

DULA: A Differentiable Ergonomics Model for Postural Optimization in Physical HRI

Amir Yazdani*, Roya Sabbagh Novin*, Andrew Merryweather*, Tucker Hermans*[†]

*University of Utah Robotics Center, Salt Lake City, UT, [†]NVIDIA, Seattle, WA

Email: amir.yazdani@utah.edu

Abstract—Ergonomics and human comfort are essential concerns in physical human-robot interaction applications. Defining an accurate and easy-to-use ergonomic assessment model stands as an important step in providing feedback for postural correction to improve operator health and comfort. In order to enable efficient computation, previously proposed automated ergonomic assessment and correction tools make approximations or simplifications to gold-standard assessment tools used by ergonomists in practice. In order to retain assessment quality, while improving computational considerations, we introduce DULA, a differentiable and continuous ergonomics model learned to replicate the popular and scientifically validated RULA assessment. We show that DULA provides assessment comparable to RULA while providing computational benefits. We highlight DULA’s strength in a demonstration of gradient-based postural optimization for a simulated teleoperation task.

I. INTRODUCTION

Autonomous Postural Optimization has received substantial attention in research with the new technologies around humans such as collaborative robots [18], smart personal trainers [4], and VR systems [8]. In these systems, the interacting agent should consider human comfort and ergonomics in its behaviour and motion planning [24]. For example, in a physical human-robot interaction (pHRI) application such as co-manipulation, one of the objectives in motion planning of the collaborative robot must be satisfying ergonomic safety for the human.

Developing a model of human comfort lies at the heart of effective postural optimization. pHRI researchers have proposed several computational models for assessing ergonomics and human comfort in terms of peripersonal space [3], muscle fatigue [17], and joint overloading [13]. In contrast ergonomists have provided simpler models which are easier for human experts to calculate by hand and as such are more common in practice. These models include the NASA TLX [6], RULA [15], REBA [7], strain index [22] and ACGIH TLV [11]. Importantly these models are supported by extensive human subject studies that validate their effectiveness on reducing ergonomic risk factors [9, 12].

Among all risk assessment tools, RULA and REBA depend most on human posture and provide quantitative scores, making them good choices for postural optimization applications. However, the discrete scores and the presence of plateaus in RULA and REBA [1] create challenges when using them in gradient-based postural optimization. Based on our experience, using the risk assessment models directly in gradient-free




Fig. 1: The structure of the DULA neural network.

optimization is time-expensive and the plateaus often prevent progress toward the global optimal solution in postural optimization. Thus, researchers in pHRI often use approximations of ergonomic assessment models in gradient-based postural optimizations; quadratic approximations [20, 1, 2] being the standard approach in the literature. However, these approximations deviate far from the scientifically validated assessments, causing doubt that they can reliably provide the same level of ergonomic benefit.

To overcome these issues, we introduce the *Differentiable Upper Limb Assessment* (DULA), a differentiable and continuous risk assessment model that is learned using a neural network to replicate the popular RULA survey tool (Fig. 1). Instead of discrete scores from 1 to 7, DULA reports the risk score as a continuous real number from 1 to 7. Furthermore, it provides the gradient of the risk with respect to each joint enabling efficient use in optimization. We compare the prediction of risk scores of DULA with RULA and provide an open-source package demonstrating how to use DULA in a gradient-based postural optimization in a simulated teleoperation.

II. RELATED WORK

Although human postural optimization has received significant attention in pHRI, only a few studies [20, 19] investigate postural improvement in teleoperation. Table I summarizes the relevant literature for postural optimization in both areas.

