



HAL
open science

Fatigue Mitigation in Human-Robot Repetitive Co-Manipulation: A Planning-Friendly Discretization of The Continuous Action Space

Aya Yaacoub, Vincent Thomas, Francis Colas, Pauline Maurice

► **To cite this version:**

Aya Yaacoub, Vincent Thomas, Francis Colas, Pauline Maurice. Fatigue Mitigation in Human-Robot Repetitive Co-Manipulation: A Planning-Friendly Discretization of The Continuous Action Space. IEEE ICRA 2024 Workshop on Innovations and Applications of Human Modelling in Physical Human-Robot Interaction, May 2024, Yokohama, Japan. . hal-04590156

HAL Id: hal-04590156

<https://hal.science/hal-04590156>

Submitted on 28 May 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Fatigue Mitigation in Human-Robot Repetitive Co-Manipulation: A Planning-Friendly Discretization of The Continuous Action Space

Aya Yaacoub¹, Vincent Thomas¹, Francis Colas¹, Pauline Maurice¹

I. INTRODUCTION

Industrial workers performing repetitive manipulation tasks are subject to work-related musculoskeletal disorders (WMSDs) which could manifest as localized joint fatigue [1] resulting from repeatedly moving the same joints. With the promise of collaborative robots (cobots) directly interacting with humans, a research direction which investigates using cobots to reduce WMSD risk has emerged. Focusing on repetitive co-manipulation tasks, some studies suggest using cobots to induce motion variability[2], recommended to reduce WMSD risk [3], in order to minimize localized fatigue [4], [5].

We consider tasks comprising many cycles where, each cycle, the cobot selects a Cartesian pose (robot action) at which it presents one object to the human. The human, then, reacts with a whole-body motion (human postural reaction) to manipulate the object. In [4], authors proposed reactively changing the object presentation pose between task cycles to the pose minimizing estimated induced muscular fatigue at the end of each cycle. However constraining the human hand pose through the object placement is not enough to determine with certainty the human postural reaction that the human is going to choose. This is due to human kinematic redundancy which allows for infinite possible postural reactions from which the human selects according to various factors including its expertise and fatigue level. But, fatigue itself is not-measurable and can only be estimated from motion measurements, which induces partial observability. Finally, the long-term effects of the cobot actions on fatigue need to be considered through planning due to the task repetitive nature. Thus, in [5], we proposed a framework representing the problem of cobot action selection for long-term fatigue mitigation as a Partially Observable Markov Decision Process (POMDP). Partially Observable Monte-Carlo Planning was used and the framework was validated on a simulated test-case. One limitation of the test-case was restricting the cobot action space to a small number of hand-picked poses instead of considering the continuous robot workspace.

For continuous-action-space POMDPs, planning is possible either through dedicated planning algorithms which discretize the action space online, or discretizing the action space offline. Among the most famous continuous-action-space POMDPs planning algorithms are Progressive Widen-

ing (PW) algorithms [6], [7]. PW approaches generate (online) a tailored set of actions for each possible system state. However refining the tailored set is required at each planning step, and can be computationally expensive, especially in a 6D Cartesian pose space. This is hardly compatible with the need for online decisions constrained by the short cycle duration in a highly repetitive task. Thus we opt for offline action space discretization.

In this work, our objective is to automatically select a set of actions of a reasonable size (to respect online planning time constraint), relevant to long-term fatigue mitigation, from the continuous action space. We seek an improved long-term fatigue mitigation compared to the simplified two-actions set tested in [5], and a better performance/time compromise compared to a large randomly sampled set representing the continuous action space. Finally, the action-set selection process should be automatic and yielding an action set robust to uncertainties surrounding the system.

