# Policies for goal directed multi-finger manipulation


**Sheldon Andrews and Paul G. Kry**

*School of Computer Science, McGill University*


## Introduction

We present a novel framework for synthesizing motions for one-handed human manipulation problems [1]. All motion is physically based. Motion capture data is not used and the exact object trajectory is never specified; we only require the final configuration of the object. Offline simulations are used to build effective control policies that adapt to real time changes in the task. Additionally, finger motions exhibit human characteristics, such as finger gaiting, and manipulations use the entire geometry of the hand.

## Methodology

We decompose each manipulation task into sequences of three phases: *approach*, *actuate*, and *release*.

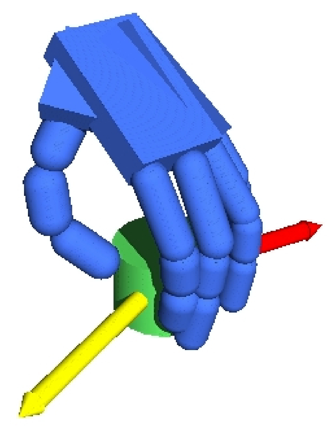

**Goal:** rotate the dial to align the red arrow with the yellow arrow.

Each state $s$ represents an optimization problem. Covariance matrix adaptation (CMA-ES) [2] is used to find suitable controller parameters by optimizing phased based objectives:

*Approach ($L_0$)*
- Finger planting

*Actuate ($L_1$)*
- Task progress
- Grasp stability
- Style

*Release ($L_2$)*
- Zero net twist
- Good recovery pose

*All phases ($L_g$)*
- Minimize torques, No back-of-hand contact

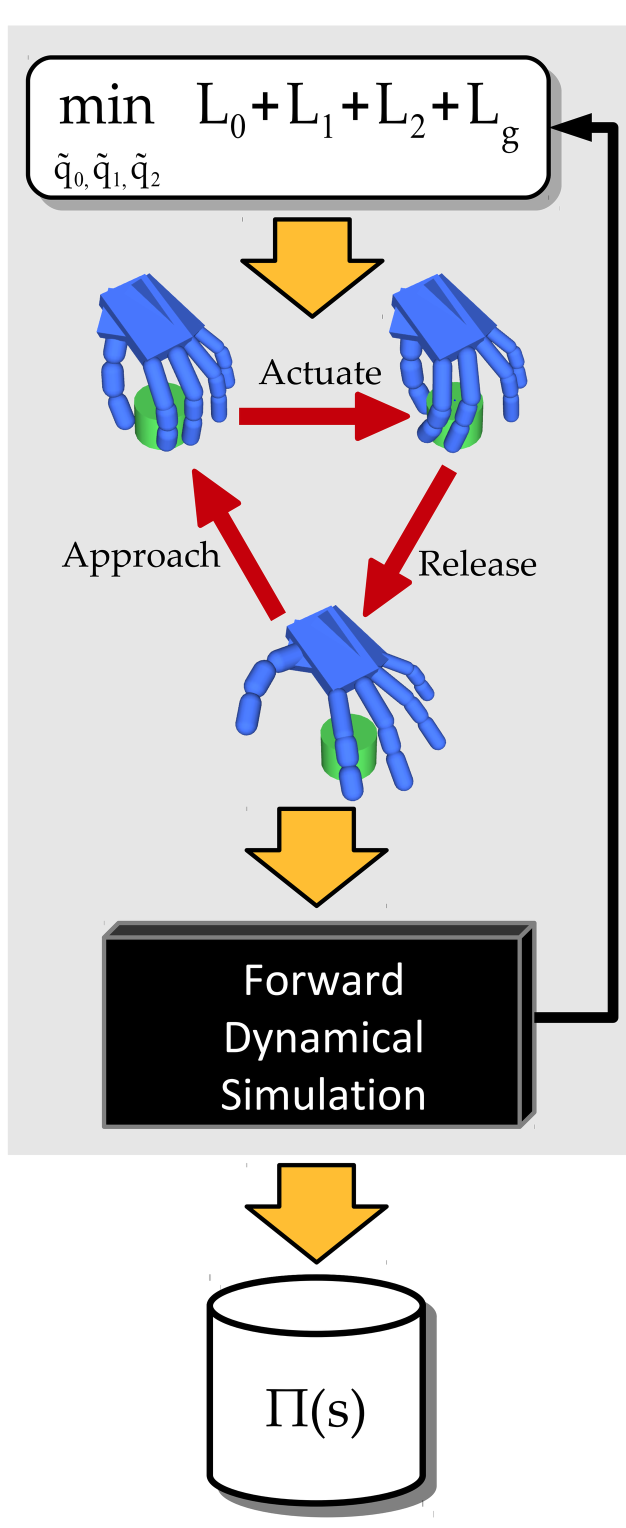


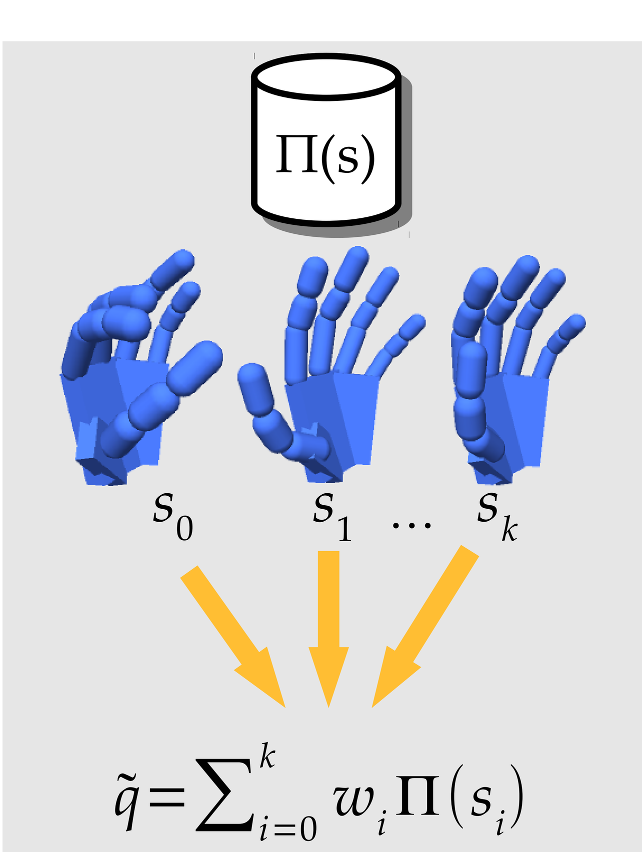


Each controller consists of three reference poses $(\tilde{q}_0, \tilde{q}_1, \tilde{q}_2)$ managed by a state machine. We evaluate the parameters by running a forward dynamical simulation over a single controller cycle. Joint torques are computed by the PD servo equation $\tau = K(\tilde{q}-q) - D\dot{q}$.

Optimized controller parameters for multiple states are used to construct a k-nearest neighbors function approximator *Π(s)*. Weighted linear interpolation is used to compute control parameters for novel states.

