Expressive Motion with X, Y and Theta: Laban Effort Features for Mobile Robots

Heather Knight, Reid Simmons, *Robotics Institute, Carnegie Mellon University*

***Abstract*— There is a saying that 95% of communication is body language, but few robot systems today make effective use of that ubiquitous channel. Motion is an essential area of social communication that will enable robots and people to collaborate naturally, develop rapport, and seamlessly share environments. The proposed work presents a principled set of motion features based on the Laban Effort system, a widespread and extensively tested acting ontology for the dynamics of "how" we enact motion. The features allow us to analyze and, in future work, generate expressive motion using position (x, y) and orientation (theta). We formulate representative features for each Effort and parameterize them on expressive motion sample trajectories collected from experts in robotics and theater. We then produce classifiers for different "manners" of moving and assess the quality of results by comparing them to the humans labeling the same set of paths on Amazon Mechanical Turk. Results indicate that the machine analysis (41.7% match between intended and classified manner) achieves similar accuracy overall compared to a human benchmark (41.2% match). We conclude that these motion features perform well for analyzing expression in low degree of freedom systems and could be used to help design more effectively expressive mobile robots.**

I. Introduction

Robots are increasingly being designed to operate in human environments. A robot that communicates its internal state can better conduct tasks, achieve rapport, and collaborate with people in settings where its intentions would otherwise be unclear. People do this all the time and are highly proficient at reading nonverbal cues, from the moment someone enters the room for a job interview, to making a snap decision about a potential romantic partner, or whether a child needs comforting.

Our goal is to make a robot's state quickly and clearly legible to interaction partners and bystanders. *Expressive motion*, which we define as a robot's ability to communicate mental state, social context and task state via body movements, can accomplish such legibility at a glance. Using expressive motion, a robot could convey that it is under a deadline by hurrying down a hallway, which might lead to fewer immediate requests for action or response, or jovially enter an office to preview having good news. Such movement-specific expressions may smooth completion of tasks involving interaction and result in a higher estimation of the robot's (social) intelligence.

*Research supported by National Science Foundation.

H. Knight and R. Simmons, Robotics Institute at Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA, USA, 15213 (phone: +1-412-268-3818, fax: +1-412-268-6436, emails: heatherbot@cmu.edu, reids@cs.cmu.edu).

Actors typically spend years training to use their bodies to express emotions, relationships, and other aspects of their characters' internal state. Our objective is to operationalize one such widely used method, the Laban Effort System [17], so that it can be used to generate legible expressive motion. Previous work has found that readable state communications result from having Laban-trained actors overlay the Laban Efforts onto flying robot trajectories [28]. The specific contribution of this work is a set of parameterized motion features that quantify the Laban Effort Vectors for mobile robots, whose motion is limited to position (x, y) and angle (theta). In other words, we want be able to analyze the Laban Efforts expressed in a mobile robot's motion without hiring a trained actor. By limiting the implementation to three degrees of freedom, we also hope to discover underlying principles connecting motion trajectories to communication of robot (or agent) state.

The ultimate goal is to apply these features to CoBot (Fig. 2), a mobile robot with omnidirectional base that autonomously completes tasks in Carnegie Mellon's Computer Science building. We chose six manners as representative samples of categories useful to current CoBot task scenarios. 1) **Emotion:** the robot might be *happy vs. sad* at the success or failure of a task. 2) **Internal State:** The robot might be *confident vs. shy* if it had explored a space, met a particular person, or completed a task many times before. 3) **Task State:** The robot might express its availability to socialize by enacting *rushed vs. lackadaisical* motion pathways.

After describing the related work, we detail our conceptual framework for specifying Laban Motion features (Section 3). Next, we collect expressive motion trajectories instructed to be performed in each of the six manners from experts in robotics and theater (Section 4). Using these sample trajectories, we evaluate the quality of our motion features by training classifiers on the human data and testing them with cross-validation (Section 5). To characterize the quality of the classification results, we give the same data to humans on Amazon Mechanical Turk and compare the human label consistency to the machine classification accuracy (Section 6). We find that our Laban features are slightly more effective than people at labeling expressive motion abstracted to x, y and theta. This implies that, for low degrees of freedom, these features can capture as much information as humans can.

II. Related Work

There is a natural human tendency to anthropomorphize even the simplest of moving shapes and machines along various channels of social expression [4][12][13][22][23]. In fact, inappropriate attribution of mental state can be used to reveal impairment in social understanding [1]. Previous researchers have also adapted theatrical methods to designing sociable robots [5][16][20]. Therefore, we believe that a properly programmed robot could fluently communicate state if we successfully quantify the relevant expressive behaviors.

