

Inverse Kinodynamics: Editing and Constraining Kinematic Approximations of Dynamic Motion

Paul G. Kry^a, Cyrus Rahgoshay^a, Amir Rabbani^a, Karan Singh^b

^a*School of Computer Science, Centre for Intelligent Machines, McGill University*

^b*Department of Computer Science, University of Toronto*

Abstract

We present inverse kinodynamics (IKD), an animator friendly kinematic work flow that both encapsulates short-lived dynamics and allows precise space-time constraints. Kinodynamics (KD), defines the system state at any given time as the result of a kinematic state in the recent past, physically simulated over a short time window to the present. KD is a well suited kinematic approximation to animated characters and other dynamic systems with dominant kinematic motion and short-lived dynamics. Given a dynamic system, we first choose an appropriate kinodynamic window size based on accelerations in the kinematic trajectory and the physical properties of the system. We then present an inverse kinodynamics (IKD) algorithm, where a kinodynamic system can precisely attain a set of animator constraints at specified times. Our approach solves the IKD problem iteratively, and is able to handle full pose or end effector constraints at both position and velocity level, as well as multiple constraints in close temporal proximity. Our approach can also be used to solve position and velocity constraints on passive systems attached to kinematically driven bodies. We describe both manual and automatic procedures for selecting the kinodynamic window size necessary to approximate the dynamic trajectory to a given accuracy. We demonstrate the convergence properties of our IKD approach, and give details of a typical work flow example that an animator would use to create an animation with our system. We show IKD to be a compelling approach to the direct kinematic control of character, with secondary dynamics via examples of skeletal dynamics and facial animation.

Keywords: inverse kinematics, secondary dynamics, key framing

1. Introduction

Physical simulation is now a robust and common approach to recreating reality in virtual worlds and is almost universally used in the animation of natural phenomena, ballistic objects, and character accessories like clothing and hair. Despite these strides, the animation of primary characters continues to be dominated by the kinematic techniques of motion capture and above all traditional keyframing. Two aspects of a primary character in particular, skeletal and facial motion, are often laboriously animated using kinematics.

We note from conversations with about half a dozen master animators that there are perhaps three chief reasons for this. First, kinematics, unencumbered by physics, provides the finest level of control necessary for animators to breathe life and personality into their characters. Second, this control is direct and history-free, in that the authored state of the character, set at any point in time, is precisely observed upon playback and its impact on the animation is localized to a neighborhood around that time. Third, animator interaction with the time-line is WYSIWYG (what you see is what you get), allowing them to scrub to various points in time and instantly observe the character state without having to playback the entire animation.

The same animators expressed the utility and importance of secondary dynamics overlaid on primarily kinematic character motion to enhance the visceral feel of their characters. Various approaches to such secondary dynamics have been proposed

in research literature [1, 2, 3], some of which are available in commercial animation software. Overlaid dynamics, unfortunately compromise the second and third reasons animators rely on pure kinematic control.

A kinematic solution incorporating secondary dynamics called *kinodynamic* skinning [4] was suggested in the context of volume preserving skin deformations. With this approach, a kinodynamic state at any time is defined as a kinematic state in the recent past, physically simulated forward to the given time. In this paper we develop this idea of kinodynamics (KD) as a history-free kinematic technique that can incorporate short-lived dynamic behavior. Note that the above usage of the term “kinodynamic”, while similar in spirit, is distinct from its use in the context of robot motion planning where it addresses planning problems where velocity and acceleration bounds must be satisfied [5].

The KD *window size* determines how far into the recent past we start a physical simulation in order to compute a KD state. We must formulate an appropriate KD window size for a given kinematic motion and physical parameters: both long enough to ensure a temporally coherent KD trajectory that captures the nuances of system dynamics, and short enough for interactive WYSIWYG computation and temporal localization of the influence of animation edits on system state. Many goal directed actions such as grasping, reaching, stepping, gesticulating, and even speaking, however, involve spatial relationships between the character and its environment, that are best specified di-

54 rectly, as targets states that the character (or parts of the char-
55 acter) must observe at given times. Techniques such as inverse
56 kinematics (IK) and space time optimization algorithmically in-
57 fer the remaining system states and animation parameters from
58 these animator specified spatio-temporal targets. However, IK
59 does not give the secondary dynamics, and space time opti-
60 mization is typically computationally expensive. Analogous to
61 these techniques, we develop an inverse kinodynamics (IKD)
62 algorithm allowing animators to prescribe position and velocity
63 constraints at specific points in time within a KD setting.

64 Kinodynamics is an interesting approach for interactive
65 character animation, where animators can continue to leverage
66 a direct history-free kinematic work flow, coupled with the ben-
67 efits of arbitrary physically simulated secondary dynamics. The
68 problem that we are solving is the inverse kinodynamics prob-
69 lem, and our solution allows an animator to easily edit kino-
70 dynamic trajectories such that desired constraints can be met.
71 Specifically, the contributions of this paper are

- 72 • an iterative inverse kinodynamics (IKD) solver with fast
73 convergence properties;
- 74 • details on how to implement our IKD solver for a wide
75 range of scenarios: constraints on full poses or hands and
76 feet in character animation, position and velocity con-
77 straints, multiple overlapping constraints, and constraints
78 on passive deformable objects and in facial animation;
- 79 • an automatic method for selecting kinodynamic time
80 windows from the physical parameters of the system and
81 acceleration bounds on the kinematic trajectory.

82 We also discuss limitations, timings, convergence rates, and we
83 describe a typical work flow example that an animator would
84 use to create an animation with our system.

85 2. Related work

86 Secondary dynamics provides a significant amount of visual
87 realism in kinematically driven animations and is an important
88 technique for animators. In the case of tissue deformations pro-
89 duced by the motion of an underlying skeleton, various methods
90 can be used to produce this motion through simulation or using
91 precomputation [1, 2, 3]. With respect to secondary dynam-
92 ics of skeletal motion, it has similarly been demonstrated that
93 tension and relaxation of the skeletal animation can be altered
94 through physically based simulation [6]. These techniques pro-
95 vide an important richness to an animation; while the style of
96 the results are controllable by adjusting the elastic parameters
97 or gains of controllers used for tracking, precise control of the
98 motion itself to satisfy given constraints or key frames is typi-
99 cally left as a separate problem.

100 In contrast to the simulations that provide secondary dy-
101 namics, it is the direct local control and WYSIWYG interface
102 provided by forward and inverse kinematics techniques that ani-
103 mators primarily use in the creation of character animation.
104 As a result, there has been a vast amount of research on inverse
105 kinematics over the years, for instance, combining direct and in-
106 verse control during editing [7], using nonlinear programming

107 [8], using priority levels to manage conflicting constraints [9],
108 or alternatively, singularity-robust inverse computations [10] or
109 damped least squares [11]. In this work our focus is on a prob-
110 lem similar to inverse kinematics, but in the new setting of kin-
111 odynamics. We leave the issue of solving inverse kinodynamics
112 in the presence of conflicting constraints for future work.

113 The rest of the related work can be categorized into two
114 groups. First, there are approaches which try to control a phys-
115 ically based simulation to have it meet some desired constraints.
116 Second, there are approaches which use kinematic editing tech-
117 niques to produce animations that meet desired constraints and
118 exhibit physical plausibility.

119 Controlling physically based simulations is a difficult prob-
120 lem. There has been a significant amount of work in this area
121 on controlling rigid bodies [12, 13], fluids [14, 15], and cloth
122 [16]. Other recent successes on controlling physically based
123 animation use gentle forces to guide an animation along a de-
124 sired trajectory, accurately achieving desired states, but also
125 allowing physical responses to perturbations [17]. Physically
126 based articulated character control has received a vast amount
127 of interest. Building on the seminal work of locomotion con-
128 trol [18], it is now possible to have, for instance, animation
129 of physically based motions that respond naturally to perturba-
130 tions [19, 20, 21], maintain balance during locomotion [22, 23],
131 and editable animations of dynamic manipulations that respect
132 the dynamic interactions between characters and objects [24].
133 Our work is very different from these approaches, and is in-
134 stead more closely related to work by Allen et al. [25], which
135 changes PD control parameters to produce skeletal animations
136 that interpolate key-frames at specific times. In our work, how-
137 ever, we keep the control parameters fixed and alter the kine-
138 matic trajectory.

139 Jain and Liu [26] show a method for interactively editing
140 interaction between physically based objects and a human. In
141 this work, it is the motion of the dynamic environment which is
142 edited through kinematic changes of a captured human motion.
143 In comparison, we focus on altering and editing a kinodynamic
144 motion with different styles (tension and relaxation) and dif-
145 ferent constraints. Directly related to the problem of authoring
146 motion, physically correct motion can be achieved by solving
147 optimizations with space-time constraints [27]. Also relevant
148 is work that uses analytic PD control trajectories for compliant
149 interpolation [28], and work on generating physically plausible
150 motion from infeasible reference motion using a dynamics filter
151 [29].