Researchers have examined ergonomics and postural optimization in three different types of physical human-robot interaction scenarios: (1) *Assistive Holding* (e.g. [14, 21, 3]), (2) *object handover* (e.g. [1, 2, 17]), and (3) *Co-Manipulation* (e.g. [23, 18, 17, 13]). In the first two tasks the postural optimization provides optimized posture for the human and joint configurations for the robot while maintaining contact with the object at the interaction interface. For the case of co-

Ergonomics Model		Analytical Models	Learned Models	Risk Assessment Tools	
Approximation Method				Quadratic	No Approximation
Optimization	gradient-based	<ul style="list-style-type: none"> ■ Peternel 2018 [18] ■ Peternel 2017 [17] □ Chen 2018 [3] ◆ Peternel 2020 [19] ■ Kim 2017 [13] 		<ul style="list-style-type: none"> ◇ Rahal 2020 [20] ■ Busch 2017 [1] ■ Busch 2018 [2] 	
	gradient-free		□ Marin 2018 [14]		■ van der Spaa 2020 [23]
Non-optimization					□ Shafti et al. [21]
Legend		<ul style="list-style-type: none"> □ physical HRI: Assistive Holding, ■ physical HRI: Handover, ■ physical HRI: Co-Manipulation ◆ Teleop: Goal-Constrained with Repositioning Postural Correction, ◇ Teleop: Goal-Constrained with Online Postural Correction 			

TABLE I: State of the art of postural optimization in pHRI and teleoperation.

manipulation, the postural optimization outputs a trajectory of optimal postures for the human and an optimal joint-space trajectory for the robot to perform the co-manipulation task.

Ergonomic assessment models provide the primary cost in postural optimization objectives. pHRI researchers have proposed many computational models to assess ergonomics and human comfort of users including peripersonal space [3], muscle fatigue [17], and joint overloading [13]. To make the optimization simpler, Marin et al. suggested the idea of a contextual ergonomics model which is a set of Gaussian process models including joint angles, moments, reaction, load and muscle activation, trained with the musculoskeletal simulation task contexts [14]. Using Gaussian process models enables search in a 2D latent space while their cost function is defined in the high-dimensional musculoskeletal space.

Some literature use approximations of the ergonomics risk assessment tools. Busch et al. proposed a differentiable surrogate of the REBA score by fitting a weighted combination of quadratic functions plus a constant for the task payload [2]. Rahal et al. suggested a quadratic approximation for RULA which is the summation of the quadratic norm of the deviation from the human neutral posture for shoulder, elbow and wrist joints angles [20]. Their approximation conceptually agrees with the qualitative idea behind RULA in which the risk score goes higher when the human deviates more from the neutral posture. Van der Spaa et al. provide the only study of directly using risk assessment tools in derivative-free postural optimization [23]. They add the REBA score to the transition cost function in an A* optimization for task and motion planning of a robot in a co-manipulation task.

III. DIFFERENTIABLE HUMAN ERGONOMICS MODEL

To build a differentiable RULA, we developed a dataset of 7.5 million upper body postures of a human model consisting of 3 joints in the torso and 7 joints in the arm. We additionally define task parameters based on the RULA worksheet—the frequency of the arm and body motions; type and maximum load on the arm and body; neck angle; and whether any legs, feet, or arms are supported. As the human range of motion is pose dependant, to ensure the validity of the postures, we used the learned pose-dependant model of posture validity provided by [10]. We developed a script for automatic RULA assessment based on the posture and tasks parameters and verified it with several ergonomists. We used this script to label the posture dataset. Since postures with labels 1, 2, 6, and 7 are not frequent in the full range of human motion, we




Fig. 2: Confusion matrix (accuracy %) for DULA vs RULA.

balanced the dataset by forcing the data generation scripts to generate enough data points with those labels. We split the dataset into 80% training and 20% testing sets.

Moreover, to learn a continuous and differentiable function for RULA, we designed a fully-connected regression-based neural network. While RULA provides discrete integer scores from 1 to 7, we choose to predict continuous labels. If we had instead performed multi-class classification based on the discrete labels the resulting model would be less useful in optimization as there is no natural choice of smooth objective to minimize or constrain ergonomic cost. This would negate our desire for a computationally useful model. Hence, we perform regression to the multi-class labels.