Previous work demonstrates that people will ascribe emotions and intention to robots exhibiting characteristic or sequentially recognizable motions. Wizard-of-Oz experiments with an 'emotive' stick [11] found its motion to impact social attributions of personhood or machine, emotional state or even social behaviors (saying hello), when performed in an interaction context. Experiments with single-axis door [14] confirm these findings. After the door opened slightly then closed, one subject reported that it indicated that the door saw them, judged them and decided not to let them in. The readability of these motion communications has also been demonstrated on a high degree of freedom PR2 robot [30], where researchers worked with a Pixar animator to expose its reactions to success or failure, and on the HERB robot during handoffs to communicate when it is trying to reach for or pass an object [29].

Finally, we know that varying particular motion characteristics on a simple robot, e.g., the Roomba, impacts human attributions toward that machine [17]. The (DE)SIRE model specifically explores speed, intensity, repetition and emotion for a humanoid form [19]. As described in the introduction, one successful set of motion features, which overlap with the studies mentioned above, has been Laban Movement Analysis. We expand upon these works by mapping their motion features to contextualized robot motions.

The idea that there might be universal features of expressive motion originates in human and animal recognition of biological motion. We can read motion expressions within and across species, where relevant for survival [10], or by simulating similar motions in ourselves [9]. In addition to the usefulness of this recognition, hormones in our bloodstream directly influence certain physical embodiments of state - freezing or running away in the face of a threat, lethargic motion when depressed [2]. Therefore, there is a biological basis for the expression of certain agent states.

Researchers have advanced our understanding of why we attribute mental state to simple representations, but do not provide a holistic framework for generating lifelike or expressive motions. Studies identify the importance of goal-directed motion for theory of mind attributions in animations of simple geometric shapes [1][6][15]. For instance, as long as robots move with or toward a readable goal, we will make attributions about their mental state, task, or social relationships. While these findings validate the plausibility of generating expressive motion, we turn to acting methodology to identify principles of motion for quantification and composition.

III. Background & General Approach

With the goal of analyzing and eventually generating expressive motion for mobile robots, we propose to translate the Effort Vectors from Laban Movement Analysis into calculable motion features. In this section, we overview Laban's Effort vectors, Time (sudden/sustained), Weight (strong/light), Space (direct/indirect) and Flow (bound/free), followed by an exploration of the eight combinations of Time, Weight and Space, also known as the Action Efforts. Next, we outline the previous technical implementations of Laban Movement Analysis and why our approach is needed. Finally, we outline our features computations, which can be calculated from position (x, y) and orientation (theta) alone.

A. Laban Effort Vectors

The Laban Effort Vectors are part of Laban Movement Analysis (LMA), a system for describing and recording human motion, developed to preserve dance choreography, much like a musical score preserves sound. Previous work has found value in applying various Laban Movement Analysis systems to humanoids [18][21][24][25][27]. LMA primarily details notation for Body, Space, Shape and Effort (our focus). While Body and Shape reference the motion of body parts, implying a humanoid form, and Space indicates where the motion is located in your kinesphere (i.e. sphere of possible motion, such as high, medium or low), a mobile robot is constrained to the floor, so we analyze Effort alone. Flying robots have used Laban Efforts to express readable robot state using actor-created trajectories [28]. Our efforts complement their findings, enabling computational features.

TABLE I. The Laban Effort System

Effort Vector	Fighting Polarity	Inducing Polarity
Time: *attitude toward time*	Sudden (abrupt)	Sustained (gradual)
Weight: *force or apparent inertia*	Strong (powerful)	Light (delicate)
Space: *attitude toward target*	Direct (single-focus)	Indirect (multi-focus)
Flow: sense of restriction	Bound (constrained)	Free (unconstrained)

The Effort system, which Laban sometimes referred to as "dynamics," attempts to relate interior intention to subtle motion characteristics such as strength and timing. Laban instructors describe these efforts as the "how" of a motion. There are many ways to walk toward a water fountain, but the combination of the efforts displayed during that path (e.g. acceleration, focus), can indicates something different about the agent's inner state (e.g. confidence, thirst). The Efforts are broken into four categories, Time, Weight, Space, and Flow (Table I), which all scale between a "fighting" and "inducing" polarity, titles indicative of the connection between these motion characteristics and internal state.

The polarity of each vector indicates the agent's attitude toward that category. For example, an agent's relaxed (sustained) attitude toward Timing might have gradual velocity transitions. For our application, manifestations of these efforts need to be adapted to a non-humanoid robot. For example, instead of using crossed arms vs. an open stance for a constrained Flow, we might limit or exaggerate the range of motion of the robot's orientation. Weight might simulate the movement pattern of a heavy robot versus one that is delicate and light.

