human posture.

- [21] L. Peternel, W. Kim, J. Babič, and A. Ajoudani, "Towards ergonomic control of human-robot co-manipulation and handover," in *IEEE-RAS 17th International Conference on Humanoid Robotics (Humanoids)*, 2017, pp. 55–60.
- [22] J. Campbell, S. Stepputtis, and H. B. Amor, "Probabilistic multimodal modeling for human-robot interaction tasks," in *Robotics: Science and Systems (RSS)*, June 2019.
- [23] A. Jain, D. Chen, D. Bansal, S. Scheele, M. Kishore, H. Sapra, D. Kent, H. Ravichandar, and S. Chernova, "Anticipatory human-robot collaboration via multi-objective trajectory optimization," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 11 052–11 057.
- [24] P. Kratzer, M. Toussaint, and J. Mainprice, "Prediction of human full-body movements with motion optimization and recurrent neural networks," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 1792–1798.
- [25] D. Berenson, J. Chestnutt, S. S. Srinivasa, J. J. Kuffner, and S. Kagami, "Pose-constrained whole-body planning using task space region chains," in *9th IEEE-RAS International Conference on Humanoid Robots*, 2009, pp. 181–187.
- [26] K. Bouyarmane and A. Kheddar, "Static multi-contact inverse problem for multiple humanoid robots and manipulated objects," in *10th IEEE-RAS International Conference on Humanoid Robots*, 2010, pp. 8–13.
- [27] M. Gienger, H. Janben, and C. Goerick, "Exploiting task intervals for whole body robot control," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2006, pp. 2484–2490.
- [28] O. Kanoun, J.-P. Laumond, and E. Yoshida, "Planning foot placements for a humanoid robot: A problem of inverse kinematics," *The International Journal of Robotics Research*, pp. 30/4/476–485, 2011.
- [29] S.-Y. Chung and O. Khatib, "Contact-consistent elastic strips for multi-contact locomotion planning of humanoid robots," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 6289–6294.
- [30] M. Toussaint, J.-S. Ha, and O. S. Oguz, "Co-optimizing robot, environment, and tool design via joint manipulation planning," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 6600–6606.
- [31] M. J. Kochenderfer, T. A. Wheeler, and K. H. Wray, *Algorithms for decision making*. MIT press, 2022.
- [32] K. Hauser and J.-C. Latombe, "Multi-modal motion planning in non-expansive spaces," *The International Journal of Robotics Research*, pp. 29/7/897–915, 2010.
- [33] A. Escande, N. Mansard, and P.-B. Wieber, "Hierarchical quadratic programming: fast online humanoid-robot motion generation," *The International Journal of Robotics Research*, pp. 33/7/1006–1028, 2014.
- [34] M. Zucker, N. Ratliff, A. D. Dragan, M. Pivtoraiko, M. Klingensmith, C. M. Dellin, J. A. Bagnell, and S. S. Srinivasa, "Chomp: Covariant hamiltonian optimization for motion planning," *The International Journal of Robotics Research*, pp. 32/9–10/1164–1193, 2013.
- [35] F. Dellaert and M. Kaess, *Factor Graphs for Robot Perception*. Foundations and Trends in Robotics, Vol. 6, 2017.
- [36] S. B. Slotine and B. Siciliano, "A general framework for managing multiple tasks in highly redundant robotic systems," in *5th International Conference on Advanced Robotics*, vol. 2, 1991, pp. 1211–1216.
- [37] C. Wang, A. Belardinelli, S. Hasler, T. Stouraitis, D. Tanneberg, and M. Gienger, "Explainable human-robot training and cooperation with augmented reality," in *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*, 2023, pp. 1–5.