# Manipulation of Pose Distributions

Mark Moll    Michael A. Erdmann

March 13, 2000

CMU-CS-00-111

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

This work was supported in part by the National Science Foundation under grant IRI-9503648.

**Keywords:** Pose distributions, parts orienting, dynamic simulation, nonprehensile manipulation

## Abstract

*For assembly tasks parts often have to be oriented before they can be put in an assembly. The results presented in this report are a component of the automated design of parts orienting devices. The focus is on orienting parts with minimal sensing and manipulation. We present a new approach to parts orienting through the manipulation of pose distributions. Through dynamic simulation we can determine the pose distribution for an object being dropped from an arbitrary height onto an arbitrary surface. By varying the drop height and the shape of the support surface we can find the initial conditions that will result in a pose distribution with minimal entropy. We are trying to uniquely orient a part with high probability just by varying the initial conditions. We will derive a condition on the pose and velocity of an object in contact with a sloped surface that will allow us to quickly determine the final resting configuration of the object. This condition can then be used to quickly compute the pose distribution. We also present simulation and experimental results that show how dynamic simulation can be used to find optimal shapes and drop heights for a given part.*

# Contents

# 1 Introduction

In our research we are trying to develop strategies to orient three-dimensional parts with minimal sensing and manipulation. That is, we would like to bring a part from an unknown position and orientation to a known orientation (but possibly unknown position) with minimal means. In general, it is not possible to orient a part completely without sensors, but it is sufficient if a particular orienting strategy can bring a part into one particular orientation with high probability. The sensing is then reduced to a binary decision; a sensor only has to detect whether the part is in the right orientation or not. If not, the part is fed back to the parts orienting device. Assuming the orienting strategy succeeds with high probability, it takes on average just a few tries. An alternative view of this type of manipulation is to consider it as manipulation of the pose distribution. The goal then is to make the pose distribution maximally skewed, thereby reducing uncertainty maximally.

Suppose a polyhedron is initially in a random configuration and the only force acting on it is gravity. We can then compute an approximation of the probability distribution function (pdf) of resting configurations. This approximation will not only depend on the geometry and mass distribution of the polyhedron, but also on the physical model (quasistatic vs. dynamic) and the coefficients of friction and restitution. To orient the part, a robot arm with a camera could detect the current orientation, pick the part up and then put it in the right orientation. This approach can be costly if a high throughput is necessary; a robot can typically orient only one part at a time and might have to re-grip to get the part from initial to desired configuration. A more common approach for small parts is to have a particular (moving) surface for a part, that can orient many parts at the same time. Examples are SONY's APOS system (Hitakawa, 1988), vibratory bowls and conveyor belts with obstacles to align the parts. In the APOS system parts are fed over a vibrating tray with extrusions such that parts will either get stuck in only one orientation, or otherwise are fed over the tray again. Vibratory bowls let parts vibrate to the top of the bowl. In the bowl are obstacles that align the parts in a certain way. For conveyor belts one can do something similar: parts are put on the belt and are aligned by obstacles (or gates) along the way. The design of APOS trays, vibratory bowls and obstacles on conveyor belts is currently still done by hand by experienced engineers who have some intuition for what could work. Still, it typically takes at least a week to design a good APOS tray or vibratory bowl.

## 1.1 Example

In this report we will discuss the use of dynamic simulation for the design of support surfaces that reduce the uncertainty of a part's resting configuration. As the support surface is changed, the pdf of resting configurations will change as well. The pdf will also vary with the initial drop position

| | Stable Poses | | | | | | |
|---|---|---|---|---|---|---|---|
| |  | 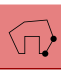 | 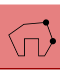 | 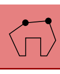 | 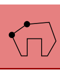 | 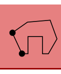 | **Entropy** |
| quasistatic approximation | 0.20 | 0.13 | 0.16 | 0.21 | 0.14 | 0.16 | 1.78 |
| dynamic, flat surface, drop height is $h = 0$ | 0.18 | 0.16 | 0.14 | 0.34 | 0.05 | 0.13 | 1.66 |
| dynamic, bowl shape is $y = 0.24x^2$, $h = 0.28$, initial hor. pos. $x_0 = -0.41$ | 0.24 | 0.03 | 0.03 | 0.50 | 0.08 | 0.15 | 1.35 |

Table 1.1: Probability distribution function of stable poses for two surfaces. The initial velocity is zero and the initial rotation is uniformly random.

above the surface. The following figure and paragraph illustrate the basic idea:

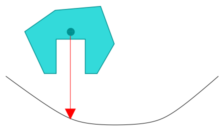

Figure 1.1: A part with an initially unknown orientation is dropped on a surface.

A part with an initially unknown orientation is released from a certain height and relative horizontal position with respect to the bowl. The only forces acting on the part are gravity and friction. We assume the bowl doesn't move. We can compute the final resting configuration for all possible initial orientations. This will give us the pdf of stable poses. The goal is to find the drop height, relative position and bowl shape that will maximally reduce uncertainty.

Table 1.1 shows three different pose distributions. Each stable pose corresponds to a set of contact points (marked by the black dots in the table). For an arbitrarily curved support surface the stable poses do not necessarily correspond to edges of the convex hull of the part. We therefore define a stable pose as a set of contact points. This means that any two poses with the same set of contact points are considered to be the same as far as the pose distribution is concerned. In our example the support surface is a parabola $y = ax^2$ with parameter $a$. Other parameters are the drop height, $h$, and the initial horizontal position of the drop location, $x_0$. We limit the surface to parabolas for illustrative purposes only; in general we would use a larger class of possible shapes (see section 5.1).

The first row in the table shows the pdf assuming quasistatic dynamics. In this case the surface is flat and the part is released in contact with the surface. The second row shows how the pdf changes if we model the dynamics. The initial conditions are the same as for the quasistatic case, yet the pdf is significantly different. The third row shows the pdf for the optimized values for $a$, $h$ and $x_0$.

The objective function over which we optimize is the entropy of the pose distribution. If $p_1, \ldots, p_n$ are the probabilities of the $n$ stable poses, then the entropy is $-\sum_{i=1}^{n} p_i \log p_i$. This function has two properties that make it a good objective function: it reaches its global minimum whenever one of the $p_i$ is 1, and its maximum for a uniform distribution. By searching the parameter space we can find the $a$, $h$ and $x_0$ that minimize the entropy. In the third row of the table the pose distribution is shown with minimal entropy[1]. The table makes it clear that even with a very simple surface we can reduce the uncertainty greatly by taking advantage of the dynamics.

[1]This is a local minimum found with simulated annealing and might not be the global minimum.

## 1.2 Outline

In section 3 we will briefly go over the physical model and in particular the way we model collisions. In section 4 we will explain the notion of capture regions and introduce an extension and relaxation of this notion in the form of so-called quasi-capture regions. Both the physical model and these quasi-capture regions allow for fast computation of the pose distribution. In section 5 we will present our simulation and experimental results. Finally, in section 6 we will discuss the results presented in this report. But first we will give an overview of related work in the next section.

# 2 Related Work

## 2.1 Parts Feeding and Orienting

One of the most comprehensive works on the design of parts feeding and assembly design is (Boothroyd et al., 1982), which describes vibratory bowls as well as non-vibratory parts feeders in detail. The APOS parts feeding system is described by Hitakawa (1988). It is part of the automatic assembly system called SMART (Sony Multi-Assembly Robot Technology). One of the strong points of the APOS system is its flexibility: by replacing the tray and fine-tuning the vibrating motion, other parts can be oriented. *How* these trays are designed and *how* to change the motion is not clear. Automating this step would increase the flexibility even further.

Berkowitz and Canny (1996, 1997) used dynamic simulation to design a sequence of gates for a vibratory bowl. They represented the effects produced by the gates as state transitions in a non-deterministic state automaton. The dynamics were simulated with Mirtich's impulse-based dynamic simulator, *Impulse* (Mirtich and Canny, 1995). Christiansen et al. (1996) used genetic algorithms to design a near-optimal sequence of gates for a given part. Optimality is defined in terms of throughput. Here, the behavior of each gate is assumed to be known. So, in a sense (Christiansen et al., 1996) is complementary to (Berkowitz and Canny, 1997): the latter focuses on modeling the behavior of gates, the former finds an optimal sequence of gates given their behavior. Akella et al. (1997) introduced a technique for orienting planar parts on a conveyor belt with a one degree-of-freedom (DOF) manipulator. Here, it was assumed that the initial orientation is known. Lynch (1999) extended this idea to 3D parts on a conveyor belt with a two DOF manipulator. Wiegley et al. (1996) presented a complete algorithm for designing passive fences to orient parts. Here, the initial orientation is unknown.