One can picture the Effort vectors as a 4-dimensional space (Fig. 1). While Flow helps connect series of motions with narrative context, e.g. feelings of freedom or constraint, the first three efforts (Time, Weight, Space) are also known as the Action Efforts. Their eight combinations (Fig. 1, right)

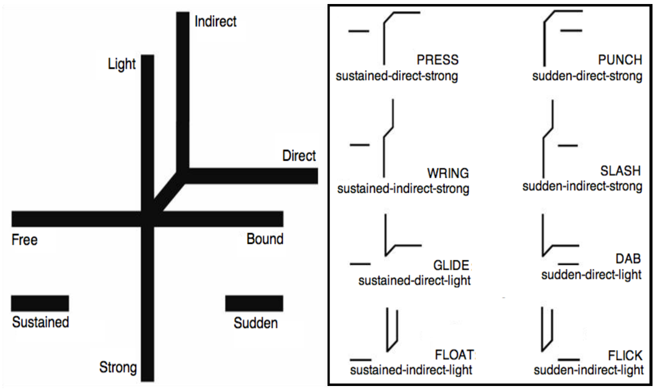

Figure 1. (left) Fully-labeled graph of Laban Effort Vector notation, (right) Sample of how the notation would be used to describe the eight action combinations of Time, Weight and Space, e.g. "dab" is sustained-direct-light and "flick" is sustained-indirect-light (adapted from [17])

map to classes of gesture (e.g. pull is an inverted press). An actor can draw from these gestures to show emotions physically, even if they have played a part many times before or are having trouble accessing those emotions. Thus they are widely used in acting training.

B. *Quantifying Laban Effort Features*

We calculate motion features based on the Laban Effort System using the limited degrees of freedom of a mobile robot, namely, X, Y and Theta. The fourth variable is time. We describe our proposed features below.

Timing: To measure how path motions scale from Sudden to Sustained, jerk (the derivative of acceleration) and the variance of the velocity provide measures of abruptness, but for a first-pass implementation, we also include speed (in x, y and the direction of motion).

Weight: To quantify the perception of force, Strong to Light, we track the acceleration patterns. Heavy objects or high-resistance environments present higher frictions, which would reflect themselves in ramping accelerations rather than steps.

Space: To rate the agent's focus on its target, from Direct to Indirect, we measure the relative angle of its orientation toward the targets, the relative angle between its motion vector and orientation. We intend to eventually include the curvilinear and rectilinear attributes of its path as well.

Flow: To quantify the agent's projection of constraint, from Bound to Free, we measure side-to-side motion in y, side-to-side motion in theta (does it just look straight ahead or all around) and orientation variance. It may also be appropriate to include temporal flow features for other kinds of robot motions, like repetition or frequency.

One could use these Laban Effort characteristics to inform various robot motion behaviors. The analysis presented here evaluates these features in classifying 'emotive traversal' data from experts in robotics or theater. Effective analysis is a necessary precursor to generation, helping validate whether such features can communicate robot state. The ability to classify motions may also help identify what robots communicate unintentionally (e.g., a robot tour guide whose acceleration patterns make it seem confrontational).

IV. Expressive Motion Sample Trajectories

We collect expressive motion sample trajectories from experts in robotics and theater using a motion-tracking table. We include both expert groups, since the former group has experience with robot motion and the latter with expressive motion, so we want to tap into the insight of both groups. Participants are instructed to move a mockup of the CoBot robot (Fig 2) from A to B and back to A in one of six manners (Fig 3). The mockup has an indicated front (an abstracted screen) and an amoeba tracking fiducial on its base. Fig. 3 (left) shows a graphic depicting the instructed exercise. In the actual experiment, the motion-tracking table had a black foam-core border overlaid on the surface, defining the elliptical region shown.

We use reacTIVision [7] software to do the tracking at 15 frames/sec, calibrated to the dimensions of the space. Data stores as (id, x, y, theta, time), in which the id is constant, x represents forward-back position between A and B, y represents side-to-side position, theta is orientation (zero is facing B), and time is the frame-number.