II. METHODS

A. Fatigue Evolution Estimation

We use the same POMDP problem formulation presented in [5]. The POMDP considers joint fatigue as a loss in torque generation capacity of a joint computed according to the joint fatigue model proposed in [8]. The evolution of torque generation capacity depends on the joint torques exerted during the human postural reaction to the cobot action. Here, we consider a human that reacts to object presentation with a whole-body motion minimizing the fatigue cost in one task cycle. To estimate joint torques for a given human whole-body motion, a physics-based Digital Human Model (DHM) simulator [9] was used. The fatigue model is uni-directional, in that agonist and antagonist fatigue values are considered separately from each joint, and calculated based on positive or negative torque values respectively. We now use "joint" to refer to a joint direction (i.e. 2 for each real joint).

B. Action Set Selection

We start by sampling a large set of actions \mathcal{A}_n (here, $n = 1500$ poses) to serve as a representative of the whole action space. In what follows, we present the automated action-space discretization method that we developed and tested.

1) *Greedy Used Actions*: Greedy action selection returns, each cycle, the action minimizing the expected human fatigue cost ($\sum_{n_{joints}} \frac{(f^j)^2}{n_{joints}}$) at the end of the cycle. Compared to other action-selection algorithms greedy is simple and fast which facilitates running a large number of tests at different initial fatigues, thereby allowing for a wide exploration of

This work was partly supported by the French ANR under Grant No. ANR-20-CE33-0004 (project ROOIBOS).

¹ The authors are with Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France. `firstname.name@loria.fr`

the fatigue space. We consider a set of initial fatigue states \mathbf{f}_0^i with $i \in \{1, 2, \dots, 78\}$, 78 being the number of joints. \mathbf{f}_0^i represents a fatigue state where joint i is half-fatigued and all other joints fully relaxed. For each \mathbf{f}_0^i , we run a greedy action-selection algorithm in a simulated 100 cycle repetitive task. The sets \mathcal{G}^i of used actions from all runs (with different initial \mathbf{f}_0^i) are kept, and $\mathcal{G} = \bigcup \mathcal{G}^i$ is considered as a candidate action set.

2) *Augmented Fatigue Space Coverage*: With \mathcal{G} alone, we are not sure to cover the entire fatigue space. To improve the performance of the action set against unexplored fatigue zones, we add to it a supplementary set \mathcal{N} . We call \mathcal{N} the set of non-dominated actions and define it as the set containing actions having the best performance for at least one joint j and one fatigue value f_0 of that joint:

$$\mathcal{N} = \bigcup_{j, f_0} \text{sample}(\{a \in \mathcal{A}_n | \forall a' \in \mathcal{A}_n, f^j(f_0, a) \leq f^j(f_0, a')\}).$$

III. RESULTS

We considered set \mathcal{A}_n (1500 actions) and the two reduced sets \mathcal{G} (22 actions) and $\mathcal{G} \cup \mathcal{N}$ (40 actions). We benchmarked the performance/time compromise of each set using different Finite-Horizon decision-making algorithms $FH(h)$ which return the action that, when applied, leads to the lowest expected discounted fatigue-cost sum over horizon h . $FH(h)$ with $h \in \{1, 2, 3\}$ were considered, $FH(1)$ actually being equivalent to greedy.

For each action set and $FH(h)$ combination, we simulated 100-cycles tasks starting from one of two initial fatigues: completely non-fatigued (b_0^c), and half fatigued in the low-back joints and non-fatigued in others (b_0^b). Figure 1 shows the accumulated fatigue reward (reward = -cost) and the average decision time in different tested scenarios.

Sets Comparison: Results show that sets \mathcal{G} and $\mathcal{G} \cup \mathcal{N}$ have equivalent performance with \mathcal{G} being faster due to the smaller number of actions. For the tested initial fatigues, no advantage was observed for including the non-dominated set. For \mathcal{A}_n , only $FH(1)$ was used since we estimated the decision time per cycle for $FH(2)$ to be around 2000s which is too long for online decision-making. \mathcal{A}_n has an equivalent performance to \mathcal{G} and $\mathcal{G} \cup \mathcal{N}$ when $FH(1)$ is used starting from b_0^c and a lower performance starting from b_0^b . In both cases, \mathcal{A}_n is slower. The lower performance of the full set \mathcal{A}_n starting from b_0^b could be due to selecting actions which seem promising in the short term but lead to degraded performance on the overall 100 cycles. **Planning Horizon:** Results also demonstrate an improved performance for \mathcal{G} and $\mathcal{G} \cup \mathcal{N}$ when used with planning for longer horizons.

