# Cost Functions for Robot Motion Style

Allan Zhou[1] and Anca D. Dragan[2]

arXiv:1809.00092v1 [cs.RO] 1 Sep 2018

*Abstract*— **We focus on autonomously generating robot motion for day to day physical tasks that is expressive of a certain style or emotion. Because we seek generalization across task instances and task types, we propose to capture style via cost functions that the robot can use to augment its nominal task cost and task constraints in a trajectory optimization process. We compare two approaches to representing such cost functions: a weighted linear combination of hand-designed features, and a neural network parameterization operating on raw trajectory input. For each cost type, we learn weights for each style from user feedback. We contrast these approaches to a nominal motion across different tasks and for different styles in a user study, and find that they both perform on par with each other, and significantly outperform the baseline. Each approach has its advantages: featurized costs require learning fewer parameters and can perform better on some styles, but neural network representations do not require expert knowledge to design features and could even learn more complex, nuanced costs than an expert can easily design.**

## I. Introduction

Our goal is to enable robots to move more expressively – communicating internal states like hesitation, or projecting personality or affect aspects like excitement or disappointment. Further, we want robots to do this *while conducting their day-to-day tasks*: we don't want them executing some prescripted motion for communication purposes, only to then go back to the same nominal robotic way when they actually do the task they are supposed to. Rather, style should go hand in hand with the task: opening up the fridge door confidently, moving in towards the juice box cautiously so as to not know over the glass bottle of milk, and happily handing it over to the person.

Motion style is an active area of research in both robotics, as well as graphics and animation. In graphics, a lot of work has focused on *motion capture style transfer* [1], [2], [3], [4]: taking a clip motion capture data in a certain style and transferring it to another clip. These approaches work well for animated characters, and are really promising for robot motions that happen *outside* of the robot's physical task, like reactions to events. But applying them to achieve style *during* the physical task is challenging, because the robot needs to maintain the constraints that the task imposes.

In robotics, work on expressive motion has focused on the design features that are predictive of style [5], [6], [7], but robots have yet to autonomously generate their motion across tasks instances and task types with style. Autonomous expressive motion generation remains largely confided to expressing intentions, not styles [8], [9].

> *Our observation is that we can capture style through a cost function that augments the robot task objective function and constraints.*

We explore generating style cost functions for manipulator arms, and leverage trajectory optimization to produce stylized motion using the same cost function across different task instances and types.

Our work is related to graphics work that learns cost functions for human locomotion styles from demonstration [10], [11]. Unfortunately though, demonstrations of stylized non-anthropomorphic manipulator arms are difficult to acquire, and features used for learning locomotion style do not transfer to manipulation.

We make the following contributions:

**Handcrafted style features.** We motion capture a dance artist performing day to day manipulation and locomotion tasks in different styles, use observations from this data to design useful features, and learn linear cost functions of these features.

**Style costs with learned features.** While hand-crafted features are useful in that they incorporate domain knowledge, they also rely on a expert to design the right features for them. We thus explore an alternative: learning a cost function represented as a neural network, that operates on the raw trajectory. We contribute an adaptation of deep comparison-based learning to this setting.

**User Study.** We compare both featurized and neural network cost functions against nominal motions for different tasks and styles. We find that users rate the cost style-optimized motion as more expressive of the intended style, and that these motions better enable them to identify the intended style.

The two approaches each have their pros and cons: on the one hand, handcrafting features can be challenging and even though it performed decently on the three styles, we already had difficulty with learning *hesitant* (this is where the neural network performed the best relative to the handcrafted features); on the other hand, neural network representations might have more limited generalization than an expertly designed cost, and tended to slightly underperform for styles that area easier to handcraft features for.

Overall, we are excited to provide an optimization-based approach to autonomously generating stylized

[1]Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, Berkeley, U.S.A. allan.zhou@berkeley.edu

[2]Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, Berkeley, U.S.A.

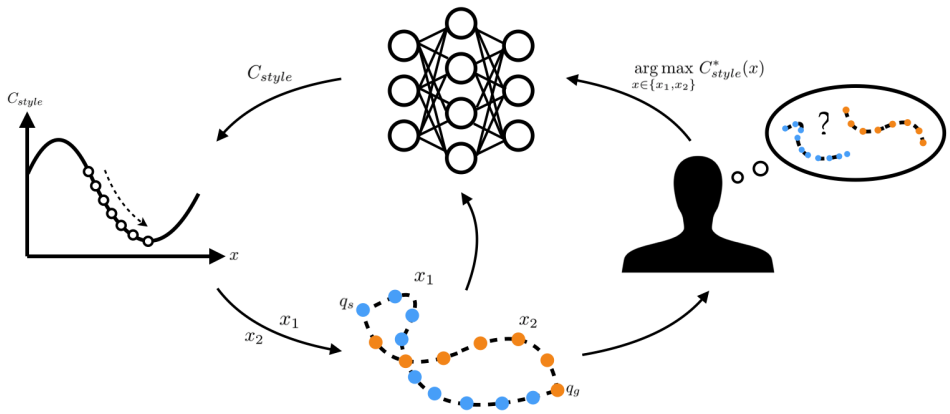


**Fig. 1:** An overview of our approach to learning style cost functions. Counter-clockwise from left: The robot generates trajectories via trajectory optimization with the current cost estimate $C_{\text{style}}$. The robot *queries* the expert with pairs of trajectories, and for each pair the expert labels the trajectory that is more expressive of the target style. The expert's labels are used to update $C_{\text{style}}$. We investigate representing $C_{\text{style}}$ both as a neural network operating on raw trajectory input, and as a weighted linear combination of hand-designed features.

motion for robot arms, along with a first attempt at comparing featurized and neural network representations of cost functions for style.

