Neurocomputing 72 (2009) 3624–3630

Contents lists available at ScienceDirect

Neurocomputing

journal homepage: www.elsevier.com/locate/neucom

Emerging motor behaviors: Learning joint coordination in articulated mobile robots ☆

Diego Pardo *, Cecilio Angulo, Sergi del Moral, Andreu Català

CETpD, Technical Research Center for Dependency Care and Autonomous Living, ESAII-UPC, Automatic Control Department, Technical University of Catalonia, Vilanova i la Geltrú, Barcelona, Spain

A R T I C L E I N F O

Available online 23 June 2009

Keywords:
Articulated mobile robots
Motor behaviors
Robot learning

A B S T R A C T

In this paper, we analyze the insights behind the common approach to the assessment of robot motor behaviors in articulated mobile structures with compromised dynamic balance. We present a new approach to this problem and a methodology that implements it for motor behaviors encapsulated in rest-to-rest motions. As well as common methods, we assume the availability of kinematic information about the solution to the task, but reference is not made to the workspace, allowing the workspace to be free of restrictions. Our control framework, based on local control policies at the joint acceleration level, attracts actuated degrees of freedom (DOFs) to the desired final configuration; meanwhile, the resulting final states of the unactuated DOFs are viewed as an indirect consequence of the profile of the policies. Dynamical systems are used as acceleration policies, providing the actuated system with convenient attractor properties. The control policies, parameterized around imposed simple primitives, are deformed by means of changes in the parameters. This modulation is optimized, by means of a stochastic algorithm, in order to control the unactuated DOFs and thus carry out the desired motor behavior.

1. Introduction

Articulated mobile robots (AMRs) are autonomous systems constructed to accomplish generic tasks. These platforms offer features that allow the generation of diverse types of motor behaviors, but also place restrictions on these behaviors. An initial example of this scenario is provided by the RoboCup competition [8], where a standard robot is used to play football. Research teams try to maximize the benefits provided by the kinematics and dynamics of the structure in order to perform motor behaviors better, such as running, kicking, heading the ball and goalkeeping. Hence, improving and generating motor behavior capacities is a research challenge. Some questions are: (i) Is the default gait of the robot efficient or fast enough? (ii) Is it able to jump? (iii) Can it lift a weight heavier than the factory-specified limit? Challenging questions arise about the optimization of default motor behaviors, the design of new motor behaviors and the overcoming of body constraints and limitations.

AMRs are often *underactuated* systems, i.e. not all of the degrees of freedom (DOFs) are actuated, and therefore the global dynamic balance of the system is constantly compromised. Additionally, they also are *redundant*, i.e. they have more DOFs than those needed for representing the position and orientation of the controlled element of the robot (the workspace). It is not evident how to control this kind of mechanism, and this is still an open area of research. Motivated by this challenge, we address a methodology to synthesize motor behaviors in AMRs in this paper, where we understand *motor behavior* as a human interpretation of the motions of a robot and their consequences (e.g. sitting, throwing an object or walking).

In order to restrict the problem, we shall focus on motor behaviors that may be encoded by the realization of rest-to-rest motions, i.e. motions defined by an initial and a final state with a velocity equal to zero. Nevertheless, many of the motor behaviors of AMRs can be understood as a consequence of rest-to-rest motions [3], for example reaching [5], throwing [10] and simple posture transitions. Moreover, cyclic and composed tasks may be decomposed into sequences of motions of this type [1].

Procedures that synthesize behaviors rely on the availability of specifications in terms of the workspace, which may include a complete path or simply the initial and final positions of the controlled element. These specifications may be derived from a direct human imposition or by path-planning methods, often offering kinematic solutions to the problem of achieving the

☆ This work was partly supported by the Spanish MEC project ADA (DPI2006-15630-C02-01).

* Corresponding author.
E-mail addresses: diego.pardo@upc.edu (D. Pardo), cecilio.angulo@upc.edu (C. Angulo), sergi.delmoral@upc.edu (S. del Moral), andreu.catala@upc.edu (A. Català).
URLS: http://www.upcnet.es/~upc26174 (D. Pardo), http://www.upcnet.es/~upc15838 (C. Angulo).

0925-2312/$ - see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.neucom.2009.01.015

desired motor behavior. Subsequently, robot control methods are required to compute torque-level actions that drive the mechanism according to the given specifications. As an inheritance from experience in the control of robot manipulators, where the element that defines the workspace is the end-effector of a joint chain with a limited motion domain, most of the approaches to robot control consider a workspace reference as input data, computed off-line or generated within the methodology of the relevant approach, starting from pure kinematic information about the solution to the task. When workspace references are followed, the torque control actions are computed with the purpose of achieving desired joint accelerations, which implies compensation of rigid-body dynamics, canceling the effects of gravity, and other nonlinearities.