152 In contrast to the work on controlling fully dynamic simu-
153 lations, we are addressing a simplified problem due to the finite
154 time window involved in simulating the state at a given time in
155 a kinodynamic trajectory. This leads to benefits in the context
156 of animation authoring, and allows for a straightforward solu-
157 tion to the inverse kinodynamic problem that we present in this
158 paper. We use correction curves with a limited temporal width
159 to solve the IKD problem. In a scenario with many constraints
160 at different times, the correction resembles a smooth displace-
161 ment function, such as the B-splines that are commonly used
162 in solving spacetime constraint problems (for instance, in the
163 motion editing work of Gleicher [30]). An important difference

164 is that our displacements are applied on the reference kinematic
165 trajectory, thus the final motion is the product of a physical simu-
166 lation as opposed to a displaced motion that satisfies physical
167 constraints. In a different approach, with similar objectives to
168 our own work, Kass and Anderson [31] propose a method for
169 including physically based secondary dynamics in a key frame
170 style editing environment through interactive solutions of space
171 time optimization problems. They focus on linear or linearized
172 space-time constraints problems, while our work, in contrast,
173 looks primarily at non-linear problems such as skeletal anima-
174 tion.

175 Specifically with respect to inverse kinematics, Boulic
176 et al. [32] include a control of the center of mass that gener-
177 ates character postures with improved plausibility by satisfy-
178 ing static balance constraints. Also within a purely kinematic
179 setting, Coleman et al. [33] create handles to edit motion ex-
180 trema of different joints clustered in time. The visual impact
181 of secondary dynamics is often captured in these temporal re-
182 lationships. Such an approach can, however, only exaggerate
183 or diminish a dynamic effect already present in the motion and
184 cannot introduce new forces and dynamic behaviors that the
185 mixing of kinematics and dynamics allows.

186 Mixing kinematics and dynamics to get the best of both
187 has promise for authoring motion in real time. For instance,
188 Nguyen et al. [34] blend kinematic animation and dynamic ani-
189 mation via a set of forces which act like puppet strings to
190 pull the character back to the kinematic trajectory. Also of
191 note is work on editing kinematic motion through momen-
192 tum and force [35], or with biomechanically inspired con-
193 straints [36]. While these different approaches use dynamic
194 principles to control accelerations and velocity, they deal with
195 dynamic systems which are not necessarily short lived, and
196 these approaches do not share our objective of a scrubbing in-
197 terface for animation editing which computes states largely in a
198 history free manner.

199 3. Overview

200 In this section, we provide an overview of our approach.
201 The animation is principally driven by a kinematic trajectory
202 $\mathbf{x}_K(t)$, typically authored and edited using traditional keyframe
203 and motion capture techniques. The kinodynamic trajectory of
204 the system $\mathbf{x}_{KD}(t)$ at a time t is the result of a physical simu-
205 lation run over a time window δ starting from an initial posi-
206 tion $\mathbf{x}_K(t - \delta)$ and velocity $\dot{\mathbf{x}}_K(t - \delta)$. The simulation uses a
207 PD (Proportional-Derivative) controller to follow the kinematic
208 trajectory, so the $\mathbf{x}_K(t)$ can be thought of as the target or desired
209 trajectory. The PD controller applies forces to the system that
210 are proportional to the difference between the set point \mathbf{x}_K
211 and the process variable \mathbf{x} . We also apply viscous damping, thus
212 the forces can be written as $K_p(\mathbf{x}_K - \mathbf{x}) - K_d \dot{\mathbf{x}}$, where the gain
213 K_p can be seen as modeling tension and relaxation, while K_d
214 controls damping. Figure 1 shows an example of KD trajec-
215 tories computed with different time windows, which can also be
216 seen in the supplementary video. Note how the history-free KD
217 trajectories capture the visual behavior of the actual dynamics
218 over a range of window sizes.

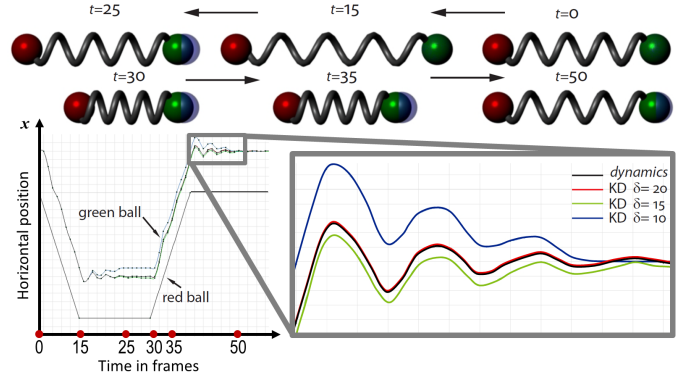


Figure 1: KD trajectories for the green ball: Kinematically the green ball is rigidly connected to the keyframed red ball, with spring dynamics overlaid. A number of frames of the KD trajectory ($\delta = 15$) are shown, with the full dynamics solution for the green ball overlaid in blue (top). KD trajectories with 3 window sizes are shown in relation to a full dynamics solution (bottom).

219 We will have kinodynamic states which deviate from the
220 kinematic trajectory because we are using a simulation with
221 control forces to generate the KD trajectory. This is desirable
222 because we want to include the effects of secondary dynamics
223 in the animation. However, there may be specific times in the
224 animation where we need constraints to be met.

225 Suppose target pose \mathbf{x}_i must be produced at time t_i . This
226 target state could be a pose in the original kinematic trajectory,
227 or something different. If the pose belongs to the original kine-
228 matic trajectory, a simple solution would be to stiffen the PD
229 control in the vicinity of the desired pose so that it is tracked
230 precisely. Note, however, that stiffness is an inherent attribute
231 of the motion’s secondary dynamics under animator control and
232 altering it to interpolate a target pose imbues the animation with
233 a different style. Instead, we iteratively compute a modification
234 to the kinematic trajectory which results in a KD state that sat-
235 isfies the constraint.

236 Example trajectories and an IKD modification are illus-
237 trated in Figure 2, where a red kinodynamic trajectory follows
238 a green kinematic trajectory (suppose it is lower due to grav-
239 ity). At left we can see an illustration of how the time window
240 for computing kinodynamic state must be long enough for any
241 impulse (smaller than a given maximum) to come sufficiently
242 close to rest that it can not be perceived (for instance, based on
243 screen pixels or a percentage of the object size). At right in
244 the figure we can see a dotted green kinematic trajectory with
245 an added bell shape correction, which produces the dotted red
246 kinodynamic trajectory satisfying the constraint at time t_i . This
247 smooth modification of the kinematic trajectory is the approach
248 we use to solve the IKD problem. We use a Gaussian shaped
249 correction curve in this work, but a variety of artist designed
250 curves that define a smooth ease-in ease-out correction can be
251 used, as described in Section 6.

252 3.1. Inverse kinodynamics

Here we formalize our core approach to solving the IKD
problem. Let $\text{SimulateKD}(\mathbf{x}_K, t_i, \delta)$ be the procedure for com-
puting \mathbf{x}_{KD_i} , the KD state at t_i for kinematic trajectory \mathbf{x}_K . We




Figure 2: An illustration of how we modify a kinematic trajectory to create a kinodynamic trajectory that satisfies the constraint that the original kinematic state be produced at time t_i .

first compute the IKD error in meeting the target as

$$\mathbf{e}_i = \mathbf{x}_i - \mathbf{x}_{KD_i}, \quad (1)$$

and from this we form bell shaped correction curves $\mathbf{e}_i \phi_i(t)$ that we add to kinematic trajectory (note that \mathbf{e}_i is a vector of same dimension as the state, and each coordinate of the state will have a bell shaped correction of a different magnitude). The bell shaped basis function $\phi_i(t)$ provides a local correction, has its peak value of 1 at t_i , and can be defined as a low degree polynomial or Gaussian. More importantly, it has a local support (a small temporal width, σ) which is selected by the artist. Conceptually, this IKD error correction introduces an additional spring force proportional to $\mathbf{e}_i \phi_i(t)$ in a small temporal neighborhood around t_i . This correction will not be sufficient, however, and our modified KD state $\tilde{\mathbf{x}}_{KD_i} = \text{SimulateKD}(\mathbf{x}_K + \mathbf{e}_i \phi_i, t_i, \delta)$ will not meet the constraint. This is because the correction did not take into account the dynamics of the system, but we can fix this by boosting the correction to account for the dynamics, assuming that the system dynamics are approximately locally linear. Letting $\mathbf{d}_i = \tilde{\mathbf{x}}_{KD_i} - \mathbf{x}_{KD_i}$, we project the error onto this initial correction result to compute the scaled correction

$$\mathbf{f}(t) = \Delta \mathbf{x}_i \phi_i(t), \quad (2)$$

where $\Delta \mathbf{x}_i = (\mathbf{e}_i \cdot \mathbf{d}_i / \|\mathbf{d}_i\|^2) \mathbf{e}_i$. Figure 3 shows a 2D illustration of how this scaling encourages good progress toward the target on each iteration. Without this linear prediction step, the convergence is significantly slower.