Goldberg (1993) showed that it is possible to orient polygonal parts with a frictionless parallel-jaw gripper without sensors. Goldberg conjectured and Chen and Ierardi (1995) proved that for every $n$-sided polygonal part, a sequence of 'squeezes' can be computed in $O(n^2)$ time that will orient it up to symmetry. The length of such a sequence is $O(n)$. These results might have analogues in three dimensions. Marigo et al. (1997) showed how to orient and position a polyhedral part by rolling it between the two hands of a parallel-jaw gripper. Here, however, infinite friction is assumed, whereas Goldberg assumed no friction.

In (Rao et al., 1995) an algorithm is described to orient polyhedral parts using so-called pivot grasps. A part is grasped with two hard finger contacts and is then free to rotate around the axis formed by the contacts. Their algorithm computes an $m \times m$ matrix of pivot grasps, where $m$ is the number of stable configurations and each entry corresponds to a transition from one stable configuration to another. In general, there will be some null entries in this matrix. In other words, it is not always possible to go from any stable configuration to any other. A vision system is used to determine a part's location and orientation. In (Gudmundsson and Goldberg, 1997) a similar system is described, where a robot arm picks up parts that pass on a conveyor belt, but here the focus is on tuning the speed of the conveyor belt. Gudmundsson and Goldberg analytically show how to maximize throughput as a function of the speed of the conveyor belt.

Erdmann and Mason (1988) developed a tray-tilting sensorless manipulator that can orient planar parts in the presence of friction. If it isn't possible to bring a part into a unique orientation, the planner would try to minimize the number of final orientations. In (Erdmann et al., 1993) it is shown how (with some simplifying assumptions) three-dimensional parts can be oriented using a tray-tilting manipulator. In particular, for polyhedral parts with $n$ faces a sequence of 'tilts' of length $O(n)$ can be found in $O(n^3)$ time. Zumel (1997) used a variation of the tray tilting idea to orient planar parts. Zumel used two actuated arms connected at a hinge to tilt parts from one arm to the other. The stable poses of the part at different angles were pre-computed. The planner then found a sequence of joint angle pairs for the two arms that would orient the part.

In recent years a lot of work has been done on programmable force fields to orient parts (Böhringer et al., 1997, 1999; Kavraki, 1997). The idea is that using some kind of force 'field' (implemented using e.g. MEMS actuator arrays) can be used to push the part in a certain orientation. Kavraki (1997) presented a vector field that induced two stable configurations for most parts. Böhringer et al. used Goldberg's algorithm (1993) to define a sequence of 'squeeze fields' to orient a part. They also gave an example how programmable vector fields can be used to simultaneously sort different parts and orient them. Luntz et al. (1997) implemented a parcel transport and manipulation system using a distributed actuator array borrowing ideas from Böhringer et al.

## 2.2 Stable Poses

To compute the stable poses of an object quasistatic dynamics is often assumed. Furthermore, usually it is assumed that the part is in contact with a flat surface and is initially at rest. Boothroyd et al. (1972) were among the first to analyze this problem. Using potential energy arguments and some simplifying assumptions they were able to get good approximations of the pdf's of some parts. They also introduced a method to get a static solution for the pdf: the probability of coming to rest on a face is simply proportional to the area of the face's projection on a sphere centered at the center of mass. The probability of an unstable face is added to the probability of the face to which it rolls. An $O(n^2)$ algorithm for $n$-sided polyhedrons, based on this idea, was implemented by Wiegley et al. (1992). Mirtich et al. (1996) improve this method by approximating some of the dynamic effects. In particular, they compute the area of the intersection of a face with a unit-area circle centered around the center of mass projected on the plane defined by that face. This is then taken as a measure of stability for that face.

Kriegman (1997) introduced the notion of a *capture region*: a region in configuration space such that any initial configuration in that region will converge to one final configuration. He also described an algorithm based on Morse theory that computes the maximal capture regions of an object. Note that this work doesn't assume quasistatic dynamics; as long as the part is initially at rest and in contact, and the dynamics in the system are dissipative, the capture regions will be correct. The capture regions will in general not cover the entire configuration space.

Much work has also been done on determining the stable orientations of assemblies (Mason et al., 1997; Trinkle et al., 1995; Mattikalli et al., 1994). Here, an object is typically in contact with several controlled rigid bodies (e.g., manipulators). Usually the only force acting on the object is gravity. In other words, this is not the same problem as determining whether we have form- or force-closure. Mason et al. (1997) noted that results in this area are also applicable to locomotion of multi-legged robots over uneven terrain.

## 2.3 Collision and Contact Analysis

Computing reaction forces for an object in contact with a surface is far from trivial. In fact, Baraff (1993) showed that deciding whether a configuration with dynamic friction is consistent is *NP*-complete (in terms of the number of contact points). Erdmann (1994) introduced the generalized friction cone, which embeds the force constraints that define the Coulomb friction cone into the part's configuration space. The possible motions have a simple geometric interpretation with this representation. Another geometric approach to analyze multiple frictional contacts was proposed by Brost and Mason (1989). Their approach is limited to two dimensions (however, the configuration space is three-dimensional). It represents forces in a dual space as points. A friction cone then reduces to a line segment in the dual space, and the dual of multiple friction cones is a convex polygon. Trinkle and Zeng (1995) developed a model to predict the quasistatic motion of a planar part in multiple contact. Their analysis yields inequalities defining regions in the space of friction coefficients for which a particular contact mode (i.e., sliding, rolling, separating or a combination thereof) is feasible. Related to this is the work of Wang and Mason (1987). They introduced an impact space, defined as all combinations of orientation and contact motion direction. Within this space one can analytically identify the areas that correspond to the different contact modes.

For rigid body collisions several models have been proposed. Many of these models are either too restrictive (e.g., Routh's model (1897) constrains the collision impulse too much) or allow physically impossible collisions (e.g., Whittaker's model (1944) can predict arbitrarily high increases of system kinetic energy). Recently, Chatterjee and Ruina (1998) proposed a new collision rule, which avoids many of these problems. Chatterjee introduced a new collision parameter (besides the coefficients of friction and restitution): the coefficient of *tangential* restitution. With this extra parameter a large part of allowable collision impulses can be accounted for, and at the same time this collision rule restricts the predicted collision impulse to the allowable part of impulse space. This is the collision rule we will use.

Instead of having algebraic laws, one could also try to model object interactions during impact. This approach is, for instance, taken by Bhatt and Koechling (1995a,b), who modeled impacts as a flow problem. While this might lead to more accurate predictions, it is obviously computationally more expensive. Also, in order to get a good approximation of the pdf of resting configurations, this level of accuracy might not be required. On the other hand, it is also possible to combine the effects of multiple collisions that happen almost instantaneously. Goyal et al. (1998a,b) studied these "clattering" motions and derived the equations of motion for this class of motions.

Given a collision model and the equations of motion, one can simulate the motion of a part. Most of the complexity in dynamic simulation is due to collision detection. Using a particular quaternion representation for orientation, Canny (1986) reduced the problem of finding the distance between polyhedrons to finding the distance between a point and a number of hyperplanes in 7 dimensions. Lin and Canny (1991) designed a fast algorithm to incrementally find the closest point between two convex polyhedra. In cases where there are a large number of collisions or with contact modes that change frequently one can simulate the dynamics using so-called impulse-based simulation (Mirtich and Canny, 1995). However, there are limits to what systems one can simulate. Under certain conditions the dynamics become chaotic (Bühler and Koditschek, 1990; Feldberg et al., 1990; Kechen, 1990). We are mostly interested in systems that are *not* chaotic, but where the dynamics can not be modeled with a quasistatic approximation. In section 5.1 a number of 'chaos plots' are shown that are very similar to the one in (Kechen, 1990).

## 2.4 Shape Design

The shape of an object and its environment imposes constraints on the possible motions of that object. Caine (1993) presented a method to visualize these motion constraints, which can be useful in the design phase of both part and manipulator. In (Krishnasamy, 1996) the mechanics of entrapment were analyzed. That is, Krishnasamy discussed conditions for a part to "get trapped" and "stay trapped" in an extrusion, in the context of the APOS parts feeder. Sanderson (1984) presented a method to characterize the uncertainty in position and orientation of a part in an assembly system. This method takes into account the shape of both part and assembly system.

In (Lynch et al., 1998) the optimal manipulator shape and motion were determined for a particular part. The problem here was not to orient the part, but to perform a certain juggler's skill (the "butterfly"). With a suitable parametrization of the shape and motion of the manipulator, a solution was found for a disk-shaped part that satisfied their motion constraints. Examples of these constraints were: (1) the part cannot break contact, and (2) the part must always be rolling. Although the analysis focused just on the juggling task, it shows that one can simulate and optimize dynamic manipulation tasks using a suitable parametrization of manipulator (or surface) shape and motion.

# 3 Physical Model

To compute the pose distribution of an object, we need to model the dynamics. Since we are interested in a whole class of pose distributions (defined by the surface shape parameters and drop location), it is important that we can compute these pose distributions for a given object quickly. That is, within the class of interest we would like to quickly find the parameter settings that result in a pose distribution with the smallest entropy. Below we will state our other assumptions regarding the physical model.

