

Anticipatory Human-Robot Collaboration via Multi-Objective Trajectory Optimization

Abhinav Jain¹, Daphne Chen¹, Dhruva Bansal¹, Sam Scheele¹, Mayank Kishore¹, Hritik Sapra¹, David Kent¹, Harish Ravichandar¹, Sonia Chernova¹

Abstract—We address the problem of adapting robot trajectories to improve safety, comfort, and efficiency in human-robot collaborative tasks. To this end, we propose *CoMOTO*, a trajectory optimization framework that utilizes stochastic motion prediction to anticipate the human’s motion and adapt the robot’s joint trajectory accordingly. We design a multi-objective cost function that simultaneously optimizes for *i*) separation distance, *ii*) visibility of the end-effector, *iii*) legibility, *iv*) efficiency, and *v*) smoothness. We evaluate CoMOTO against three existing methods for robot trajectory generation when in close proximity to humans. Our experimental results indicate that our approach consistently outperforms existing methods over a combined set of safety, comfort, and efficiency metrics.

I. INTRODUCTION

With the advances in robotics automation, robots are more frequently working in close proximity to humans. For instance, in manufacturing, humans and robots can work together to assemble components, and in household and assistive robotics, robots can provide physical assistance to humans. Thus, there is a growing need for robots to effectively and safely interact with humans in close proximity.

A key challenge in robot-human close-proximity interaction is the generation of robot trajectories that are safe, i.e. they do not physically harm the human, and are comfortable, i.e. the human is able to interpret and anticipate the robot’s behavior. Collaborative robotic systems can create safe trajectories with frequent monitoring and replanning [1], [2], but at the cost of efficiency. Anticipatory methods that use predictions of human motion can instead be used to generate safer trajectories using learned models of human motion.

Several different factors define safety and comfort of a robot’s trajectory. While a trajectory may be safe for a nearby human, it might not be comfortable. The visibility of the robot’s end effector in the peripheral vision of the human can increase comfort. Inference of the robot’s intent through partial observation of its trajectory can also greatly increase comfort. Finally, sudden and unexpected robot behavior can be a major source for discomfort. Effective robot trajectories in human-robot collaboration must take into account all of these factors. Further, user experience factors are not




Fig. 1: Effective human-robot collaboration requires both anticipation and simultaneous optimization of a variety of costs associated with safety, efficiency, and comfort.

necessarily complimentary to trajectory efficiency, as previously shown for robot-human handover tasks [3]. Thus, collaborative robot trajectory generation must effectively *balance* efficiency, safety, and comfort factors.

In this paper, we address the problem of adapting robot trajectories for human-robot collaborative environments with an overall goal of improving human safety and comfort, as well as increasing task efficiency. We define the default trajectory executed by the robot for a particular task as the nominal trajectory. Given some observation of human motion directly preceding execution of the nominal trajectory, we use a prediction of the human’s motion to adapt the nominal trajectory for improved safety and comfort.

Figure 1 shows an overview of our approach. We combine multiple objective functions to satisfy several factors of human comfort in addition to safety. We use time-sampled stochastic predictions of human motion $\hat{x}^h(t)$ to generate objective functions $C_i(\xi^r, \hat{x}^h)$, using the uncertainty of the prediction to generate appropriately conservative trajectories. We call our approach Collaborative Multi-Objective Trajectory Optimization – *CoMOTO*.

In order to evaluate CoMOTO, we define several metrics that incorporate key factors in safety and comfort of humans in collaborative environments. We perform experiments in three collaborative picking test cases, and compare the results against established baselines. Our results show that

*This work was supported in part by the Army Research Lab under Grant W911NF-17-2-0181 (DCIST CRA).

¹The authors are with the Institute of Robotics and Intelligent Machines, Georgia Institute of Technology, Atlanta, GA, USA {jain, daphne.chen, dbansal36, scheele, mkishore3, hritiksapra, dekent, harish.ravichandar, chernova}@gatech.edu

CoMOTO performs consistently well for all of our metrics across different close-proximity collaborative picking scenarios, while the baselines are able to perform well in only a particular metric or for only a specific test scenario.