Using $\mathbf{x}_K + \mathbf{f}$, the process is repeated, until the system state converges to within a numerical threshold of \mathbf{x}_i at t_i . That is, we find the new kinodynamic state at t_i , compute the error \mathbf{e}_i , the modified kinodynamic state using $\mathbf{x}_K + \mathbf{f} + \mathbf{e}_i \phi_i$, the correction result \mathbf{d}_i , and finally an update to the correction function

$$\Delta \mathbf{x}_i \leftarrow \Delta \mathbf{x}_i + (\mathbf{e}_i \cdot \mathbf{d}_i / \|\mathbf{d}_i\|^2) \mathbf{e}_i. \quad (3)$$

4. KD animation and IKD scenarios

In this paper, we look at a number of scenarios that can largely be described as either pose constraints (as described




Figure 3: An illustration for a 2D state of two steps of the IKD iteration for a constraint at time t_i . At bottom left, \mathbf{x}_{KD_0} is the initial KD state at time t_i , which is far from target \mathbf{x}_i . A correction based on \mathbf{e} results in the modified KD state $\tilde{\mathbf{x}}_{KD_1}$, which does not take into account the system dynamics. We project the error onto $\mathbf{d}_i = \tilde{\mathbf{x}}_{KD_1} - \mathbf{x}_{KD_0}$ to determine a scaling of the correction that would produce a KD state as close as possible to \mathbf{x}_i assuming linear system dynamics. Using the scaled correction (Equation 3), we produce the new KD state \mathbf{x}_{KD_1} , and repeat the process until $\|\mathbf{e}\|$ falls below a threshold.

above) or end effector constraints (Section 4.1). For instance, we may want a kinodynamic skeletal animation of a dance to produce some key poses, or a kinodynamic skeletal animation of a punch that actually hits the desired target at a specific time. Alternatively, another scenario which is important to consider is the case where we drive a deformable mesh animation to follow a target mesh animation. In contrast to joint angles, this case involves a state vector formed by the Cartesian position of vertices in the mesh. In Section 4.5, we show how this approach can be used for facial animation.

We note that the blending of the correction can be done in a number of ways. If we only have one position constraint to satisfy in the entire animation, then it would be possible to naively apply a constant offset to the kinematic trajectory in order to meet the constraint at time t_i . Typically we will have several constraints at different times, so we only make a local edit to the desired trajectory (see Section 4.2). Any of a number of smoothly shaped curves with compact support will serve this purpose, as discussed in Section 6. The shape and width of the correction basis functions are an important artist control, much like setting ease-in ease-out properties in a key frame animation.

4.1. Skeletal animation end effector IKD

In the case of an articulated character, the state \mathbf{x} is a set of joint angles, and the simulation uses a PD controller to follow the kinematic trajectory. The gains of the controller set the level of tension or relaxation of the character [6].

When editing a skeletal motion, we may wish to set constraints on the entire pose, as described above, but it is also important that we are able to constrain only part of the state at

290 the time of a contact event, for instance, a point on a hand or
 291 foot (we will call such points *end effectors*). Suppose the end
 292 effector position of an articulated character is given by $p(\mathbf{x})$,
 293 and that it must reach position p_i at time t_i . In this case, we
 294 have the constraint $p(\mathbf{x}_{KD_i}) = p_i$, and we use an inverse kinematics
 295 solution to map the end effector error to an error in the
 296 state.

297 Figure 4 shows an example of how we solve the IKD prob-
 298 lem of punching a target. While the motion in Figure 4(a) hits
 299 the target at the desired time, we change the motion style by ad-
 300 justing the tracking gains of the physically simulated character
 301 shown in orange, to produce the more relaxed KD motion show
 302 in Figure 4(b). This relaxed motion fails to hit the target, but we
 303 can solve an inverse kinematics problem to adjust the joints of
 304 our relaxed character so that the end effector does hit the target.
 305 This IK solution pose is shown in dark blue in Figure 4(c). We
 306 could make a purely kinematic fix to our KD trajectory by sim-
 307 ply layering this IK solution on top our KD trajectory, using a
 308 bell shaped curve to slowly ease the correction in and out. How-
 309 ever, this does not respect the relaxed dynamics of the character
 310 (see the accompanying video). Instead, we modify the kinematic
 311 trajectory used to produce the kinodynamic animation.
 312 By editing the kinematic trajectory, we produce a natural look-
 313 ing motion that exhibits a relaxed style with a follow through
 314 motion. This modification is shown in Figure 4(d) and (e) for
 315 two bell shaped correction curves of different widths.

The algorithm iterates as described in the previous section,
 using an update to the correction curve that is based on an in-
 verse kinematics solution,

$$e_i = \text{SolveIK}(\mathbf{x}_{KD_i}, p_i) \quad (4)$$

where SolveIK computes a state displacement e_i such that
 $p(\mathbf{x}_{KD_i} + e_i) = p_i$. Note that the correction function up-
 date must be modified to use the end effector error, $\Delta p_i =$
 $p_i - p(\mathbf{x}_{KD_i})$, instead of the state displacement e_i . The update
 becomes

$$\Delta \mathbf{x}_i \leftarrow \Delta \mathbf{x}_i + (\Delta p_i \cdot d_i / \|d_i\|^2) e_i. \quad (5)$$

316 where $d_i = p(\tilde{\mathbf{x}}_{KD_i}) - p(\mathbf{x}_{KD_i})$.

317 4.2. Multiple constraints

318 The process of designing an animation typically involves
 319 setting multiple constraints at different times throughout the an-
 320 imation. If these events are sufficiently far apart, we can treat
 321 each as an independent IKD problem. However, constraints
 322 in close temporal proximity may need to be solved simultane-
 323 ously. For instance, we might want an animation of a drummer
 324 to hit a set of cymbals at precise times in quick succession (see
 325 also Figure 7-b for an example that involves hitting targets on
 326 a control panel). Note that two end effectors, for instance two
 327 hands, each with a constraint at the *same* time, can be solved
 328 with the technique in Section 4.1 by letting the IK problem in-
 329 volve two end effectors; it is only when constraints are at dif-
 330 ferent times that we need a different approach.

331 Temporally coupled constraints can happen in a variety of
 332 ways. If the bell shaped correction curve necessary to satisfy
 333 one constraint modifies kinematic states that fall within the time

334 window used to simulate the KD state at a another constraint,
 335 then the solution of the latter constraint will depend on the
 336 solution of the former. This dependence can be one way, or
 337 both ways, depending on the temporal width of the bell shaped
 338 curves used for each constraint, and the KD time window size
 339 used for the kinodynamic simulation.

340 While we may typically be able to solve most constraints
 341 independently, if the KD time window is large and the tempo-
 342 ral width of the correction curve is small then we may need to
 343 solve the constraints in order from earliest to latest. Here, we
 344 specifically consider the harder case where the temporal width
 345 is large enough to cause the constraints to be temporally cou-
 346 pled, requiring the constraints to be solved simultaneously. In
 347 our multi-constraint examples, we typically choose correction
 348 function widths that provide an ease-in trajectory with a dura-
 349 tion of approximately one or two seconds, so constraints that
 350 fall within one or two seconds of one another will need to be
 351 addressed simultaneously.

The correction f that we must add to the kinematic state to
 satisfy a number of constraints can now be seen as an interpo-
 lation function. That is, f interpolates a set of corrections $\Delta \mathbf{x}_i$ at
 t_i , for $i = 1..N$ where N is the number of constraints. We imple-
 ment this interpolated correction function using a sum of basis
 functions,

$$f(t) = \sum_i^N \lambda_i \phi_i(t), \quad (6)$$

where the basis function coefficients λ_i are computed by solv-
 ing a linear system of equations,

$$f(t_i) = \Delta \mathbf{x}_i, \text{ for } i = 1..N. \quad (7)$$

352 Note that the coefficients λ_i are vectors with the same dimen-
 353 sion as f , that is, the dimension of the state.

354 At each iteration, the multi-constraint IKD solver must pro-
 355 duce an appropriate update to each $\Delta \mathbf{x}_i$. In Section 3.1, we ef-
 356 fectively computed numerical partial derivatives with respect to
 357 the bell shaped basis magnitudes, and found a least squares so-
 358 lution for the desired update using a projection (computed with
 359 a dot product). In the case of two constraints i and j we have d_i
 360 influenced by an adjustment for constraint j , but we avoid the
 361 expense of computing these relationships by fixing only one
 362 constraint at a time. Thus we have an inner loop that consists of
 363 computing the KD state at t_i , the error e_i , the updated KD state
 364 for trajectory $\mathbf{x}_K + f + e_i \phi_i$, the update for $\Delta \mathbf{x}_i$ (using Equation 3
 365 or Equation 5), and then finally we recompute the interpolation
 366 function weights. This approach, similar to Gauss Seidel iter-
 367 ation, works well because the effect of ϕ_i on \mathbf{x}_{KD_i} is typically
 368 much larger than at \mathbf{x}_{KD_j} .

369 The technique we use to solve the multi-constraint IKD
 370 problem is summarized in Algorithm 1, and consists of a nested
 371 loop of adjusting the correction f to fix each of the violated con-
 372 straints, until all constraints are sufficiently satisfied or a maxi-
 373 mum number of iterations is reached.

374 Solving for the basis function coefficients λ_i is fast. To
 375 solve the interpolation function, we can compute an LU decom-
 376 position, from which we can find λ_i using a back solve. Re-
 377 peated solves of the interpolation function can be done quickly




Figure 4: Illustration of how IKD is used to produce an animation of a relaxed character that punches a target. (a) shows the motion capture at the time of contact in both wire-frame and solid orange. (b) the solid orange character shows the KD state of the relaxed character, which fails to reach the target at the time of contact. (c) inverse kinematics produces the pose of the character in dark blue. (d) iteratively computing the error and modifying the kinematic trajectory produces a KD state which hits the target (orange). Here the modified motion capture pose is shown in wire-frame. (e) shows the result of using a smaller temporal width for the bell shaped correction curve, which results in more of an upper cut.