## Results

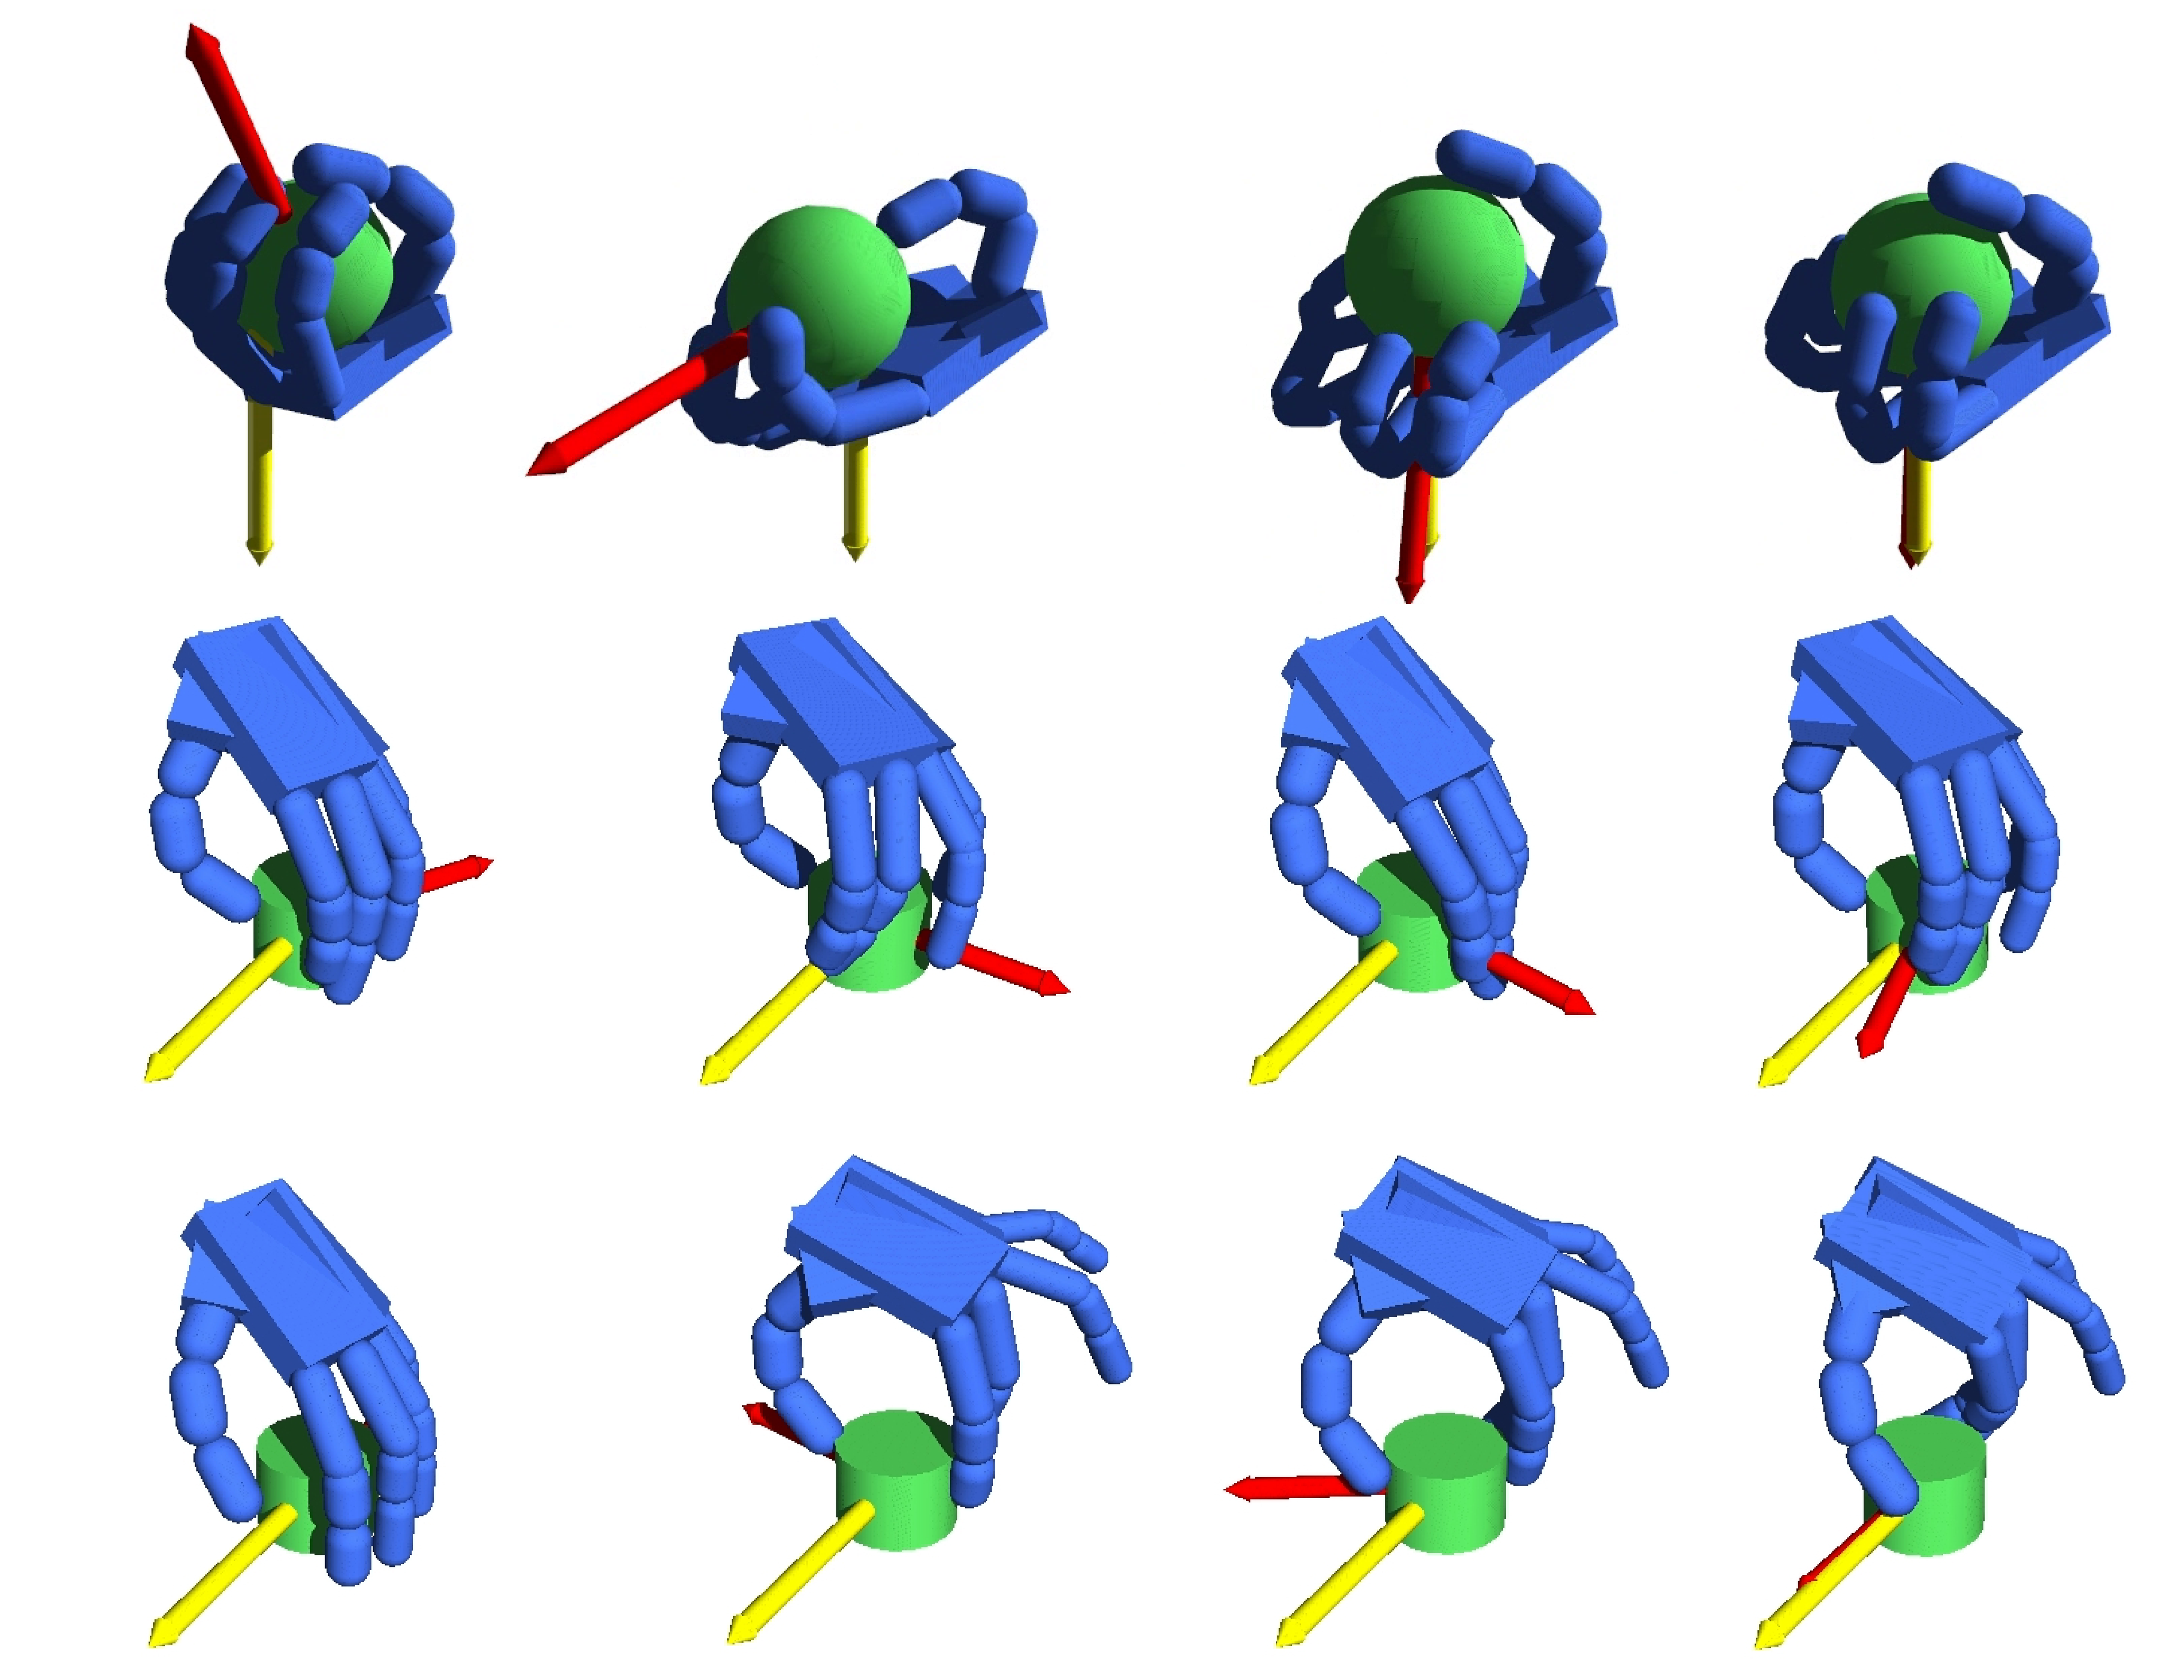

Top row: Ball-in-hand example. Middle and bottom rows: Two variations for a dial turning example. The middle row enforces a short phase duration; the bottom row uses a longer phase duration, and only the thumb, index, and middle fingers are selected for participation.

