

Kinematic Retargeting with Potential Fields: Enabling a Unified Social Behavior Repository

Tarik Tosun
Princeton University
875 Frist Center
Princeton, NJ 08544
302-547-4196

ttosun@princeton.edu

Ross Mead
University of Southern California
3710 McClintock Ave., RTH 423
Los Angeles, CA 90089-0781
213-740-6245

rossmead@usc.edu

Maja J Matarić
University of Southern California
3710 McClintock Ave., RTH 407
Los Angeles, CA 90089-0781
213-740-4520

mataric@usc.edu

ABSTRACT

In this paper, we describe a method for *kinematic retargeting*, the porting of configurations and motions from one embodiment (e.g., human, robot, or animated character) to another. This approach is implemented through an optimization process involving potential fields. Unlike platform-specific approaches, this approach has the capability to retarget motion between arbitrary embodiments. That portability enables the development of a unified social behavior repository for human-robot interaction.

Categories and Subject Descriptors

I.2.9 [Artificial Intelligence]: Robotics – *kinematics and dynamics*.

General Terms

Algorithms, Performance, Human Factors, Standardization, Theory

Keywords

Retargeting, kinematics, unified social gesture repository, socially assistive robotics, embodiment, cross-platform

1. INTRODUCTION AND MOTIVATION

Socially assistive robotics (SAR) is a growing field in which robots engage in non-contact, co-present social interactions with people with special needs [1]. For example, preliminary research exists indicating the effectiveness of robot use in the promotion of social skills in children with autism, exercises to retain or improve memory in patients with Alzheimer’s disease, and activities of daily living to increase range of motion in post-stroke rehabilitation [2]. For each interaction scenario, a set of communicative and task-specific embodied robot behaviors must be explicitly defined; these behaviors are often platform-specific. A unified repository of experimentally validated behaviors, coupled with a method for automatically adapting those behaviors to specific robot platforms, would facilitate SAR research. A critical challenge in the creation of such a unified behavior repository is the method for the automatic porting of behaviors from one embodiment to another, a process termed *retargeting*. Existing retargeting systems are specific to a small subset of platforms [3]. To alleviate this and enable the development of a unified social behavior repository, we propose a novel potential-field-based approach to the automatic retargeting of kinematic configurations from one embodiment (e.g., human, robot, or animated character) to another.

Copyright is held by the author/owner(s).
HRI’12, March 5–8, 2012, Boston, Massachusetts, USA.
ACM 978-1-4503-1063-9/12/03.

2. APPROACH

2.1 Potential Field Generation

Our approach to kinematic retargeting defines a Laplacian potential field in which the “skeleton” (kinematic chain) of the *source platform* is the surface of minimum potential, and then configures the pose of the *target platform* to minimize the field integral over its own skeleton. Laplacian fields mathematically describe electrostatic fields, so this process can be conceptualized intuitively as imbuing the two skeletons with opposite electrical charges and then allowing the target skeleton to “snap” to the rigid source skeleton in the pose that minimizes the potential of the system.

2.2 Pose Optimization Process

The optimization process involves evaluation of the line integral of the source potential over the skeleton of the target. This integral is a scalar function of the target’s pose, parameterized by its joint angles. For a target platform with n degrees of freedom, this becomes an n -dimensional optimization problem with domain bounds defined by the joint angle limits of the target platform. This function typically has many local minima; gradient-based algorithms tend to succeed in the vicinity of the global minimum, but fail when the initial target pose is far from the optimal pose. In these cases, a gradient-free algorithm (such as genetic algorithms, Monte Carlo methods, or particle-based approaches) can be used. In practice, the source and target chains are often of different lengths, in which case they are scaled to the same total length for calculation purposes during the optimization process.

3. IMPLEMENTATION AND RESULTS

The foundations of a cross-platform, rapid retargeting system have been implemented in MatLab using the described potential-field-based approach. In this section, we show that our method captures the source embodiment’s pose in a way that inverse kinematics often cannot. We illustrate the capability of our method in retargeting poses to target embodiments with both lower and higher degrees of freedom than the source embodiment. We then show how end-effector position may be incorporated into the optimization process to produce more effective gesture retargeting solutions. Finally, we discuss fitness criteria we have developed for evaluation of goodness-of-fit between two poses.

3.1 Inverse Kinematics vs. Chain Retargeting

The retargeting problem cannot be adequately solved using inverse kinematics (IK) over the source end-effector position. There are often multiple IK solutions for a given end-effector position and, because IK does not explicitly capture the configuration of the entire kinematic chain, the configuration of the target chain will not necessarily be similar to that of the source chain, as depicted in Figure 1 (in solid red).

3.2 Dissimilar Degrees of Freedom (DOF)

The hallmark of the approach is its ability to retarget between kinematic chains with different degrees of freedom. Figure 1 illustrates a 3-DOF source chain (in dashed black) compared to both a 2-DOF target chain (in solid blue) and a 4-DOF target chain (in solid green). In contrast to the IK solution (in solid red), both of these retargeted solutions closely match the source pose; however they do not necessarily match the end-effector position.