Let $\rho$ be the radius of gyration and $\theta$ be the relative orientation of the contact point with respect to the center of mass. Then it will turn out to be useful to model the dynamics using the generalized coordinates $(x, y, q)$, where $(x, y)$ is the position of the center of mass and $q = \rho\theta$ represents the orientation of the object. Using these particular generalized coordinates some equations are greatly simplified. For example, the kinetic energy can then be written as

$$KE = \tfrac{1}{2}m\left(v_x^2 + v_y^2 + \rho^2\omega^2\right) = \tfrac{1}{2}m\|\mathbf{v}\|^2.$$

## 3.1 Collisions

Collisions are modeled using Chatterjee's collision rule (Chatterjee and Ruina, 1998). This rule computes a collisional impulse that guarantees a non-negative dissipation of kinetic energy, non-interpenetration between colliding objects and a non-negative normal impulse. Furthermore, the collisional impulse is restricted to lie on the surface or inside the friction cone. Since it is an algebraic collision rule, the post-collision velocity can be computed quickly. The collisional impulse is a linear combination of two base impulses:

$$\mathbf{p}_1 = -\frac{n^T v_{pre} n}{n^T M^{-1} n},$$
$$\mathbf{p}_2 = -M v_{pre},$$

where $n$ is the contact normal, $M$ is the local mass matrix at the contact point and $v_{pre}$ is the relative pre-collision velocity in the work space between the objects at the contact point. In the general case $M$ will depend on the mass properties of both objects. However, if one of the objects is considered immobile (i.e., has infinite mass), it will only depend on the pose and mass distribution of the other object. In our case we consider the surface to be immobile. How to compute mass matrices in the general case is described in e.g. (Smith, 1991). The 'candidate' collisional impulse is defined as

$$\hat{\mathbf{p}} = (1+e)\mathbf{p}_1 + (1+e_t)(\mathbf{p}_2 - \mathbf{p}_1),$$

where $e$ is the coefficient of restitution and $e_t$ the coefficient of tangential restitution. This new parameter $e_t$ gives us an extra parameter to model collisions without violating any physical constraints. The coefficient of tangential restitution can range from $-1$ to 1. In case the candidate impulse lies outside the (real space) friction cone, we project the candidate impulse onto the surface of the friction cone. The collisional impulse then becomes

$$\begin{aligned} \mathbf{p} &= (1+e)\mathbf{p}_1 + \lambda(\mathbf{p}_2 - \mathbf{p}_1), \quad \text{where} \\ \lambda &= \frac{\mu(1+e)n^T\mathbf{p}_1}{\|\mathbf{p}_2 - (n^T\mathbf{p}_2)n\| - \mu n^T(\mathbf{p}_2 - \mathbf{p}_1)}. \end{aligned}$$

As usual, $\mu$ denotes the coefficient of (dry sliding) friction. We assume that the coefficient of static friction is equal to $\mu$. If the candidate impulse lies on or inside the friction cone, then the collisional impulse $\mathbf{p}$ is simply equal to $\hat{\mathbf{p}}$. The collisional impulse in generalized coordinates is then equal to

$$\begin{aligned} \mathbf{p}_g &= \begin{pmatrix} \mathbf{p} \\ \frac{\mathbf{r}\times\mathbf{p}}{\rho} \end{pmatrix} = F\mathbf{p}, \quad \text{where} \\ \mathbf{r} &= R\begin{pmatrix} \cos\theta \\ \sin\theta \end{pmatrix} \\ F &= \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ -\frac{R}{\rho}\sin\theta & \frac{R}{\rho}\cos\theta \end{pmatrix}. \end{aligned}$$

Note that $F^T$ is just a transformation of generalized coordinates to world-space coordinates. Using the method described in (Smith, 1991) the mass matrix for our generalized coordinates simplifies to $M = m(F^TF)^{-1}$, where $m$ is the mass of the object and $M$ is $2\times 2$ matrix.

Let the pre-collision velocity (in configuration space) be $\mathbf{v}_i$. We can then write the post-collision velocity (in configuration space) as

$$\mathbf{v}_f = \mathbf{v}_i + \frac{F\mathbf{p}}{m}.$$

Between collisions the only force acting on the part is gravity.

## 3.2 Rolling Contact

Rolling contact is modeled as a compound pendulum (see e.g. (Symon, 1971, p. 216) for details). The differential equation that describes the rotation around the contact point is

$$\ddot{\theta} = \frac{gr_c}{\rho_c^2}\sin\theta,$$

where $r_c$ is the distance between the center of mass and the contact point and $\rho_c$ is the radius of gyration about the contact point. It can be shown that $\rho_c^2 = \rho^2 + r_c^2$. We can numerically solve this differential equation and check whether the part can maintain rolling contact.

## 3.3 Sliding Contact

We assume that friction is sufficiently high so that a part cannot slide for an infinite amount of time. Furthermore, we assume that once the object starts sliding, it will not change contact points anymore (i.e., the resting configuration remains the same as far as the pose distribution is concerned). In general, this assumption will not be true. However, on a surface with a constant slope one can find the coefficient of friction such that once the part starts sliding, it will only slow down and come to rest without changing contact points. For a concave surface our assumption is slightly harder to justify. Consider the following example:

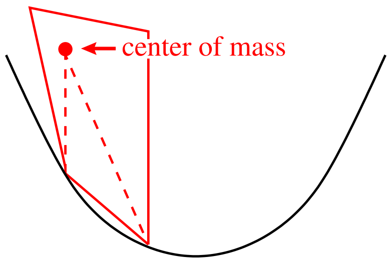


Figure 3.1: Sliding changes contact points

The figure above shows a polygon in contact with a concave surface. If the polygon starts sliding down, its center of mass will no longer be between the contact points. The polygon will therefore rotate counterclockwise.

Since the part is released from a certain height it will bounce first. For numerical simulation, we can therefore treat the motion of the part as a sequence of bounces and rolling motions that 'converge' to sliding. This means that our assumption will only be violated when the velocity in the contact normal direction is exactly zero. A similar approach has been used successfully before by Mirtich and Canny (1995).

## 3.4 Summary of Assumptions

- Both surface and object are rigid bodies. This means we can use an algebraic collision law. Even though this assumption is not true for our experimental setup as discussed in section 5.2, the results following from this assumption are still realistic.
- Collisions follow Chatterjee's collision rule.
- The coefficients of dry sliding friction and static friction are the same.
- Rolling is modeled as a compound pendulum, where the object rotates around the contact point.
- Friction is sufficiently high so that a part cannot slide for an infinite amount of time.
- Once the object starts sliding, it will not change contact points anymore.

# 4 Analytic Results

## 4.1 Quasi-Capture Intuition

In our efforts to analyze pose distributions in a dynamic environment, we have been working on a generalization of so-called 'capture regions' (Kriegman, 1997) that we have termed *quasi-capture regions*. Specifically, for a part in contact with a sloped surface, we would like to determine whether it is captured, i.e., whether the part will converge to the closest stable pose. For simplicity, let the surface be a tilted plane.

**Definition 1.** *Let a* pose *be defined as a point in configuration space such that the part is in contact with the surface.*

We assume that friction is sufficiently high so that a part cannot slide for an infinite amount of time. In general capture depends on the whole surface and everything that happens after the current state, but the friction assumption and our definition of pose allow us to define quasi-capture of the part in terms of local state. The closest stable pose can be defined as follows:

**Definition 2.** *We define a* stable pose *to be a pose such that there is force balance when only gravity and contact forces are acting on the part. The* closest stable pose *is the stable pose found by following the gradient of the potential energy function (using e.g. gradient descent) from the current pose.*

We can now define quasi-capture regions:

**Definition 3.** *A* quasi-capture region *is the largest possible region in configuration phase space such that (a) all configurations in this region have the same closest stable pose and (b) no configuration in a quasi-capture region has enough (kinetic and potential) energy to leave this region either with a rolling motion or one collision-free motion.*

Ideally these quasi-capture regions would induce a partition of configuration phase space, so that for each point in phase space we would immediately know what its final resting configuration is. Of course, this is not the case in general, since with a sufficiently large velocity an object can reach any stable pose. But if we restrict the velocity to be small to begin with, then we are able to quickly determine the pose distribution. It has been our experience that without the use of quasi-capture regions a lot of computation time is spent on the final part/surface interactions (e.g, clattering motions) before the part reaches a stable pose. In other words, *with our analytic results it is possible to avoid computing a potentially large number of collisions.*

In section 3.3 we gave an example where sliding would change the final resting pose. The numerical approach we gave there showed that our assumption is still usable. A different approach would be to shrink the quasi-capture regions by some amount to allow for some sliding. The exact amount can be computed numerically.

In our analysis we have focused on the two dimensional case. To illustrate the notion of capture, we will start with another example. Consider a rod of length $l$ with center of mass at distance $R$ from each vertex. One can visualize this as a disk with radius $R$ and uniform mass, but with contact points only at the ends of the rod (see figure 4.1).