II. RELATED WORKS

Human-robot collaborative manipulation is a well-studied topic. Several works focus on reactive systems for safety in collaborative environments. Lasota *et al.* propose a reactive speed control system for collaborative robots [1]. Their system monitors human pose to measure the separation distance between the human and the robot, scales the robot's execution speed by the separation distance, and stops execution completely below a specified separation threshold. Dumonteil *et al.* propose a similar reactive approach for collaborative robots in industrial applications using a state machine [2], by reactively replanning trajectories to avoid potential collisions. Reactive systems based on unadapted trajectories are task inefficient due to repeated replanning, and so we instead focus on using human motion prediction to generate initially safer trajectories.

Several prior works have utilized human motion prediction in order to proactively adapt robot trajectories. Mainprice and Berenson propose a prediction based planning framework that generates swept volumes for collision detection based on human motion prediction using a Gaussian mixture model (GMM), and interleaves planning and execution to update the prediction [4]. However, since their system selects an updated predicted human trajectory at each iteration, their framework requires constant replanning. Fishman *et al.* address the problem of coordinated human-robot collaboration, specifically in a handover task [5], using a joint optimal control model to simultaneously plan the robot's behavior and predict the humans' behavior by inferring human goals. Stouraitis *et al.* develop a method that involves estimating a human partner's policy to optimize trajectories for dyadic collaborative manipulation, where a human and robot work together to manipulate a single large object [6]. Huang and Mutlu present an anticipatory control method based on inference from human gaze [7], highlighting the task efficiency benefits of using predictive and anticipatory planning methods. Maeda *et al.* use early human action recognition to initiate a corresponding robot response [8]. Maeda *et al.* also present a Probabilistic Movement Primitive framework for learning a mixture model of human-robot interaction primitives, used to identify human tasks as well as to coordinate robot movement with the observed human movement. These works either focus on improving task efficiency through predicting human intent, or use human motion prediction to create safe trajectories. Our work builds on these ideas by using human motion prediction to improve human comfort factors as well.

Our objective of combining several factors of safety and comfort for trajectory optimization is similar to Mainprice *et al.*'s work [9], which considers distance for safety, and visibility and reachability for comfort in robot handover

tasks. However, their work assumes the human will remain still during the robot's trajectory execution, and it does not consider the pose of the human's arm. We extend collaborative multi-objective trajectory optimization to account for a moving human, with an articulated human model.

Human comfort factors beyond collision avoidance are important considerations in human-robot collaborative environments [10]. Dragan *et al.* propose legibility and predictability of robot motion to a human observer [11], [12]. A legible motion is one from which an observer can quickly and confidently infer the motion's goal after only partial observation, and predictable motion is the most expected motion to reach a goal. Stulp *et al.* present legibility as a task-specific behavior that can be learned rather than a general characteristic of a trajectory [13]. Medina *et al.* emphasize the importance of smoothness for robot-human handover trajectories [14]. In order to account for multiple factors that affect human comfort, our work considers legibility, predictability (through efficient execution), and smoothness.

III. COLLABORATIVE TRAJECTORY OPTIMIZATION

Our framework, CoMOTO, uses stochastic human motion prediction to calculate an objective function, composed of a set of costs relevant to close-proximity interaction, which is minimized using a trajectory optimization framework. Specifically, the trajectory generation pipeline consists of a brief 1 second observation period of the human's motion, which is used to predict the remaining trajectory of the human (see Section V-A for details). The predicted trajectory is then used as input to calculate a set of costs, including separation distance, visibility, legibility, deviation from a nominal trajectory, and smoothness, that account for the stochastic nature of the prediction. The costs themselves are detailed in Section IV. We formulate an objective function using a weighted combination of the costs, which is then minimized to generate a robot trajectory using TrajOpt [15], although other trajectory optimization works, such as [16], [17], and cost based planning algorithms, such as [18], can be used instead. The generated trajectory is then executed concurrently with the remainder of the human's motion.

While sensor-based reactive stops offer absolute collision prevention, they reduce efficiency with frequent interruptions requiring replanning. Instead, we directly address safety during planning. CoMOTO leverages motion prediction to anticipate human motion and generate trajectories that are inherently safer, thereby reducing the need for replanning and increasing task efficiency.

