A Formal Approach to the Automatic Generation of Ballet Phrases

Amy LaViers, Yushan Chen, Calin Belta, and Magnus Egerstedt

Abstract— **This paper presents a novel example where formal methods can be used to generate a type of human behavior. Drawing inspiration from classical ballet, poses are cast as discrete states and movements as the transitions between these states. Using Linear Temporal Logic (LTL), we are able to further constrain the set of possible sequences through the transition system and thus prevent it from evolving through a sequence of states that is physically impossible or aesthetically undesirable. Thus, a given movement style is encoded in the availability of transitions at each state, and the dynamics of a complex physical trajectory are abstracted as a system which moves between these states. Our overarching objective is to facilitate subtle degrees of control over systems through a useful parameterization for human movement. Such subtleties are required, for example, by humanoid robots to interact in and analyze a social and aesthetically driven world.**

I. INTRODUCTION

Understanding human behavior is a goal in many disciplines: behavioral psychology, neuroscience, robotics, computer vision, artificial intelligence, athletics, and dance, to name a few. Thus, this paper may be understood in terms of a larger dialogue about how to represent human movement with a concise parameterization which is comprised of two components: recognition and interpretation. Such a parameterization would provide a baseline for observational data in psychology and neuroscience; allow robots to recognize and imitate human behavior; give new insight to the phenomena of human athletic achievements; and give tools for analysis of dance sequences.

Formal methods for verification can play a surprising role in this discourse. We will ensure that our system behaves according to certain guidelines in a framework where verification methods become an integral part of the system model itself. Here, our system model is one describing human movement and our specifications control this model, preventing it from entering disallowed and undesirable states. This allows us to produce sequences of movement in a particular style with particular emotive effects.

Namely, in this paper we draw inspiration from the performing art of dance where the particular style and certain emotive interpretation of movement sequences are tightly coupled; however, for this to be possible we need new tools for representing human movement. From the artistic perspective, this induces quite an imbalance: a musician works with notes arranged in octaves, chords, and scales; an artist paints with a palette of colors known to span the space of light that is perceivable to humans; and writers knit nouns and verbs into clauses, sentences, and paragraphs. Movements are certainly classified and named (pirouettes and spiral 3-step turns), and some words from these other artistic disciplines have filtered into the dance world to describe conglomerate structures of motion (phrases and movements). Certainly there is organization in artistic movement - but can we as engineers use analytical tools to justify any of this organization?

Such a justification would cascade easily into the field of robot recognition and imitation by providing an internal model of behavior for such systems as humanoid robots. The emotive content found in the highly formalized movements of a specific genre of dance is also present in the casual, pedestrian gestures of everyday body language. An abstraction based in such aesthetic principles could one day provide a link between sensing and actuation that is plastic and redundant, as human behavior. As humanoid robots are expected to operate in dynamic environments amid social human beings, deterministic rules at each junction of decision-making often conflict due to an unforeseen order of environmental events. A robust and socially-aware resolution should guide such a system to choose from one of many solutions to a given problem based both on hard, task-related constraints and encouraged movement principles.

A number of human motions have been successfully encoded using dynamic motion primitives [6], [16] and labeled with tasks, such as reaching, drawing, and arguably walking. These primitives, or *movemes* [4], are designed to produce rich and complex human-like motions through systematic, temporal composition. Traditionally, these primitives are obtained from empirical data, e.g., collected using motion capturing devices, that is segmented (often by hand, [3], but with progress towards automatic segmentation, [8]) into appropriate motion chunks and stored in a motion library [16]. But full fledged behaviors do not follow predictable, continuous trajectories as humans constantly make discrete decisions and may abruptly change behavioral modes, and there is no clear method for stitching these chunks of stereotyped trajectories together. Hence, a drawback of this representation is that it cannot interpret long sequences of movement.

Rather than basing the models on such patterns mined from empirical data, in this paper, we draw inspiration from a highly constrained set of motion patterns, namely, those found in classical ballet. The formal principals of movement

Amy LaViers and Magnus Egerstedt are with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA, E-mail: alaviers@gatech.edu, magnus@gatech.edu

Yushan Chen is with the Department of Electrical and Computer Engineering, Boston University, Boston, MA 02215, USA, E-mail: yushanc@bu.edu

Calin Belta is with the Department of Mechanical Engineering, Division of Systems Engineering, and the Bioinformatics Graduate Program, Boston University, MA 02215, USA, E-mail: cbelta@bu.edu

organization in basic ballet can be distilled into a few simple rules, making this an excellent candidate movement genre since the execution of these rules produces such highly sophisticated and complex motions. Using the ability to *produce* movement as an initial metric for successful representation, we develop machinery to generate coherent movement phrases - in the style of classical ballet. That machinery comes in two general parts: modeling and control.

On the part of modeling, we extract key poses for the experience of classical ballet and dub them the states of a transition system. Transitions between these states model the movements between poses. Beyond consisting of formal movements as discussed above, ballet simplifies the process of instantiating such a system further with its canonical warm-up routine, the *barre*. The motions in this warm-up constitute basic building blocks of ballet as they are the primary tool used to warm and train muscle patterns at the start of every ballet class. Thus, our transition system will, in some sense, span the allowable trajectories of basic ballet technique (in terms of discrete, untimed events or transitions). This is the subject of [15] and is reviewed for our context in Sec. II.

Yet, how do we select among these many available pose sequences for one(s) that are both physically feasible and aesthetically meaningful? This problem requires discrete, non-physical control methods which can interface with the continuous trajectories of systems (such as the human body) and provide discrete decision-making power. In this paper as in [14], we use linear temporal logic (LTL) statements to provide formal and verifiable system specifications which limit our system output to that of our desired behavior (here, sequences in the style of classical ballet) and thus, begin to enumerate fundamental rules and secondary, aesthetic principles which govern this instantiation of human behavior. This is the subject of Sec. III.

We leverage the framework developed in [14] and [7] to bring the power of our model to life in Sec. IV. More specifically, outputs that satisfy the hard constraints and encouraged movement qualities are generated by enforcing temporal logic specifications and assigning rewards to desired states and transitions; these specifications guide the selection of the output sequence (which may not be unique). We provide a concluding impression of our work in Sec. V.