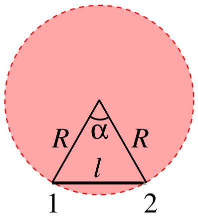


Figure 4.1: A rod with an off-center center of mass.

Note that the endpoints of the rod are numbered. We will refer to these endpoints later. Let the 'side' of the rod where the center of mass is above the rod be the *high energy side*, and the other side be the *low energy side*. We can then define that the rod is 'on' the high energy side if and only if the center of mass is between and above the endpoints of the rod. Suppose the rod is in contact with a flat, horizontal surface. For the rod to make a transition from one side to the other, it will have to rotate, either by rolling or by bouncing. At some point during the transition the center of mass will pass over the contact point. Its potential energy at that point will always be greater or equal than the potential energy at the start of the transition. Hence, to make that transition the rod has to have a minimum amount of kinetic energy. This can be written more formally as

$$\tfrac{1}{2}m\|\mathbf{v}\|^2 > -mg\Delta h \tag{4.1}$$

Figure 4.2 illustrates this.

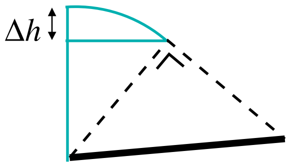


Figure 4.2: Capture condition for a rod.

For a polygonal object in contact with a surface with constant slope we will derive in section 4.2 a lower bound on the norm of the velocity such that for all velocities below that bound the part will be quasi-captured. As we vary the position of the center of mass with respect to the rod endpoints, the slope of the support surface and the drop height, the bound for the capture velocity will change. This bound will also depend on the relative orientation of the contact point with respect to the center of mass.

For a sloped surface the capture condition is not as simple as for the horizontal surface. By bouncing and rolling down the slope, the rod can increase its kinetic energy. We have derived an upper bound on how far the rod can bounce. This gives an upper bound on the increase of kinetic energy. So the capture condition can now be stated as: the current kinetic energy plus the maximum gain in kinetic energy has to be less than the energy required to rotate to the other side. To guarantee that the rod is indeed captured, we have to make sure that the maximum gain in kinetic energy is less than the decrease in kinetic energy due to a collision. There are some additional complicating factors. For instance, a change in orientation can increase the kinetic energy, but to rotate to the other side the rod has to rotate back, undoing the gain in kinetic energy.

## 4.2 Quasi-Capture Velocity

What we will prove is a sufficient condition on the pose and velocity of the rod such that it is quasi-captured on the titled plane. The condition will be of the following form: if the current kinetic energy plus the maximal increase in kinetic energy is less than some bound, the rod is quasi-captured. This bound depends on the current orientation, the current velocity, the slope of the surface and the geometry of the rod. Because of the way we have set up our generalized coordinates, the kinetic energy is $\frac{1}{2}m\|\mathbf{v}\|^2$. In other words, the mass is just a constant scalar. Without loss of generality we can assume $m = 2$. That way the kinetic energy is simply $\|\mathbf{v}\|^2$. We will write $v$ for $\|\mathbf{v}\|$.

**Definition 4.** *Let a bounce be defined as the flight path between two impacts.*

The closest distance between the rod and the slope during one bounce can be described by

$$d(t) = \tfrac{1}{2}g(\cos\phi)t^2 + (v_y\cos\phi + v_x\sin\phi)t - d_{\theta(t)}, \tag{4.2}$$

where $v_x$ and $v_y$ are the translational components of the velocity and $d_{\theta(t)}$ is a component that depends on the orientation. Let the rod be in contact at $t = 0$. Then $d(0) = 0$ (and therefore $d_{\theta(0)} = 0$). Let $\hat{t}$ be the smallest positive solution to $d(t) = 0$. The change in height is then $\Delta h = \frac{1}{2}g\hat{t}^2 + v_y\hat{t}$, so that the change in $v^2$ is $\Delta v^2 = 2g\Delta h = g^2\hat{t}^2 + 2v_yg\hat{t}$. To find the maximum $\Delta v^2$ for all velocity vectors of length $v$ we can parametrize the translational velocity as $v_x = v\cos\xi$ and $v_y = v\sin\xi$, and maximize over $\xi$. This ignores the rotational component of the velocity, but the following lemma shows that for a certain value of $d_{\theta(\hat{t})}$ the resulting solution for $\Delta v^2$ is an upper bound for the true maximal increase of $v^2$.

**Definition 5.** *Let the* ideal orientation *be defined as the orientation where the rod is parallel to the surface and the center of mass is below the rod.*

**Lemma 1.** *We can always increase the rod's kinetic energy after a bounce by allowing it to rotate around the center of mass 'for free' (i.e., without using energy) to the ideal orientation (ignoring penetrations of the surface) and then letting it continue to fall while maintaining this orientation. However, if the rod is already in the ideal orientation after the bounce, its kinetic energy cannot be increased.*

*Proof.* One can easily verify that rotating around the center of mass to the ideal orientation of the bounce maximizes distance between the rod and the surface. This distance will always be greater than or equal to 0. If we allow the rod to continue to fall until it hits the surface, its kinetic energy will increase. □

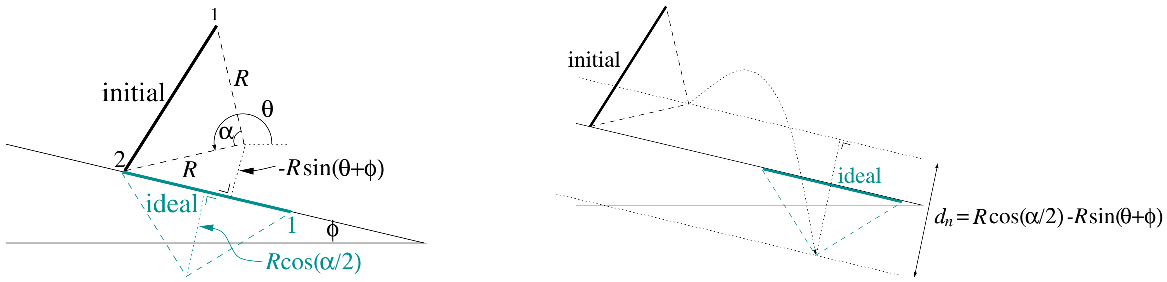


(a) Change in distance between the center of mass and the surface in poses with the initial and ideal orientation

(b) Trajectory of the center of mass during a bounce

Figure 4.3: Increase in kinetic energy when rotating to the ideal orientation

From this lemma it follows that by assuming the rod rotates to the ideal orientation the increase in kinetic energy due to one bounce is an upper bound on the true increase of kinetic energy. With this lemma computing the next contact point is a lot easier. Let $\theta$ be the relative orientation of the contact point at $t = 0$. $\theta = 0$ corresponds to the contact point being to the right of the center of mass. The signed distance from the center of mass to the surface at $t = 0$ is then $-R\sin(\theta+\phi)$, as shown in figure 4.3(a). One can easily verify that in the ideal orientation the relative orientation of endpoint 1 is $\frac{\pi}{2} - \frac{\alpha}{2} - \phi$. Let $\hat{\theta}$ be equal to this relative orientation. In the pose where the rod is in contact with the surface and has the ideal orientation the signed distance from the center of mass to the surface is $-R\sin(\hat{\theta}+\phi) = -R\cos\frac{\alpha}{2}$. So in total the center of mass travels a distance $R(\cos\frac{\alpha}{2} - \sin(\theta+\phi))$ in the direction normal to the surface during one bounce. Let $d_n$ be equal to this distance. To solve for the time of impact we can treat the rod as a point mass centered at the center of mass and replace $d_{\theta(\hat{t})}$ in equation 4.2 with $-d_n$. Equation 4.2 is then simply a paraboloid in $t$. The distance function now measures the distance between the center of mass and the dotted line parallel to the surface shown in figure 4.3(b). This approach is not limited to the case where our new orientation is the ideal orientation. Suppose an oracle would tell us that the new orientation is $\tilde{\theta}$. Then we can solve for the time of impact by substituting $R(\sin(\tilde{\theta}+\phi) - \sin(\theta+\phi))$ for $d_{\theta(\hat{t})}$ in equation 4.2.

The following lemma gives a bound on the velocity needed to *roll* to the other side.

**Lemma 2.** *If the rod is in rolling contact, then to roll to the other side the following condition has to hold:* $v^2 > -2gR(1+\sin\theta+(\mathrm{sign}(\cos\theta)-1)\sin\frac{\alpha}{2}\sin\phi)$. *We assume* $0 \le \phi < \frac{\pi}{2}$.

*Proof.* We can distinguish several cases: endpoint 1 of the rod or endpoint 2 can be in contact with the slope, and the rod can be on the low or high energy side. We will prove the case where endpoint 1 is in contact and the rod is on the high energy side. The proof for the other cases is analogous. The case under consideration is shown in figure 4.4(a). To roll counterclockwise over to the left side, $v^2 > -2gh_1$. The distance $h_1$ is simply equal to $R(1+\sin\theta)$. If the rod rolls clockwise over to two-point contact and continues to roll over endpoint 2, the rod gains kinetic energy because the second contact point is lower than the first contact point. This gain is proportional to $h_2$.