IV. TRAJECTORY ADAPTATION COSTS

Our objective function is split into several costs that cover different elements of safety and comfort in a collaborate environment. Each cost is a function of time parameterized robot joint trajectory $\xi^r(t)$ and predicted human motion $\hat{x}^h(t) \sim N(\mu^h(t), \Sigma^h(t))$ for $t \in [0, T]$.

A. Distance Cost

Distance between the human and the robot is the most critical factor in safe collaborative manipulation. Thus we formulate a cost that penalizes lower separation distances between the human and the robot. The cost is further scaled by the covariance of the prediction, with higher covariance resulting in a higher cost, resulting in more conservative trajectories when the predicted motion has higher uncertainty. The cost is formulated as follows:

$$C_{dist}(\xi^r, \hat{x}^h) = \sum_{t=0}^T \sum_i \sum_j \frac{1}{d_{i,j}(t)^T (\Sigma_i^h(t))^{-1} d_{i,j}(t)} \quad (1)$$

$$d_{i,j}(t) = \mu_i^h(t) - p_j^r(t)$$

where $\mu_i^h(t)$ and $\Sigma_i^h(t)$ are the mean and covariance of the predicted 3D position of the i^{th} human joint at time t , and $p_j^r(t)$ is the 3D position of the j^{th} robot joint at time t .

B. Visibility Cost

During trajectory execution, visibility of the robot's end effector is an essential factor for human comfort [19]. If the robot is out of the field of view of the human, the human may be distracted and try to locate it, thus decreasing both human comfort and task efficiency. This is a basic human instinct for safety against unpredictable moving objects. The visibility cost penalizes the end effector for being farther from the human's gaze.

We define the visibility cost as the angle between the predicted human gaze and the line between the position of the robot end effector and the human's head. We define the predicted human gaze as the line from the predicted position of the human head to the position of the object with which the human is interacting. The cost is scaled inversely to the variance of the prediction of the human head pose.

$$C_{vis}(\xi^r, \hat{x}^h) = \sum_{t=0}^T \frac{\angle(O, \mu_{head}^h(t), p_{ee}^r(t))}{\sigma_{head}^h(t)} \quad (2)$$

where O is the 3D position of the object with which the human is interacting, $\mu_{head}^h(t)$ and $\sigma_{head}^h(t)$ are the mean and variance of the predicted 3D position of the human head, and $p_{ee}^r(t)$ is the 3D end-effector position at time t .

C. Legibility Cost

The robot's motion must be legible, that is, it must convey its intent through its trajectory. Dragan *et al.* define a legible robot trajectory as one from which the user can quickly and confidently infer the task goal after only partial trajectory execution [11]. We choose to implement a legibility cost in order to improve the human's ability to understand the robot's intent. We replicate the Legibility cost from [12].

$$C_{legibility}(\xi^r) = \frac{\sum_t P(G|\xi_{S \rightarrow Q_t}^r) f(t)}{\sum_t f(t)} \quad (3)$$

$$P(G|\xi_{S \rightarrow Q}^r) = \frac{\exp(-C(\xi_{S \rightarrow Q}^r) - C(\xi_{Q \rightarrow G}^{*r}))}{\exp(-C(\xi_{S \rightarrow G}^{*r}))} \quad (4)$$

$\xi_{A \rightarrow B}^r$ is the trajectory from configuration A to configuration B . S denotes the robot's start configuration, G denotes its goal configuration, and Q_t denotes its configuration at time t . $f(t)$ is a weighing function that increases cost of legibility towards the beginning of the trajectory. The optimal trajectory ξ^{*r} is a linear trajectory in the Cartesian space. $C(\xi)$ is the length of the trajectory in Cartesian space. Dragan *et al.* include a regularizer term $\lambda C(\xi)$ in order to prevent excessively long trajectories, which we exclude from our cost as that requirement is met by the Nominal Trajectory Cost.

D. Nominal Trajectory Cost

The nominal trajectory is the default trajectory executed by the robot without any adaptation. This trajectory is calculated using a collision cost and a joint velocity cost in TrajOpt. The nominal trajectory can be viewed as one that optimizes smoothness, collision avoidance (with objects), and efficiency in the *absence* of a human. While the costs defined thus far focus solely on the human, our nominal trajectory cost brings balance to the overall cost function, and acts as a regularizer to preserve efficiency.