## II. Approach

In trajectory optimization, we can formulate a motion planning problem as a constrained optimization problem:

$$\begin{aligned} \underset{x}{\text{minimize}} \quad & C(x) \\ \text{subject to} \quad & h_i(x) = 0,\ i = 1, \ldots, m. \\ & g_j(x) \leq 0,\ j = 1, \ldots, n. \end{aligned} \tag{1}$$

The trajectory $x$ is represented by a sequence of waypoints, where each waypoint is a particular configuration. If the robot has $D$ degrees of freedom and the trajectory has $T$ waypoints, then $x$ is a $D \times T$ matrix where $x[t]$ is a single waypoint. The constraints $h_i$ and $g_i$ ensure that $x$ completes the task of moving from the start to goal while avoiding collisions. $C$ is a cost function that is often designed to encourage minimum-length paths. One such cost is the sum of squared differences of configurations:

$$C_{ssd}(x) = \sum_{t=1}^{T-1} \|(x[t+1] - x[t])\|^2 \tag{2}$$

Problems of the form Eqn. (1) can be solved locally by using trajectory optimizers [12], [13].

Given a desired style such as *hesitant*, our goal is to find $C_{\text{style}}$ such that doing trajectory optimization on the objective:

$$C(x) = \underbrace{C_{\text{style}}(x)}_{\text{Style term}} + \lambda \underbrace{C_{ssd}(x)}_{\text{Task term}} \tag{3}$$

generates trajectories that complete the task in a hesitant style. The task term encourages the robot to complete the task with a reasonably efficient and smooth trajectory, while the style term encourages the motion to be in the desired style.

### A. Featurized Costs

We first approached the problem by hand-designing trajectory features that are relevant to a style, and then expressing the cost as a weighted linear combination of them:

$$C^F_{\text{style}}(x) = w^T \phi(x) \tag{4}$$

where $\phi$ are the features and $w$ are the learned weights.

We identified trajectory features pertaining to three styles: *happy*, *sad*, and *hesitant*. For *happy* and *sad*, we began by studying motion capture data of an actress who performed different tasks in specified styles. Across different tasks, we noticed her tendency to "dip" her head for the *sad* style. Fig. 2 illustrates this pattern for both a handover and place task. As we are working with a non-humanoid manipulator arm, we naturally decided to focus on features of the end effector as the robot's analog to a human's head.

Additionally, the actress tended to keep her arms close to the torso for the *sad* style, while extending them further out for the *happy* style. From these observations, we chose to define:

- $f_r$: the average horizontal distance from end effector to base (radius)
- $f_h$: the average end effector z-coordinate (height)
- $f_o$: the average angle between the vertical, or positive z-axis, and the direction the end effector is pointing (orientation)

We can then express our featurized style costs for *happy* and *sad*

$$C^F_{\text{style}}(x) = [w_r, w_h, w_o]^T \begin{bmatrix} f_r \\ f_h \\ f_o \end{bmatrix} \tag{5}$$

For *hesitant*, we included the same three features $f_r, f_h, f_o$, and added additional features relevant to timing and motion speed. During execution the waypoints are equally spaced in time, so distances between waypoints determine the relative speed of the robot as it moves. Let

$$f_{v_i} = \|x[i+1] - x[i]\| \tag{6}$$

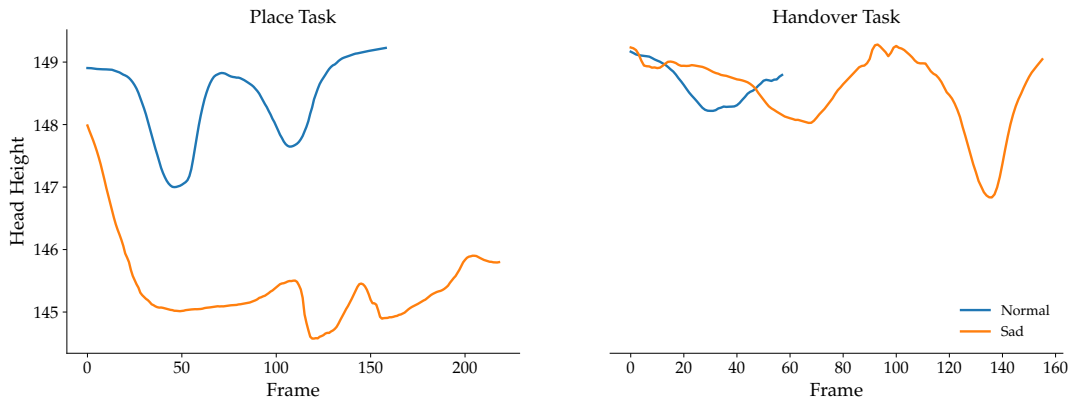


**Fig. 2:** The head height over time of the actress for handover and place tasks, as recorded by motion capture.

A large value of $f_{v_i}$ means the robot will move *quickly* between waypoints $x[i]$ and $x[i+1]$, while a small value means the robot will move slowly between those waypoints. Naturally, we will call our $f_{v_i}$ the *velocity* features. We then define our featurized *hesitant* cost in terms of the end effector and velocity features:

$$C^F_{\text{style}}(x) = \left[w_r, w_h, w_o, w_{v_1}, \ldots, w_{v_{T-1}}\right]^T \begin{bmatrix} f_r \\ f_h \\ f_o \\ f_{v_1} \\ \cdots \\ f_{v_{T-1}} \end{bmatrix} \tag{7}$$

### B. Neural Network Costs

Work in deep reinforcement learning has shown some success in using neural networks to learn more complicated cost functions without painstakingly designing complicated features by hand [14], [15], [16]. We can parameterize our cost with a Multi-layer Perceptron (MLP) that takes in a waypoint and forward kinematics information as input, and outputs a vector $y_t$ for each waypoint of the trajectory. We also provide the neural network with velocity and timing information, using the predecessor waypoint $x[t-1]$ and waypoint number $t$. Similar to [15], our style cost functions are then of the form

$$C^{NN}_{\text{style}}(x) = \sum_{t=1}^{T-1} ||y_t||^2 \tag{8}$$

The MLP we chose to parameterize $C^{NN}_{\text{style}}$ has two hidden layers with *tanh* activations and 42 and 21 units, respectively. The output layer $y_t$ has 21 units is linear (no activation). We apply Dropout [17] after the first two layers during training.

### C. Learning from Preferences

*1) Overview:* Both the featurized cost $C^F_{\text{style}}$ and neural network cost $C^{NN}_{\text{style}}$ are parameterized by weights which we will learn from the human's style cost. There has been a line of previous work focused on learning cost functions from humans through demonstrations [18], [19], [15], [16]. Although it may seem natural to obtain demonstrations of styled motion from humans, demonstrations of styled motion in robots that may be non-humanoid is more difficult.

A useful alternative to learning from demonstrations is learning costs from preferences, which has been explored for both linear combinations of features [20], as well as for cases where the cost is parameterized by neural networks [21]. In this setup, we repeatedly generate pairs of trajectories $(x_A, x_B)$, and query the human to pick the one they think is better.

In our setting, we are interested in learning a cost for a style, where task completion is taken care of separately (by the $C_{ssd}$ term and the trajectory optimization constraints). If we were learning a cost for, say, the *sad* style, we would query the human to select the trajectory they think looks more sad.

We assume that the probability the human selects a trajectory decreases exponentially with cost [22], so we predict the probability they prefer $x_A$ as:

$$P(x_A > x_B) = \frac{exp(-C_{\text{style}}(x_A))}{exp(-C_{\text{style}}(x_A)) + exp(-C_{\text{style}}(x_B))} \tag{9}$$

The human's actual selection serves as the label to the prediction. Define $I(x)$ as 1 if $x$ is the human selects $x$, and 0 otherwise. The cross entropy loss for this sample is then:

$$L = -I(x_A) \log P(x_A > x_B) - I(x_B) \log P(x_B > x_A) \tag{10}$$

For both $C^F_{\text{style}}$ and $C^{NN}_{\text{style}}$, we can use gradient-based optimization to update our weights and minimize Eqn. (10).

In our implementation, we generate comparisons in small batches which are sent to the human for labeling. Once a batch of comparisons is labeled, we update the weights using by minimizing the mean loss over the batch. Fig. 1 shows a visual overview of our approach.

*2) Generating Trajectories:* Given a $C_{\text{style}}$, we form an optimization problem as in Eqn. (1), using $C_{\text{style}}$ to build the objective according to Eqn. (3). We use TrajOpt [12] to find a locally optimal solution $x_0$ to this optimization problem.

Assuming $C_{\text{style}}$ is fully learned, we are only interested in the solution trajectory $x_0$. We can time