Figure 2. The motion-tracking table during a pre-study, Two views of study Mockup and a photo of the CoBot robot, future recipient of this expressive motion work

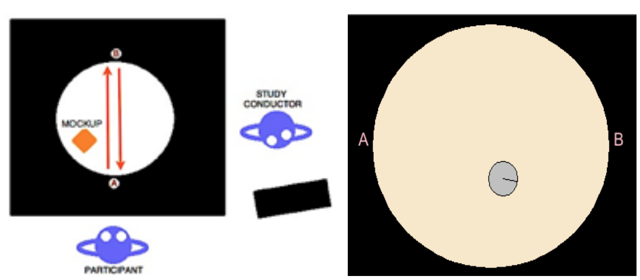

Figure 3. (left) Graphic depicting data collection, Participant stands while conductor watching screen, task is to move mockup from A-B-A (right) screenshot of rendered trajectory video given to MTurk participants

One goal of this work is to demonstrate whether our features will be generalizable to expressive motions. Utilizing oppositional pairings heightens the likelihood of high contrast data that shows the potential of algorithmic motions to be as identifiable as handcrafted animations. We evaluate three oppositional pairs, *happy vs. sad* (emotion), *confident* *vs. shy* (internal state), and *rushed vs. lackadaisical* (task state), as all are relevant to CoBot task scenarios.

We base our analysis on 16 participants (8 female and 8 male) with high self-reported expertise in either Robotics or Theater. The average participant age was 28 (std. dev. 7.4 years), and the average years in their specialty was 7.7 (std. dev. 5.1 years). Participants demonstrated each manner twice. We use only the second trial as they had had more time to refine their solution.

The study procedure began with a film depicting CoBot. Next, the study conductor presented the mockup and motion tracking table (Fig 2), which had markings indicating positions A and B (Fig 3). The study conductor then instructed the participant to move the mockup from A-to-B-to-A in a particular manner. Subjects enacted the six manners in randomized order within oppositional pairings: happy/sad, confident/shy, rushed/ lackadaisical. The ordering follows from our pilot in which participants found affective states easier to conceptualize. They ran through each set of six manners twice.

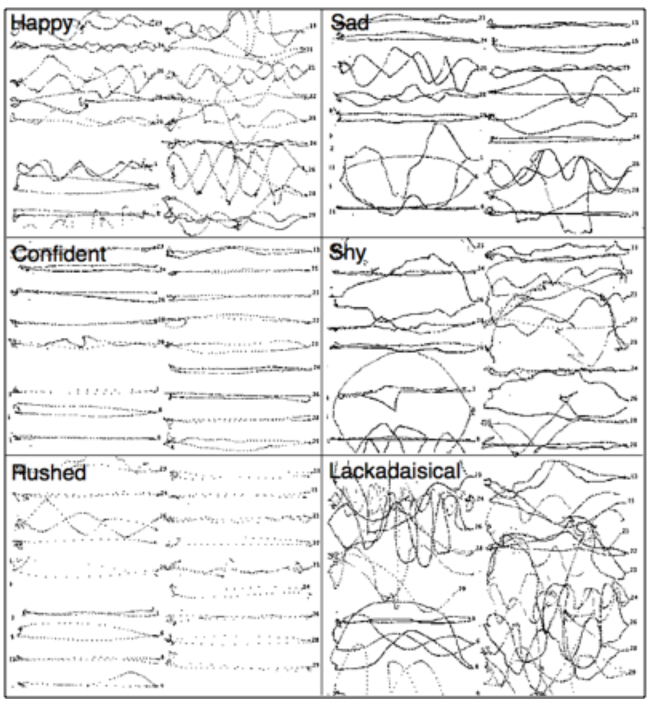

Figure 4. Expressive Motion Trajectories. Each manner rectangle includes two columns of A-B-A path renderings from a bird's eye view.

Participants described the resulting trajectories (Fig. 4) as having explicit motion characteristics. They also engaged in a great deal of storytelling and human references to come up with their expressive motion solutions. Inspection of the participant trajectories reveals certain commonalities in the solutions. Confident and Rushed are mostly straight line paths. Happy frequently undulates. Lackadaisical meanders widely.

Note that the display format of Fig. 4 reveals or obscures different motion characteristics. Widely spaced dots result from rapid motion (contrast Rushed to Sad), and deviating paths and curves are easily noticeable. Orientation, overlapping paths (e.g. hesitating or backing up), and small motions, such as the turn characteristics at point B, are better found in the data file.

V. Path Classification with Laban Features

We use the collected expressive motion trajectories to evaluate our Laban features, using cross validation to test the classification results on the oppositional pairings (75% accuracy) and across all manners (41.7%). We have emotive traversal samples of each of the 16 participants for each of the six manners, thus we analyze 96 data files. To correct for our sampling frequency (15frames/sec), we smooth our data at half the sampling frequency. We also truncate the start and end of the data files to active motion sequences.