The nominal trajectory cost penalizes deviation from the nominal trajectory. The cost is calculated as a sum of Cartesian distances of the end effector between the nominal trajectory and the adapted trajectory at each timestep:

$$C_{nominal}(\xi^r) = \sum_{t=0}^T \|p_{nom}^r(t) - p_{ee}^r(t)\| \quad (5)$$

where $p_{nom}^r(t)$ is the position of the end effector at time t in the nominal trajectory.

E. Smoothness Cost

Smooth robot motion is a necessary component for a comfortable collaborative environment. Several dynamical quantities can be minimized across the trajectory to generate smooth motion. Prior trajectory optimization frameworks such as [16] use sum of squared velocities of the robot as a smoothing cost. However, in order to better decrease jerkiness of adapted trajectories as well as to even out speed across the execution of the trajectory, we use the sum of squared acceleration of the robot as follows:

$$C_{smooth}(\xi^r) = \sum_{t=0}^{T-2} \left(\left\| \frac{d^2}{dt^2} \xi(t) \right\|^2 \right) \quad (6)$$

F. Multi-Objective Optimization

The final objective function is the sum of all the above costs. The overall optimization problem is given by

$$\{\xi^{r*}\}_{t=0}^T = \arg \min_x \sum_{i \in \mathcal{C}} \alpha_i C_i(\xi^r, \hat{x}^h) \quad (7)$$

$$\text{s.t. } \xi^{r*}(T) = \xi_g^r \quad (8)$$




Fig. 2: Simulated test environment showing the KUKA robot and the keypoints of the human skeleton.

where $\{\xi^{r^*}\}_{t=0}^T$ denotes the optimized robot joint trajectory, $\mathcal{C} = \{dist, vis, legibility, nominal, smooth\}$, α_i represents the pre-specified weights associated with the i^{th} cost, and ξ_g^r denotes the desired goal location.

We note that the costs used in the objective function do not necessarily incentivize the same behavior. For instance, minimizing the distance cost will push the robot trajectory away from the human. On the other hand, the visibility cost will work to pull the trajectory closer to the human. Nevertheless, each cost function considers an important aspect of the interaction. CoMOTO thus attempts to find optimized trajectories that effectively trade-off different costs according to the specified weighting factors.

V. EXPERIMENTS AND RESULTS

We perform a series of experiments on three different test cases involving a human and a robot working in a collaborative environment in order to evaluate the performance of our approach. We define four key metrics to evaluate CoMOTO and compare our results against several baselines. The human motion predictions are generated in Matlab. All trajectory optimization is performed using TrajOpt [15]. The experiments are run on a KUKA LBR iiwa R820 Robot in a ROS Gazebo simulation [20], a visualization of which is shown in Figure 2. The coefficients α_i are chosen empirically for optimized performance.

We present three test cases involving close human-robot collaboration, categorized by the behavior of the human:

- **Stationary:** Stationary human with robot reaching for an object.
- **Reaching-far:** Human and robot reaching for distant objects.
- **Reaching-near:** Human and robot reaching for closely-positioned objects.

For each test case, we use 5 unique human trajectories and a unique nominal robot trajectories. We note that the trajectories within each test case represent different relative initial positions for the human and the robot.

A. Human Motion Prediction

We use [21] as the framework for stochastic human motion prediction. The provided code includes ground truth trajectories that are split into training and testing datasets. The trajectories are 3D positions of the human’s right arm, 2 to 3 seconds long, recorded at 100Hz. The GMM model is

trained on 100 trajectories of a human reaching for an object and 100 trajectories of a still human with arm stretched out. Since the dataset only contains recorded trajectories for the right arm (shoulder, elbow, wrist and palm), the remaining human skeleton consisting of the neck, head, torso and left arm is extrapolated using fixed offsets. The same offsets are applied to the mean of the prediction of the right shoulder to generate the predictions of the remaining joints. The covariances for the remaining skeleton are identical to that of the right shoulder.