II. A Discrete Model for Behavior

A concept central to ballet's doctrine is that the barre trains and safely warms the muscle groups critical to the correct execution of the freestanding, full-fledged movements that comprise the second portion of class and performances. Hence, these canonical exercises contain the poses and allowable trajectories through them that construct the more rich and expressive remaining vocabulary of ballet. The term "barre" refers to the physical hand-railing, or bar, that dancers hold on to in order to balance during the warm-up. Exercises typically focus on one side of the body and are repeated twice in order to work both sides of the body. The working leg is the leg that is away from the bar and is more active than the other (standing or supporting) leg during a given iteration of the exercise. For now, we limit our focus to the positions of the working leg in the body's coronal plane.

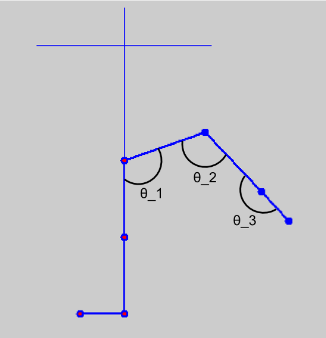

Fig. 1. The discrete states are interpreted as poses corresponding to three joint angles: hip, knee, and ankle and are restricted to the body's coronal plane. Amy don't forget to insert the ballerina photo with the same pose and modify the caption accordingly

As such, using the notion of a "working leg," we define 10 states which correspond to poses constructed from a triplet of joint angles, the hip, leg, and ankle configuration for this working leg, as seen in Figure 1. These poses represent shapes critical to the experience of ballet. They are chosen from goal positions at the barre; as such, they are highly recognizable snapshots from the vocabulary of ballet and are found in more complex movements used for choreography.

The state transitions are given by events modeled as the movements from the barre exercises. These movements are listed in the table below, together with the transition labels (assigned according to the the first four letters of the name of the movement from which the event was derived). Additionally, we distinguish two transitions for each movement listed in the table using a subscript to indicate an *in* and *out* variant. The variants stem from the fact that each movement has a goal end pose; during a movement sequence, the dancer system is either on its way out to the goal pose or on its way back in, to a previous state.

Movement	Transition Label
plié	plie
relevé	rele
battement tendu	tend
degajé	dega
coupé	coup
frappé	frap
grand battement	gran
possé	poss
battement	batt
développé	deve

Each of the above movements has a starting and ending pose. Our transition system connects these movements via the appropriate poses in Fig. 2. This system represents allowable movement for one leg only. That is, the resulting language, i.e., the set of all trajectories that start at the initial state produce feasible barre routines. Some strings might be somewhat unusual (perhaps with movements repeated a strange number of times), but they will certainly be recognizable as being in the style of classical ballet.

Formally, the transition system in Fig. 2 that models the motion of the right leg of a ballerina during a simple ballet barre is defined as

$$T_R = (Q_R, q_{0_R}, \rightarrow_R, \Pi_R, h_R), \tag{1}$$

where

(i) $Q_R = \{q_1^R, \ldots, q_{10}^R\}$ is the finite set of states;
(ii) $q_{0_R} = q_2^R$ is the initial state representing the initial pose;
(iii) $\rightarrow_R \subseteq Q_R \times Q_R$ is the reflexive transition relation (i.e., each state has a self transition);
(iv) $\Pi_R = \{p_1, \ldots, p_9\} \cup \{\text{Roffground}, \text{Rcoronal}\}$ is a finite set of atomic propositions;
(v) $h_R : Q_R \mapsto 2^{\Pi_R}$ is a satisfaction (output) map, where state q_i^R satisfies the set $h_R(q_i^R)$ of propositions from Π_R as shown in Fig. 2.

Correspondingly, we define the transition system that models the motions of the left leg transition system to be:

$$T_L = (Q_L, q_{0_L}, \rightarrow_L, \Pi_L, h_L), \tag{2}$$

where each component is defined as in Eq. 1; that is, items (i) - (v) are identical for the left leg's transition system with "R" replaced with "L" as necessary. In particular, $\Pi_L = \{p_1, \ldots, p_9\} \cup \{\text{Loffground}, \text{Lcoronal}\}$.

The atomic propositions are statements which are either true or false about every state of our system; in the next section we will use the power of temporal logic to watch the evolution of our system in terms of these statements of particular interest. The pose propositions $\{p_1, \ldots, p_9\}$ for each leg are satisfied only when that leg is in the corresponding pose. The additional proposition Roffground represents "the right leg of the robot is off the ground," while Rcoronal represents "the right leg is in a coronal extension away from the body" - the additional propositions for the left leg have similar meaning.

In order to demonstrate how a sample path through the system works, consider for example a développé; this movement is found both in barre exercises and more complex ballet movement phrases. A développé is the action when the working leg's foot is moved to the ankle, then the knee, then extends from the body so that the leg is parallel to the floor. However, lifting the foot to the ankle or knee (without, for example, any extension to follow) are allowable movements called coupé and possé, respectively. Thus, to keep our transitions (and trajectories) uniquely defined, we model this as three separate events for the working leg: $coup_o$, $poss_o$, $deve_o$. Next, the dancer performs a closing movement where the foot remains extended from the body and the leg is lowered till the foot is returned to the starting stance. This is modeled as the event $batt_i$ - the transition from pose 8 directly to pose 2 with a label that corresponds to the *in*-trajectory of a battement (a simpler movement that looks like a high straight-legged kick). The events $coup_i$, $poss_i$, and $deve_i$ are also defined, that is, the reverse pose transitions are allowed and used for more complex movements. For the stationary, supporting leg, the transition is simply a repeated self-loop event which corresponds to no movement. This leg may be in state 1, 2, or 3 as we will see in Sec. IV.

Our next task is to compose two such systems in a way that ensures the resulting system produces physically and aesthetically allowable movement for two legs. We first compose these systems using a synchronous product; this composition is liberal and naive because it incorporates every available joint state and transition (some of which are no longer physically possible and/or aesthetically desirable) without taking into consideration whether the composed system is still appropriate. In the next section, we whittle away which of these joint states and transitions the system will inhabit.