One can easily verify that for two-point contact the relative orientations of contact points 1 and 2 are $\frac{3\pi}{2} - \frac{\alpha}{2} - \phi$ and $\frac{3\pi}{2} + \frac{\alpha}{2} - \phi$, respectively. The bound for rolling over endpoint 2 is

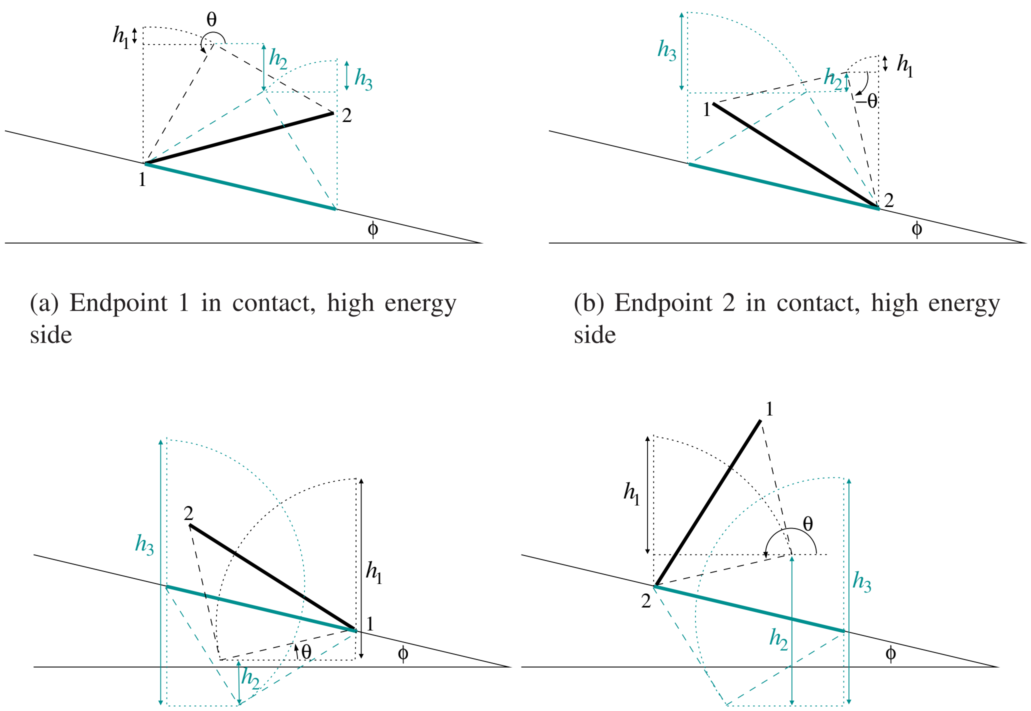


(a) Endpoint 1 in contact, high energy side

(b) Endpoint 2 in contact, high energy side

(c) Endpoint 1 in contact, low energy side

(d) Endpoint 2 in contact, low energy side

Figure 4.4: Capture condition for rotation

therefore

$$
\begin{aligned}
v^2 &> -2g(h_3 - h_2)\\
&= -2gR(1+\sin(\tfrac{3\pi}{2}+\tfrac{\alpha}{2}-\phi)-\sin(\tfrac{3\pi}{2}-\tfrac{\alpha}{2}-\phi)+\sin\theta)\\
&= -2gR(1+\sin\theta-\cos(\tfrac{\alpha}{2}-\phi)+\cos(\tfrac{\alpha}{2}+\phi))\\
&= -2gR(1+\sin\theta-2\sin\tfrac{\alpha}{2}\sin\phi).
\end{aligned}
$$

If the center of mass is to the left of the contact point the last term will change sign. We can therefore combine the two bounds (one for rotating clockwise, and for rotating counterclockwise) into this bound:

$$
\begin{aligned}
v^2 &> \min\left(-2gR(1+\sin\theta),\ -2gR(1+\sin\theta+2\,\mathrm{sign}(\cos\theta)\sin\tfrac{\alpha}{2}\sin\phi)\right)\\
&= -2gR\left(1+\sin\theta+(\mathrm{sign}(\cos\theta)-1)\sin\tfrac{\alpha}{2}\sin\phi\right). \qquad (4.3)
\end{aligned}
$$

□

**Theorem 1.** *The rod with a velocity vector of length $v$ and in contact with the surface is in a quasi-capture region if the following condition holds:*

$$v^2 + \tfrac{2v\cos\xi\sin\phi}{\cos^2\phi}\left(v\sin(\xi+\phi)+\sqrt{v^2\sin^2(\xi+\phi)-2gd_n\cos\phi}\right) - 2g\left(\tfrac{d_n}{\cos\phi}+R\varepsilon\right) \leq -2gR\left(1+\cos(\tfrac{\alpha}{2}+\phi)\right),$$

*where $\xi$ is the direction of the velocity vector that will result in the largest increase of kinetic energy, $d_n = R(\cos\frac{\alpha}{2} - \sin(\theta+\phi))$ and $\varepsilon = \cos(\frac{\alpha}{2}+\phi) - \frac{\cos(\alpha/2)}{\cos\phi} + \max\left(\tan\phi,\ 2\sin\frac{\alpha}{2}\sin\phi\right)$.*

*Proof.* The path of the center of mass during a bounce that increases the kinetic energy maximally is described by $\frac{1}{2}gt^2\cos\phi + v(\sin\xi\cos\phi+\cos\xi\sin\phi)t + d_n = 0$. The smallest positive solution of this equation is

$$\begin{aligned}\hat{t} &= \tfrac{-v(\sin\xi\cos\phi+\cos\xi\sin\phi)-\sqrt{v^2(\sin\xi\cos\phi+\cos\xi\sin\phi)^2-2gd_n\cos\phi}}{g\cos\phi}\\ &= \tfrac{-v\sin(\xi+\phi)-\sqrt{v^2\sin^2(\xi+\phi)-2gd_n\cos\phi}}{g\cos\phi}.\end{aligned} \tag{4.4}$$

The maximum change in $v^2$ is then bounded by

$$\begin{aligned}\Delta v^2 &= 2g\Delta h\\ &\leq 2g(\tfrac{1}{2}g\hat{t}^2 + v(\sin\xi)\hat{t})\\ &= \tfrac{1}{\cos^2\phi}\left(2v^2\sin^2(\xi+\phi) - 2gd_n\cos\phi + 2v\sin(\xi+\phi)\sqrt{v^2\sin^2(\xi+\phi)-2gd_n\cos\phi}\right)\\ &\quad - \tfrac{2v\sin\xi}{\cos\phi}\left(v\sin(\xi+\phi)+\sqrt{v^2\sin^2(\xi+\phi)-2gd_n\cos\phi}\right)\\ &= \tfrac{2v\cos\xi\sin\phi}{\cos^2\phi}\left(v\sin(\xi+\phi)+\sqrt{v^2\sin^2(\xi+\phi)-2gd_n\cos\phi}\right) - \tfrac{2gd_n}{\cos\phi}.\end{aligned} \tag{4.5}$$

After one bounce the orientation is assumed to be such that rod is parallel to the surface and the center of mass is below the rod, as this will result in the largest increase in kinetic energy according to lemma 1. This means that endpoint 1's relative orientation is equal to $\hat{\theta} = \frac{\pi}{2} - \frac{\alpha}{2} - \phi$. Substituting this value in equation 4.3 of lemma 2 gives $-2gR(1+\sin\hat{\theta})$. In other words, if the kinetic energy after the bounce is less than $-2gR(1+\sin\hat{\theta})$ and the rod is in the ideal orientation, the rod cannot roll to the other side.

We can combine the two bounds to obtain a sufficient condition to determine whether the rod can rotate to the other side *if its new orientation after one bounce is equal to the ideal orientation.* Unfortunately this condition does not imply a similar condition for the general case where the new orientation is not necessarily equal to the ideal orientation.