For each experiment, the ground truth human trajectory is split into two. The first 100 samples (1 second) are used as the observation. A prediction of the remaining human motion is generated based on the observation of those 100 samples. Subsequently, CoMOTO optimizes the robot trajectory such that it is synchronized with the anticipated human motion. We assume that the human remains stationary after reaching the goal.

B. Baselines

We evaluate CoMOTO, against the following baselines:

Nominal: the non-adapted nominal trajectory generated by TrajOpt using common costs and constraints, including collision, joint velocity, and joint target constraints. We include the nominal trajectory alone to show how our approach improves this trajectory’s performance with respect to the full set of metrics described in Section V-C.

Speed-Adjusted: the nominal trajectory executed with real-time speed adjustment based on human-robot separation distance as described in [1].

Legible: the legible motion optimization algorithm of [12]. The baseline implementation uses a legibility cost identical to the one used in our approach. The optimal trajectory is again a linear trajectory in the Cartesian space.

Distant+Visible: local path optimization using the method presented in [9], optimizing for costs based on human-robot separation distance and human visibility. To provide a direct multi-objective optimization comparison to our approach, we use their cost-based optimization to adapt the nominal trajectory, rather than a path generated by a T-RRT planner.


C. Evaluation Metrics

We measure the performance of each algorithm according to the following metrics:


Separation distance (Dst.): percentage of the trajectory where the separation distance between the robot and the human exceeds 20cm.

End effector visibility (Vis.): percentage of the trajectory where the robot’s end effector is within the human’s 160° field of view. When calculating this metric, we assume the human is looking at their target object.


Legibility (Leg.): legibility of the robot’s end effector motion to the human observer, calculated as described in [11].



(a) Performance comparison for the stationary scenario.



(b) Performance comparison for the reaching-far scenario.



(c) Performance comparison for the reaching-near scenario.

Fig. 3: Comparison plots between CoMOTO and each baseline with respect to all evaluation metrics (mean ± 1 SD)

Deviation from the nominal trajectory (Nom.): sum of squared distance between the adapted trajectory and the nominal trajectory.

We note that we utilize the ground truth (as opposed to the predicted) human trajectories to compute the metrics.

D. Results and Discussion

The performance of all algorithms on all metrics can be found in Table I. We ran a one-way analysis of variance (ANOVA) with correlated samples for each metric to determine whether differences in our measurements were statistically significant. Where ANOVA showed a significant difference of trajectory optimization approach on any of our metrics at $p < 0.05$, we conducted post tests between the approaches with Tukey’s HSD test. For brevity, we show significant differences in Table I only for approaches that performed significantly better than *all* other approaches. We also provide visual comparisons of our approach to each baseline to better show the comparison over all metrics at once in the radar plots of Figure 3. We break down the results individually for each test case below.

1) *Stationary:* We first consider the scenario where the robot must pick an object with a stationary human present in the workspace. The results for this test case are shown in

Table I and Figure 3a. Since the motion prediction model was trained on stationary trajectories as well as moving trajectories, the predictions are observed to have minimal motion. With a stationary human, CoMOTO and the Dist+Vis baseline perform comparably on both distance and visibility, with no significant difference. CoMOTO outperforms all other approaches in legibility.

2) *Reaching-far:* We next consider the scenario where the human and robot are concurrently reaching for different objects. The results for this test case are shown in Table I and Figure 3b. The anticipatory nature of CoMOTO allows it to outperform the Dist+Vis baseline in separation distance and come a close second in visibility. CoMOTO also outperforms all other approaches in legibility.

3) *Reaching-near:* Finally, we consider the scenario where the human and robot are concurrently reaching for objects that are close together, providing a situation where high separation distance cannot be maintained. As can be seen in Table I and Figure 3c, all methods perform poorly in the distance metric. However, CoMOTO is able to outperform other methods in distance while still maintaining good performance in the other metrics. The Speed-Adjusted baseline is not always able to complete execution of the trajectory since the goal position of the robot may be within

TABLE I: Each approach’s mean ± 1 SD for all metrics and scenarios, with the best-performing result shown in bold. Note that lower values are preferred for the Nominal metric. We denote statistically significant improvements over all other methods with * ($p < 0.05$) or ** ($p < 0.01$).