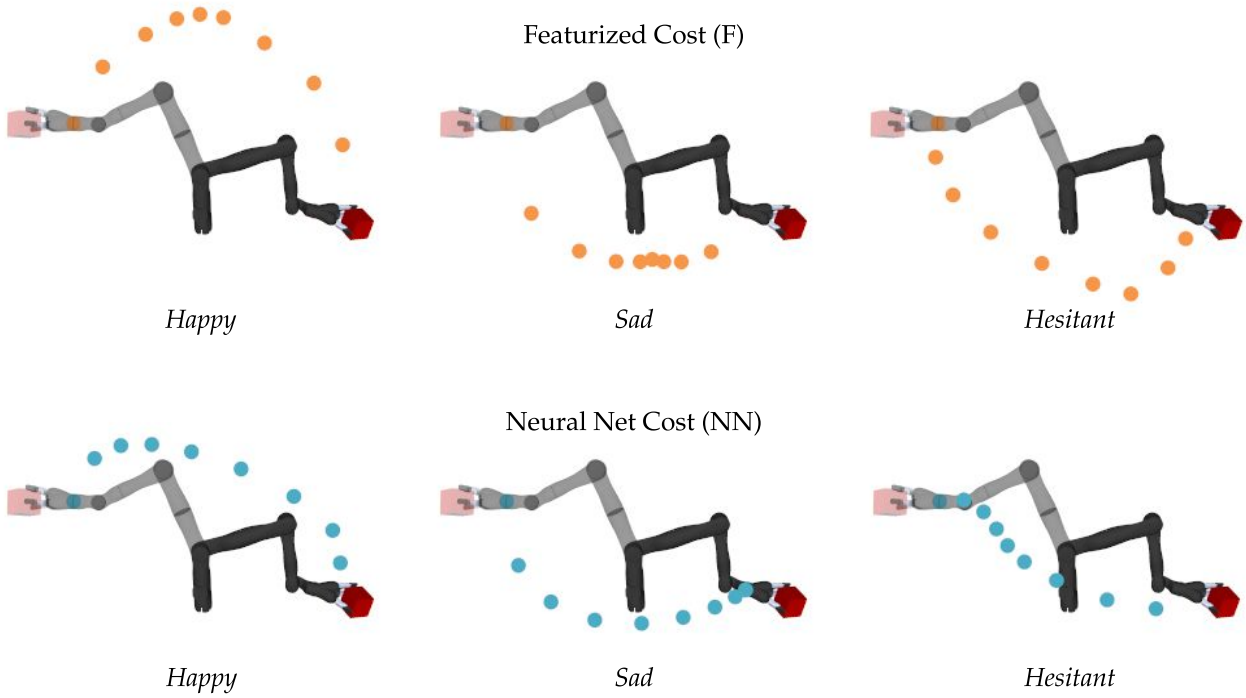


**Fig. 3:** Output of optimizing the featurized and neural network costs in the place task, for each style.

the trajectory (spacing waypoints equally in time) and execute it on the robot.

During the learning process, however, we want to produce pairs of trajectories $(x_A, x_B)$ for the human to compare and label. Moreover, if we only query the human with trajectories that are optimal with respect to the current $C_{\text{style}}$ estimate, the human will have to compare trajectories that all look very similar and we may not adequately explore the space of possible trajectory styles. We introduce *exploration* in the learning process by creating additional trajectories which are random variations of $x_0$. In particular we can create a new trajectory of the form $x' = x + \Delta$, where $\Delta$ is a small, smooth perturbation. Repeating this process we create many variations of $x_0$, which we use to form query pairs $(x_A, x_B)$.

To create the smooth perturbation $\Delta$, we start with a small random perturbation $\delta \in \mathbb{R}^D$ to only a single randomly selected waypoint. If we think of a trajectory $x$ as a $D \times T$ dimensional vector, with all the waypoints concatenated together, then:

$$\Delta_0 = [0 \quad \cdots \quad 0 \quad \delta \quad 0 \quad \cdots \quad 0]^T \tag{11}$$

$$\Delta = \beta A^{-1} \Delta_0 \tag{12}$$

$$x' = x_0 + \Delta \tag{13}$$

where

$$A = \begin{bmatrix} 2I & -I & 0 & \cdots & 0 \\ -I & 2I & \ddots & \ddots & \vdots \\ 0 & \ddots & \ddots & \ddots & 0 \\ \vdots & \ddots & \ddots & 2I & -I \\ 0 & \cdots & 0 & -I & 2I \end{bmatrix} \tag{14}$$

The effect of Eqn. (12) is to smooth out the single-waypoint perturbation $\delta$ to multiple waypoints, and $\beta$ is simply a scalar coefficient to ensure that the size of $\delta$ does not change.

### D. Training Details

For each of the three styles, we trained a featurized cost $C^F_{\text{style}}$ per Sec. II-A and a neural net cost $C^{NN}_{\text{style}}$ as described in Sec. II-B. For both cost function types, we trained with 16 human labels for the *happy* and *sad* styles and 15 labels for the *hesitant* style.

For the neural network cost $C^{NN}_{\text{style}}$ in particular, we observed better performance when we augmented the data with rotations about the z-axis. Specifically, suppose the human gives their preference $I(x)$ for a pair of trajectories $(x_A, x_B)$. Since the robot is upright for all tasks, we assume that the preference does not change if both trajectories are rotated about the z-axis by the same angle $\theta$. This means for a single pair of trajectories and the preference label we can generate multiple training data points with randomly selected $\theta$:

$$(Rot_z(x_A, \theta), Rot_z(x_B, \theta), I(x)) \tag{15}$$

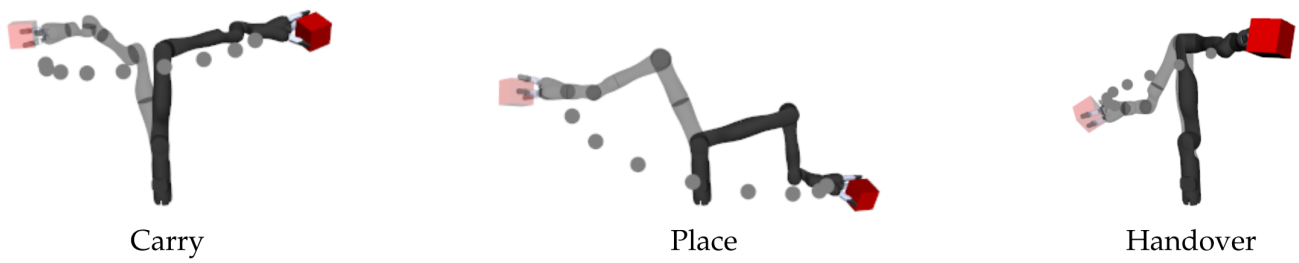


**Fig. 4:** Trajectories generated using sum of squared differences cost ($C_{ssd}$) to illustrate each of the 3 tasks.

We found that this augmentation helped prevent the neural network from over-fitting to certain start and goal pairs. Note that for $C^F_{\text{style}}$, the features $\phi$ are already invariant to these rotations and so this augmentation does not have any effect.

**Featurized cost analysis.** For the featurized costs $C^F_{\text{style}}$, we can manually inspect the learned weights on each feature in order to analyze behavior. For the *happy* and *sad* styles we learned weights on the end effector radius, height, and orientation features ($f_r, f_h, f_o$) as described in Eqn. (5). The final learned weights for *happy* and *sad* were:

$$w^{happy} = [0.03, -0.79, 0.38]^T \tag{16}$$

$$w^{sad} = [0.97, 0.42, -0.50]^T \tag{17}$$

As we expected, $w^{happy}$ puts a negative weight on the end effector height feature, rewarding trajectories that move higher. It puts a positive weight on the orientation feature: since the orientation feature is the angle to the vertical, a positive weight penalizes the robot if its end effector "dips" downward. Correspondingly, $w^{sad}$ penalizes end effector height and rewards the end effector for dipping down.

Meanwhile, $w^{sad}$ puts a large positive weight on the radius feature, encouraging the arm to stay closer to the torso and appearing "withdrawn." However, $w^{happy}$ not have the expected corresponding behavior, putting only a small positive weight on the radius feature.

The learned weights $w^{hesitant}$ act similarly to $w_{sad}$ on the end effector features. Interestingly, $w^{hesitant}$ puts negative weights on the velocity features $f_{v_i}$ for $i <= 4$, and positive weights on $f_{v_i}$ for $i > 4$. This rewards the robot for moving more quickly for the first part of the trajectory, and then penalizes the robot for moving quickly in the second part. The overall effect is that our featurized cost for *hesitant* learns to encourage trajectories that slow down closer to the goal.

## III. Experiments

We run two users studies to evaluate and contrast the featurized and neural network cost functions trained in Sec. II-D We compared both types of cost functions against the neutral non-stylized baseline.

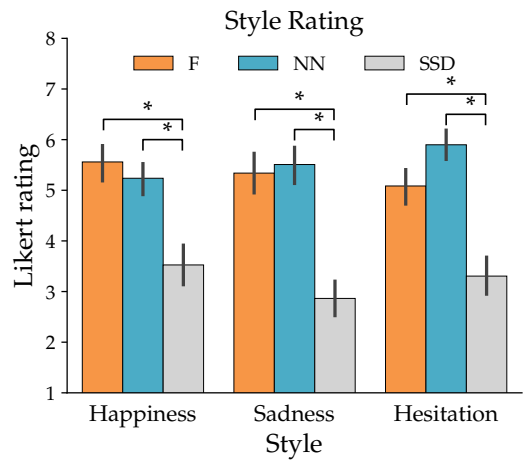


**Fig. 5:** Average *style rating* of trajectories produced using each type of cost function: featurized (F), neural network (NN), and sum of squared differences (SSD). The *style rating* is the participant's Likert rating on the style the cost function was trained for. For example, if a trajectory was generated with a cost trained for *happy*, the *style rating* is the Likert rating for "Happiness."

### A. Experimental Design

*1) Manipulated Factors:* We manipulate the *cost function type* that the robot uses with three levels: featurized (F), learned neural network (NN), and sum of squared differences (SSD).

We apply each cost function type to three different styles: *happy*, *sad*, and *hesitant*. For each style, we planned trajectories for three different tasks: a carry task, a place task, and a handover task (see Fig. 4). The task's start and goal configurations were held out during training. Throughout our experiments we generated trajectories with 10 waypoints ($T = 10$). The trajectories generated using each of the learned cost functions in the place task is shown in Fig. 3.

As a simple baseline, we also generated trajectories for each task using the sum of squared differences cost, $C_{ssd}$.

### B. Dependent Measures

We measure how effective each cost type is at producing motion that has that style. We measure effectiveness in two ways, in two separate user studies.

**Study 1: Style Rating.** The first study collects ratings of how much the generated motion made the robot look like the intended style. Participants saw groups of three

motions for the same task and style, each produced with a different cost type (featurized, neural network, or SSD).

First, the participants observed each trajectory and responded to the free response question "What style or emotion would you attribute to this robot?"