More formally, the synchronous product of the two transition systems T_L and T_R, denoted as $T_L \otimes T_R$, is defined as

$$T_P = (Q_P, q_{0_P}, \rightarrow_P, \Pi_P, h_P), \tag{3}$$

where

(i) $Q_P \subseteq Q_L \times Q_R$;
(ii) $q_{0_P} = (q_{0_L}, q_{0_R})$;
(iii) $\rightarrow_P \subseteq Q_P \times Q_P$ is defined by $(q, q') \in \rightarrow_P$ if and only if $q \neq q'$, $(q_L, q_L') \in \rightarrow_L$ and $(q_R, q_R') \in \rightarrow_R$, where $q = (q_L, q_R)$ and $q' = (q_L', q_R')$;
(iv) $\Pi_P = \Pi_L \cup \Pi_R \cup \{\mathit{Spose}\}$;
(v) $h_P : Q_P \mapsto 2^{\Pi_P}$ is defined by *(1) Spose* $\in h_P(q_L, q_R)$ if and only if $q_L = q_R$, *(2)* for all $\pi \in \Pi_L \cup \Pi_R$, $\pi \in h_P(q_L, q_R)$ if and only if $\pi \in h_L(q_L)$ or $\pi \in h_R(q_R)$.

The synchronous product allows for both synchronous and asynchronous transitions of the original single leg systems since the reflexive transitions defined in Eq. 1 (iii) establish self-loops at each state. Hence, the composed system captures all possible transitions that can appear in either of the single leg systems. For the composed system, we add another atomic proposition, "Spose," which is satisfied by all joint-states at which both legs have the same configuration or pose.

III. LINEAR TEMPORAL LOGIC FOR BEHAVIOR SPECIFICATION

Recently, there has been an increasing interest in developing computational frameworks that enable rich specification languages for robotics. In particular, temporal logics, such as Linear Temporal Logic (LTL) and Computation Tree Logic (CTL) have been suggested as motion specification languages [23], [18], [10], [13], [7]. The use of such logics allows for a large spectrum of specifications which include: choice of a goal ("go to either A or B"); convergence to a region ("reach A eventually and stay there for all future times"); visiting targets sequentially ("reach A, and then B, and then C"); surveillance ("reach A and then B infinitely often"); the satisfaction of more complicated temporal and logic conditions about the reachability of regions of interest ("Never go to A. Don't go to B unless C or D were visited"). Such

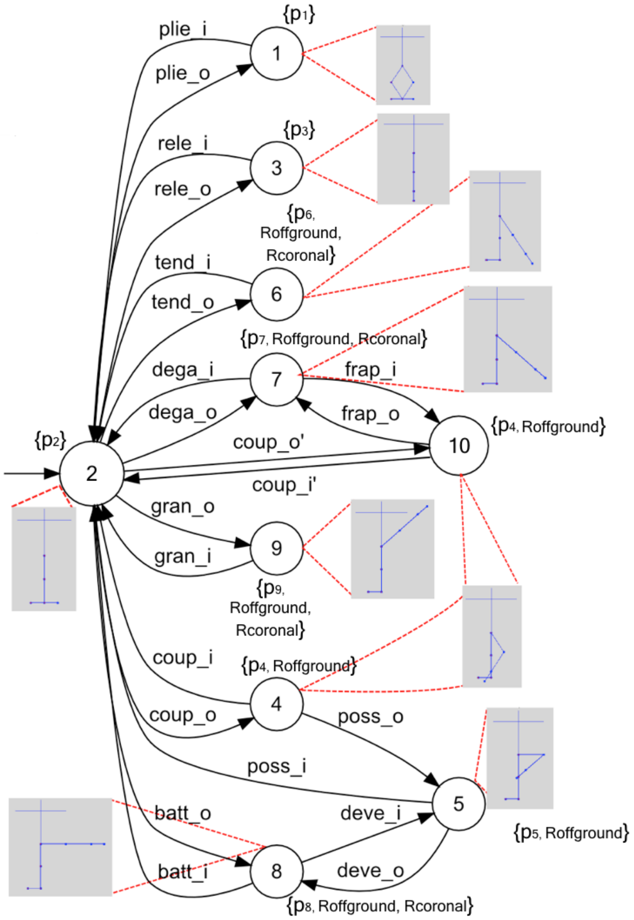

Fig. 2. A transition system which models the working leg (right leg) of a dancer during a ballet barre exercise. The atomic propositions $\{p_1, \ldots, p_9\}$ correspond to poses defined by three joint angles: hip, knee, and ankle. Images of the poses are shown for clarity. For simplicity, we used i to denote state q_i^R and we omitted the self-loops at each state. Note that states 10 and 4 satisfy the same propositions, i.e., correspond to the same pose. We differentiate these two related states based on whether the motion of the leg is to remain low (below the hips) or high (at or above hip level) before returning to a neutral state and beginning the next movement.

robot motion planning and control objectives are achieved based on algorithms inspired from model checking [5] and temporal logic games [17].

Here, we use LTL to express robotic tasks that include an aesthetic component, e.g., "go to goal with grace" or, more specifically, "move in the quick, staccato style of allégro ballet." These tasks have an objective component, such as "take ten steps," and also incorporate subjective qualities like intention and aesthetics, such as "make ten movements that give the impression of being happy." Thus, in this paper, through a particular approach to modeling and formal synthesis, we begin to quantify subjective qualities, which are a significant missing link in human-like robot interaction, by scripting them in the established language of LTL.

LTL formulas are built from a set of atomic propositions Π, standard Boolean operators $\neg$ (negation), $\vee$ (disjunction), $\wedge$ (conjunction), and temporal operators $\mathbf{X}$ (next), $\mathcal{U}$ (until), $\mathbf{F}$ (eventually), $\mathbf{G}$ (always). The semantics of LTL formulas are given over infinite words generated by transition systems, such as T_R, T_L, and T_P defined in Sec. II. For simplicity, in what follows, we will denote a general transition system and its set of propositions by T and Π, respectively. A word of T is an infinite sequence over the power set of Π that satisfies the transition relation of T. For example, $\{p_2\}\{p_7, Roffground, Rcoronal\}\{p_4, Roffground\}\{p_2\}\ldots$ is the word generated by a run of T_R that starts at q_2^R, goes to q_7^R, next to q_{10}^R, and then keeps looping among these three states.