Next, we calculate our Laban Effort Features across all paths, made up of initial implementations of the four efforts. We quantify the agent's attitude toward **Time** by velocity characteristics, such as means and variances. We quantify the apparent **Weight** via the path's acceleration characteristics. We quantify the agents sense of goal (**Space**) by tracking the percent of time it spends oriented toward point B (starting orientation), point A (ending orientation), versus towards a side. Finally, we quantify **Flow**, the apparent boundaries imposed on the motion, or lack thereof, via measure of side-to-side motion off the path and range of orientation.

TABLE II. Significant Laban Effort Features in Distinguishing Oppositional Pairs, via ANOVA analysis (N=96)

Pairing	TIMING					
	ave-y-vel	ave-x-vel	ave-speed	var-y-vel	var-x-vel	var-speed
happy/sad	0.0020*	<.0001*	<.0001*	0.0094*	0.0021*	<.0001*
confident/shy	0.7793	.0096*	.0098*	0.3776	0.0653	0.1223
rushed/lackad	0.0067*	<.0001*	<.0001*	0.0359*	<.0001*	<.0001*
	STRENGTH					
	ave-y-acc	ave-x-acc	ave-ddt	var-y-acc	var-x-acc	var-ddt
happy/sad	0.0007*	<.0001*	<.0001*	0.0140*	0.0036*	0.0006*
confident/shy	0.3856	0.1465	0.1433	0.277	0.4438	0.7701
rushed/lackad	0.0655	<.0001*	<.0001*	0.0804	0.0004*	0.0045*
	SPACE			**FLOW**		
	%-b-orient	%-a-orient	side-orient	y-dist	theta-dist	var-y-pos
happy/sad	0.0781	0.3757	0.4171	0.0521	0.2093	0.7795
confident/shy	0.4569	0.2523	0.1177	0.1336	0.2547	0.0427*
rushed/lackad	0.1128	0.9645	0.0250*	0.0004*	0.0101*	0.0025*

*statistically significant predictor of manner based on Anova analysis

An ANOVA analysis (Table II) of the motion features finds 16 of the 18 features to be useful in distinguishing one or more of the oppositional pairs (e.g. rushed/lackadaisical). As might be expected, happy/sad and rushed/lackadaisical find clear distinctions in all of the Timing and most of the Strength features. Using Weight, we can now interpret the lower acceleration characteristics of sadness as a lack of force or energy that contrasts to the relatively energetic motions in happiness. Confident/shy is best distinguished by Flow characteristics, likely because most confident paths adhere closely to the boundaries of the path, while shy incorporates more curvature. We may need to improve our Space features, as only one of the three relative orientation features showed significance for the high-contrast oppositional pair data. It is also possible that orientation will become more important and telling during the presence of an interaction.

We calculate machine classification results for each oppositional pairing, then across all data. The diametric qualities of the former pairing makes it more likely that we will be able to answer the Boolean query; is it possible for expressive motion in X, Y and Theta to communicate state?

The latter gives a measure of the potential impact of these features in more complex state space (e.g, perhaps happy but definitely lackadaisical), such as we hope to use with mobile robots in future work.

The average oppositional pair classification accuracy is 75% on the test sets and 100% on the training data. The average classification accuracy for happy/sad was 75%, while confident/shy was 78.5% and rushed/lackadaisical was 71.9%. There are 32 path samples for each oppositional pair (1/3 of the total data), which we divided into eight subsets for cross-validation. Each time, we perform a discriminant analysis using all but one subset, then test the solution on the reserved paths.

The overall classification accuracy across all six manners, again with 8-fold cross-validation (over all 96 samples), is 41.7%, performing best for happy (62.5%), sad (56.3%) and lackadaisical (50%), and worst for rushed (18.8%), as in Table V. While rushed might seem like an intuitive manner, similar paths were present in almost all manners, as the confusion matrix shows. The slow speed characteristics of sad likely aided classification, while happy paths had consistent features in acceleration. Confident and Shy fell in the middle (both 31.3%). To further improve the within-manner results, we could refine our training set to use only the paths found to be most readable to people.

VI. Contrasting Human and Machine Labels

Ultimately, human capacity to read social expression will be the benchmark by which we judge machine capabilities. Thus, to understand the quality of the classification results, we contrast machine classification results to people's ability to label the same set of trajectory samples.

First, we render the expressive motion videos (Fig. 3, right) directly from the collected data files, representing the mockup as a circle with a line to indicate orientation and tracing the boundaries of the motion area. The relative sizes of grey circle to mockup and motion area to motion-tracking table are proportional to the original setup.