We then asked participants to rate on a 7-point Likert scale the happiness, sadness, and hesitation of each motion they observed, regardless of whether that trajectory was produced using the corresponding style cost (to avoid biasing the participants with the question). For the *style rating* we only took the Likert rating that matched the cost function. For example, if a trajectory was produced using a style cost for *happy*, then the *style rating* is the participant's rating for "Happiness," as opposed to "Sadness" or "Hesitation."

We also asked a forced-choice question for each style. For instance for *happy*, we asked to choose the "most happy" motion between three trajectories generated using either the featurized cost, the neural net cost, or SSD.

**Study 2: Correct Identification.** The second study tests whether participants can identify the correct style from distractor styles. In the second study, for each task we presented the *happy*, *sad*, and *hesitant* trajectories generated using either the featurized cost or the neural net cost. The trajectory generated with sum of squared differences was also presented. The participants were asked to rank each of the four trajectories from "most ___" to "least ___" in each of the three styles.

### C. Subject Allocation

In the first study, both the cost function type and style were within subjects, while the task was between subjects. Each subject saw trajectories generated using all cost types for each of the 3 styles, but only in a single task.

In the second study, cost function type, style, and task were all within subjects.

We recruited 60 (20 per task) participants for the first study and 30 participants for the second study using Amazon's Mechanical Turk (AMT). Every participant spoke fluent English and had a minimum 95% approval rating on the AMT platform.

### D. Hypothesis

*We hypothesize that cost type has a significant effect on style rating and correct identification: both of the learned costs should outperform the nominal cost. We do not know which of the two will perform best: the neural network does not have the benefit of hand-designed features, but also has the capacity to learn useful features on its own that we might have not thought about.*

### E. Analysis

**Study 1: Style Ratings.** Fig. 5 plots the results of the *style rating* from the first study. We analyzed the style ratings using a fully factorial repeated measures ANOVA with cost type, style, and task as factors and

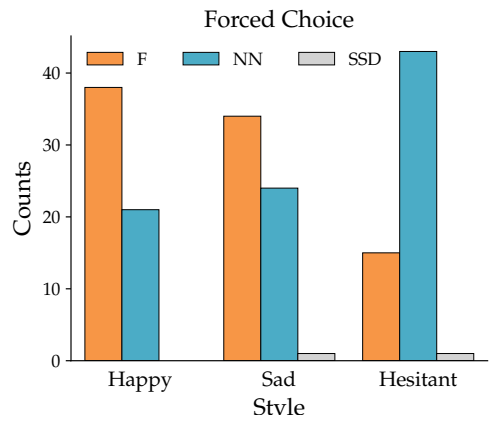


**Fig. 6:** Participants compared trajectories in a certain style generated using either a featurized (F), neural net (NN), or sum of squared differences (SSD) cost function. They then chose the trajectory they thought most expressed that style. This plot shows how often trajectories generated using each cost were chosen.

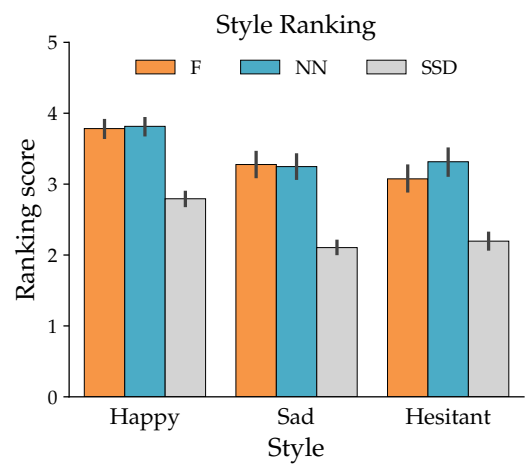


**Fig. 7:** The ranking score of each cost function type, which measures how successfully participants could correctly identify the style of trajectories generated using each cost function type. A higher number means more accurate identification.

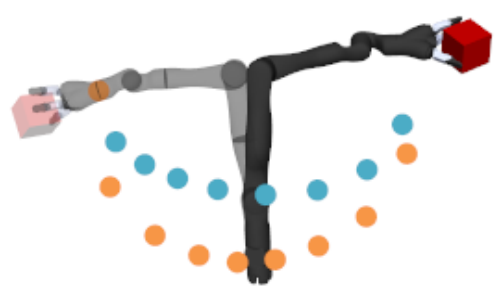

**Fig. 8:** The end effector path of *sad* trajectories generated with the neural network (blue) cost and the featurized (orange) cost, for the carry task.