3.3 End-Effector Terms (EET)

As discussed in Section 2.2, our retargeting approach has been formulated as the optimization of a scalar fitness function. As described, this function has only one term—the line integral of potential over the skeleton of the target embodiment. However, the approach is extensible and terms may be added to this fitness function to tailor the method to a particular application; for example, research indicates that end-effector position carries great meaning in socially expressive gestures [4]. For the purpose of retargeting social gestures in a unified behavior repository, a term is included to penalize solutions that place the target end-effector far from the source end-effector. This is illustrated in Figure 2, in which an EET solution (in solid magenta) has been retargeted with a fitness function that incorporates an end-effector term. This solution matches the source pose closely, and places the target end-effector position closer to the source end-effector position than the non-EET solution (in solid green).

3.4 Fitness Criteria

The potential field approach can quantify a scalar metric of “goodness-of-fit” provided by the fitness function. Intuitively, a definitive best possible value for this fitness function is obtained for any source chain by computing its value when the source chain is retargeted to itself. A fitness ratio is then computed for any solution by dividing the solution’s fitness function value by the best possible value (Table 1).

Table 1. Values of pose fitness term (Pose Fit) and end-effector fitness term (EET Fit) for each of the kinematic chains (with respective degrees of freedom) in Figures 1 & 2.

Chain	Source	IK	w/o EET	w/o EET	w/ EET
DOF	3	4	2	4	4
Pose Fit	1.0000	0.5210	0.9344	0.9876	0.9470
EET Fit	1.0000	1.0000	0.9890	0.9289	0.9804

4. DISCUSSION AND ONGOING WORK

We have presented a novel potential-field based approach to kinematic retargeting; this enables the development of a unified social gesture repository, which we save as future work. This approach has been implemented in a simulated MatLab environment and is currently being developed for physical robot platforms. Our preliminary investigations have shown the utility of the approach over inverse kinematic solutions and with varying

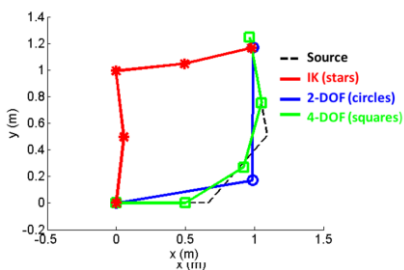


Figure 1. IK and retargeted (2-DOF and 4-DOF) solutions from a source platform.

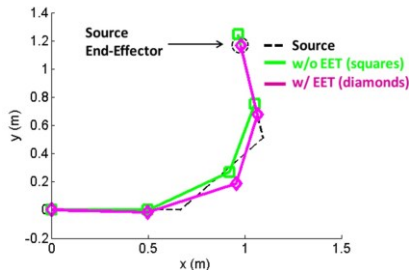


Figure 2. Retargeted solutions with and without end-effector terms (EET).

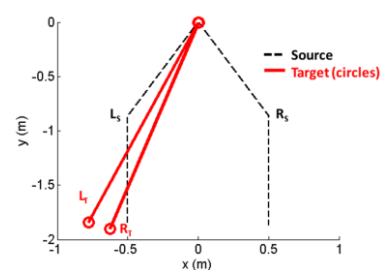


Figure 3. Interacting potential fields are problematic in kinematic trees.

degrees of freedom between the source and target embodiments. We have also illustrated the extensibility of the approach through the incorporation of end-effector terms, which are particularly useful for socially expressive gesture in human-robot interaction and socially assistive robotics [4]. Initial metrics have also been established for goodness-of-fit between source and target configurations. We now discuss extensions of this work that will address motion and whole-body retargeting (kinematic trees).

4.1 Motion Retargeting

To retarget motion, the source can be sampled at some frequency during the motion to generate pose frames that are individually retargeted from the source platform to the target platform. Motion retargeting typically involves small changes in pose from one frame to the next, and can be accomplished with a gradient-based approach. However, this requires real-time performance, which is not achieved in our current implementation; we are investigating indexing techniques (dynamic programming or lookup tables) to meet this real-time constraint.

4.2 Kinematic Tree Retargeting

Retargeting an entire skeleton at once is challenging. The skeleton is represented as a kinematic tree (a collection of branching kinematic chains, or “limbs”), and its potential field has a distinct region of low potential for each limb. Problems arise when a target limb falls into the potential field produced by another source limb (see Figure 3). This will be alleviated by retargeting each limb individually in a divide-and-conquer manner.

5. ACKNOWLEDGMENTS

This work was sponsored in part by the NSF Research Experience for Undergraduates program and an NSF Graduate Research Fellowship. The authors would like to thank David Feil-Seifer for his ideas regarding end-effector terms, and Amin Atrash for his advice regarding optimization techniques.

6. REFERENCES

- [1] Feil-Seifer, D.F. and Matarić, M.J. 2005. Defining socially assistive robotics. In *International Conference on Rehabilitation Robotics* (June 2005), Chicago, IL, 465-468.
- [2] Tapus, A., Matarić, M.J., and Scassellati, B. 2007. The grand challenges in socially assistive robotics. In *IEEE Robotics and Automation Magazine*, 14, 1 (2007), 35-42.
- [3] Pollard, N., Hodgins, J., Riley, M., and Atkeson, C. 2002. Adapting human motion for the control of a humanoid robot. In *The Proceedings of the International Conference on Robotics and Automation*, 1390-1397.
- [4] Mataric, M.J. and Pomplun, M. 1998. Fixation behavior in observation and imitation of human movement. In *Cognitive Brain Research*, 7, 2 (1998), 191-202.