Next, we use "master workers" on Amazon Mechanical Turk (MTurk) for two classes of video categorization. MTurk is a useful resource for crowdsourcing such tasks [31]. Master workers have high reliability ratings gained from accurately completing previous MTurk tasks over an extended training period. We present workers with two kinds of categorization tasks: in the first, they must label a video as one of two manners (an oppositional pairing like rushed/lackadaisical); in the second, participants label each video as one of the six manners.

To mirror our oppositional pair classification results, we present participants videos by oppositional pairing (e.g. happy/sad), asking them to assign each video one of two labels. There are 44 samples in each oppositional pair subset, including all 32 path samples of the oppositional pairing being tested (16 of each label) and eight randomly selected videos from the other two subsets. The number of workers doing the labeling varies with the supply and demand of MTurk at the time, but we had 4-5 workers per oppositional pairing (see breakdown in Table III), where two distinct master workers label each path. As an indication of the reliability of the worker labels, we provide data on their agreement in Table III, ranging from 73% to 82%.

TABLE III. MTurk Oppositional Pair Label reliability*

happy/sad (5 workers)	**confident/shy** (5)	**lackad./rushed** (4)
agreed on 73% (32/44) disagreed 27% (12/44)	agreed 82% (36/44) disagreed 18% (8/44)	agreed on 77% (34/44) disagreed 23% (10/44)

*two distinct worker labels for each trajectory

Overall readability within oppositional pairs was ~80% across the three pairings, only slightly higher than machine classification of the same (section V). For each oppositional pair, we analyze labels of all paths intended to be in the manners of that pairing, i.e., 64 samples for each pairing, made up of 32 samples of each manner. The overall readability of Happy vs. Sad was 79.7%, while Confident vs. Shy was 78.1% and Rushed vs. Lackadaisical was 81.3%.

The oppositional pairs are intentionally high-contrast, and their strong readability scores confirm the differences classification also found, confirming the possibility of using motion features to distinguish between states via expressive motion. A more challenging test for our motion features comes when trying to recognize to which of six categories a trajectory belongs, thus we continue with another MTurk categorization set in which workers label each video as one of six categories (happy, sad, rushed, lackadaisical, confident, shy).

There were 96 paths, each labeled by four distinct master workers with 18 workers overall. Because of the way the categorization application is set up on MTurk, that meant running the experiment two times, with two distinct worker labels per video in each set (see Table IV). As an indication of the reliability of the category label, workers agreed on about one third of the labels for each set. Chance would have each label occur 12.5% of the time, which indicates that the expressive motion paths do provide information about the category, but that they are not all equally good at communicating state.

TABLE IV. MTurk All- Manner Label reliability

first set - 9 workers	**second set - 9 workers**
agreed on 36% (35/96) disagreed 64% (61/96)	agreed on 31% (30/96) disagreed 69% (66/96)

The categories over which the worker labels most frequently disagreed (Confident and Rushed 21/192, Happy and Lackadaisical 16/192, Sad and Shy 15/192), may indicate correlations between manners. Of the 32 intended samples of Confident and 32 intended samples of Rushed, 21 were labeled both "Confident" and "Rushed." Regardless of the initial quality of the expressive paths, this suggests overlaps between the motion characteristics we associate with them, for example, both confident and rushed may present as “pushing” across the space, as the Press in Fig. 1.

Readability across all manners is 41.2%. We calculate this from the frequency with which the worker label matches the

TABLE V. MACHINE CLASSIFICATION CONFUSION MATRIX* . ROWS ARE INTENDED PATH CATEGORIES AND COLUMNS ARE PREDICITIONS (N=96)

	Conf predict	Happy	Lackad.	Rushed	Sad	Shy	**Accuracy**
Conf. path	**5**	0	5	1	3	2	**31.3%**
Happy	0	**10**	1	3	0	2	**62.5%**
Lackad.	2	0	**8**	2	3	1	**50.0%**
Rushed	1	6	3	**3**	1	2	**18.8%**
Sad	4	0	2	1	**9**	0	**56.3%**
Shy	4	5	0	2	0	**5**	**31.3%**

*41.7% accuracy overall

TABLE VI. MTURK CONFUSION MATRIX*. ROWS ARE INTENDED PATH CATEGORIES AND COLUMNS ARE HUMAN LABELS (N=386)

	Conf. (label)	Happy	Lackadaisical	Rushed	Sad	Shy	**Readability**
Confident (path)	**23**	8	7	17	2	7	**35.9%**
Happy	13	**23**	11	9	4	4	**35.9%**
Lackadaisical	5	10	**22**	1	10	16	**34.4%**
Rushed	17	6	0	**38**	0	3	**59.4%**
Sad	11	2	10	0	**36**	5	**56.3%**
Shy	6	5	19	3	15	**16**	**25.0%**

*41.2% readability overall

intended expressive manner, given each of 96 videos is labeled four times (384 samples). This readability score is significantly higher than chance, which would be 12.5%. The full data is in Table VI.