A word satisfies an LTL formula ϕ if ϕ is true at the first position of the word; $\mathbf{X}\phi$ states that at the next state, an LTL formula ϕ is true; $\mathbf{F}\phi$ means that ϕ eventually becomes true in the word; $\mathbf{G}\phi$ means that ϕ is true at all positions of the word; $\phi_1 \mathcal{U} \phi_2$ means ϕ_2 eventually becomes true and ϕ_1 is true until this happens. More expressivity can be achieved by combining the above temporal and Boolean operators (several examples are given later in the paper). For example, the word given above satisfies the following LTL formulas: $\mathbf{X}p_7$ - p_7 is true at the next state, $\mathbf{F}p_4$ - p_4 will be true eventually, $\mathbf{GF}(p_4 \wedge Roffground)$ - p_4 and $Roffground$ are always eventually true.

An LTL formula can be represented in an automata-theoretic setting as a *Büchi Automaton*. These automata accept infinite length words which satisfy the corresponding LTL formula. For each LTL formula ϕ over Π, there exists a Büchi automaton that accepts exactly the language over 2^Π satisfying ϕ. There exist off-the-shelf algorithms for translation of LTL formulas to Büchi automata [9], [19].

Since the one leg transition system does not contain all the information about the physical capabilities of the robot (T_R does not tell us anything about forces required for jumping but instead accepts the correct sequence of leg positions during the jump), it is entirely possible that the product T_P accepts runs that are not physically possible to execute and that are not within the range of our goal aesthetic. Hence, we formulate LTL specifications (in the form of Büchi automata) that enable our system to make discrete decisions about viable trajectories (that is, accept or reject various runs through the transition system) - where "viable" is defined in terms of both physical and aesthetic constraints.

Specifically, we want to prevent the robot from executing any physically infeasible runs. In addition, we are interested in applying aesthetic conditions to the accepted runs of T_P in order to make them adhere to our chosen dance style. For example, we may disallow a list of two-legged body poses which are perhaps considered ugly as judged by the metric of ballet; often, these are asymmetrical poses or poses which cannot be seen from the audience's distant perspective. Even given a system which only produces movements in the style of ballet, we may further restrict our output so as to only produce a specific type of movement phrase within the genre that, for example, is typified by more frequent use of certain movements.

Thus, to define our problem, we assume that our system is required to satisfy the physical constraints of a bipedal geometry and the aesthetic requirements of basic, classical ballet. We consider two types of specifications to express the restrictions: (1) *hard* specifications and (2) *soft* specifications. A hard specification incorporates all the physical constraints and aesthetic requirements that the robotic system must satisfy while a soft specification captures certain additional aesthetic requirements that the robotic system is encouraged to achieve.

We use LTL formulas to represent the hard specification. The physical restrictions of the robot and the aesthetic requirements of ballet can be easily translated to LTL formulas as in the following two examples:

1) "stand up infinitely many times" can be converted to

$$\mathbf{G}\ \mathbf{F}\ (\text{Spose} \wedge p_2)$$

where Spose means both legs are in the same pose and p_2 is the standing up pose.

2) "always bend knees and then straighten them before having two legs off ground" can be converted to

$$\mathbf{G}\ (\neg((p_1 \wedge \text{Spose}) \wedge \mathbf{X}\ (p_2 \wedge \text{Spose})) \quad \rightarrow \\ \mathbf{X}\ \mathbf{X}\ (\neg(\text{Loffground} \wedge \text{Roffground})))$$

where p_1 is the bent knee pose and p_2 is the standing pose. This is a sensible example as rarely will a dancer, or any person, jump without first bending their knees in order to use this large joint for extra force.

The soft specifications tweak the viable output sequence by specifying aesthetic conditions that the system is encouraged (instead of forced) to achieve. We define our first type of soft specification as in [14], denoted by $\mathcal{S}$, as a collection of subsets of Π_P: $\mathcal{S} \subset 2^{\Pi_P}$. The collection $\mathcal{S}$ models our desire for the robot to prefer satisfying certain propositions more than others. Here, we will also define a second type of soft specification that encourages movements - transitions between poses.

For example, a specific type of ballet phrase, known as *allégro*, is typified by movements that produce a quick and upbeat dynamic quality. One movement that is common in an allégro phrase is entrechat (a little jumpy move where the dancer beats his or her feet back and forth). In our system, an infinitely repeated entrechat would correspond to the word

$$\{p_2\}\{p_1, \text{Spose}\}\{p_2, \text{Spose}\}\{p_3, \text{Spose}\}\{p_2, \text{Spose}\}\{p_1, \text{Spose}\}\{p_2\}...,$$

which is generated by a run of T_P

$$(q_2, q_2)(q_1, q_1)(q_2, q_2)(q_3, q_3)(q_2, q_2)(q_1, q_1)(q_2, q_2)....$$

Consequently, we want to encourage the system to select the pose that satisfies $\{p_3, \text{Spose}\}$ (i.e., state (q_3, q_3)) more often for a higher appearance frequency of entrechat. In this example, the soft specification is $\mathcal{S} = \{p_3, \text{Spose}\}$, and we will encourage the system to visit state (q_3, q_3) more frequently in the output sequence.

To achieve the soft specification, we introduce rewards (positive real numbers) to the states indicated by $\mathcal{S}$. Alternatively, if we want to encourage specific movements, which we think of as transitions between poses, we assign rewards to transitions. We make the natural assumption that, at each time instant, the system can "foresee" only a few steps ahead (analogous to a dancer that is performing a free, unchoreographed solo who is planning just a few steps ahead at a time). The rewards are assigned in a time varying fashion and the appearance and disappearance of rewards and their values are randomized.

Upon visiting a state or taking a transition with a reward, the system collects the corresponding reward. We aim to maximize the collected rewards; hence, the system is encouraged to visit some states or take some transitions more frequently than others. Specifically, we satisfy the LTL formula and maximize the collected reward. Since the system produces infinite trajectories, it does not make sense to look for a run that maximizes the total collected rewards (which can be infinite). Rather, we design a (local, real-time) receding-horizon controller and find a run that maximizes rewards collected based on local information obtained at the current state.

To achieve this, we use an approach similar to that in [7], where a receding-horizon controller was designed. More specifically, we use a measure of progress towards satisfying the formula. If the controller is designed to always increase this progress as defined by our measure, then we can show that the LTL formula is satisfied. The proposed approach relies on 1) the construction of the product automaton between the transition system and the Büchi Automaton corresponding to ϕ and 2) an assignment of a suitable cost to each state of the resulting product automaton. This cost assignment is computed off-line once and then is used on-line with the real-time controller. The cost assignment is designed so that when used in conjunction with our proposed control strategy, an accepting state on the product automaton will be reached in a finite number of transitions. If this is repeated infinitely many times, the acceptance condition of the product automaton is enforced.