Consider the case $v = 0^+$. Substituting this value in equation 4.5 and expanding the definition of $d_n$ shows that the maximum increase in kinetic energy is then

$$-\frac{2gd_n}{\cos\phi} = -\frac{2gR(\sin(\hat{\theta}+\phi)-\sin(\theta+\phi))}{\cos\phi}. \tag{4.6}$$

Therefore, when $v = 0^+$ and the relative orientation of the contact point after the bounce is equal to ideal orientation the quasi-capture constraint is

$$-2gR\frac{\sin(\hat{\theta}+\phi)-\sin(\theta+\phi)}{\cos\phi} \leq -2gR(1+\sin\hat{\theta}). \tag{4.7}$$

That is, if an upper bound on the kinetic energy after one bounce is less than the energy needed to rotate to the other side, the rod will not be able to rotate to the other side. Now suppose the new orientation is *not* equal to the ideal orientation. Then the increase of kinetic energy will be less, but the energy required to roll to the other side will be less, too. Let $\tilde{\theta}$ be the relative orientation of the contact point after the bounce. Equation 4.7 is of the form $f(\hat{\theta}) \leq g(\hat{\theta})$, where $f(\cdot)$ computes the kinetic energy after one bounce for a given new orientation and $g(\cdot)$ computes the energy needed to roll to the other side for a given orientation[1]. Unfortunately, this bound does not imply $\forall\tilde{\theta}. f(\tilde{\theta}) \leq g(\tilde{\theta})$. From the 'oracle argument' on page 18 it follows that $f(\tilde{\theta})$ is indeed an upper bound on the maximum increase of the kinetic energy. Substituting $\tilde{\theta}$ in lemma 2 shows that $g(\tilde{\theta})$ is a lower bound on the kinetic energy needed to roll to the other side. We would like to determine the smallest possible $\varepsilon$ such that

$$f(\hat{\theta}) - 2gR\varepsilon \leq g(\hat{\theta}) \qquad \Rightarrow \qquad \forall\tilde{\theta}. f(\tilde{\theta}) \leq g(\tilde{\theta}).$$

It is not hard to see $\varepsilon$ has to be equal to $\max_{\tilde{\theta}}(g(\hat{\theta}) - g(\tilde{\theta}) - f(\hat{\theta}) + f(\tilde{\theta}))/(-2gR)$. The difference between $f(\hat{\theta})$ and $f(\tilde{\theta})$ is

$$-2gR\tfrac{(\sin(\hat{\theta}+\phi)-\sin(\theta+\phi))}{\cos\phi} + 2gR\tfrac{(\sin(\tilde{\theta}+\phi)-\sin(\theta+\phi))}{\cos\phi} = -2gR\tfrac{\sin(\hat{\theta}+\phi)-\sin(\tilde{\theta}+\phi)}{\cos\phi}.$$

Similarly, the difference between $g(\hat{\theta})$ and $g(\tilde{\theta})$ is

$$\begin{aligned} &-2gR(1+\sin\hat{\theta}) + 2gR\left(1+\sin\tilde{\theta}+(\operatorname{sign}(\cos\tilde{\theta})-1)\sin\tfrac{\alpha}{2}\sin\phi\right) \\ &\qquad = -2gR\left(\sin\hat{\theta}-\sin\tilde{\theta}-(\operatorname{sign}(\cos\tilde{\theta})-1)\sin\tfrac{\alpha}{2}\sin\phi\right). \end{aligned}$$

The correction $\varepsilon$ is therefore

$$\varepsilon = \max_{\tilde{\theta}}(\sin\hat{\theta}-\sin\tilde{\theta}-(\operatorname{sign}(\cos\tilde{\theta})-1)\sin\tfrac{\alpha}{2}\sin\phi - \tfrac{\sin(\hat{\theta}+\phi)-\sin(\tilde{\theta}+\phi)}{\cos\phi})$$

By differentiating the expression inside $\max(\cdot)$ with respect to $\tilde{\theta}$ we find that there is a local maximum at $\tilde{\theta} = 0$. Other local maxima occur when $\tilde{\theta}$ approaches $-\frac{\pi}{2}$ from below or $\frac{\pi}{2}$ from above. The correction $\varepsilon$ therefore simplifies to

$$\begin{aligned} \varepsilon &= \max\left(\sin\hat{\theta} - \tfrac{\sin(\hat{\theta}+\phi)-\sin\phi}{\cos\phi},\ \sin\hat{\theta}+2\sin\tfrac{\alpha}{2}\sin\phi - \tfrac{\sin(\hat{\theta}+\phi)}{\cos\phi}\right) \\ &= \cos(\tfrac{\alpha}{2}+\phi) - \tfrac{\cos(\alpha/2)}{\cos\phi} + \max\left(\tan\phi,\ 2\sin\tfrac{\alpha}{2}\sin\phi\right). \end{aligned}$$

For $v \neq 0^+$ the difference between $f(\hat{\theta})$ and $f(\tilde{\theta})$ is even larger and $g(\tilde{\theta})$ does not depend on $v$, so the value for $\varepsilon$ is an upper bound for all $v$. Combining all the bounds we arrive at the following

[1]Analagous to expression 4.3, $g(\tilde{\theta})$ equals $-2gR\left(1+\sin\tilde{\theta}+(\operatorname{sign}(\cos\tilde{\theta})-1)\sin\frac{\alpha}{2}\sin\phi\right)$.

quasi-capture condition

$$v^2 + \tfrac{2v\cos\xi\sin\phi}{\cos^2\phi}\left(v\sin(\xi+\phi)+\sqrt{v^2\sin^2(\xi+\phi)-2gd_n\cos\phi}\right)-2g\left(\tfrac{d_n}{\cos\phi}+R\varepsilon\right) \le -2gR\left(1+\cos(\tfrac{\alpha}{2}+\phi)\right).$$

□

Note that for $\phi = 0$ this bound reduces to $v^2 \le -2gR(1+\sin\theta)$. In other words, this bound is as tight as possible when the surface is horizontal.

For an arbitrary $d_n$ it is not possible to compute the optimal $\xi$ analytically. Fortunately, we *can* analytically solve for $\xi$ if we assume that the bounce consists of pure translation. The resulting $\xi$ can be used as an approximation. To find this approximation we substitute $d_{\theta(t)} = 0$ in the distance function (equation 4.2):

$$d(t) = \frac{1}{2}g(\cos\phi)t^2 + (v_y\cos\phi + v_x\sin\phi)t.$$

The positive solution to this is $t = -\frac{2(v_y+v_x\tan\phi)}{g}$. The change in height for the center of mass is therefore

$$\begin{aligned}\Delta h &= \frac{1}{2}gt^2 + v_y t\\ &= \tfrac{2}{g}v_x\tan\phi(v_y+v_x\tan\phi)\\ &= \tfrac{2}{g}v^2\cos\xi\tan\phi(\sin\xi+\cos\xi\tan\phi)\\ &= \tfrac{1}{g}v^2\tan\phi\left(\tan\phi(1+\cos(2\xi))+\sin(2\xi)\right).\end{aligned}$$

To determine the most negative change in height, we differentiate with respect to $\xi$ and find its roots

$$\frac{d}{d\xi}\Delta h = \tfrac{1}{g}v^2\tan\phi(\cos(2\xi)-\tan\phi\sin(2\xi)) = 0 \qquad\Rightarrow\qquad \cot(2\xi)=\tan\phi.$$

By checking the second derivative we can verify that this is indeed a minimum for $0 \le \phi < \frac{\pi}{2}$. We can rewrite the solution for $\xi$ as

$$\begin{aligned}\cos\xi &= \frac{\cos\phi}{\sqrt{2}\sqrt{1-\sin\phi}}\\ \sin\xi &= \frac{\sqrt{1-\sin\phi}}{\sqrt{2}}.\end{aligned}$$

Substituting these values in equation 4.4, we find that the approximation for the bound for $\Delta v^2$ then simplifies to

$$\Delta v^2 \le -\frac{2gd_n}{\cos\phi} + \frac{v^2\sin\phi}{1-\sin\phi}\left(1+\sqrt{1-4d_n g(1-\sin\phi)/(v^2\cos\phi)}\right)$$

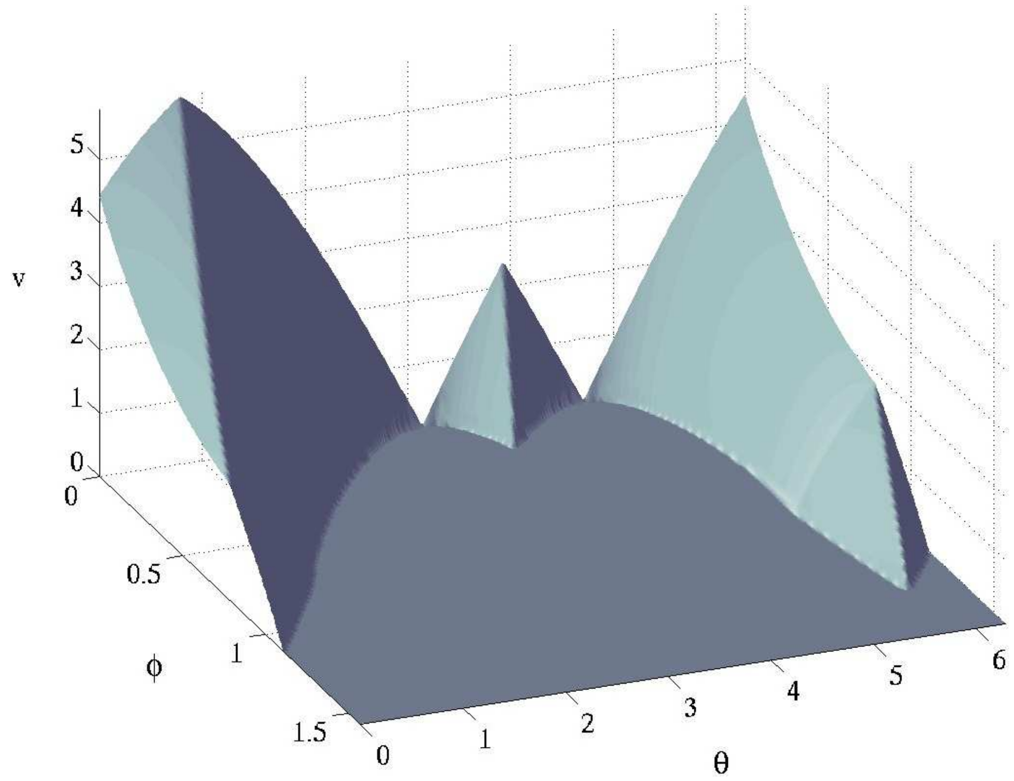


Figure 4.5: Quasi-capture velocity as a function of the slope of the surface and the orientation of the rod

The relative error in this approximation depends on $\phi$, $d_n$, $v$ and $g$ and can be computed numerically. Somewhat surprisingly, the relative error appears to be constant in $v$, $d_n$ and $g$. The relative error does vary significantly with $\phi$, but is still fairly small (on the order of $10^{-2}$).