The structure of the neural network is shown in Fig. 1. It includes 4 hidden layers with ReLU activation function. We found that a network with 124 units for the first three layers and 7 units for the last hidden layer worked best. We train the network using the standard mean squared error loss function for 2000 epochs using a learning rate of 0.001. We used 5-fold cross-validation to find the optimal network parameters. This results in a model with 99.73% accuracy. We round our continuous output to the nearest integer when reporting accuracy. Figure 2 shows the confusion matrix for the learned DULA model. The lowest diagonal element is 99.38% which




Fig. 3: Postural optimization categories in teleoperation.




Fig. 4: Teleoperation simulation environment in ROS. The green skeleton visualizes the suggested optimal posture and the white skeleton shows the simulated human.

shows the high accuracy of the learned model across all ranges.

IV. POSTURAL OPTIMIZATION USING DULA

We use the learned DULA model in a gradient-based postural optimization for a simple teleoperation scenario in which a human interacts with a leader robot to remotely control a follower robot. As the human performs the task, the intelligent teleoperation system estimates the human posture using e.g. marker-based or markerless posture estimation, or directly from the leader robot using the method proposed in [25]. Then, it performs risk assessment to obtain the ergonomics risk score using the standard RULA model. The postural optimization algorithm uses DULA to find the optimal posture that results in the same hand pose at the interaction point between the human and the leader robot, while having the minimum risk of injuries and then provides the user with the online optimal postural correction to move towards while completing the task.

The human operator can correct the posture while performing the task without pausing. This approach is beneficial for goal-constrained teleoperation tasks such as pick-and-place in which the goal position for placing the object is defined, but, the path toward the goal point is not constrained.

We define online postural optimization for teleoperation as:

$$\begin{aligned} \mathbf{q}_t^* &= \underset{\mathbf{q}_t}{\operatorname{argmin}} \operatorname{DULA}(\mathbf{q}_t) \\ \text{s.t. } & \|\mathbf{x}_t - \Phi(\mathbf{q}_t)\|_{\Sigma}^2 < \epsilon \\ & \mathbf{q}_t \in \text{Range of Motion} \end{aligned} \quad (1)$$

where \mathbf{q}_t^* and \mathbf{q}_t are optimal posture and posture at time t , respectively, \mathbf{x}_t is the observed pose of the hand measured by the leader robot, Φ is the forward kinematics of the human, and Σ is the weight vector for position and orientation elements.

It is important to note that the optimal posture \mathbf{q}_t^* from Eq. (1) for each time step is then suggested to the human to

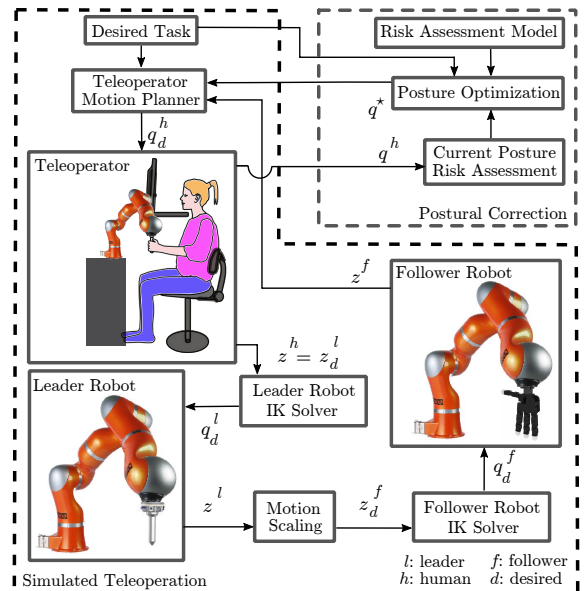


Fig. 5: Overview of teleoperation simulation.

move towards. The human can refuse or accept, and try to apply it as much as possible while completing the task.