IV. GENERATING BALLET BEHAVIORS

In this section, we show that our solution generates movement phrases, which we represent by a series of poses, within the genre of classical ballet (satisfying the physical restrictions of bipedal geometry and the aesthetic requirements of ballet) and that satisfies additional constraints of specified style. We present two case studies: one demonstrating the basic difference between our hard and soft specifications and the second demonstrating greater faculty of our soft specifications by producing more subtle style variation in the output sequences.

The hard specification is given according to a collection of physical and aesthetic rules which translate to LTL formulas as following:

1) We first ensure that no physically infeasible joint-poses are entered by the system. Such poses include (i) those

poses which imply impossible jumping positions, i.e., state (q_6, q_7), (ii) poses impossible to hold because they are neither jumping poses nor feasible standing poses, i.e., both legs attempting to touch the ankle of the other state (q_4, q_4), and (iii) poses that the geometry of the biped's hips will not easily allow and thus form awkward poses which are ugly in the eyes of traditional ballet choreography, i.e., state (q_1, q_3). Thus, a set of poses are disallowed:

(i) *"always avoid both legs off ground except poses in which two legs are in the same position, and poses satisfying $p_7 \wedge p_{10}$ or $p_5 \wedge p_8$"*

$$\mathbf{G}\ \neg(\text{Roffground} \wedge \text{Loffground} \wedge \\ (\neg\text{Spose}) \wedge \neg(p_7 \wedge p_{10}) \wedge \neg(p_5 \wedge p_8))$$

(ii) *"always avoid both legs in poses p_4, p_6, and p_{10}"*

$$\mathbf{G}\ \neg(\text{Spose} \wedge (p_4 \vee p_6 \vee p_{10}))$$

(iii) *"always avoid both legs in poses satisfying $p_1 \wedge p_2$, $p_1 \wedge p_3$, or $p_2 \wedge p_3$"*

$$\mathbf{G}\ \neg((p_1 \wedge p_2) \vee (p_1 \wedge p_3) \vee (p_2 \wedge p_3))$$

2) Next, we prevent the system from switching to pose which, given a current pose, would cause a biped to fall over. We restrict transitions from states that have only one leg supporting the body; these include (i) jumping from a leg in relevé - this would require just the toe joint of one leg to lift the entire body weight into the air, i.e., (q_3, q_8) to (q_8, q_5), (ii) extending the supporting leg when it is in plié, i.e., (q_1, q_5) to (q_8, q_5), and (iii) extending the supporting leg when it is on flat, i.e., (q_2, q_5) to (q_8, q_5). Thus, a collection of sequences of poses are disallowed:

   (i) *"when one leg is in pose p_1 and the other leg is already off ground, always avoid lifting both legs off ground"*

$$\mathbf{G}\ (\neg((p_1 \wedge (\text{Roffground} \vee \text{Loffground})) \\ \rightarrow \mathbf{X}\ (\text{Roffground} \wedge \text{Loffground})))$$

   (ii) *"when the right leg is in pose p_2 and the left leg is currently off ground, always avoid performing coronal extensions using the right leg without putting down the left leg first"*

$$\mathbf{G}\ ((\text{p}_2 \wedge \text{Roffground}) \\ \rightarrow \mathbf{X}\ \neg(\text{Roffground} \wedge \text{Lcoronal}))$$

   (iii) *"when the left leg is in pose p_2 and the right leg is currently off ground, always avoid performing coronal extensions using the left leg without putting down the right leg first"*

$$\mathbf{G}\ ((\text{p}_2 \wedge \text{Loffground}) \\ \rightarrow \mathbf{X}\ \neg(\text{Loffground} \wedge \text{Rcoronal}))$$

3) Finally, we ensure that the system puts both legs standing on the ground once in a while. This is phrased as a requirement to visit the standing pose, state (q_2, q_2), infinitely often:

   "visit the pose where both legs are in pose p_2 infinitely many times"

$$\mathbf{G}\ \mathbf{F}\ (\text{Spose}\ \wedge p_2)$$

The hard specifications ensure that the robot only selects sequences which are in line with physical and aesthetic restraints. Now, we will use our soft specifications and our optimal controller to produce sequences which satisfy more subtle style specifications. Namely, we can specify certain poses and/or movements (which we can think of as the transitions between poses or a series of desired poses[1]) that we want the output sequence to favor. This results in an increased frequency of these poses or movements in the output sequence without breaking the rules encoded in the hard specifications.

The type of ballet that we first reproduce is one in the style of *allégro*. Allégro is typified by quick and upbeat dynamic quality; while this type of quality is not directly a part of our discrete model, we can produce runs which contain the movements (series of poses) that ballet dancers use to produce such a movement style. Namely, there are three moves we want to see in the output: *entrechat*, *assemblé*, and *jeté*.

In Sec. III we explained the soft specification that would result in a greater number of entrechats in the output sequence. Note that entrechat is a symmetrical move of the body. Assemblé is a jump which is asymmetrical about the body. In our system, this would also correspond to two words of T_P:

$$\binom{p_2}{\text{Spose}}\begin{pmatrix} p_1 & , \text{Roffground} \\ p_7 & , \text{Rcoronal} \end{pmatrix}\binom{p_2}{\text{Spose}}\binom{p_3}{\text{Spose}}\binom{p_2}{\text{Spose}}\binom{p_1}{\text{Spose}}\binom{p_2}{\text{Spose}}$$

or

$$\binom{p_2}{\text{Spose}}\begin{pmatrix} p_1 & , \text{Loffground} \\ p_7 & , \text{Lcoronal} \end{pmatrix}\binom{p_2}{\text{Spose}}\binom{p_3}{\text{Spose}}\binom{p_2}{\text{Spose}}\binom{p_1}{\text{Spose}}\binom{p_2}{\text{Spose}}.$$

Hence, for a output sequences with greater frequency of entrechat and assemblé, we encourage the system to enter states satisfying $\binom{p_3}{\text{Spose}}$ (i.e., state (q_3, q_3)) more often for the higher appearance frequency of entrechat, and to enter states satisfying $\begin{pmatrix} p_1 & , \text{Roffground} \\ p_7 & , \text{Rcoronal} \end{pmatrix}$ or $\begin{pmatrix} p_1 & , \text{Loffground} \\ p_7 & , \text{Lcoronal} \end{pmatrix}$(i.e., states (q_1, q_7) and (q_7, q_1)) more often for a higher appearance frequency of assemblé.