	Approach	Dst.(%)	Vis.(%)	Leg.	Nom. (m^2)
Stationary	CoMOTO	94.8 \pm 0.4	79.6 \pm 5.8	59.5 \pm 0.9**	7.0 \pm 1.4
	Nominal	83.3 \pm 15.2	65.6 \pm 9.3	39.7 \pm 0.0	<i>n/a</i>
	Speed-Adj	83.3 \pm 15.2	64.0 \pm 11.7	36.5 \pm 6.5	0.4 \pm 0.8*
	Legible	91.2 \pm 6.2	62.6 \pm 9.4	47.1 \pm 0.0	2.2 \pm 0.0
	Dist+Vis	91.4 \pm 3.8	76.4 \pm 6.4	40.2 \pm 1.2	6.4 \pm 0.5
Reaching-far	CoMOTO	94.8 \pm 0.2	70.0 \pm 3.7	58.9 \pm 1.8**	6.4 \pm 1.1
	Nominal	74.6 \pm 0.1	57.9 \pm 4.0	39.7 \pm 0.0	<i>n/a</i>
	Speed-Adj	74.6 \pm 13.6	56.2 \pm 6.5	36.3 \pm 6.4	0.5 \pm 0.8**
	Legible	90.3 \pm 5.7	54.7 \pm 4.6	47.1 \pm 0.0	2.2 \pm 0.0
	Dist+Vis	90.3 \pm 3.2	72.3 \pm 2.5	40.1 \pm 1.6	6.4 \pm 0.3
Reaching-near	CoMOTO	62.2 \pm 5.5	100.0 \pm 0.0	45.7 \pm 5.8	2.9 \pm 0.7
	Nominal	51.5 \pm 5.7	100.0 \pm 0.0	12.5 \pm 0.0	<i>n/a</i>
	Speed-Adj	51.6 \pm 5.6	100.0 \pm 0.0	4.1 \pm 3.5	1.5 \pm 0.7
	Legible	48.2 \pm 13.7	100.0 \pm 0.0	42.0 \pm 0.0	3.6 \pm 0.0
	Dist+Vis	57.8 \pm 6.2	100.0 \pm 0.0	0.0 \pm 1.8	2.6 \pm 1.0

its stopping threshold of 6 cm.

Summary: Our experiments demonstrate that CoMOTO is a highly adaptable framework that can optimize trajectories to improve factors of safety, comfort, and efficiency across different operating scenarios. CoMOTO significantly improves the performance in separation distance, visibility and legibility over the Nominal baseline. We note that, across all three scenarios, CoMOTO scores consistently better than or equivalent to all baselines both in terms of the distance metric and the visibility metric (see Table I). This observation suggests that our approach outperforms all the baselines in terms of maintaining a safe and comfortable distance from the human while *simultaneously* making sure that the robot’s end effector is visible to the moving human. As one would expect, the Dist+Vis baseline results in slightly better visibility than CoMOTO only in the reaching-far scenario. CoMOTO consistently scores the highest in the legibility metric as well. While CoMOTO doesn’t excel in the nominal metric, its performance is comparable to other metrics. Overall, CoMOTO is able to perform well across all metrics while the baselines perform well only in specific metrics or scenarios.

VI. CONCLUSION

We have presented CoMOTO, a multi-objective trajectory optimization framework for anticipatory human-robot collaboration. Our results demonstrate that existing methods tend to specialize and excel on either individual safety, comfort or efficiency metrics, or in particular collaborative scenarios. In comparison, since it simultaneously optimizes for various costs, CoMOTO performs comparably or better across the full set of metrics and scenarios. Future work will explore reactive replanning to tackle unexpected human movements and address computational bottlenecks to enable real-time implementation.