In conclusion, we have presented an approach to improve fatigue mitigation in repetitive co-manipulation tasks. The proposed approach automatically reduces a naively sampled large action set into smaller sets which allow for improved fatigue mitigation when combined with planning while respecting online-planning time constraints. Further verification is needed to test the robustness of our approach to uncertainties around the human reaction and initial fatigue, and to unexplored fatigue states.

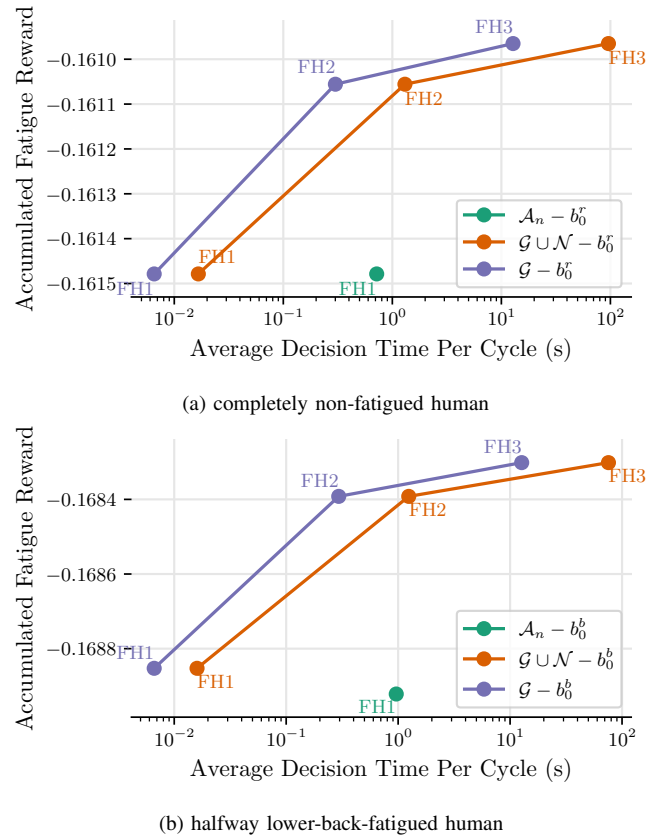


Fig. 1: Accumulated fatigue-reward vs average decision-time plots for different tested set-algorithm combinations.

Keywords: WMSD, Fatigue Mitigation, Cobots, Long-term, Continuous-Action-Space


REFERENCES

- [1] European Agency for Safety and Health at Work., “E-fact 9 - Work-related musculoskeletal disorders (MSDs): an introduction | Safety and health at work EU-OSHA,” 2007.
- [2] M. Zolotas, R. Luo, S. Bazzi, D. Saha, K. Mabulu, K. Kloeckl, and T. Padir, “Imposing motion variability for ergonomic human-robot collaboration,” *IJSE Transactions on Occupational Ergonomics and Human Factors*, vol. 0, no. 0, pp. 1–12, 2024.
- [3] D. Srinivasan and S. E. Mathiassen, “Motor variability in occupational health and performance,” *Clinical Biomechanics*, vol. 27, no. 10, pp. 979–993, 2012.
- [4] L. Peternel, C. Fang, N. Tsagarakis, and A. Ajoudani, “A selective muscle fatigue management approach to ergonomic human-robot co-manipulation,” *Robotics and Computer-Integrated Manufacturing*, vol. 58, pp. 69–79, 2019.
- [5] A. Yaacoub, V. Thomas, F. Colas, and P. Maurice, “A probabilistic model for robot decision making to mitigate human fatigue in repetitive co-manipulation tasks,” *IEEE Robotics and Automation Letters*, pp. 1–8, 2023.
- [6] G. Chaslot, M. Winands, J. Uiterwijk, H. van den Herik, and B. Bouzy, “Progressive strategies for Monte-Carlo tree search,” *Information Sciences* 2007, vol. 10, pp. 655–661, 2007.
- [7] M. Hoerger, H. Kurniawati, D. Kroese, and N. Ye, “Adaptive Discretization using Voronoi Trees for Continuous-Action POMDPs,” 2022.
- [8] L. Ma, W. Zhang, D. Chablat, F. Bennis, and F. Guillaume, “Multi-objective optimisation method for posture prediction and analysis with consideration of fatigue effect and its application case,” *Computers & Industrial Engineering*, vol. 57, no. 4, pp. 1235–1246, 2009.
- [9] P. Maurice, V. Padois, Y. Measson, and P. Bidaud, “Digital human modeling for collaborative robotics,” in *DHM and Posturography*. Elsevier, 2019, pp. 771–779.

