Learning Controller Preconditions for Physics-Based Character Animation

Petros Faloutsos, Michiel van de Panne, and Demetri Terzopoulos

University of Toronto
Department of Computer Science
10 King's College Road
Toronto, ON M5S 3G4, Canada

An ambitious goal in the area of physics-based computer animation within the domain of computer graphics is the creation of virtual actors that autonomously synthesize realistic human motions and possess a broad repertoire of lifelike motor skills. To this end, the control of dynamic, anthropomorphic figures subject to gravity and ground contact forces remains a difficult open problem. While a divide-and-conquer strategy is clearly prudent in emulating the enormous variety of controlled motions that humans and other animals may perform, little effort has been directed at how the resulting control solutions may be integrated to yield significantly broader functionalities. We propose a simple yet effective framework for composing basic motor controllers in order to achieve broader control competencies, such as the ability to stand up autonomously in several different ways after sustaining a fall. Through controller composition, we illustrate the potential of building a "dynamic virtual stuntman" that can execute some of the sorts of dangerous stunts seen in action movies. The challenge here lies in developing appropriate control strategies for specific actions and ways of integrating these strategies into a coherent whole.

An important technical contribution of our motor controller composition framework is an explicit model of "preconditions"; i.e, those regions of a dynamic figure's state space within which a given motor controller is applicable and expected to work properly. In particular, we investigate the provocative problem of determining preconditions automatically and introduce a promising solution which employs support vector machine (SVM) learning theory. Our novel technique systematically learns appropriate preconditions by stochastically sampling controller performance in practice. Using performance data sampled from a controller, we train a nonlinear binary-output SVM classifier to predict the success or failure of the controller. Hence, the trained SVM classifier represents the preconditions of the controller. In our application, SVM training can take hours (on a Pentium III processor), but this is done off-line. Once trained, the SVM classifier can be interrogated in milliseconds to evaluate the preconditions of a controller during on-line operation.

We demonstrate successful SVM learning of preconditions for families of composable controllers for articulated figures whose physical parameters reflect anthropometric data consistent with a fully-fleshed adult male. One family of controllers is for a 37 degree-of-freedom (DOF) 3D articulated figure. A second family of controllers is for a comparable 16 DOF 2D articulated figure. While the 3D figure illustrates the ultimate promise of the technique, the easier control problem associated with the 2D skeleton allows for more rapid prototyping of larger families of controllers and more careful analysis of their operation. As has been recognized in the robotics literature, the control of broadly skilled motion repertoires remains very much an open problem even for 2D articulated figures.

We have engineered a software system that implements our controller composition framework. Our system is built on top of DANCE, a portable, extensible object-oriented modeling and animation system that provides a versatile plug-in mechanism (DANCE was developed at the University of Toronto and is freely available for non-commercial use). Individual controllers and the composition method are implemented as DANCE plug-ins. Using this software, we have produced a variety of physics-based animation results.