REFERENCES

- [1] P. A. Lasota, G. F. Rossano, and J. A. Shah, “Toward safe close-proximity human-robot interaction with standard industrial robots,” in *IEEE International Conference on Automation Science and Engineering (CASE)*. IEEE, 2014, pp. 339–344.
- [2] G. Dumonteil, G. Manfredi, M. Devy, A. Confetti, and D. Sidobre, “Reactive planning on a collaborative robot for industrial applications,” in *12th International Conference on Informatics in Control, Automation and Robotics (ICINCO)*, vol. 2. IEEE, 2015, pp. 450–457.
- [3] C.-M. Huang, M. Cakmak, and B. Mutlu, “Adaptive coordination strategies for human-robot handovers,” in *Robotics: Science and Systems*, vol. 11. Rome, Italy, 2015.
- [4] J. Mainprice and D. Berenson, “Human-robot collaborative manipulation planning using early prediction of human motion,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013, pp. 299–306.
- [5] A. Fishman, C. Paxton, W. Yang, N. Ratliff, and D. Fox, “Trajectory optimization for coordinated human-robot collaboration,” *arXiv preprint arXiv:1910.04339*, 2019.
- [6] T. Stouraitis, I. Chatz Nikolaidis, M. Gienger, and S. Vijayakumar, “Dyadic collaborative manipulation through hybrid trajectory optimization,” in *CoRL*, 2018, pp. 869–878.
- [7] C.-M. Huang and B. Mutlu, “Anticipatory robot control for efficient human-robot collaboration,” in *11th ACM/IEEE international conference on human-robot interaction (HRI)*. IEEE, 2016, pp. 83–90.
- [8] G. Maeda, A. Maloo, M. Ewerton, R. Lioutikov, and J. Peters, “Anticipative interaction primitives for human-robot collaboration,” in *AAAI Fall Symposium Series*, 2016.
- [9] J. Mainprice, E. A. Sisbot, L. Jaillet, J. Cortés, R. Alami, and T. Siméon, “Planning human-aware motions using a sampling-based costmap planner,” in *IEEE International Conference on Robotics and Automation*. IEEE, 2011, pp. 5012–5017.
- [10] P. A. Lasota, T. Fong, J. A. Shah, *et al.*, “A survey of methods for safe human-robot interaction,” *Foundations and Trends® in Robotics*, vol. 5, no. 4, pp. 261–349, 2017.
- [11] A. Dragan, K. Lee, and S. Srinivasa, “Legibility and predictability of robot motion,” in *Proceedings of Human-Robot Interaction*, March 2013.
- [12] A. Dragan and S. Srinivasa, “Generating legible motion,” in *Robotics: Science and Systems*, 2013.
- [13] F. Stulp, J. Grizou, B. Busch, and M. Lopes, “Facilitating intention prediction for humans by optimizing robot motions,” in *IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2015, pp. 1249–1255.
- [14] J. R. Medina, F. Duvallat, M. Karnam, and A. Billard, “A human-inspired controller for fluid human-robot handovers,” in *IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2016, pp. 324–331.
- [15] J. Schulman, Y. Duan, J. Ho, A. Lee, I. Awwal, H. Bradlow, J. Pan, S. Patil, K. Goldberg, and P. Abbeel, “Motion planning with sequential convex optimization and convex collision checking,” *The International Journal of Robotics Research*, vol. 33, no. 9, pp. 1251–1270, 2014.
- [16] 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*, vol. 32, no. 9–10, pp. 1164–1193, 2013.
- [17] M. Kalakrishnan, S. Chitta, E. Theodorou, P. Pastor, and S. Schaal, “Stomp: Stochastic trajectory optimization for motion planning,” in *IEEE international conference on robotics and automation*. IEEE, 2011, pp. 4569–4574.
- [18] L. Jaillet, J. Cortés, and T. Siméon, “Transition-based rrt for path planning in continuous cost spaces,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2008, pp. 2145–2150.
- [19] E. A. Sisbot, L. F. Marin-Urias, R. Alami, and T. Simeon, “Spatial reasoning for human-robot interaction,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2007.
- [20] C. Hennersperger, B. Fuerst, S. Virga, O. Zettinig, B. Frisch, T. Neff, and N. Navab, “Towards mri-based autonomous robotic us acquisitions: a first feasibility study,” *IEEE transactions on medical imaging*, vol. 36, no. 2, pp. 538–548, 2017.
- [21] R. Luo, R. Hayne, and D. Berenson, “Unsupervised early prediction of human reaching for human-robot collaboration in shared workspaces,” *Autonomous Robots*, vol. 42, no. 3, pp. 631–648, 2018.