Some of the approaches described in [9] consider the kinematics and rigid-body dynamics of the robot in the generation of the workspace reference, but this is usually done to establish sufficient conditions for the accomplishment of the behavior, rather than to take advantage of the kinematics and rigid-body dynamics. An example of this is evidenced in the gait control system used in the ASIMO humanoid (which has one of the best humanoid gaits developed so far), where control forces are computed to maintain balance stability during gait execution, i.e. the effects of gravity are canceled while suitable accelerations are imposed to accomplish the motion; consequently, the energy consumption is more than 15 times (scaled) the amount required during human gait [18]. However, it has been demonstrated that during human gait, not only are the dynamic effects of gravity not always canceled but also they are actually employed [14]. It seems that the current strategies to carry out a given motor behavior are well-suited to obtaining a particular solution of the problem. Thus, the space of behavior solutions is narrowed by the approach used rather than by the capacities of the robot.

However, some results using new perspectives show evidence of alternative solutions, ones that favor the execution of the motion and expand the capacities of the robot. For instance, results in [19] show that the given factory-maximum payload of an industrial manipulator (a 6-DOF PUMA-762) can be greatly increased by exploring new zones of the solution space with suitable control policies. The approach used was the formulation of a parameterized optimal control problem, where body dynamics and time ranges were stated as restrictions. Torque-level actions were found such that the payload lifted by the manipulator was much more (six times) than the load reachable by the default aggregation of path planning, workspace reference and torque control. Surprisingly, contrary to standard procedures, the resulting trajectories included singularities, letting the robot rest the payload against its structure on its way to the goal. Along the same lines, a similar result was later presented in [15], where a simple manipulator (2D, 3-DOF) accomplished a weightlifting behavior, avoiding workspace restrictions in the formulation. Besides maximizing the payload lifted, the results included quite different workspace trajectories that accomplished the same behavior.

The key attribute in both approaches was the direct connection between the desired behavior and the torque commands, i.e. the workspace requirements were almost null, leaving the system free to be modulated in order to fulfill the behavior, i.e. lift a defined weight. Both approaches use optimization as the main route; nevertheless, the analytical solution in [19] implies a detailed formulation of the problem and its restrictions, which is perfectly viable for manipulators in structured environments, but this is not the case for AMRs. On the other hand, the solution given in [15] is not analytical but numerical; it searches in the solution space using a learning algorithm, i.e. a numerical optimization of policy parameters by means of iterative evaluation of experiences. Nevertheless, its control framework, based on the coordination of lower-level PID controllers, cannot be directly extrapolated to more complex problems.

Recently, the attention given to the use of learning as a paradigm to exploit the capacity of robots has being growing. The latest publications on learning of motions by robots [7] revolve around early results on imitation [6], where the initial solution in the workspace is directly guided by a human, and afterwards the robot joints are controlled by parameterized policies that are intended to accomplish the behavior. The type of the functions used as control policies is that of dynamical systems (DSs). The optimal parameters of the policy are found using reinforcement learning (RL) [17] algorithms. Extensive work on RL algorithms suitable for computing robot control policies has been presented in [13].

In the methodology presented in this paper, we assume the availability of kinematic information equivalent to the initial and final states of the desired behavior. In contrast to the imitation approach, a reference in the workspace is not specified. Our control framework, based on local control policies at the joint acceleration level, attracts actuated DOFs to the desired final configuration; meanwhile, the resulting final states of the unactuated DOFs are viewed as a consequence of the actuated acceleration profiles. DSs are used as acceleration controllers, providing the system with these attractor properties. Additionally, the control policies are parameterized around imposed simple primitives, which may by deformed by means of changes in the parameters in order to obtain complex accelerations.

Subsequently, we present an example that provides a qualitative description of the type of problems that this paper addresses. The *standing-up* behavior illustrates those motor behaviors of underactuated systems in which dynamic balance is compromised. Fig. 1 shows the initial and final states for this behavior. Note that the behavior is enclosed by a motion where the initial and final velocities are equal to zero. The robot starts in a lying-down posture and should stand up, ending up as shown in Fig. 1b. However, gravity and other nonlinearities can influence the behavior in such a way that the robot ends up in a different state (see Fig. 1c). The achievement of desired values for the actuated DOFs is not enough for the desired behavior to be the result.

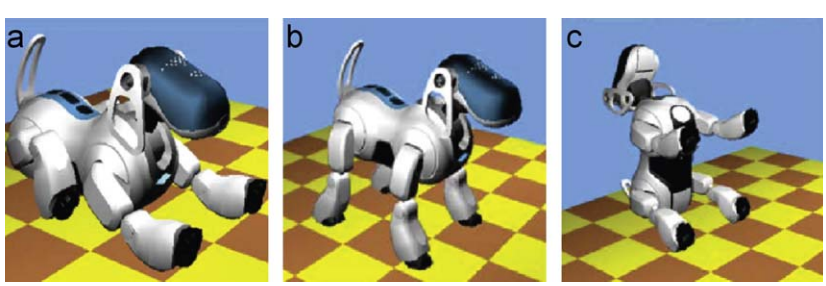

Fig. 1. (a) Initial state of the robot for the standing-up behavior. (b) Desired final configuration. (c) Undesired final configuration, where motor behavior has failed.