Algorithm 1 Inverse Kinodynamics Multi-Constraint Solve

Input: constraints \mathbf{x}_i or p_i at t_i , for $i = 1..N$, δ

Output: state correction curve f

```

1:  $itr \leftarrow 0$ 
2:  $E \leftarrow \infty$ 
3:  $\Delta \mathbf{x}_i \leftarrow 0$ , for  $i = 1..N$ 
4:  $f \leftarrow \text{SolveInterpolation}(\Delta \mathbf{x})$ 
5: while  $itr++ < \text{maximum}$  and  $E > \text{threshold}$  do
6:   for  $i = 1 \rightarrow N$  do
7:      $\mathbf{x}_{KD_i} \leftarrow \text{SimulateKD}(\mathbf{x}_K + f, t_i, \delta)$ 
8:      $\mathbf{e}_i \leftarrow \text{compute using Equation 1 or 4}$ 
9:      $\tilde{\mathbf{x}}_{KD_i} \leftarrow \text{SimulateKD}(\mathbf{x}_K + f + \mathbf{e}_i \phi_i, t_i, \delta)$ 
10:     $\Delta \mathbf{x}_i \leftarrow \text{compute using Equation 3 or 5}$ 
11:     $f \leftarrow \text{SolveInterpolation}(\Delta \mathbf{x})$ 
12:   end for
13:    $E \leftarrow 0$ 
14:   for  $i = 1 \rightarrow N$  do
15:      $\mathbf{x}_{KD_i} \leftarrow \text{SimulateKD}(\mathbf{x}_K + f, t_i, \delta)$ 
16:      $E \leftarrow E + \|p(\mathbf{x}_{KD_i}) - p_i\|$  or  $\|\mathbf{x}_i - \mathbf{x}_{KD_i}\|$ 
17:   end for
18: end while

```

378 because we can reuse the same decomposition (the basis func-
 379 tions and their centers do not change).

380 Note that the inner loop update could skip an update for a
 381 given constraint if its contribution to the error was known to be
 382 small. However, the size of this error can only be verified by
 383 recomputing the KD state as it is influenced by other changes
 384 to f . The computation of \mathbf{x}_{KD_i} is the bulk of the cost.

385 4.3. Constraining velocities

386 When constraining a pose or an end effector position, we
 387 might also want to set constraints on velocities. For instance,
 388 we may want the hand of a character to tap the surface of a

389 stationary object without penetrating the surface. The hand
 390 end effector must satisfy both position and velocity constraints,
 391 meaning it must reach the target at the time of contact and have
 392 zero velocity at that time. Alternatively, the desired velocity
 393 can follow the original velocity of the animation, or can be set
 394 to achieve a different velocity at the time of the constraint.

395 We can solve the IKD problem for constrained velocities in
 396 a similar manner to the position problem, and likewise solve
 397 for simultaneously constrained position and velocity. Again,
 398 the IKD solution comes from layering a correction overtop of
 399 the kinematic trajectory.

Suppose that at time t_i we have desired state velocity $\dot{\mathbf{x}}_i$
 (or alternatively, a desired end effector velocity \dot{p}_i). Instead of
 adding a bell shaped curve to change the velocity, we will add
 a wiggle to change the velocity $\dot{\mathbf{x}}_K(t_i)$ without changing $\mathbf{x}_K(t_i)$.
 We use the derivative of the bell shaped position correction ba-
 sis function as a basis function for setting the derivative,

$$\psi_i(t) = \frac{\partial}{\partial t} \phi_i(t), \quad (8)$$

400 though this function could likewise be selected by the animator.

For simplicity, suppose we are dealing with a set of N pairs
 of constraints, that is, constraints on both position and velocity
 at times t_i , for $i = 1..N$. To deal with position and velocity con-
 straints in close proximity we use an interpolation of the correc-
 tions $\Delta \mathbf{x}_i$ with velocities $\Delta \dot{\mathbf{x}}_i$ necessary to correct the kinematic
 trajectory. Thus, the interpolation function has the form

$$f(t) = \sum_i^N (\lambda_i \phi_i(t) + \beta_i \psi_i(t)). \quad (9)$$

Again, the basis function coefficients λ and β can be found by
 solving the system of $2N$ linear equations for each dimension
 of the state, given by the required corrections and correction

velocities:

$$\mathbf{f}(t_i) = \Delta \mathbf{x}_i, \text{ for } i = 1..N, \quad (10)$$

$$\frac{\partial \mathbf{f}(t_i)}{\partial t} = \Delta \dot{\mathbf{x}}_i, \text{ for } i = 1..N. \quad (11)$$

It is important to observe that we update the desired velocity correction $\Delta \dot{\mathbf{x}}_j$ by comparing the desired velocity $\dot{\mathbf{x}}$ with the velocity of the KD trajectory. The velocity of the dynamic simulation which produces $\mathbf{x}_{KD}(t)$ does *not* give us this KD velocity (it is not the dynamic simulation velocity that we want to control). Instead, we must approximate this KD state velocity from successive frames of the KD state,

$$\dot{\mathbf{x}}_{KD}(t_i) \approx \frac{1}{h}(\mathbf{x}_{KD}(t_i) - \mathbf{x}_{KD}(t_i - h)). \quad (12)$$

We measure the difference to set the velocity error $\dot{\mathbf{e}}_i$, with which we compute a new KD state, and ultimately find an update to the required velocity correction $\Delta \dot{\mathbf{x}}_i$, using a computation similar to Equation 3.

In the above example, we are considering a target velocity on the entire state. If instead our constraint is only on the end effector of a skeleton, then the approach is slightly different. In this case, we compute the approximate KD end effector velocity,

$$\dot{p}_{KD}(t_i) \approx \frac{1}{h}(p(\mathbf{x}_{KD}(t_i)) - p(\mathbf{x}_{KD}(t_i - h))). \quad (13)$$

The difference between this velocity and the artist requested end effector velocity \dot{p}_i is then mapped to a state error,

$$\dot{\mathbf{e}}_i = J^+(\dot{p}_i - \dot{p}_{KD}(t_i)). \quad (14)$$

where J^+ is a pseudoinverse of the end effector Jacobian $J = \partial p / \partial \mathbf{x}$ evaluated at pose $\mathbf{x}_{KD}(t_i)$. Again, this error is used to update the required velocity correction $\Delta \dot{\mathbf{x}}_i$, and the process is repeated until our IKD algorithm has converged or we reach a maximum number of iterations.

4.4. Passive deformable object IKD

The previous sections present solutions that deal with constraints on skeletal motion. While this plays an important role in character animation, animators may wish to have more control over passive deformable objects attached to kinematically driven bodies. Controlling secondary dynamics of such deformations can be tricky since the motion can only be indirectly edited by changing the kinematic motion that drives the secondary dynamics. Fortunately, the same techniques can be applied to solve IKD problems on deformable objects.

For example, consider the scenario shown in Figure 7-d, where a character with a floppy hat must walk through a door without the hat hitting any part of the door frame. In this case, the hat would hit the top of door frame at time t_i . We can use IKD to set a constraint that the tip of the hat must be at a position just below the door frame at time t_i . In our example, the hat is a passive elastic system rigidly attached to the head of the character, and the character motion is purely kinematic (and it must be altered to change the trajectory of the tip of the hat).

IKD can be used to control the tip of the hat using the skeletal IKD technique described in Section 4.1, with a small adjustment. At each iteration of the IKD algorithm, we fix the end effector position. In this example it is the tip of the hat, based on its kinodynamic position at time t_i , even though it is not fixed with respect to the character’s head. The IKD solution involves a change in the character’s posture, allowing the kinodynamic position of the tip of the hat to miss the door frame (see the accompanying video).

4.5. Dynamic blend shape IKD

The IKD problem for passive deformable objects discussed in Section 4.4 can be applied to elastic tissue deformation in a variety of scenarios. Particularly in the context of facial animation, deformation is tediously authored by animators by keyframing linearly blend shape targets [37]. Overlaying secondary jiggle and other dynamic nuance currently comes at the cost of letting dynamics have the “final word” on the animation, with no guarantees of hitting certain expressions. IKD allows one to overlay this desired secondary dynamics in a kinematic setting and further specify critical poses as target shapes to be precisely interpolated, independent of the kinematically authored blend shape animation. A loosened facial animation can also be kept in sync with the environment (like taking a puff from a cigarette or sip from a glass) or an audio track by adding checkpoints from the kinematic trajectory as IKD targets, so the final facial trajectory has a limited deviation from the kinematic input. We implement IKD as a deformation that tracks control points on a shape using springs and dampers as in work by Müller et al. [38]. Figure 5 shows examples of pose constraints applied to a kinodynamic trajectory for two different characters.

The inverse kinodynamic solution follows the same algorithm presented above, with the correction update following Equation 3, which is very easy to compute as we simply need the difference between the kinodynamic state and the target. The techniques for dealing with multiple constraints and velocity constraints are likewise similar to those describe above.

5. Time window selection

Setting the size of the KD time window δ has an important influence on the quality and cost of the kinodynamic trajectory. We want a small window to make it cost effective to simulate kinodynamic states on the fly, but the window also needs to be long enough to produce the desired secondary dynamics effects.