In addition to goal-constrained teleoperation, we also defined and formulated the use of DULA in postural optimization for path-constrained teleoperation (e.g. turning a valve where the path to follow is constrained based on the diameter of the valve) and trajectory-constrained teleoperation (e.g. arc welding where the operator should follow a velocity profile in addition to the welding path). Figure 3 summarizes our formulation for different types of teleoperation tasks and their corresponding postural optimization approaches. We focus on online postural correction in this paper; analyzing the other problem formulations is ongoing work.

We used a sequential quadratic programming (SQP) [16] solver from SciPy, bounded on the range of motion, to solve the nonlinear optimization in Eq. (1), and calculated DULA gradients using automatic differentiation in PyTorch.


V. SIMULATED TELEOPERATION ENVIRONMENT

We developed an open-source simulator for postural correction in teleoperation using ROS. It includes a human seated on a stool, and two 7-DOF KUKA LWR-4 robots as leader and follower robots as shown in Fig. 4. We model our simulated human operator to behave like a human in two ways: (1) physically controlling the teleoperation task and (2) accepting or rejecting the recommended postural corrections.


We model this as an optimal motion planning framework with re-planning that finds a human joint trajectory that controls the follower robot for the desired task while moving toward the optimal ergonomic posture:

$$\tau_{t \rightarrow H}^{h*} = \underset{\tau_{t \rightarrow H}^h}{\operatorname{argmin}} \sum_{t=t}^H \|\mathbf{x}_g^f - \mathbf{x}_t^f\|_{\Sigma}^2 + \alpha \|\mathbf{q}_t^h - \mathbf{q}_t^{h*}\|_2^2 \quad (2)$$

where $\tau_{t \rightarrow H}^f$ is the trajectory of human posture from time t to the time of the horizon H , \mathbf{x}_g^f is the goal pose of the follower robot, and $^f \mathbf{x}_t$ is the pose of the follower robot at time t .



A. Gradient-free postural optimization



B. Gradient-based postural optimization

Fig. 6: Comparison of gradient-free and gradient-based postural correction on a teleoperation task. Lower scores are better.

\mathbf{q}_t and \mathbf{q}_t^* are the current posture and optimal posture of the human from the postural optimization at time t , respectively. Here, $0 \leq \alpha \leq 1$ is a scalar number that models the postural correction acceptance and the effort of the human operator towards applying the postural correction. Details of the motion planning for the human and the robots are presented in Fig. 5. We note that this model is likely a simplification of true human teleoperation behavior. We do not advocate for its use over human subject studies. Instead, we propose it as a useful tool for systematically exploring new algorithms in human safety assessment and improvement.

As a comparison to DULA, we directly use the standard RULA in a gradient-free postural optimization using the cross-entropy method [5]. Figure 6(A) shows the postural correction using this approach in a simple teleoperation task. It presents the RULA scores for calculated optimal posture (green), human posture without postural correction (orange), and the human posture after applying the postural correction (blue) according to the human control and acceptance model during the task. The plot shows that the risk is reduced after the simulated human applies the suggested optimal postural correction. Also the task completion time increases without a significant decrease in risk.

Figure 6(B) shows the postural correction for the same task using gradient-based optimization with DULA. We can see that gradient-based approach provides a smoother motion with respect to the RULA risk score and avoids going through postures with risk higher than 4. It also results in shorter task




Fig. 7: comparing the optimal posture and corrected posture between gradient-free optimization using RULA and gradient-based optimization using DULA. Lower scores are better. An iteration of the gradient-free method takes 2.7 minutes, where the gradient-based method takes only tenths of a second.

completion time than the gradient-free approach.

Figure 7 provides more information on comparing gradient-free and gradient-based postural optimization. It shows that the median risk scores of target optimal postures from the gradient-free approach is lower. However, the postures of the simulated human are more comfortable after applying the optimal posture calculated from the gradient-based optimization. We believe it is due to the smoother optimal postures that has been suggested to the simulated human from the gradient-based method. Moreover, the gradient-based approach is much faster. Each iteration of the gradient-free approach using 10000 samples takes 2.7 minutes to solve, while the gradient-based approach takes only tenths of a second.