Examining the manner readabilities one by one, we find that some are more difficult to successfully represent than others, ranging from 25% to 59% accuracy. Shy appears to be the most challenging to represent (only 25% readable), while Rushed and Sad may have been the easiest (59.4 and 56.3% readable). Confident, Happy and Lackadaisical were in the middle (35.9, 35.9 and 34.4% readable), with just over one-third labeled with their intended category. It is also interesting to contrast what the machine classification versus MTurk workers found clearer or more difficult. Although the overall classification and readability scores were similar, the within-manner differences mean that there is still more for the machine to learn.

The thespian (40.1% readable) and roboticist paths (42.2% readable) also had similar overall readability scores but, again, each group had different readability scores within manners. The thespians created significantly more readable Lackadaisical, Sad and Shy paths while the roboticists created significantly more readable Confident and Rushed. This may be because the former trajectories benefit from more nuanced exploration, while the second set were more readable as straightforward instantiations. Both groups were similarly successful at expressing Happy.

If we compare the overall readability scores to the machine classification results based on Laban Features, we find that the human labels are just about as accurate as the calculated label. This indicates that the Laban features did well in helping distinguish expressive motion paths computationally.

VII. CONCLUSIONS

As collaborative robots move into everyday life, the need for algorithms enabling their acceptance becomes critical. By adhering to the types of expressive motions that people already use and innately understand, we expect to produce robots that are more legible, and hence more acceptable, to the general, untrained populace. As a first step in that direction, we conclude that Laban Effort Features prove useful in distinguishing expressive motions computationally, supporting their potential for use in systems for analyzing and generating expressive motions with X, Y and Theta.

Using these results, we can also track the effort feature characteristics of current mobile robot systems, gaining tools for evaluating the unintentional communications they already project to humans in their environment. For example, CoBot frequently requests help in pressing the elevator button, as it has no arms, but is often ignored. If it has been requesting assistance without orienting toward someone (indirect), while moving at a smoothly accelerating (light) and measured pace (sustained), people may subconsciously feel less special, and infer the task to be low-priority, therefore being less likely to help. By making sure the robot orients toward people (direct) in a more forceful way (strong) while maintaining sustained Timing so as not to seem aggressive, people may be more inclined to help or respond to its query.

In limiting our implementation to three degrees of freedom, we have already begun to discover principles connecting motion trajectories to certain communications of robot (or agent) state. For example, we have found that the clearest communications of shy involve direct paths with hesitations and that timing characteristics are one of the most significant features separating the different manner paths.

We will continue to evaluate and extend our algorithms on the CoBot robot, a system deployed daily on autonomous tasks at Carnegie Mellon. By exploring motion analysis and, eventually, expressive motion generation on a real-world system, we will be able to assess the effect and effectiveness of incorporating the expressive motions discussed in this paper to robot tasks, both in terms of task performance and human attributions.

To do so, we would like to extend our current feature set to contextualized expressive behaviors. We believe this will be particularly relevant to generating robot motions in dynamic human environments. If Laban features are our letters, we can use a higher-level framework to help design situated robot motion behaviors. One possibility is Anne Bogart and Tina Landau's nine Viewpoints [3]. The system specifies contextualized social and environmental behaviors, like the temporal reactions we have to seeing someone unexpectedly, or the influence of room dimensions on the overall shape of our path.

By handing off this set of parameterized Laban Effort System motion features to the research community, we can continue to collectively explore and extend our understanding of this space. These features could be used on any mobile robot without the need for special hardware, and could therefore be immediately and widely evaluated. For example, autonomous car researchers might find value in the projection of friendliness or hurry at a four-way intersection. Moreover, if expressive motion is already communicatory in just position (x, y) and angle (theta), imagine its potential in systems with higher degrees of freedom and multi-modal communication.

Acknowledgment

Special thanks to the National Science Foundation for funding this research.