There are two variants of allégro, *grand allégro* and *petit allégro*. We intend to reproduce *grand allégro*, which is distinguished from petit allégro by its large sized movements. We use jeté (a split-leg leap, represented by the proposition $\begin{pmatrix} p_8 & , \text{Roffground} & , \text{Rcoronal} \\ \text{Spose}, & \text{Loffground} & , \text{Lcoronal} \end{pmatrix}$ in our system) to characterize our outputs for grand allégro; namely, we enforce recurrence of a pose for *grand jeté* in grand allégro.

Hence, the soft specification can be summarized as

$$\mathcal{S} = \left\{ \begin{matrix} \binom{p_3}{\text{Spose}} & \begin{pmatrix} p_1 & , \text{Loffground} \\ p_7 & , \text{Lcoronal} \end{pmatrix} \\ \begin{pmatrix} p_1 & , \text{Roffground} \\ p_7 & , \text{Rcoronal} \end{pmatrix} & \begin{pmatrix} p_8 & , \text{Roffground} & , \text{Rcoronal} \\ \text{Spose}, & \text{Loffground} & , \text{Lcoronal} \end{pmatrix} \end{matrix} \right\}. \quad (4)$$

---

[1] The relationship between movements and poses is also a function of the structure of our transition system. Because transitions are sparse, encouraging a goal pose may be enough to encourage a specific movement or transition; however, sometimes, if a pose has multiple transitions leading away from it, we need to encourage transitions directly.

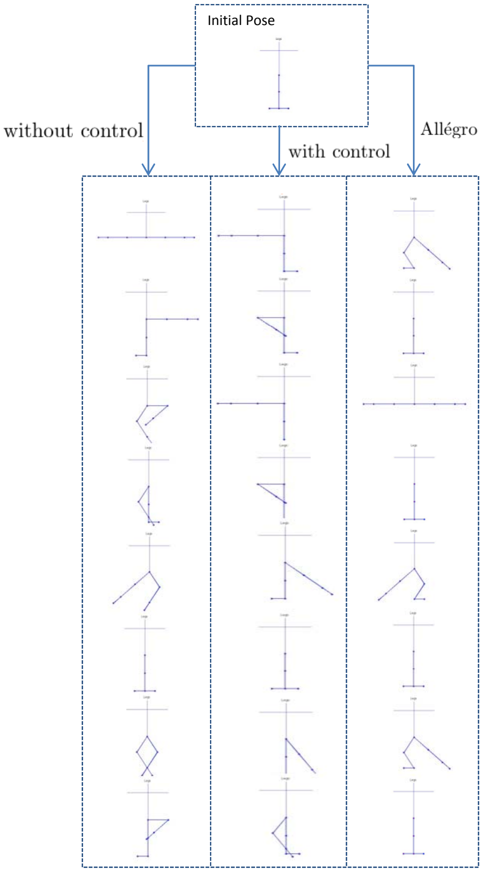

Fig. 3. Three sample sequences demonstrate the results of our control method. The left sequence is an example of a nonphysical (and thus also unaesthetic in some sense) series of poses. The middle sequence is an example of a series of poses endowed with only the hard specifications. Hence, the physical restrictions and aesthetic requirements are satisfied; however, the propositions included in Eqn. 4 do not appear frequently. The right-hand sequence is an example of a series of poses endowed with both the hard and soft specifications (thus is it an allégro phrase). We can see the high occurrence of states (p_1, p_7), (p_7, p_1), and (p_8, p_8). Note that state (p_3, p_3) also occurs frequently in the produced sequence (52 steps in total), but not in the first 9 steps that are shown here. We plot the occurrence rates of these states in Fig. 4. [14]

The dance produced by our method is animated using a MATLAB script; snapshots of three illustrative sample cases are provided in Fig. 3. These animations have been evaluated by a trained eye and found to be a reasonable initial model of ballet technique. Clearly, significant changes take place between the distinct cases of systems that we animated. In Fig. 4, we show the occurrence rates of the states satisfying the soft specifications in different sample cases.

Next, we provide a case study that demonstrates how small changes in style specifications can lead to a sequence with a new emotive effect. Ultimately, it is this type of coupling between choreographic structure and emotive effect that this framework aims to crystallize. This case study demonstrates the effect of changing just one aspect of a specification.

We encourage the movement called développé, also discussed in Sec. II, as it is part of the barre warm-up as well as a movement often found in performance phrases. Like most of the barre exercises it focuses on one half of the dancer's body performed with what is often called the working leg. The dancer draws the leg up the shin bone to the knee and then stretches the leg out and away from the body leaving it high in the air. Thus, the movement does not specify the configuration of the other leg (the supporting leg). In fact, the supporting leg may either be in pose 1 (in plié), 2 (on flat), or 3 (on relevé) - a choice left up to the choreographer.

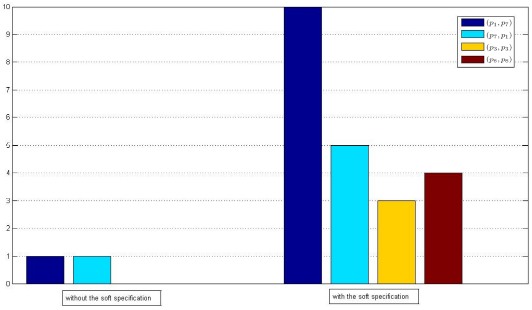

Fig. 4. Occurrence rates of the states satisfying the soft specification for different sample paths show how this specification can encourage certain behaviors of the system. Each sample path contains 52 steps. The red and blue bars show the case with and without the soft specification, respectively. In the case without the encouragement, states (p_3, p_3) and (p_8, p_8) are not visited at all and (p_1, p_7) and (p_7, p_1) are visited only once. In comparison, these states are visited more frequently in case with encouragement from target states. Hence, the soft specifications achieve a differentiated behavior which we call the allégro case as these specifications derive from this style of ballet. [14]

While many aspects of performance - most notably facial expression and dynamic quality which are excluded from this model - affect the emotive quality of a movement, the choice of the standing leg position distinctly changes the quality of the movement. Namely, if the développé occurs when the standing leg is in pose 2 (or "on flat"), the movement does not induce a change in level of the dancer's body whereas if the standing leg is in pose 1 or 3, a sharp depression or elevation of the dancer's body occurs.