VI. CONCLUSION AND FUTURE DIRECTIONS

We introduced DULA, a differentiable and continuous ergonomics model to assess human upper body posture. DULA learned to replicate the non-differentiable RULA using a neural network. The proposed model is 99.73% accurate and computationally designed for efficient use in postural optimization for pHRI and other related applications. We also introduced a framework for postural optimization in teleoperation using DULA and presented a demo task. The results reveal postural optimization using DULA lowers the risk score for goal-constrained teleop. The trained DULA model, data, algorithm, and demo are available as an open-source package¹.

There are several directions for future work. As a first step, we plan to follow the same procedure for REBA a whole body assessment tool. Additionally, we intend to conduct a human subject study to evaluate our postural optimization and correction approach. We wish to compare different means of feedback for the posture correction to the human operator including visual, auditory, and haptic feedback.

VII. ACKNOWLEDMENT

This work is supported by DARPA under grant N66001-19-2-4035.

¹<https://sites.google.com/view/differentiable-ergonomics>

REFERENCES

- [1] Baptiste Busch, Guilherme Maeda, Yoan Mollard, Marie Demangeat, and Manuel Lopes. Postural optimization for an ergonomic human-robot interaction. In *Intl. Conf. on Intelligent Robots and Systems*, pages 2778–2785, 2017.
- [2] Baptiste Busch, Marc Toussaint, and Manuel Lopes. Planning ergonomic sequences of actions in human-robot interaction. In *Intl. Conf. on Robotics and Automation*, pages 1916–1923, 2018.
- [3] Lipeng Chen, Luis FC Figueredo, and Mehmet R Dogar. Planning for muscular and peripersonal-space comfort during human-robot forceful collaboration. In *IEEE-RAS Intl. Conf. on Humanoid Robotics*, pages 1–8, 2018.
- [4] Steven Chen and Richard R Yang. Pose trainer: correcting exercise posture using pose estimation. *arXiv preprint arXiv:2006.11718*, 2020.
- [5] Pieter-Tjerk De Boer, Dirk P Kroese, Shie Mannor, and Reuven Y Rubinstein. A tutorial on the cross-entropy method. *Annals of operations research*, 134(1):19–67, 2005.
- [6] Sandra G Hart and Lowell E Staveland. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. In *Advances in psychology*, volume 52, pages 139–183, 1988.
- [7] Sue Hignett and Lynn McAtamney. Rapid entire body assessment (REBA). *Applied ergonomics*, 31(2):201–205, 2000.
- [8] Thuong N Hoang, Martin Reinoso, Frank Vetere, and Egemen Tanin. Onebody: remote posture guidance system using first person view in virtual environment. In *Proceedings of the 9th Nordic Conf. on Human-Computer Interaction*, pages 1–10, 2016.
- [9] Syazwan Aizat Ismail, Shamsul Bahri Mohd Tamrin, Mohd Rafee Baharudin, Mohamad Azhar Mohd Noor, Muhamad Hanafiah Juni, Juliana Jalaludin, and Zailina Hashim. Evaluation of two ergonomics intervention programs in reducing ergonomic risk factors of musculoskeletal disorder among school children. *Res J Med Sci*, 4(1):1–10, 2010.
- [10] Yifeng Jiang and C Karen Liu. Data-driven approach to simulating realistic human joint constraints. In *Intl. Conf. on Robotics and Automation*, pages 1098–1103, 2018.