References

[1] F. Abella, F. Happéb and U. Fritha. Do triangles play tricks? Attribution of mental states to animated shapes in normal and abnormal development. Cognitive Development, Vol. 15:1, pp. 1–16, January–March 2000
[2] J. Becker, S. Breedlove and D Crews. Behavioral Endocinology. MIT Press, Cambridge, 2003
[3] A. Bogart and T. Landau. The Viewpoints Book: A practical guide to Viewpoints and Composition. Theater Communications Group, 2005
[4] C. Breazeal. Designing Sociable Robots. MIT Press, 2004
[5] A. Bruce, et al. Robot Improv: Using drama to create believable agents. In Proceedings of the International Conference of the Association for the Advancement of Artificial Intelligence, 2000.
[6] F. Castelli, et al. Movement and Mind: A Functional Imaging Study of Perception and Interpretation of Complex Intentional Movement Patterns. NeuroImage, Vol 12: 314-325, 2000
[7] J. Chen, W. Lin, K. Tsai and S. Dai. Analysis and Evaluation of Human Movement based on Laban Movement Analysis. Tamkang Journal of Science and Engineering, Vol. 14, No. 3, pp. 255-264, 2011
[8] A. Engel, et al. How moving objects become animated: The human mirror neuron system assimilates non-biological movement patterns. Social Neuroscience, 3:3-4, 2008
[9] V. Gazzola, et al. The anthropomorphic brain: the mirror neuron system responds to human and robotic actions. Neuroimage 35.4 1674-1684, 2007
[10] M. Hauser, N. Chomsky and W. Fitch. The Faculty of Language: What Is It, Who Has It, and How Did It Evolve? Science: 298 (5598), 1569-1579, November 22, 2002
[11] J. Harris and E. Sharlin, Exploring emotive actuation and its role in human-robot interaction. In Proceedings International Conference on Human-Robot Interaction, 2010
[12] F. Heider and M. Simmel. An Experimental Study of Apparent Behavior. The American Journal of Psychology, Vol. 57(2), 1944
[13] F. Heider. Social perception and phenomenal causality. Psychological Review, Vol 51(6): 358-374, 1944
[14] W. Ju and L. Takayama. Approachability: How people interpret automatic door movement as gesture. Int'l Journal of Design, 2009
[15] H. Kelley. The processes of causal attribution. American Psychologist, Vol 28(2): 107-128, Feb, 1973
[16] H. Knight. Eight lessons learned about non-verbal interactions through robot theater. In Proceedings International Conference on Social Robotics, Amsterdam Netherlands, November 2011
[17] R. Laban. Modern Educational Dance. Macdonald & Evans, 1963
[18] A. LaViers and M. Egerstedt. Style Based Robotic Motion. In Proceedings of the American Control Conference, 2012
[19] A. Lim. Design and Implementation of Emotions for Humanoid Robots based on the Modality-independent DESIRE Model. Master's Thesis at Kyoto University, 2012
[20] D. Lu and W. Smart. Human Robot Interaction as Theatre. In IEEE International Symposium on Robot and Human Communication, 2011
[21] M. Masuda, S. Kato and H. Itoh. Laban-Based Motion Rendering for Emotional Expression of Human Form Robots. Chapter in Knowledge Management and Acquisition for Smart Systems and Services Lecture Notes in Computer Science, Volume 6232: 49-60, 2010.
[22] B. Mutlu and J. Forlizzi. Robots in organizations: the role of workflow, social, and environmental factors in human-robot interaction. Proceedings Int'l Conf on Human-Robot Interaction, 2008
[23] Picard, R. Affective computing. MIT Press, Cambridge, MA, 1997
[24] J. Rett and J. Dias. Human-robot interface with anticipatory characteristics based on Laban Movement Analysis and Bayesian models. In IEEE 10th Int'l Conf on Rehabilitation Robotics, 2007
[25] J. Rett, L. Santos, and J. Dias. Laban movement analysis for multi-ocular systems. In Proceedings IEEE/RSJ International Conference on Intelligent Robots and Systems, 2008.
[26] M. Saerbeck and C. Bartneck. Perception of affect elicited by robot motion. In Proc. Int'l Conference on Human-Robot Interaction, 2010
[27] L. Santos and J. Dias. Human-Robot Interaction: Invariant 3-D Features for Laban Movement Analysis Shape Component. In Proceedings of Int'l Conference on Robotics and Applications, 2009
[28] M. Sharma, et al. Communicating affect via flight path: exploring use of the laban effort system for designing affective locomotion paths. In Proceedings Int'l Conference on Human-Robot Interaction, 2013
[29] K. Strabala, M. Lee, A. Dragan, J. Forlizzi, S. Srinivasa, M. Cakmak, and V. Micelli Towards Seamless Human-Robot Handovers. Journal of Human-Robot Interaction, 2013
[30] L. Takayama, D. Dooley and W. Ju. Expressing thought: Improving robot readability with animation principles. ACM/IEEE Proceedings International Conference on Human-Robot Interaction, 2011
[31] A. Kittur, E. H. Chi, and B. Suh. "Crowdsourcing user studies with Mechanical Turk." Proceedings of the SIGCHI conference on human factors in computing systems. ACM, 2008.