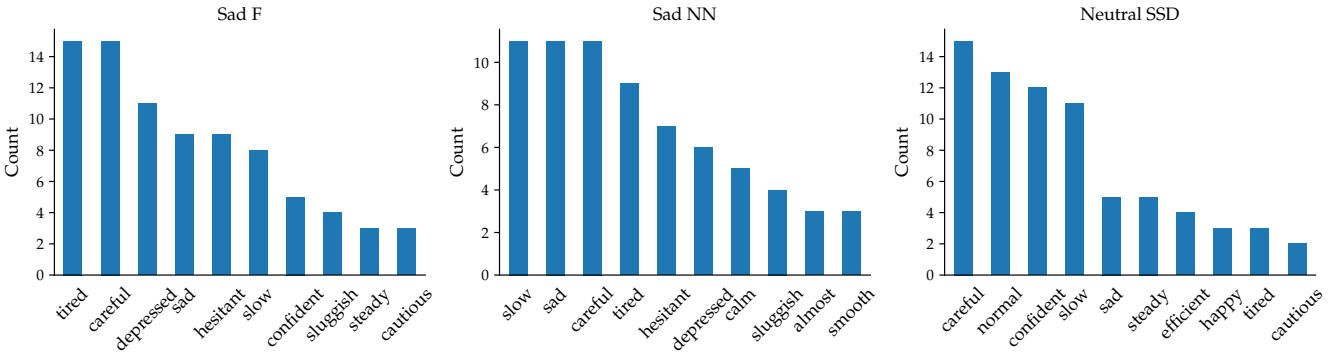


**Fig. 9:** Words that participants used to describe robot motion produced with the *sad* featurized (F) and neural net (NN) costs, as well as the sum of squared differences (SSD) cost.

user id as a random effect. We found a significant main effect for cost type ($F(2, 448) = 160.96, p < .0001$), but also interaction effects with the style ($F(4, 448) = 3.97, p < .01$), task ($F(4, 448) = 2.45, p = .05$), and the three-way interaction was also significant ($F(8, 448) = 2.03, p = .04$). We followed up with Tukey HSD posthocs, and saw that the only significant differences show that both learned costs performed better than the nominal SSD baseline, across tasks and styles.

There was no conclusive difference between the neural net and the featurized costs. However, in the forced choice section the featurized cost was preferred for the *happy* and *sad* styles, while the neural net cost was preferred for hesitant (see Fig. 6).

Fig. 8 plots a comparison of trajectories generated using each cost type *sad* in the carry task: the end effector dips much lower in the trajectory generated using the featurized cost, compared to the one generated using the neural network cost. A corresponding effect is observed for the *happy* case. Not surprisingly, in simple styles like *happy* and *sad*, learning weights on a few hand-designed features outperforms the neural network cost.

The neural net seems to perform best on the hesitant style, judging by the forced choices results. Indeed, looking at Fig. 3, we see that for this style the neural net cost produces a more sophisticated motion that is slow at first and faster after – this behavior is more nuanced than what the corresponding motion produced with the featurized cost.

To analyze the responses to the free response questions, we split the responses into individual words and then removed common "stop words" such as "the," "it," "a," etc. We also removed the words "robot," "robots," and "video" as participants commonly used them to reference what they were seeing but they are not relevant to our analysis. After filtering, we plotted the most commonly used words in a histogram. The histogram for the *sad* style is shown in Fig. 9. As we would expect from the *Style Rating* results, the responses to the trajectories generated by both the neural network and featurized style costs were fairly similar. For both cases, in addition to "sad" and the obvious descriptor "slow," responses commonly referenced styles that are visually similar to *sad*, such as "tired" and "depressed." Meanwhile, the responses for trajectories generated by $C_{ssd}$ are very different from responses to the other two types. For example, a common descriptor for these trajectories is "normal," which is not commonly used for the other trajectories.

The pattern for the *happy* and *sad* styles is the same: responses to trajectories produced with the neural network or featurized costs are very similar to each other, while the responses to the trajectories produced by $C_{ssd}$ are different to the other two.

**Study 2: Correct Identification.** We flipped the ranking data so that higher is better, then analyzed it in the same way as the style rating. We found a significant main effect for cost type ($F(2, 1082) = 300.75$, $p < .0001$), with the Tukey HSD posthoc again supporting that SSD has worse performance. Style also had a significant main effect ($F(2, 1080) = 75.83$, $p < .0001$), with the posthoc finding happy to be easier to identify. The task was also significant, ($F(2, 1088) = 3.04, p = .05$), where the correct style in the handover task was easier to identify than in the carry task.

There were not significant differences between the neural net and the featurized costs. These results echo the subjective user ratings.

## IV. Discussion

**Summary.** We approached the problem of generating styled motion in robotic manipulation tasks by learning two different kinds of styled cost functions. First, we learned a linear cost function of hand-designed features, then we learned a neural network cost on raw trajectory input. We trained both types of cost functions by utilizing human preferences.

We ran two experiments to compare the performance of these two methods, and the results of the experiments showed that both methods performed significantly better than neutral trajectories at expressing the desired style. They also showed some advantages for the neural network cost in terms of expressing more complicated costs, such as hesitant, while the featurized cost matched or beat the neural network cost in simple styles such as *happy* or *sad*.

**Limitations and Future Work.** This work touches only a small part of the problem of generating styled motion in robots. We focused on three styles in this paper, but both approaches described in the paper could be generalized to completely different styles.

As we saw in the experiments, our featurized costs use a rather limited set of features which could limit their expressivity in some cases. They could be made more expressive by more carefully considering the style at hand and designing more complicated features.