- [11] Jay M Kapellusch, Arun Garg, Kurt T Hegmann, Matthew S Thiese, and Elizabeth J Malloy. The Strain Index and ACGIH TLV for HAL: risk of trigger digit in the WISTAH prospective cohort. *Human factors*, 56(1):98–111, 2014.
- [12] Zahra Khodabakhshi, Seyed Amin Saadatmand, Mehrdad Anbarian, and Rashid Heydari Moghadam. An ergonomic assessment of musculoskeletal disorders risk among the computer users by rula technique and effects of an eight-week corrective exercises program on reduction of musculoskeletal pain. *Iranian Journal of Ergonomics*, 2(3):44–56, 2014.
- [13] Wansoo Kim, Jinoh Lee, Luka Peternel, Nikos Tsagarakis, and Arash Ajoudani. Anticipatory robot assistance for the prevention of human static joint overloading in human–robot collaboration. *IEEE Robotics and Automation Letters*, 3(1):68–75, 2017.
- [14] Antonio Gonzales Marin, Mohammad S Shourijeh, Pavel E Galibarov, Michael Damsgaard, Lars Fritzsche, and Freek Stulp. Optimizing contextual ergonomics models in human-robot interaction. In *Intl. Conf. on Intelligent Robots and Systems*, pages 1–9, 2018.
- [15] L. McAtamney and E. N. Corlett. RULA: a survey method for the investigation of work-related upper limb disorders. *Applied ergonomics*, 24(2):91–99, 1993.
- [16] Jorge Nocedal and Stephen Wright. *Numerical optimization*. Springer, 2006.
- [17] L. Peternel, W. Kim, J. Babič, and A. Ajoudani. Towards ergonomic control of human-robot co-manipulation and handover. In *IEEE-RAS Intl. Conf. on Humanoid Robotics*, pages 55–60, 2017.
- [18] L. Peternel, N. Tsagarakis, D. Caldwell, and A. Ajoudani. Robot adaptation to human physical fatigue in human-robot co-manipulation. *Autonomous Robots*, pages 1–11, 2018.
- [19] Luka Peternel, Cheng Fang, Marco Laghi, Antonio Bichi, Nikos Tsagarakis, and Arash Ajoudani. Human arm posture optimisation in bilateral teleoperation through interface reconfiguration. In *IEEE RAS/EMBS Intl. Conf. for Biomedical Robotics and Biomechanics*, 2020.
- [20] Rahaf Rahal, Giulia Matarese, Marco Gabiccini, Alessio Artoni, Domenico Prattichizzo, Paolo Robuffo Giordano, and Claudio Pacchierotti. Caring about the human operator: haptic shared control for enhanced user comfort in robotic telemanipulation. *IEEE Tran. on Haptics*, 13(1):197–203, 2020.
- [21] Ali Shafti, Ahmad Ataka, B Urbistondo Lazpita, Ali Shiva, Helge A Wurdemann, and Kaspar Althoefer. Real-time robot-assisted ergonomics. In *Intl. Conf. on Robotics and Automation*, pages 1975–1981, 2019.
- [22] J Steven Moore and Arun Garg. The strain index: a proposed method to analyze jobs for risk of distal upper extremity disorders. *American Industrial Hygiene Association Journal*, 56(5):443–458, 1995.
- [23] Linda van der Spaa, Michael Gienger, Tamas Bates, and Jens Kober. Predicting and Optimizing Ergonomics in Physical Human-Robot Cooperation Tasks. In *Intl. Conf. on Robotics and Automation*, pages 1799–1805. IEEE, 2020.
- [24] Amir Yazdani and Roya Sabbagh Novin. Posture estimation and optimization in ergonomically intelligent teleoperation systems. In *Companion of the 2021 ACM/IEEE International Conference on Human-Robot Interaction*, pages 604–606, 2021.
- [25] Amir Yazdani, Roya Sabbagh Novin, Andrew Merryweather, and Tucker Hermans. Is The Leader Robot an Adequate Sensor for Posture Estimation and Ergonomic Assessment of A Human Teleoperator? In *IEEE Intl. Conf. on Automation Science and Engineering*, 2021.