The theorem above shows a sufficient condition on the velocity and pose of the rod such that it cannot rotate to the other side during one bounce. But suppose there is a sequence of bounces, each of them increasing the kinetic energy. It is possible that the rod satisfies the quasi-capture condition, but is still able to rotate to the other side in more than one bounce. Thus, the theorem by itself is not enough to guarantee that the rod will converge to its closest stable orientation. In the analysis above we have ignored the dissipation of kinetic energy during collisions. If in the case the capture condition is true the dissipation of kinetic energy is larger than the increase due to the bounce, the rod will indeed be captured after an arbitrary number of bounces. To make sure this is the case the coefficients of restitution can not be too large.

In figure 4.5 the quasi-capture velocity is plotted as a function of the slope of the surface and the orientation of the rod. The slope ranges from 0 to $\frac{\pi}{2}$ and the orientation ranges from 0 to $2\pi$. Note that the orientation of the rod is not the same as the relative orientation of the contact point. However, for each combination of $\phi$ and $\theta$ the relative orientation of the contact point can be easily computed. The other relevant parameter values for this plot are: $R = 1$, $g = -9.81$ and $\alpha = \frac{\pi}{2}$. The little bump in the middle corresponds to the rod being captured on the high-energy side. The bigger bumps on the left and right correspond to being captured on the low-energy side.

# 5 Simulations and Experiments

## 5.1 Dynamic Simulation

To numerically compute the pose distribution of parts, we have written two dynamic simulators. One is based is on David Baraff's Coriolis simulator (Baraff, 1991, 1993), which can simulate the motions of polyhedral rigid bodies. Coriolis takes care of the physical modelling. Our simulator then computes pose distributions for different (parametrized) support surfaces and different initial conditions.

Our simulator uses simulated annealing to optimize over the surface parameters and drop location with respect to the surface. The objective function is to minimize the entropy of the pose distribution. Initially the sampling of orientations of the object is rather coarse, so that the resulting pose distribution is not very accurate. But as the simulator is searching, the simulated annealing algorithm is restarted with an increased sample size and the best current solution as initial guess. This way we can quickly determine the potentially most interesting parameter values and refine them later. Our implementation is based on the one given in (Press et al., 1992, pp. 444–455).

Surfaces are parametrized using wavelets (Strang, 1989; Daubechies, 1993). Wavelet transforms are similar to the fast Fourier transform, but unlike the fast Fourier transform basis functions (sines and cosines) wavelet basis functions are localized in space. This localization gives us greater flexibility in modeling different surfaces compared with the fast Fourier transform or, say, polynomials. There are many classes of wavelet basis functions. We are using the Daubechies wavelet filters (Daubechies, 1993) and in particular the implementation as given in (Press et al., 1992, pp. 591–606). To reduce an arbitrary surface to a small number of coefficients we first discretize the function describing the surface. We then perform a wavelet transform and keep the largest components (in magnitude) in the transform to represent the surface. When we minimize the entropy, we optimize over these components. We can either keep the smaller components of our initial wavelet transform around or set them to zero.

Development of a second simulator was started, because Coriolis had some limitations. In particular, the collision model could not be changed and we wanted to experiment with Chatterjee's collision model (Chatterjee and Ruina, 1998). The second simulator also allowed us to optimize for our specific dynamics model. In our model there is only one moving object, and the only forces acting on it are gravity and friction. Currently, the simulator only handles two dimensional objects, but in the future it might be extended to handle three dimensions as well. It uses the analytic results from the previous section to stop simulating the motion of the part once it is captured.

Using the simulator we can compute the quasi-capture regions for the rod. Figures 5.1(a)–(e) show the quasi-capture regions for the low energy side after one through five bounces, respectively. The dark areas correspond to initial orientations and initial velocities that result in the rod being