5.1. Manual time window selection

We find that it is easy to select a reasonable window by hand. Given fixed gains for the PD controller, this can be done by simulating the physical system in response to an impulse and visually selecting the time at which vibrations are no longer visible. Instead of an isolated impulse, we typically focus on a maximum acceleration or an abrupt change in the kinematic trajectory defined by a motion capture clip or a keyframe animation. We alternate between adjusting the KD time window

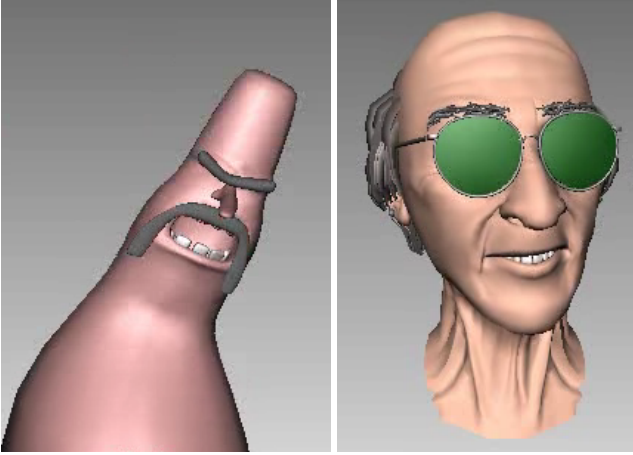


Figure 5: Two facial animation IKD examples (see accompanying video). Left, a temporal pose constraint produces a head tilt. Right, a temporal pose constraint produces a smile.

481 size, and scrubbing back and forth on the time line to observe
 482 the results. We stop adjusting the window size when we are sat-
 483 isfied that we have selected the smallest window that does not
 484 prematurely truncate damped vibrations caused by the maxi-
 485 mum acceleration in the kinematic trajectory.

486 5.2. Automatic time window selection for 1D systems

The KD time window can also be set automatically based on a computation that takes into account both the maximum acceleration of the kinematic animation and the physical properties of the kynodynamic system. Here we first consider the one dimensional example shown in Figure 1 (we will extend this to more complex examples in the next section). The system consists of a particle of mass m attached to a damped spring, the equation of motion is

$$m\ddot{x} + c\dot{x} + kx = 0, \quad (15)$$

487 and suppose the initial position be zero, $x_0 = 0$, which is like-
 488 wise the equilibrium position of this system.

Let J_{\max} be the maximum impulse that our particle experiences when following a given kinematic trajectory. We can compute this from the maximum acceleration in the kinematic trajectory as $J_{\max} = \ddot{x}_{\max}hm$ where h is the simulation step size. We want to know how long it will take for the system in Equation 15 to come to rest if activated by this maximum impulse. Here, we define the time at which the system comes to rest as the time, after which the trajectory will always remain within some small distance ε from the equilibrium position. We can solve this by first obtaining an analytic solution for the trajectory [39] using the maximum impulse to define an initial condition on the velocity $\dot{x}_0 = J_{\max}/m$. We are typically interested in mass, stiffness, and damping settings that result in an underdamped system, in which case $c^2 < 4mk$, giving a solution of

the form

$$x(t) = Ae^{\gamma t} \sin(\omega t), \quad (16)$$

$$\gamma = -\frac{c}{2m}, \quad (17)$$

$$\omega = \frac{\sqrt{4mk - c^2}}{2m}, \quad (18)$$

$$A = \frac{\dot{x}_0}{\omega}, \quad (19)$$

where γ is the decay rate, ω is the frequency, and A is the amplitude. The function $Ae^{\gamma t}$ provides a bound on the magnitude of the simulation. Given a minimum perceptual magnitude ε , we can choose our time window δ as the time when the system appears to have come to rest, that is, by solving $Ae^{\gamma\delta} = \varepsilon$, obtaining

$$\delta = \frac{1}{\gamma} \ln\left(\frac{\varepsilon}{A}\right). \quad (20)$$

When the system is over damped, the solution can be written as a difference of exponentials,

$$x(t) = iA \left(e^{(\gamma-i\omega)t} - e^{(\gamma+i\omega)t} \right) \leq iAe^{(\gamma-i\omega)t}. \quad (21)$$

The function is always positive because the velocity initial condition is positive, and thus it can be bounded by the exponential with positive coefficient (which also has the slower decay due to a smaller exponent). The time window can again be computed easily with a logarithm

$$\delta = \frac{1}{\gamma - i\omega} \ln\left(\frac{\varepsilon}{iA}\right). \quad (22)$$

489 However, if the system is critically damped, then the trajectory
 490 has the form $x(t) = \dot{x}_0 e^{\gamma t}$, and the time window should be se-
 491 lected as the larger root of $\dot{x}_0 e^{\gamma t} - \varepsilon = 0$. Alternatively, we
 492 can create an approximate exponential bound by avoiding the
 493 critically damped case though a small adjustment of the system
 494 parameters.

495 5.3. Case study: time windows for multidimensional systems

496 While the one dimensional example above is useful for un-
 497 derstanding how we can choose a KD time window, we are in-
 498 terested in more complicated systems with multiple degrees of
 499 freedom, for instance, our floppy hat example described in Sec-
 500 tion 4.4. This system can be analyzed in a similar way, pro-
 501 vided that a linearized version of the system yields a reasonable
 502 prediction of the behavior. This is typical in many elastic sys-
 503 tems when deformations are small due to small forces or large
 504 stiffness. Using linearized forces at the equilibrium pose of the
 505 system, the idea is to diagonalize the system to break the prob-
 506 lem up into a set of independent oscillators [40, 2]. While we
 507 specifically look at elastic deformable objects as a case study,
 508 the same approach can be used for our examples of PD control
 509 driven skeletal animation by computing the vibration modes of
 510 the skeletal system [41, 42].

Consider, for instance, a 3D multidimensional deformable elastic system such as a finite element model, written in terms

of N nodal displacements $u \in \mathbb{R}^{3N}$, and recall that some nodes of this system are rigidly attached to a frame, or a “bone”, that is driven by a kinematic animation. Assuming for now that the bone is stationary, the equations of motion can be written as

$$M\ddot{u} + C\dot{u} + Ku = 0 \quad (23)$$

where M is the mass matrix, K is the stiffness computed as the force gradient at the equilibrium pose, and $C = (\alpha M + \beta K)$ is the damping matrix that we assume is Rayleigh. The eigenvalue decomposition of $M^{-1}K$ gives the modal vibration basis matrix U , which provides a change of coordinates $u = Uq$, and allows for the system to be written as a set of independent oscillators,

$$m_i\ddot{q}_i + c_i\dot{q}_i + k_iq_i = 0. \quad i = 1..n, \quad (24)$$

where $c_i = (\alpha m_i + \beta k_i)$. With initial conditions $q(0) = 0$ and $\dot{q}(0)$ nonzero, the solution in the under-damped case is

$$q_i(t) = A_i e^{\gamma_i t} \sin(\omega_i t), \quad (25)$$

$$\gamma_i = -\frac{c_i}{2m_i}, \quad (26)$$

$$\omega_i = \frac{\sqrt{4m_i k_i - c_i^2}}{2m_i}, \quad (27)$$

$$A_i = \frac{\dot{q}(0)}{\omega_i}. \quad (28)$$

511 The exponential term decays at a minimum rate given by α ,
 512 and higher frequency modes will have faster decays when β
 513 is nonzero. Because of this, we focus on a maximum activation
 514 (through an acceleration of the bone) of the lowest frequency
 515 mode, $i = 1$, in order to compute the time window. That is, in-
 516 stead of solving for the smallest δ such that $\sqrt{u^T u} < \epsilon$ for all
 517 $t > \delta$, we assume that $u(t)$ will approximately equal $U_1 q_1(t)$
 518 near our solution because the contribution of all the higher fre-
 519 quency modes will have long since become insignificant due to
 520 their much faster decay rates. Note that $\sqrt{q_1^T U_1^T U_1 q_1} = q_1$ be-
 521 cause the eigenvector U_1 is unit length, thus we solve $A_1 e^{\delta \gamma_1} =$
 522 ϵ for the time window size δ . Because the problem has been
 523 reduced to one dimension, we can use the same technique as
 524 in the previous section to solve for the time window in over-
 525 damped and critically damped cases too.

526 This leaves two points to clarify. First, the selection of the
 527 perceptual threshold ϵ , and second, the largest activation of this
 528 lowest mode that will let us specify the initial condition $\dot{q}_1(0)$.

529 We use the Euclidean norm to measure displacement from
 530 the rest pose, and it can be difficult to set the perceptual thresh-
 531 old because a model can be discretized with different numbers
 532 of nodes. We account for this by setting the threshold as the
 533 value where all nodes are displaced 0.1% of the object diameter
 534 d , or $\epsilon = 0.001d \sqrt{N}$.