Further investigation is also needed to test the preference based learning system. The styles we investigated required relatively few human responses to train the neural network cost, but a more difficult style might require more iterations of the training process and would better test the effectiveness of the query generation process. The query generation process itself could potentially be improved to increase the efficiency of the learning process.

## ACKNOWLEDGMENT

This work was supported in part by the A.F.O.S.R., the Wagner foundation, and H.K.U.S.T. We thank the members of the InterACT lab for providing helpful discussion and feedback.

## References

[1] L. Torresani, P. Hackney, and C. Bregler, "Learning motion style synthesis from perceptual observations," in *Advances in Neural Information Processing Systems*, 2007, pp. 1393–1400.

[2] S. Xia, C. Wang, J. Chai, and J. Hodgins, "Realtime style transfer for unlabeled heterogeneous human motion," *ACM Transactions on Graphics (TOG)*, vol. 34, no. 4, p. 119, 2015.

[3] M. E. Yumer and N. J. Mitra, "Spectral style transfer for human motion between independent actions," *ACM Transactions on Graphics (TOG)*, vol. 35, no. 4, p. 137, 2016.

[4] D. Holden, J. Saito, and T. Komura, "A deep learning framework for character motion synthesis and editing," *ACM Transactions on Graphics (TOG)*, vol. 35, no. 4, p. 138, 2016.

[5] H. Knight and R. Simmons, "Expressive motion with x, y and theta: Laban effort features for mobile robots," in *Robot and Human Interactive Communication, 2014 RO-MAN: The 23rd IEEE International Symposium on*. IEEE, 2014, pp. 267–273.

[6] D. Szafir, B. Mutlu, and T. Fong, "Communication of intent in assistive free flyers," in *Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction*. ACM, 2014, pp. 358–365.

[7] M. Sharma, D. Hildebrandt, G. Newman, J. E. Young, and R. Eskicioglu, "Communicating affect via flight path exploring use of the laban effort system for designing affective locomotion paths," in *Human-Robot Interaction (HRI), 2013 8th ACM/IEEE International Conference on*. IEEE, 2013, pp. 293–300.

[8] M. J. Gielniak and A. L. Thomaz, "Generating anticipation in robot motion," in *RO-MAN, 2011 IEEE*. IEEE, 2011, pp. 449–454.

[9] A. D. Dragan, K. C. Lee, and S. S. Srinivasa, "Legibility and predictability of robot motion," in *Human-Robot Interaction (HRI), 2013 8th ACM/IEEE International Conference on*. IEEE, 2013, pp. 301–308.

[10] C. K. Liu, A. Hertzmann, and Z. Popović, "Learning physics-based motion style with nonlinear inverse optimization," *ACM Transactions on Graphics (TOG)*, vol. 24, no. 3, pp. 1071–1081, 2005.

[11] S. J. Lee and Z. Popović, "Learning behavior styles with inverse reinforcement learning," in *ACM Transactions on Graphics (TOG)*, vol. 29, no. 4. ACM, 2010, p. 122.

[12] J. Schulman, J. Ho, A. X. Lee, I. Awwal, H. Bradlow, and P. Abbeel, "Finding locally optimal, collision-free trajectories with sequential convex optimization." in *Robotics: science and systems*, vol. 9, no. 1, 2013, pp. 1–10.

[13] N. Ratliff, M. Zucker, J. A. Bagnell, and S. Srinivasa, "Chomp: Gradient optimization techniques for efficient motion planning," in *Robotics and Automation, 2009. ICRA'09. IEEE International Conference on*. IEEE, 2009, pp. 489–494.

[14] M. Wulfmeier, P. Ondruska, and I. Posner, "Maximum entropy deep inverse reinforcement learning," *arXiv preprint arXiv:1507.04888*, 2015.

[15] C. Finn, S. Levine, and P. Abbeel, "Guided cost learning: Deep inverse optimal control via policy optimization," in *International Conference on Machine Learning*, 2016, pp. 49–58.

[16] J. Ho and S. Ermon, "Generative adversarial imitation learning," in *Advances in Neural Information Processing Systems*, 2016, pp. 4565–4573.

[17] N. Srivastava, G. E. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting." *Journal of machine learning research*, vol. 15, no. 1, pp. 1929–1958, 2014.

[18] A. Y. Ng, S. J. Russell, *et al.*, "Algorithms for inverse reinforcement learning." in *Icml*, 2000, pp. 663–670.

[19] B. D. Ziebart, A. L. Maas, J. A. Bagnell, and A. K. Dey, "Maximum entropy inverse reinforcement learning." in *AAAI*, vol. 8. Chicago, IL, USA, 2008, pp. 1433–1438.

[20] D. Sadigh, A. D. Dragan, S. Sastry, and S. A. Seshia, "Active preference-based learning of reward functions," in *Robotics: Science and Systems (RSS)*, 2017.

[21] P. Christiano, J. Leike, T. B. Brown, M. Martic, S. Legg, and D. Amodei, "Deep reinforcement learning from human preferences," *arXiv preprint arXiv:1706.03741*, 2017.

[22] R. A. Bradley and M. E. Terry, "Rank analysis of incomplete block designs: I. the method of paired comparisons," *Biometrika*, vol. 39, no. 3/4, pp. 324–345, 1952.