Changing levels is a well-known tool in choreography: if an entire phrase of movement occurs on the same level, it may be perceived as boring or monotone. Adding elevation can be thought of as adding high notes in a musical score - which typically bring about images which are light and happy. On the other hand, low notes on the bass clef in music imbue a more somber, mysterious tone to the score. While dance cannot boast the tool of musical clefs, which arrange each movement in an ascending order, there is certainly a similar effect for changes in elevation (particularly when we have also abstracted away other performance qualities and theatrical elements like stage design, costumes, and lighting) [1], [11].

Hence, a sequence containing many développés in plié will emote a distinctly different mood than one with, say, more développés on relevé. This mood is likely to be deeper and more mysterious in contrast to one that is lighter and more open. The ability of the machinery presented in this

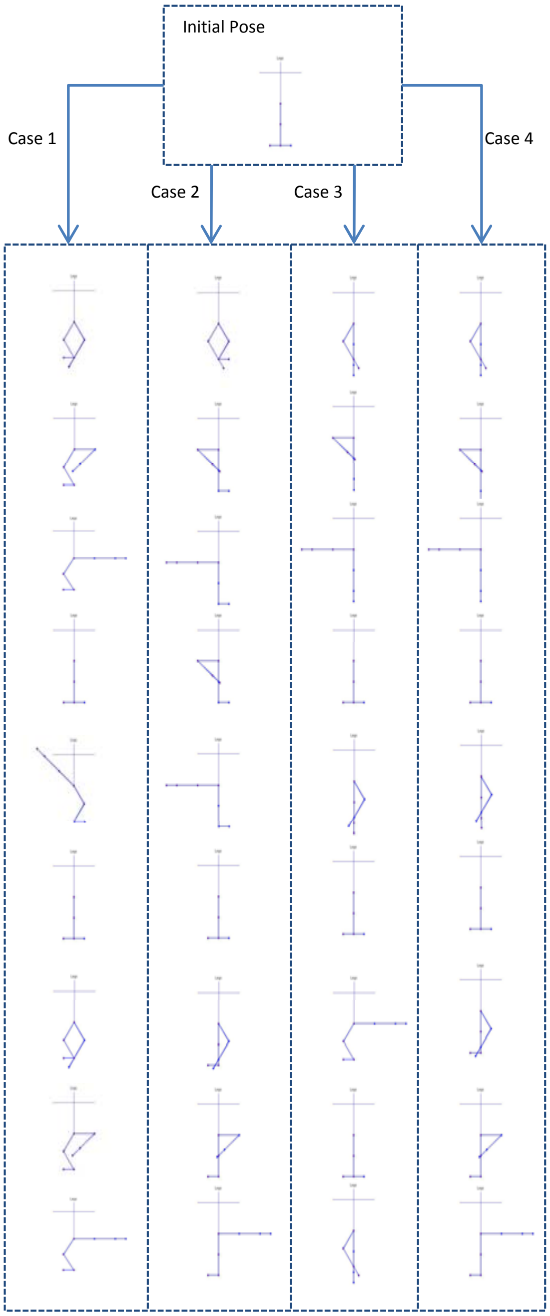

Fig. 5. Four sequences demonstrate how a simple change in specification produces great variation in the output sequence. The sequence with the most general specification is Case 4, the rightmost sequence; here the développé is encouraged without regard to the standing leg's position. The first three cases restrict the position of the standing leg to produce specific emotive effects. The leftmost sequence, Case 1, encourages more développés in plié. This produces a languid, perhaps mysterious, effect. Next, Case 2 encourages développé on flat. The resulting sequence contains fewer changes in the level of the dancer's body which leads to a more matter-of-fact, even blunt, presentation of the leg. On the other hand, the elevated développés on relevé in Case 3 are a practical choice to give a light, ethereal emotive effect to the output sequence.

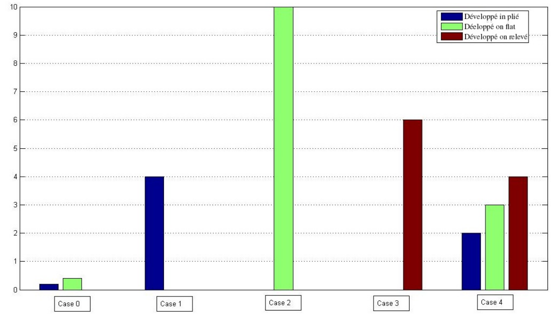

Fig. 6. Occurrence rates of the transitions satisfying the soft specification for different sample paths show how this specification can encourage certain behaviors of the system. Each sample path contains 52 steps. Case 0 shows the average of 10 runs without any soft specification, showing the tendency of the system without any encouragement from the soft specifications. The occurrences of développé in plié, on flat, and on relevé are shown for the remaining four cases in blue, green, and red, respectively.

paper to produce such sequences is demonstrated in Fig. 5 and further analyzed in Fig. 6. In these cases, we have added weights directly to our composed transition system. These soft specifications encourage specific movements to occur and are not phrased in LTL statements like the hard specifications - instead we employ our receding-horizon controller to collect the rewards on transitions and thus produce a desired output sequence.

V. TOWARDS A METHOD FOR STYLE SPECIFICATION AND COMPARISON

An emerging philosophy in several disciplines is that dynamical equations in terms of external parameters such as joint angles are an inherently poor choice of coordinates for parameterizing human motion. Although these external quantities are easy to measure, we present three examples where they have been demonstrated to break down. (1) Recent results in neuroscience ([12]) indicate that the motor cortex (the part of the brain that controls animal movements) is organized in terms of behavioral actions, not body parts and joint angles. (2) It is an emerging practice in dance education to give students corrections in terms of actions, not body part placement [2]. The philosophy is that, in general, humans do not have control of individual body parts. As a result, corrections that identify misplaced body parts result in weird, undesirable movement patterns that the dancer essentially invents in attempt to align something which he or she cannot control directly. (3) Biophysicists studying the movements of *C. Elegans*, a microscopic worm that is one of the most well-studied model organisms in biology, provided the first quantitative explanation of its movement patterns by phrasing their analysis in terms of four body poses. These body poses, and their linear combinations, were shown to account for 95% of the worm's behavior [20]. Viewing the worm's movement in this so-called shape space ([21], [22]) allowed the researchers to, for the first time, accurately predict the motion of the worm due to external stimuli. Thus, it is perhaps an internal parameterization, such as body position, that will provide a more simple, useful model for human movement.