To define the largest activation of the first modal vibration, we follow the work of James and Pai [2] in specifying the effect of a 6 dimensional spatial velocity $\phi = (\omega^T, v^T)^T$ on the nodal velocities. With node i at position r_i , the velocity of the nodal




Figure 6: Left, the finite element model cubes in which a christmas hat is embedded. Right, a number of superimposed frames from the simulation showing the vibration.

points can be written $\dot{u} = \Gamma \phi$ where

$$\Gamma = \begin{pmatrix} -[r_1] & I \\ \vdots & \\ -[r_N] & I \end{pmatrix}, \quad (29)$$

and $[r_i]$ denotes the skew symmetric matrix that computes the cross product $r_i \times$. The inverse mode matrix lets us compute the modal velocities due to the spatial velocity, $\dot{q} = U^{-1} \Gamma \phi$, and we are specifically interested in the first mode. Notice that the six values in the top row of $U^{-1} \Gamma$ allow us to determine how large \dot{q}_1 can be for a given ϕ . Let $h_{\omega_{max}}$ be the maximum among the absolute values of the first three components, and similarly $h_{v_{max}}$ for the last three components. If we bound the magnitude of the largest angular and linear acceleration of the kinematic bone, then we can define scalars ω_{max} and v_{max} as the magnitude of largest angular and linear velocity that can be obtained in a single time step when starting from zero. Finally, we can set the initial condition of the first mode velocity as

$$\dot{q}_1(0) = h_{\omega_{max}} \omega_{max} + h_{v_{max}} v_{max}. \quad (30)$$

535 This gives a conservative setting of the initial velocity of the
 536 first mode that corresponds to the maximum accelerations we
 537 are willing to have in our kinematic animation.

538 Figure 6 shows a finite element model hat, for which we
 539 compute a time window $\delta = 0.568$ seconds with the approach
 540 described above where we let $\omega_{max} = 4$ radians per second and
 541 $v_{max} = 2$ meters per second.

542 6. Results and discussion

543 In this section we discuss a number of aspects of our IKD
 544 solution. We discuss convergence of the IKD solver in Sec-
 545 tion 6.1, the choice of correction curves in Section 6.2, the lim-
 546 itations of artist designed contact constraints in Section 6.3,
 547 issues in setting KD time windows in Section 6.4, and the
 548 specifics of our prototype implementation and timings in Sec-
 549 tion 6.5.

550 Figure 7 shows snapshots of examples that highlight the dif-
 551 ferent scenarios demonstrating the IKD techniques presented in
 552 Sections 4 and 5. Please refer to the accompanying video to
 553 view the full animations. The video also includes a work flow
 554 example demonstrating the interactivity of our system for skele-
 555 tal animation IKD, which is also described in Section 6.6.




Figure 7: From left to right, (a) punch, (b) control panel, (c) YMCA’s “C” Pose , (d) passive deformable hat, (e) position constraint for grasp

556 6.1. Convergence

557 It is important to discuss the convergence of our IKD algorithm.
 558 If there are many abrupt motions in the kinematic
 559 trajectory, then the resulting simulation could be chaotic. As
 560 such, we might not expect a small change in the joint angles to
 561 produce a predictable result, even if we smoothly and slowly
 562 blended in and out of this desired trajectory. While we do not
 563 assume linear dynamics, we do assume that the function map-
 564 ping \mathbf{x}_K to \mathbf{x}_{KD} is smooth “enough” (as is the case for all of our
 565 example systems) to allow us to use a linear approximation of
 566 the dynamics in our IKD solver.

567 While the convergence rate of our IKD algorithm depends
 568 on the actual scenario, Figure 8 compares the convergence rates
 569 achieved using different temporal widths of the correction func-
 570 tion, for the target punching example from Figure 4. While con-
 571 vergence can be slower when very small temporal widths are
 572 used, the number of iterations can be reduced by damping the
 573 correction adjustment computed in Equation 3 or 5 (see curves
 574 marked *scaled* in Figure 8).

575 We find that IKD converges quickly under a wide variety
 576 of constraints. Figure 9 shows the error in cm after 6 iterations
 577 for varying target positions in the punch example. The error is
 578 small in all cases, except for target constraints which are phys-
 579 ically out of reach of the character.

580 We have noticed that our IKD algorithm can fail to make
 581 further progress once the error falls below 10^{-2} cm. We believe
 582 this is because we are using the Open Dynamics Engine (ODE)
 583 to compute the simulations that produce our KD states. While
 584 repeating simulations using the same initial conditions should
 585 produce the same results, aggressive optimizations within ODE
 586 make use of randomization. This does not present a problem as
 587 the error in end effector placement is significantly smaller than
 588 the overall size of the articulated character.

589 6.2. Correction curves

590 The shape of the correction curve we use to modify the kine-
 591 matic motion directly affects the motion which is produced. We
 592 use Gaussian shaped curves in our examples because they are
 593 simple and smooth. We effectively treat them as if they have
 594 compact support, and could easily use any other ease-in-ease-
 595 out curve of a desired shape and support, and we leave the se-
 596 lection of this curve to the animator. That is, the width of the




Figure 8: IKD convergence rates for the punch scenario using a bell shaped correction function with big temporal width (1 sec), and a small temporal width (0.5 sec), with error measured in cm, and error threshold 10^{-2} cm. IKD convergence can be slower when a small temporal width is used, but this can be improved slightly by damping the correction adjustment by a small amount, for instance, 0.8 in the examples labeled *scaled* in the legend.

597 Gaussian is selected by the animator; a wide curve will pro-
 598 duce a smooth anticipatory motion, while a short curve will
 599 produce a motion that abruptly moves to meet the constraint
 600 with a larger acceleration (and in turn, produces a larger follow
 601 through). While we only look at symmetric curves, any smooth
 602 artist created ease-in ease-out curve can be used. For instance,
 603 a non-symmetric correction curve can be designed to create a
 604 quick reaction followed by a slow return to the unmodified tra-
 605 jectory.

606 6.3. Contact constraints

607 While we are adding constraints to deal with contact, we re-
 608 quire the artist to specify these constraints. Although contacts
 609 may naturally happen in the dynamic simulations that produce
 610 our kinodynamic states, we will only have a “memory” of con-
 611 tacts that happen in the KD time window. For instance, we
 612 cannot correctly handle a braid of hair which is normally at rest
 613 down the back of a character but flips over a shoulder with the





Figure 9: IKD error in cm after 6 iterations for varying target locations (cm) in the punch scenario. The error consistently falls to less than 0.1 mm after 6 iterations, except when the target is nearly out of reach.

614 turn of a head. As such, these cases are comfortably handled as
 615 pure skeletally driven dynamics, but we hope to address such
 616 scenarios in the future by analysing collision events to adap-
 617 tively vary the KD window size.

618 6.4. KD time windows

619 In the discussion of time window selection in Section 5, we
 620 noted that the maximum acceleration in the kinematic trajectory
 621 will influence the size of the time windows (large accelerations
 622 will require longer time windows in order for oscillations pro-
 623 duced by these accelerations to become imperceptible). This is
 624 also true for the altered trajectory which includes the correction
 625 to solve a given IKD problem. We are using smooth bell shaped
 626 curves to add this displacement, so generally the accelerations
 627 due to the correction will be small. But if we set the temporal
 628 width of this curve to be small, then the IKD solution will need
 629 to involve a very large displacement to the kinematic trajectory
 630 to force the dynamic trajectory to the desired target, thus re-
 631 quiring larger time windows for computing a KD state. While
 632 we impose no explicit restrictions on the physical simulation of
 633 characters, our approach is largely suited to well-conditioned
 634 and continuous simulations.

635 The time window does not need to be the same for the entire
 636 time line and can be smaller when there are only small accel-
 637 erations in the kinematic trajectory. Figure 10 shows measure-
 638 ments we have made of the Euclidean error of the KD trajectory
 639 compared to a normal simulation, for a variety of time windows.
 640 The system shown here consists of a floppy elastic hat attached
 641 to the head of an articulated character driven by motion capture
 642 as shown in Figure 7-d. The stiffness and damping parameters
 643 were set quite low, requiring a relatively large time window of 4
 644 seconds before the KD trajectory to be imperceptibly different

645 from simulation for the accelerations in this example kinematic
 646 trajectory. However, the plot shows that there are moments dur-
 647 ing the animation where a much smaller KD time window can
 648 be used and still result in an imperceptible amount of error. This
 649 example illustrate the possible benefits from modulating the KD
 650 windows size across an animation, and is a problem we leave
 651 for future work.

652 6.5. Implementation and timings

653 Here we present timing details of the different implemen-
 654 tations we have tested. Note that the skeletal motion examples
 655 make use of motion capture that was recorded at 100 Hz, and
 656 we also produce KD states and IKD solutions based on this
 657 same frame rate as opposed to resampling and simulating at the
 658 final video rate. In contrast, the facial animation examples have
 659 a 25 Hz simulation and frame rate.

660 The IKD Skeletal animation examples were generated with
 661 our Java implementation which uses ODE (Open Dynamics En-
 662 gine) to simulate the forward dynamics. On an Intel Core i7
 663 3.2 GHz processor, the KD state takes roughly 0.01 s to gener-
 664 ate at each frame with a time window of 0.3 s, which allows for
 665 interactive scrubbing of the time line. Because the main cost
 666 of the IKD solution is the computation of KD states, an IKD
 667 solution for the punch scenario in Figure 4 can be computed in
 668 under half a second (that is, less than 50 iterations are required
 669 as shown in Figure 8).

670 IKD has also been implemented as a Maya 2011 deformer
 671 for control point shapes that track a kinematic trajectory using a
 672 spring and damper simulation. On an Intel i7 1.87 GHz proces-
 673 sor, the model at left in Figure 5 (approximately 1200 vertices)
 674 can be scrubbed at about 8 frames per second when using a
 675 KD time window of 1 s. We note that a large portion of the
 676 computation time is external to the KD algorithm, and a better
 677 optimized implementation would be possible by making use of
 678 the Maya plug-in application programming interface.