In the context of fatigue mitigation in

Human-Robot Co-manipulation:

We propose an automatic method to extract, from the continuous action space, a small set of robot actions for planning.

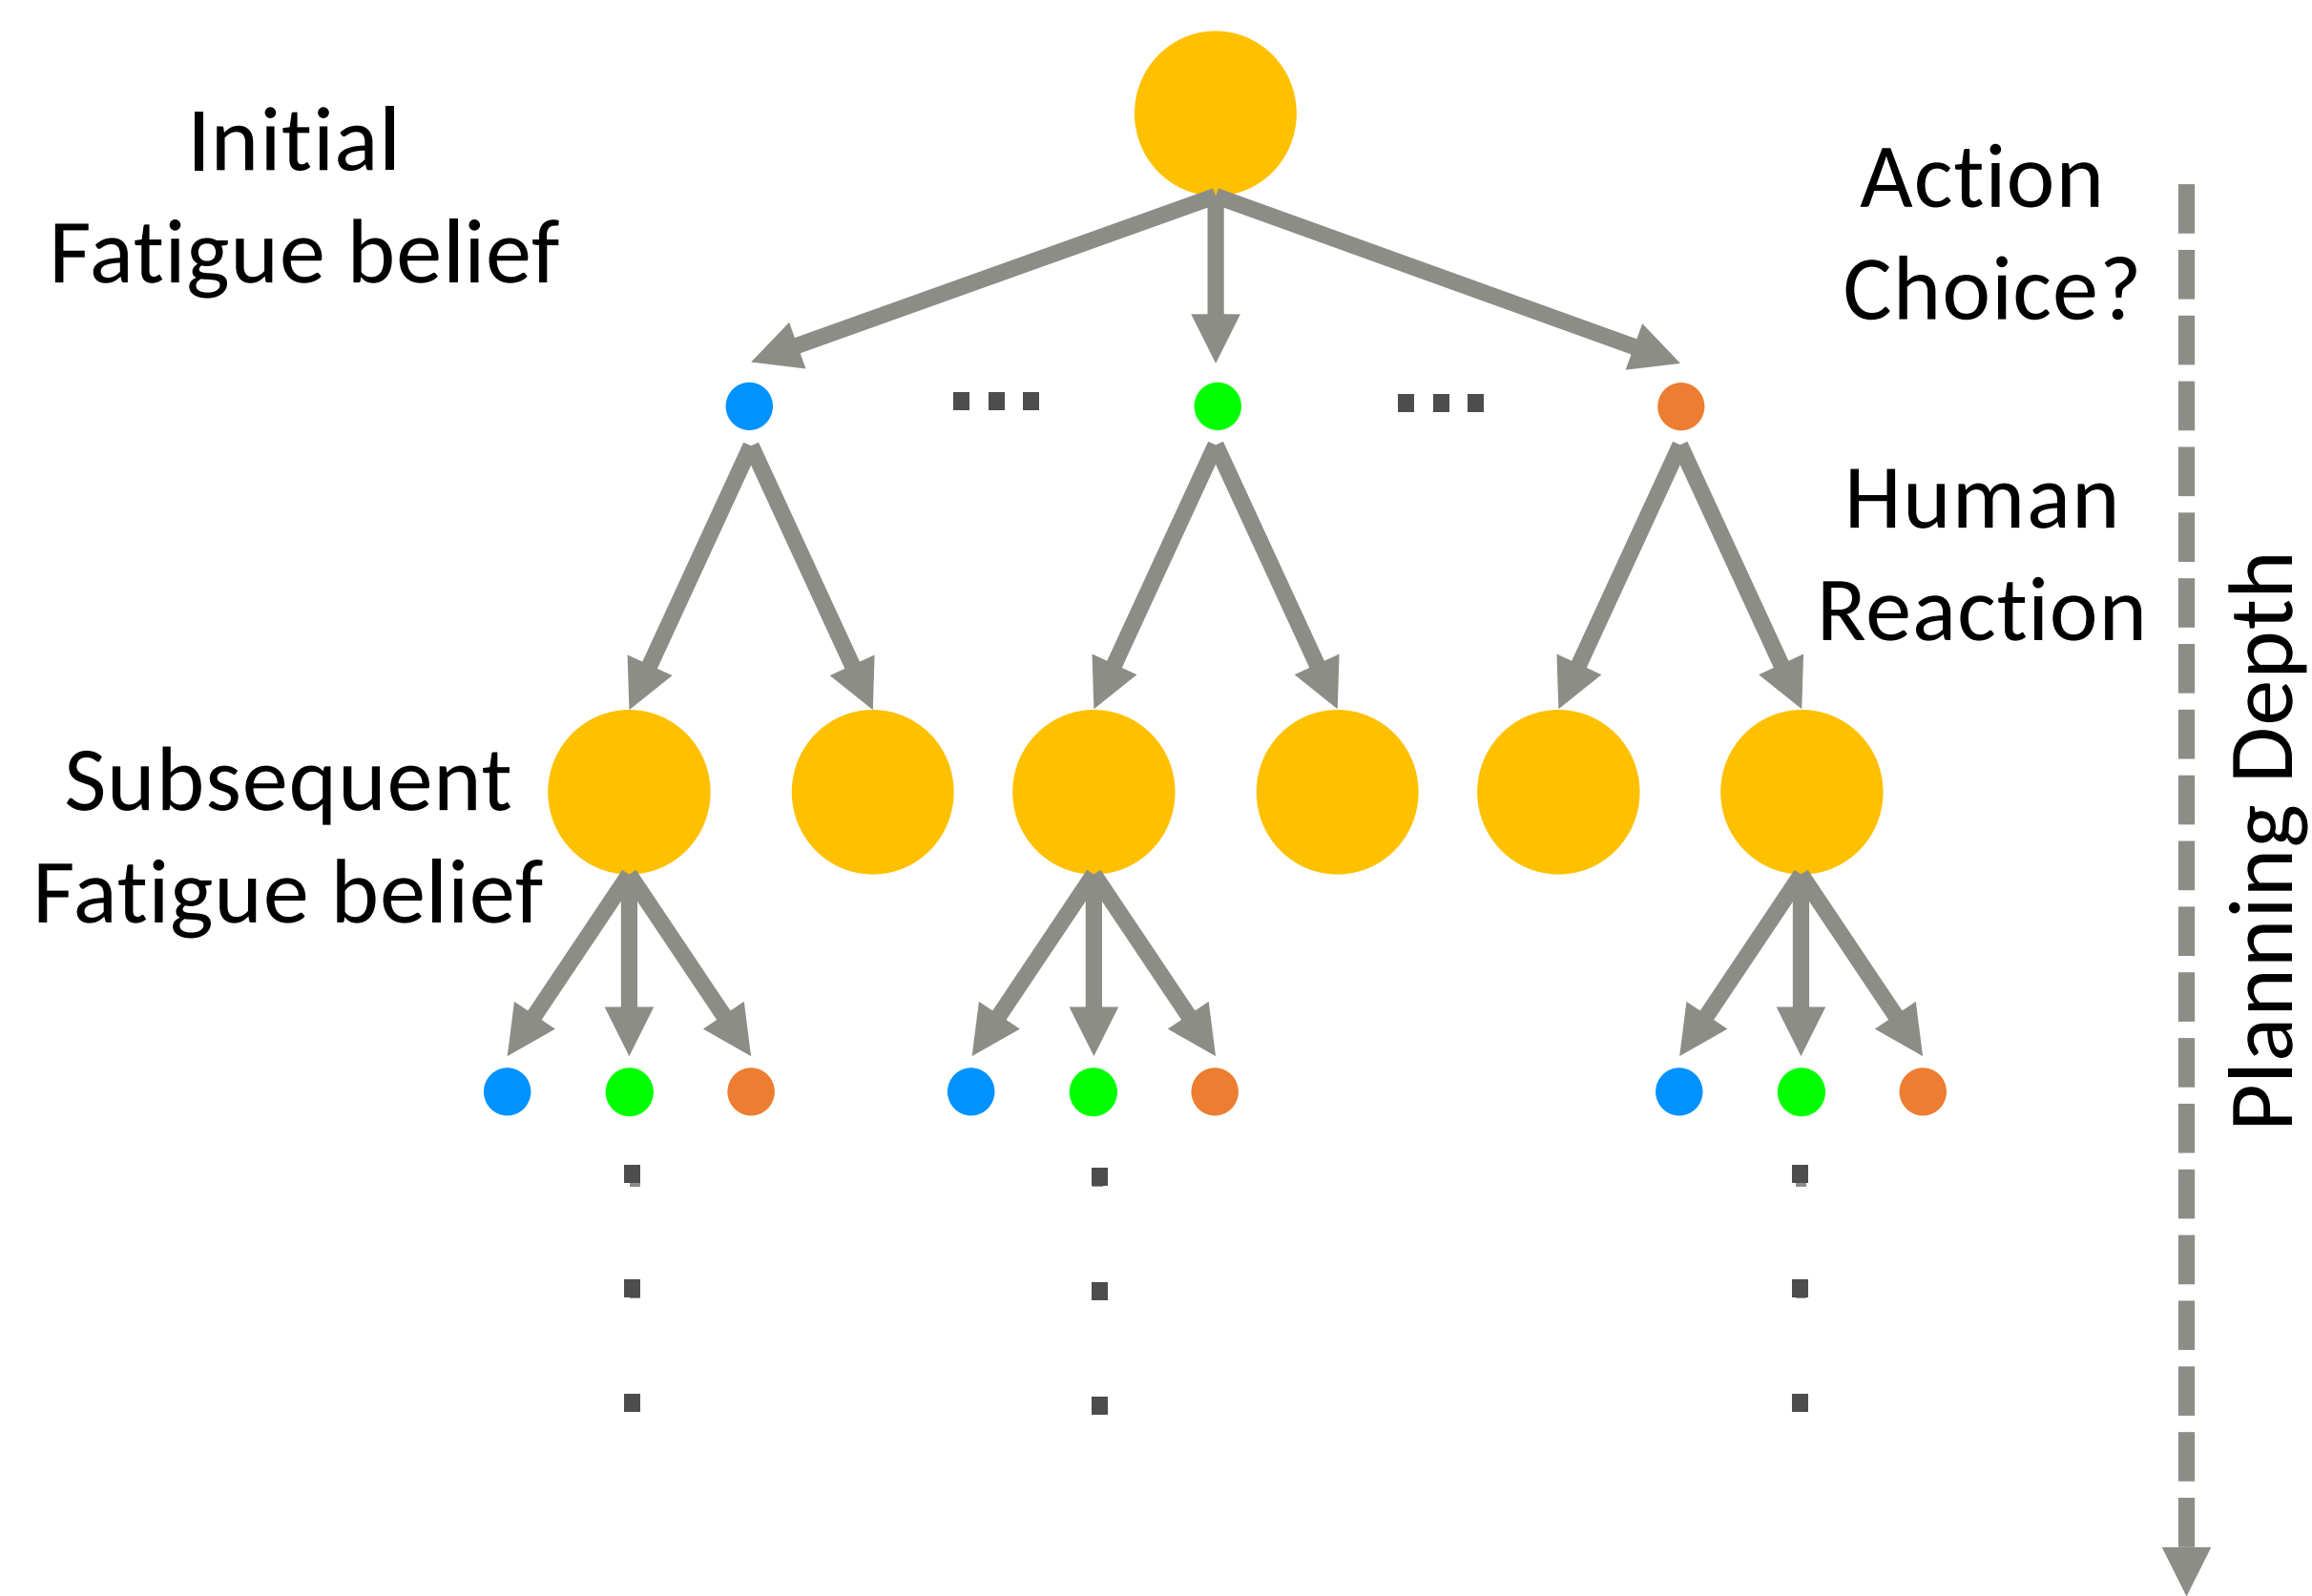


Aya Yaacoub, Vincent Thomas, Francis Colas, Pauline Maurice
Université de Lorraine, CNRS, Inria, LORIA, Nancy, France

- Industrial workers performing repetitive manipulation tasks are at risk of developing work-related musculoskeletal disorders indicated at early stages by localized joint fatigue.
- In a scenario where a cobot presents the worker with an object to co-manipulate, previous work has proposed optimizing the object pose to minimize localized joint fatigue after each object presentation [1].