## Analysis

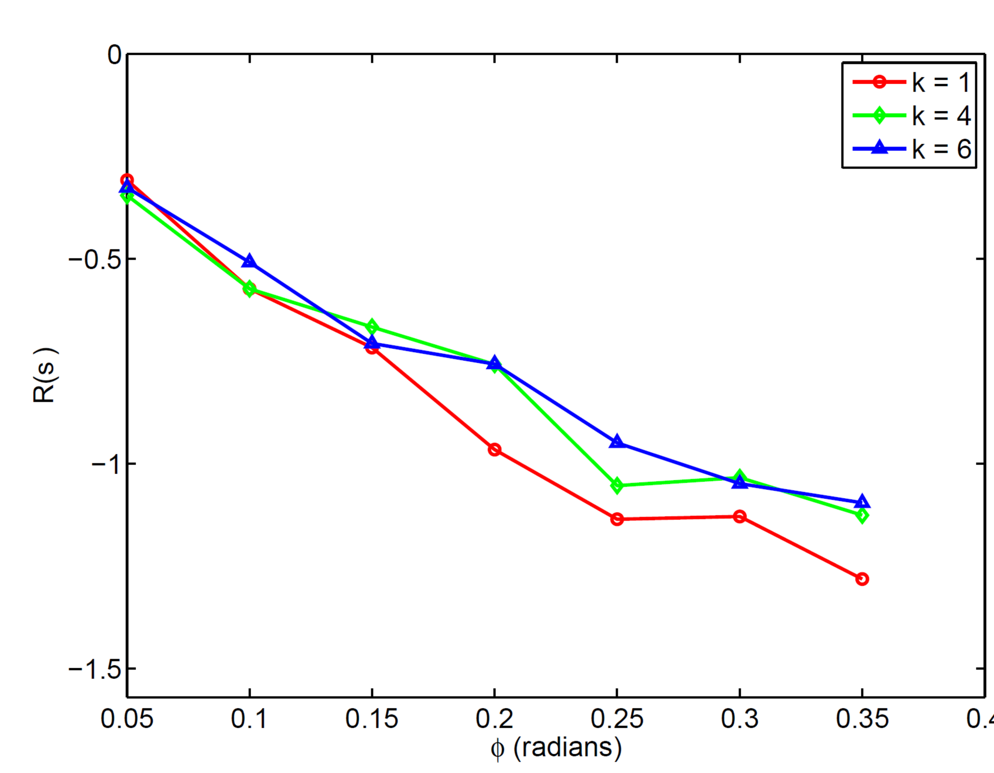


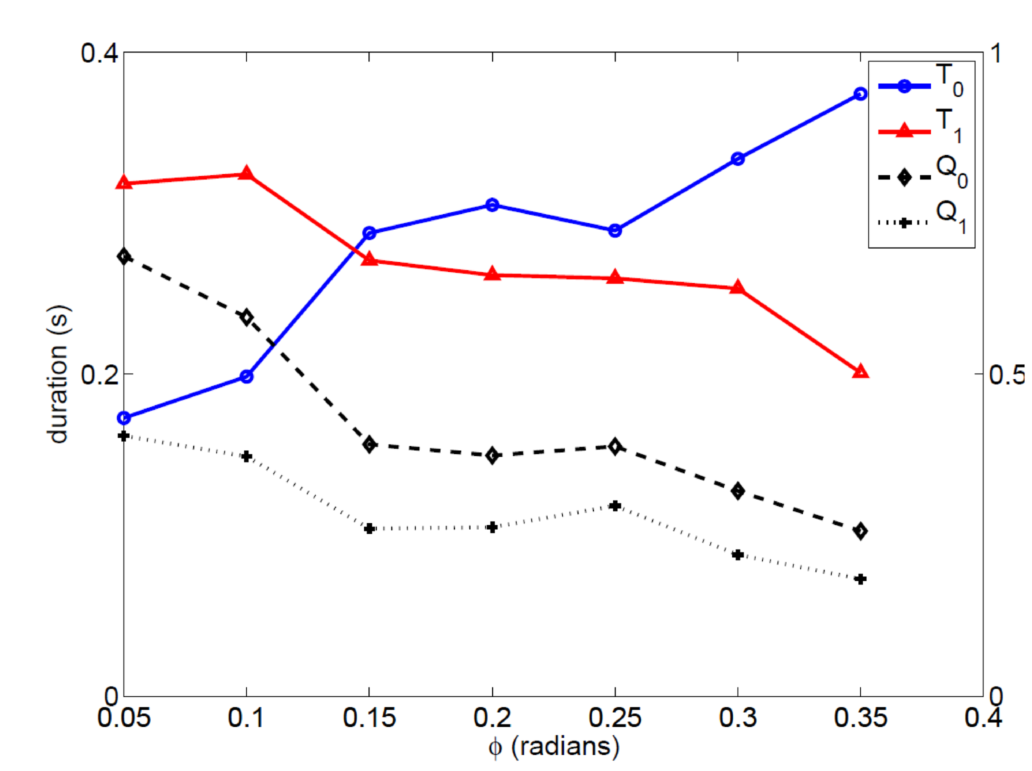


We examine the effect of joint perturbations $\phi$ on the effectiveness of learned controllers. Task progress is measured by the reward *R(s)* at the end of a controller cycle (left). A shortened duration for the actuation phase ($T_1$) suggests that progress is related to the grasp quality during the approach ($Q_0$) and actuation phase ($Q_1$) (right). Grasp quality is computed similarly to [3].

## Conclusion and future work

Our framework for generating human grasping motions creates plausible, coordinated motions that adapt to task changes in real-time. We do not assume a pre-defined trajectory or require motion capture data. In future work we will investigate the use of linear feedback control to improve the robustness of the policies by tracking features relevant for successful human manipulation.

## References


[1] Andrews, S. and Kry, P.G. (2013). Goal Directed Multi-Finger Manipulation: Control Policies and Analysis. *Computers & Graphics.*

[2] Hansen, N. (2006) The CMA evolution strategy: A comparing review. *Advances in Estimation of Distribution Algorithms.*

[3] Klein, C. and Baho, B.E. (1987) Dexterity measures for the design and control of kinematically redundant manipulators. *The International Journal of Robotics Research.*



GRAND NSERC CRSNG McGill