679 6.6. Interface and work flow example

680 A work flow example helps explain how animators can use
 681 IKD to create a desired animation and the types of controls they
 682 have for editing the animation. Figure 11 shows a snap shot
 683 of our Java implementation interface as it is being used during
 684 the example of creating an animation of a character punching a
 685 target.

686 The interface has a time frame slider that the animators can
 687 scrub back and forth to get to a desired frame of the anima-
 688 tion, as well as three buttons for playing the animation forward,
 689 backward, and stopping playback. With these controls, the ani-
 690 mator can play or scrub the animation to observe the character
 691 with secondary dynamics that exist in the kinodynamic trajec-
 692 tory. The interface allows various parameters to be adjusted,
 693 such as the physical parameters of the character, the kinody-
 694 namic window size, as well as the temporal width of the IKD
 695 correction function. Finally, there is a button that can be pressed
 696 to compute the IKD solution for any constraints that are set in
 697 the time line.

698 In our work flow example, the animator starts by loading
 699 an animation of a character throwing a punch. In this case, the

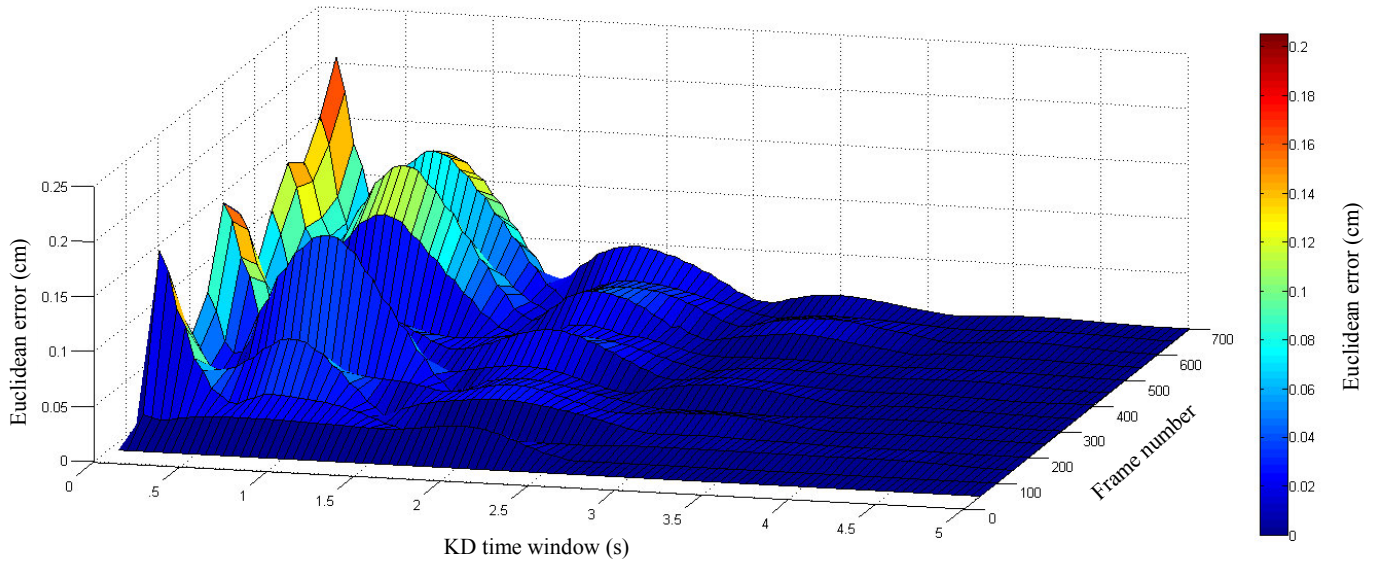


Figure 10: Euclidean error for the KD state of the hat example shown in Figure 7-d. The plots shows the error, a comparison between the KD state and the ground truth simulation, for different KD time window settings and across different frames of the animation. Notice that the KD time window necessary to meet a given level of error varies across the 7 seconds of animation (700 frames at 100 Hz).

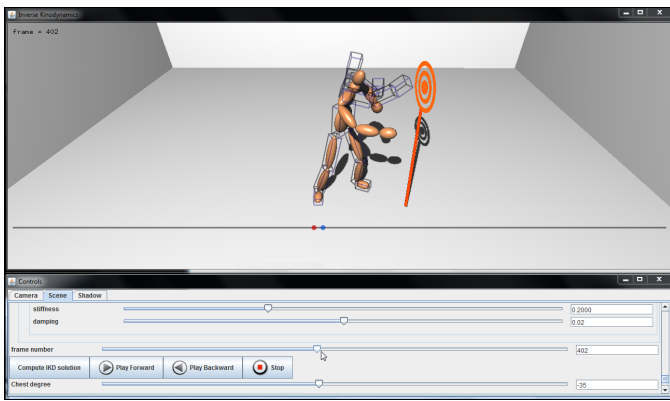


Figure 11: Work flow framework showing tools for animators to scrub back and forth in animation and change animation parameters.

700 animation is a motion capture clip, though it could likewise be
 701 a kinematic trajectory produced with key frames. In Figure 11,
 702 the main window shows the kinematic pose with a wire cube
 703 style, while the KD state of the character is rendered with orange
 704 ellipsoids. The animator can interactively adjust the stiffness
 705 of the joints and the kinodynamic window size and observe
 706 the influence of these changes at the current pose. The animator
 707 can also immediately see the influence of these changes on the
 708 entire animation by scrubbing back and forth in the time line. In
 709 our example, the animator lowers the stiffness of the character
 710 which gives a drunken boxing style to the animation. Modifications
 711 to the kinematic trajectory can also be made, and the results
 712 viewed immediately. Our work flow example also includes
 713 a modification of the bend in the back to have the drunken boxer
 714 stand straight, or stoop further. Again, the results of these mod-
 715 ifications can be seen immediately at the current pose, or viewed
 716 in the KD animation by scrubbing back and forth. This high-




Figure 12: (a) shows a KD state (orange) failing to punch a target. (b) is a modified motion capture pose (wire-frame) that produces a KD state which hits the target. The previous solution fails to punch a modified target position (c). A satisfied punch can be achieved by recomputing the IKD solution for the new target (d).

717 lights the WYSIWYG editing interface that we provide to the
 718 animator.

719 Once the animator has adjusted tension, relaxation, or other
 720 aspects of the style of the motion, there will still be other con-
 721 straints that need to be satisfied. In our work flow example,
 722 the animator would like to have the character punch a target at
 723 the desired time. Suppose the motion capture involves a punch
 724 at a target p_i at time t_i . The KD trajectory with the tension
 725 and relaxation parameters set to produce a drunken boxing style
 726 does not satisfy the target as shown in Figure 12-a. Setting the
 727 time line to the specific frame (t_i) and observing the error in the
 728 punch, the animator can use IKD to satisfy the constraint, produ-
 729 cing the solution shown in Figure 12-b. If a different target
 730 is desired, for instance, a lower target as shown in Figure 12-c,
 731 then the animator can change the constraint and use IKD again
 732 to satisfy the constraint as show in in Figure 12-d.

733 The ability to see the effect of edits in real-time is what
 734 makes the interface a powerful tool for animators. Besides char-
 735 acter animation, we also include deformable objects in our sim-
 736 ulation and control their pose, as discussed in Section 4.4.

737 7. Conclusions

738 While the IKD interface itself is not a contribution of our
739 work, a demonstration of the Maya implementation to a few
740 keyframe animators was positively received. From a work flow
741 standpoint, the animators felt they would have to consciously
742 omit keyframing dynamic nuances but this would be a welcome
743 change allowing them focus on the primary motion. For the
744 approach to be used in practice they expressed a need for in-
745 terface tools that make the addition and management of IKD
746 targets user friendly. Our current implementation, while inter-
747 active for skeletal animation, is only interactive for face blend
748 shapes with around 1000 control vertices. The vectorizable na-
749 ture of our algorithm, however, makes it a good candidate for
750 a faster GPU implementation. In future work we would like
751 to address the coupling of kinodynamic trajectories with fully
752 dynamic environments via adaptive kinodynamic window sizes
753 that are aware of collision events and other discontinuities in a
754 full physical simulation. In summary, we propose the concept
755 of Inverse Kinodynamics and present a first algorithm which
756 opens up new possibilities for editing traditional keyframe ani-
757 mations that are augmented with secondary dynamics.

758 Acknowledgments

759 This work was supported by NSERC, and GRAND NCE
760 Project MOTION, and originated from discussions at the Bel-
761 lairs Research Institute during the 2011 Workshop on Computer
762 Animation.