The extension of this philosophy into a system with well-defined inputs and outputs presents interesting questions for systems theory, and as has been presented in this paper, systems theory allows a new articulation of the creative process involved in choreographing human movement. We see the success of this model in terms of the previous paragraph as twofold 1) we phrase the states of our system in terms of natural body poses and 2) we phrase the evolution of these states in terms of the most meaningful facts about the system via LTL.

Namely, using this body position parameterization, we have specified a grammar for body positions in ballet movements restricted to the coronal plane. Furthermore, our specification is guaranteed to follow certain principles, both physical and aesthetic, through the expressive power of LTL. This specification may lead to systems (humanoid or otherwise) which behave in a way that is natural for its given context. A deterministic program can limit the flexibility of a system's ability to cope and adapt to its environment. When implemented on a robotic system, a behavior, as outlined here, provides one or more states (and perhaps a preferred state) that the system may enter given its current state.

A novel venue for movement and style analysis is the second contribution of this model. The movement analysis required to produce a system capable of automatically generating movement phrases in the style of classical ballet results in a quantitative phrasing of the rules that govern this somewhat curious example of human behavior. A quantitative survey of specific dance styles (between different genres and choreographers) would bolster and perhaps corroborate years of qualitative dance study which hold that specific movement patterns evoke very different aesthetic and emotive effects in dance choreography.

In summary, ballet is a highly ordered behavior of a truly complex biological system whose attributes have important analogs in systems theory that warrant quantitative study. By formulating aesthetic style from a systems theoretic perspective and, thus, resolving the attributes of human movement which typify and comprise stylized movement, we are beginning to define a metric for a previously abstract concept. Furthermore, the structure of the aesthetic movement explored here provides an interesting challenge for control theory, namely that of discrete event systems and their composition.

REFERENCES

[1] Susan Au. *Ballet and Modern Dance*. Thames and Hudson New York, 2nd edition, 2002.

[2] I. Bartenieff and D. Lewis. *Body movement: Coping with the environment*. Routledge, 1980.

[3] E. Bradley and J. Stuart. Using chaos to generate variations on movement sequences. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 8:800, 1998.

[4] C. Bregler. Learning and recognizing human dynamics in video sequences. In *Computer Vision and Pattern Recognition, 1997. Proceedings., 1997 IEEE Computer Society Conference on*, pages 568–574. IEEE, 2002.

[5] E. M. M. Clarke, D. Peled, and O. Grumberg. *Model checking*. MIT Press, 1999.

[6] D. Del Vecchio, R.M. Murray, and P. Perona. Decomposition of human motion into dynamics-based primitives with application to drawing tasks* 1. *Automatica*, 39(12):2085–2098, 2003.

[7] Xu Chu Ding, Calin Belta, and Christos G. Cassandras. Receding horizon surveillance with temporal logic specifications. In *the 49th IEEE Conference on Decision and Control*, 2010 (to appear).

[8] C. Fanti, L. Zelnik-Manor, and P. Perona. Hybrid models for human motion recognition. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 1, pages 1166–1173. IEEE, 2005.

[9] P. Gastin and D. Oddoux. Fast LTL to Buchi automata translation. *Lecture Notes in Computer Science*, pages 53–65, 2001.

[10] H. Kress Gazit, G. Fainekos, and G. J. Pappas. Where's waldo? sensor-based temporal logic motion planning. In *IEEE Conference on Robotics and Automation*, Rome, Italy, 2007.

[11] R. Gottlieb. *George Balanchine: the ballet maker*. Eminent Lives, 2004.

[12] M.S.A. Graziano and T.N. Aflalo. Mapping behavioral repertoire onto the cortex. *Neuron*, 56(2):239–251, 2007.

[13] M. Kloetzer and C. Belta. A fully automated framework for control of linear systems from temporal logic specifications. *IEEE Transactions on Automatic Control*, 53(1):287–297, 2008.

[14] Amy LaViers, Yushan Chen, Calin Belta, and Magnus Egerstedt. A formal approach to the automatic generation of ballet motions. *ACM/IEEE Second International Conference on Cyber-Physical Systems*, 2011.

[15] Amy LaViers and Magnus Egerstedt. The ballet automaton: A formal model for human motion. *(under review) American Control Conference, Proceedings of the 2011*, 2011.

[16] B. Peng, G. Qian, and S. Rajko. View-invariant full-body gesture recognition from video. In *Pattern Recognition, 2008. ICPR 2008. 19th International Conference on*, pages 1–5. IEEE, 2009.

[17] N. Piterman, A. Pnueli, and Y. Saar. Synthesis of reactive(1) designs. In *VMCAI*, pages 364–380, Charleston, SC, 2006.

[18] M. M. Quottrup, T. Bak, and R. Izadi-Zamanabadi. Multi-robot motion planning: A timed automata approach. In *Proceedings of the 2004 IEEE International Conference on Robotics and Automation*, New Orleans, LA, April 2004.

[19] F. Somenzi and R. Bloem. Efficient büchi automata from ltl formulae. *Twelfth Conference on Computer Aided Verification (CAV'00)*, pages 248–263, 2000.

[20] G.J. Stephens, B. Johnson-Kerner, W. Bialek, and W.S. Ryu. Dimensionality and dynamics in the behavior of C. elegans. *PLoS Comput Biol*, 4(4):e1000028, 2008.

[21] G.J. Stephens, B. Johnson-Kerner, W. Bialek, W.S. Ryu, and E. Warrant. From Modes to Movement in the Behavior of Caenorhabditis elegans. *PloS one*, 5(11):462–465, 2010.

[22] G.J. Stephens, W.S. Ryu, and W. Bialek. The emergence of stereotyped behaviors in C. elegans. *Bulletin of the American Physical Society*, 55, 2010.

[23] T. Wongpiromsarn, U. Topcu, and R. M. Murray. Receding horizon control for temporal logic specifications. In *Proc. International Conference on Hybrid Systems: Computation and Control*, 2010.