- In [2], we considered a task with many short-duration cycles in which the cobot presents the human with an object at some pose and the human reacts with a whole-body motion. We validated that considering the consequences of the robot actions over some future horizon (higher than 1), through planning, can improve upon greedy fatigue minimization.

Problem

Planning Cost $\propto (|Actions| \cdot |Reactions|)^{depth}$



- To allow for a bigger horizon (planning depth), it is important to reduce the size of the cobot's action set.

Requirements

- Relevant to the objective of long-term fatigue minimization
- Good performance/time compromise for online planning
- Automatically obtained, adaptable extraction
- Useful in many possible human fatigue-state scenarios

Action-Set Selection

1 Random discretization of a large action set to serve as a representative of the continuous action space: $\mathcal{A}_n, n = 1500$.


2 Refining the action set:

a **Greedy-Used:**

- each task cycle (object presentation):

$$\text{belief state } b \xrightarrow{\text{Greedy}} a^* = \arg \max_{a \in \mathcal{A}_n} \text{fatigueCost}(a, b)$$


- for every joint j , starting from a state where only j is half-fatigued (others not fatigued), run greedy in a simulation :




- keep action set $\mathcal{G}^j : \{ \bullet, \bullet, \bullet, \bullet \}$

The greedily used actions set is : $\mathcal{G} = \bigcup \mathcal{G}^j$

b **Non-Dominated:**



Results



- Sets were tested with Finite Horizon (FH) action-selection returning the action leading to the best expected accumulated sum of rewards in h cycles.
- \mathcal{G} and $\mathcal{G} \cup \mathcal{N}$ had equivalent performances with \mathcal{G} being faster. Both sets improved in performance as a higher planning horizon was considered.
- Further testing is needed to check the adaptability of the sets to uncertainty in the fatigue state and different possible fatigue scenarios

References

- [1] L. Peternel, C. Fang, N. Tsagarakis and A. Ajoudani, "A selective muscle fatigue management approach to ergonomic human-robot co-manipulation", Robot. Comput.-Integr. Manuf, 2019.
[2] A. Yaacoub, V. Thomas, F. Colas and P. Maurice, "A Probabilistic Model for Cobot Decision Making to Mitigate Human Fatigue in Repetitive Co-Manipulation Tasks," in IEEE RA-L, 2023