763 References

764 [1] Capell S, Green S, Curless B, Duchamp T, Popović Z. Interactive
765 skeleton-driven dynamic deformations. In: Proceedings of the 29th an-
766 nual conference on Computer graphics and interactive techniques. SIG-
767 GRAPH '02. ISBN 1-58113-521-1; 2002, p. 586–93.
768 [2] James DL, Pai DK. Dyr: dynamic response textures for real time de-
769 formation simulation with graphics hardware. In: SIGGRAPH. 2002, p.
770 582–5.
771 [3] Kim J, Pollard NS. Fast simulation of skeleton-driven deformable body
772 characters. *ACM Trans Graph* 2011;30:121:1–121:19.
773 [4] Angelidis A, Singh K. Kinodynamic skinning using volume-
774 preserving deformations. In: Proceedings of the 2007 ACM SIG-
775 GRAPH/Eurographics symposium on Computer animation. SCA '07;
776 Eurographics Association. ISBN 978-1-59593-624-4; 2007, p. 129–40.
777 [5] Donald B, Xavier P, Canny J, Reif J. Kinodynamic motion planning. *J*
778 *ACM* 1993;40:1048–66.
779 [6] Neff M, Eugene F. Modeling tension and relaxation for computer anima-
780 tion. In: Proceedings of the 2002 ACM SIGGRAPH/Eurographics sym-
781 posium on Computer animation. SCA '02; New York, NY, USA: ACM.
782 ISBN 1-58113-573-4; 2002, p. 81–8.
783 [7] Boulic R, Thalmann D. Combined direct and inverse kinematic con-
784 trol for articulated figure motion editing. *Computer Graphics For-*
785 *um* 1992;11(4):189–202. doi:10.1111/1467-8659.1140189. URL
786 <http://dx.doi.org/10.1111/1467-8659.1140189>.
787 [8] Zhao J, Badler NI. Inverse kinematics positioning using nonlinear
788 programming for highly articulated figures. *ACM Trans Graph*
789 1994;13(4):313–36.
790 [9] Baerlocher P, Boulic R. An inverse kinematics architecture enforcing an
791 arbitrary number of strict priority levels. *Vis Comput* 2004;20(6):402–17.
792 [10] Yamane K, Nakamura Y. Natural motion animation through constrain-
793 ing and deconstraining at will. *IEEE Transactions on Visualization and*
794 *Computer Graphics* 2003;9(3):352–60.

795 [11] Buss SR, Kim JS. Selectively damped least squares for inverse kinemat-
796 ics. *Journal of Graphics Tools* 2004;10:37–49.
797 [12] Popović J, Seitz SM, Erdmann M, Popović Z, Witkin A. Interactive ma-
798 nipulation of rigid body simulations. In: Proceedings of the 27th an-
799 nual conference on Computer graphics and interactive techniques. SIG-
800 GRAPH '00. ISBN 1-58113-208-5; 2000, p. 209–17.
801 [13] Popović J, Seitz SM, Erdmann M. Motion sketching for control of rigid-
802 body simulations. *ACM Trans Graph* 2003;22:1034–54.
803 [14] Treuille A, McNamara A, Popović Z, Stam J. Keyframe control of smoke
804 simulations. *ACM Trans Graph* 2003;22:716–23.
805 [15] McNamara A, Treuille A, Popović Z, Stam J. Fluid control using the
806 adjoint method. *ACM Trans Graph* 2004;23:449–56.
807 [16] Wojtan C, Mucha PJ, Turk G. Keyframe control of complex particle
808 systems using the adjoint method. In: SCA '06: Proceedings of the
809 2006 ACM SIGGRAPH/Eurographics symposium on Computer anima-
810 tion. Eurographics Association. ISBN 3-905673-34-7; 2006, p. 15–23.
811 [17] Barbič J, Popović J. Real-time control of physically based simulations
812 using gentle forces. *ACM Trans on Graphics (SIGGRAPH Asia 2008)*
813 2008;27(5):163:1–163:10.
814 [18] Raibert MH, Hodgins JK. Animation of dynamic legged locomotion.
815 *SIGGRAPH Comput Graph* 1991;25:349–58.
816 [19] Zordan VB, Hodgins JK. Motion capture-driven simulations that hit and
817 react. In: Proceedings of the 2002 ACM SIGGRAPH/Eurographics sym-
818 posium on Computer animation. SCA '02; New York, NY, USA: ACM.
819 ISBN 1-58113-573-4; 2002, p. 89–96.
820 [20] Yin K, Cline MB, Pai DK. Motion perturbation based on simple neuro-
821 motor control models. In: Proceedings of the 11th Pacific Conference on
822 Computer Graphics and Applications. PG '03; Washington, DC, USA:
823 IEEE Computer Society. ISBN 0-7695-2028-6; 2003, p. 445.
824 [21] Zordan VB, Majkowska A, Chiu B, Fast M. Dynamic response for motion
825 capture animation. *ACM Trans Graph* 2005;24:697–701.
826 [22] Yin K, Loken K, van de Panne M. Simbicon: simple biped locomotion
827 control. *ACM Transactions on Graphics* 2007;26(3).
828 [23] Coros S, Beaudoin P, van de Panne M. Generalized biped walking con-
829 trol. In: ACM SIGGRAPH 2010 papers. SIGGRAPH '10; New York,
830 NY, USA: ACM; 2010, p. 1–9.
831 [24] Abe Y, Popović J. Interactive animation of dynamic manipulation. In:
832 Proceedings of the 2006 ACM SIGGRAPH/Eurographics symposium on
833 Computer animation. SCA '06; Aire-la-Ville, Switzerland, Switzerland:
834 Eurographics Association. ISBN 3-905673-34-7; 2006, p. 195–204.
835 [25] Allen BF, Chu D, Shapiro A, Faloutsos P. On the beat!: Timing and ten-
836 sion for dynamic characters. In: SCA '07: Proceedings of the 2007 ACM
837 SIGGRAPH/Eurographics symposium on Computer animation. Euro-
838 graphics Association. ISBN 978-1-59593-624-4; 2007, p. 239–47.
839 [26] Jain S, Liu CK. Interactive synthesis of human-object interaction. In:
840 ACM SIGGRAPH/Eurographics Symposium on Computer Animation
841 (SCA). ACM. ISBN 978-1-60558-610-6; 2009, p. 47–53.
842 [27] Witkin A, Kass M. Spacetime constraints. *SIGGRAPH Comput Graph*
843 1988;22:159–68.
844 [28] Allen BF, Neff M, Faloutsos P. Analytic proportional-derivative control
845 for precise and compliant motion. In: Robotics and Automation (ICRA),
846 2011 IEEE International Conference on. 2011, p. 6039–44.
847 [29] Yamane K, Nakamura Y. Dynamics filter - concept and implementa-
848 tion of online motion generator for human figures. *IEEE Transactions on*
849 *Robotics and Automation* 2003;19(3):421–32.
850 [30] Gleicher M. Motion editing with spacetime constraints. In: Proceedings
851 of the 1997 symposium on Interactive 3D graphics. I3D '97; New York,
852 NY, USA: ACM; 1997, p. 139–49.
853 [31] Kass M, Anderson J. Animating oscillatory motion with overlap: wiggly
854 splines. *ACM Trans Graph* 2008;27:1–8.
855 [32] Boulic R, Mas R, Thalmann D. A robust approach for the control
856 of the center of mass with inverse kinetics. *Computers & graphics*
857 1996;20(5):693–701.
858 [33] Coleman P, Bibliowicz J, Singh K, Gleicher M. Staggered poses:
859 a character motion representation for detail-preserving editing of pose
860 and coordinated timing. In: Proceedings of the 2008 ACM SIG-
861 GRAPH/Eurographics Symposium on Computer Animation. SCA '08;
862 Eurographics Association. ISBN 978-3-905674-10-1; 2008, p. 137–46.
863 [34] Nguyen N, Wheatland N, Brown D, Parise B, Liu CK, Zordan V. Perfor-
864 mance capture with physical interaction. In: Proceedings of the 2010
865 ACM SIGGRAPH/Eurographics Symposium on Computer Animation.

- 866 SCA '10; Eurographics Association; 2010, p. 189–95.
- 867 [35] Sok KW, Yamane K, Lee J, Hodgins J. Editing dynamic human motions via momentum and force. In: Proceedings of the 2010 ACM SIGGRAPH/Eurographics Symposium on Computer Animation. SCA '10; 868 Aire-la-Ville, Switzerland, Switzerland: Eurographics Association; 2010, 869 p. 11–20.
- 870
- 871
- 872 [36] Lockwood N, Singh K. Biomechanically-inspired motion path editing. In: Proceedings of the 2011 ACM SIGGRAPH/Eurographics Symposium 873 on Computer Animation. SCA '11; New York, NY, USA: ACM. ISBN 874 978-1-4503-0923-3; 2011, p. 267–76.
- 875
- 876 [37] Pighin F, Hecker J, Lischinski D, Szeliski R, Salesin DH. Synthesizing realistic facial expressions from photographs. In: Proceedings of the 877 25th annual conference on Computer graphics and interactive techniques. SIGGRAPH '98. ISBN 0-89791-999-8; 1998, p. 75–84.
- 878
- 879
- 880 [38] Müller M, Heidelberger B, Teschner M, Gross M. Meshless deformations based on shape matching. *ACM Trans Graph* 2005;24:471–8.
- 881
- 882 [39] Meirovitch L. Analytical methods in vibrations. Macmillan series in advanced mathematics and theoretical physics; Macmillan; 1967.
- 883
- 884 [40] Bathe KJ. Finite Element Procedures in Engineering Analysis. Prentice-Hall, Inc.; 1982.
- 885
- 886 [41] Kry P, Revéret L, Faure F, Cani MP. Modal locomotion: animating virtual characters with natural vibrations. *Computer Graphics Forum* 887 2009;28(2):289–98. Special Issue: Eurographics 2009.
- 888
- 889 [42] Nunes RF, Cavalcante-Neto JB, Vidal CA, Kry PG, Zordan VB. Using natural vibrations to guide control for locomotion. In: Proceedings of the 890 ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games. I3D '12; New York, NY, USA: ACM; 2012, p. 87–94.
- 891
- 892