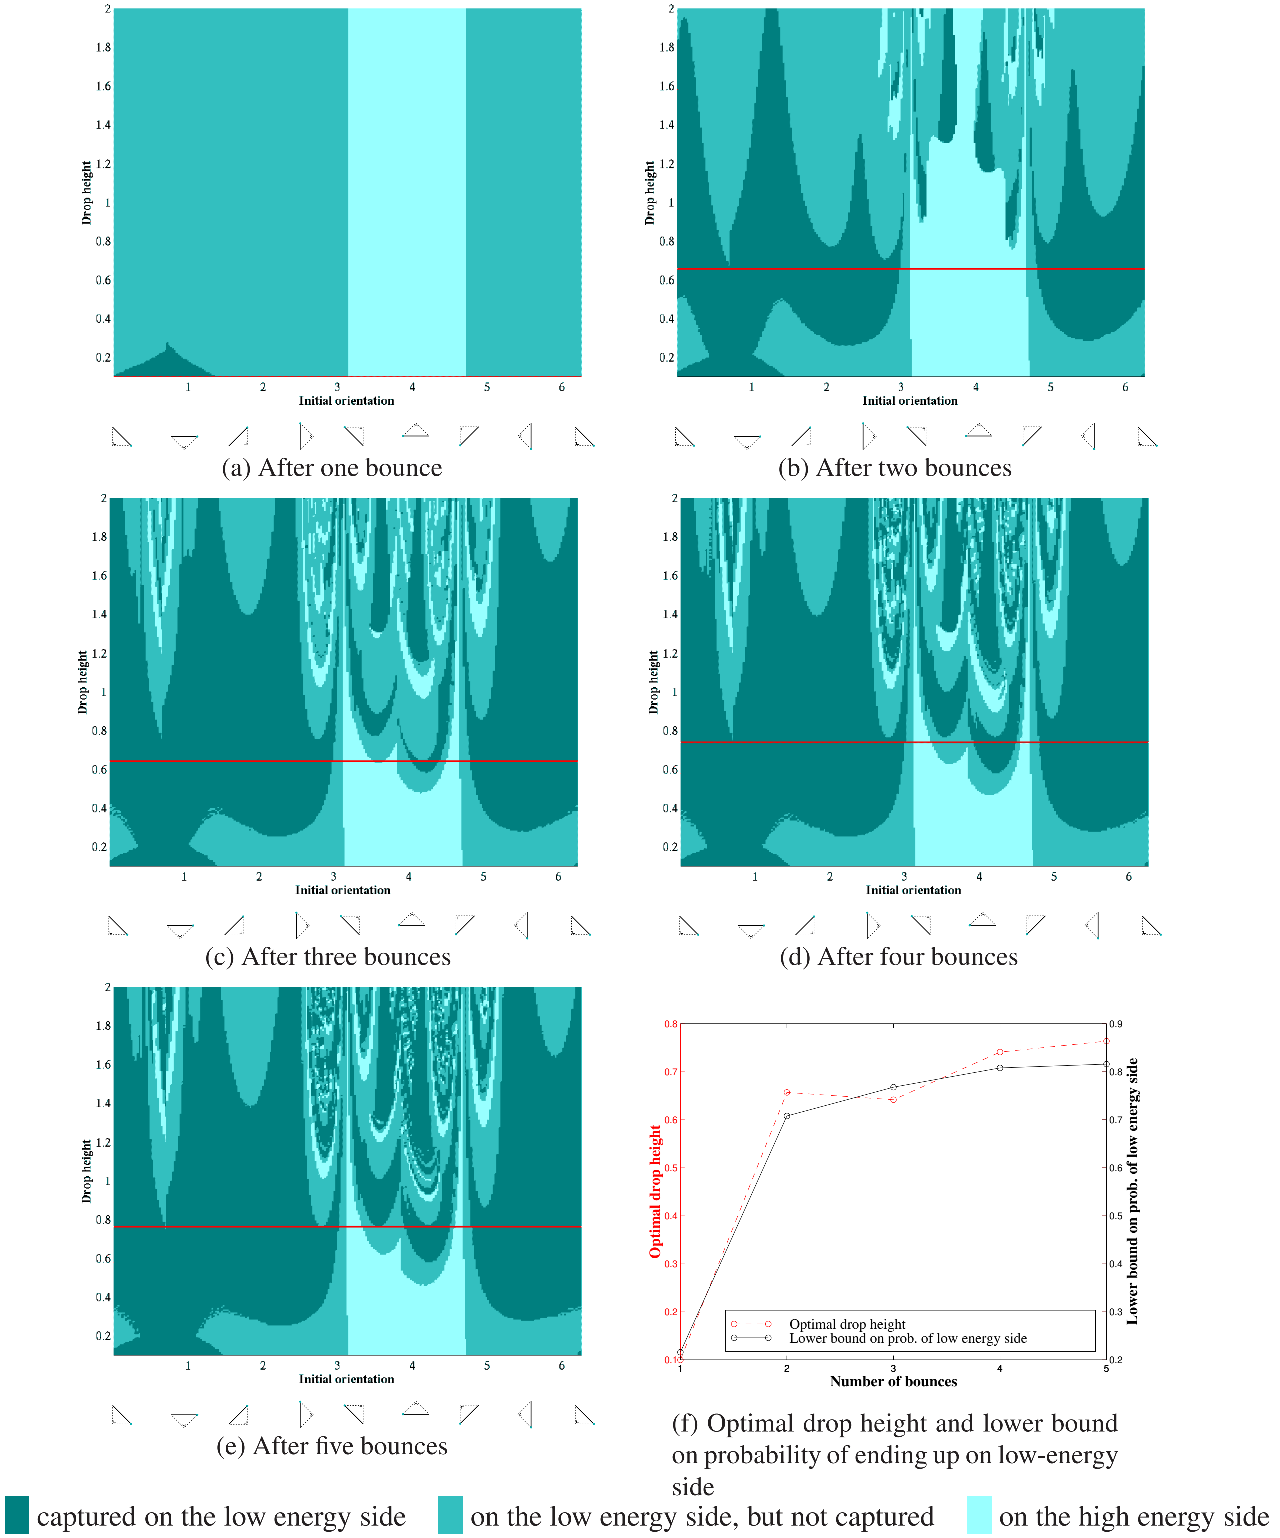


(a) After one bounce

(b) After two bounces

(c) After three bounces

(d) After four bounces

(e) After five bounces

(f) Optimal drop height and lower bound on probability of ending up on low-energy side

Figure 5.1: **Quasi-capture regions for the low energy side of the rod on a 7$^{\text{o}}$ slope.** The rod parameters are $\alpha = \pi/2$, $R = 0.05$, gravity, $g$, equals -1.20, the coefficient of friction, $\mu$, is 5 and the restitution parameters are: $e = 0.1$, $e_t = -0.2$.

quasi-captured. The zero orientation is defined as the orientation where endpoint 1 is to the right of the center of mass. The triangles below the X-axes show the pose of the rod corresponding to the orientation at that point of the X-axes.

Let the optimal drop height be defined as the drop height that maximizes the probability of ending up on the low energy side. Then dropping the rod with uniformly random initial orientation from the optimal drop height will reduce uncertainty about its orientation maximally (unless there exists a drop height that will result in an even higher probability for the high energy side). In figures 5.1(a)–(e) the drop height that results in the maximum probability of ending up on the low energy side is marked by a horizontal line. After each successive bounce this drop height is likely to be a better approximation of the optimal drop height. In figure 5.1(f) the approximate optimal drop height and lower bound on the probability of ending up on the low-energy side after one through five bounces is shown. One thing to note is that both the optimal drop height and the lower bound on the probability of ending up on low-energy side rapidly converge. This seems to suggest that after only a small number of bounces we could make a reasonable estimate of the optimal drop height and uncertainty reduction. Further study is needed to find out if this is true in general.

## 5.2 2D Results

To verify the simulations we also performed some experiments. Our experimental setup was as follows. We used an air table to effectively create a two-dimensional world. By varying the slope of the air table we can vary the gravity. At the bottom of the slope is the surface on which the object will be dropped. The angle $\phi$ of the surface in the plane defined by the air table can, of course, be varied.

The rod of the previous section has been implemented as a plastic disk with two metal pins sticking out from the top at an equal distance from the center of the disk. When released from the top of the air table the disk can slide under the surface and will only collide at the pins. Experimentally we determined the pose distribution of the rod for different values for $g$, $h$ and $\phi$ by determining the final stable pose for 72 equally spaced initial orientations. Our simulation and experimental results of some tests have been summarized in table 5.1. The rows marked with an asterisk have been used to estimate the moment of inertia of the rod and the coefficients of friction and restitution. The estimated values for these parameters are: $e = 0.404$, $e_t = -0.136$, $\rho = 0.0376$ and $\mu = 4.71$. Note that for a low drop height and a horizontal surface (row 13 in table 5.1) the pdf is equal to a quasistatic approximation, as one would expect. More surprisingly, we see that the probability of ending up on the low-energy side can be changed to approximately 0.95 by setting $g$, $h$ and $\phi$ to appropriate values. In other words, we can reduce the uncertainty almost completely.

One can identify several error sources for the differences between the simulation and experimental results. First, there are measurement errors in the experiments: in some cases slight changes in the initial conditions will change the side on which the rod will end up. Second, since the simulations are run with finite precision, it is possible that numerical errors affect the results. Finally, the physical model is not perfect. In particular, the rigid body assumption is just false. The surface on which the rod lands is coated with a thin layer of foam to create a high-damping, rough surface. This is done to prevent the rod from colliding with the sides of the air table.

| | gravity $(m/s^2)$ | drop height $(mm)$ | slope | Prob. low energy side: Sim. | Prob. low energy side: Exp. | |
|---|---|---|---|---|---|---|
| | -0.68 | 58 | $20^o$ | 0.85 | 0.94 | |
| | -0.68 | 122 | $20^o$ | 0.90 | 0.94 | |
| $*$ | -0.68 | 186 | $20^o$ | 0.91 | 0.93 | |
| | -0.68 | 246 | $20^o$ | 0.93 | 0.96 | |
| | -1.5 | 58 | $20^o$ | 0.85 | 0.93 | |
| $*$ | -1.5 | 122 | $20^o$ | 0.90 | 0.92 | |
| | -1.5 | 186 | $20^o$ | 0.91 | 0.97 | |
| | -1.5 | 246 | $20^o$ | 0.93 | 0.97 | |
| | -2.6 | 58 | $20^o$ | 0.85 | 0.94 | |
| | -2.6 | 122 | $20^o$ | 0.90 | 0.93 | |
| | -2.6 | 186 | $20^o$ | 0.91 | 0.93 | |
| $*$ | -2.6 | 246 | $20^o$ | 0.91 | 0.94 | |
| | -2.6 | 76 | $0^o$ | 0.75 | 0.75 | $\longleftarrow$ quasi-static case |
| | -2.6 | 156 | $0^o$ | 0.88 | 0.83 | |
| $*$ | -2.6 | 220 | $0^o$ | 0.92 | 0.85 | |
| | -2.6 | 284 | $0^o$ | 0.87 | 0.89 | |

Table 5.1: Simulation and experimental results for the rod. Shown are the probabilities of ending up on the low-energy side for different values for $g$, $h$ and $\phi$. The drop height is measured from the center of the disk to the surface.

## 5.3 3D Results

We have not generalized our analytic results to three dimensions yet, but we can still use our optimization technique to find a good surface and drop height for a given object. For the dynamic simulation we rely now on Baraff's Coriolis simulator. Figure 5.2(a) shows an orange insulator cap[1] at rest on flat, horizontal surface. The contact points are marked by the little spheres. In figure 5.2(b) the bowl resulting from the simulated annealing search process is shown. The initial shape is a paraboloid: $f(x,y) = (x^2+y^2)/20$. This shape is reduced to a triangulation of a $8 \times 8$ regular grid. The part is always released on the left-hand side of the bowl.

We optimized over the four largest wavelet coefficients of the initial shape and the drop height. The search for the optimal bowl and drop height is visualized in figure 5.3. The five-dimensional parameter space is projected onto a two-dimensional space using Principal Component Analysis (Jolliffe, 1986). Each point corresponds to a bowl shape evaluation, i.e., for each point a pose distribution is computed. The size of each point is proportional to the sample size used to determine the pose distribution. Computing a pose distribution by taking 600 samples takes about 40 minutes on a 500 MHz Pentium III. The surface in figure 5.3 is a cubic interpolation between the points. The dark areas correspond to areas of low entropy. Notice that most of the points are in or near a dark area.

Table 5.2 compares the simulation results with experimental data from (Goldberg et al., 1999).

[1] This object has been used before as an example in (Goldberg et al., 1999; Kriegman, 1997; Rao et al., 1995).

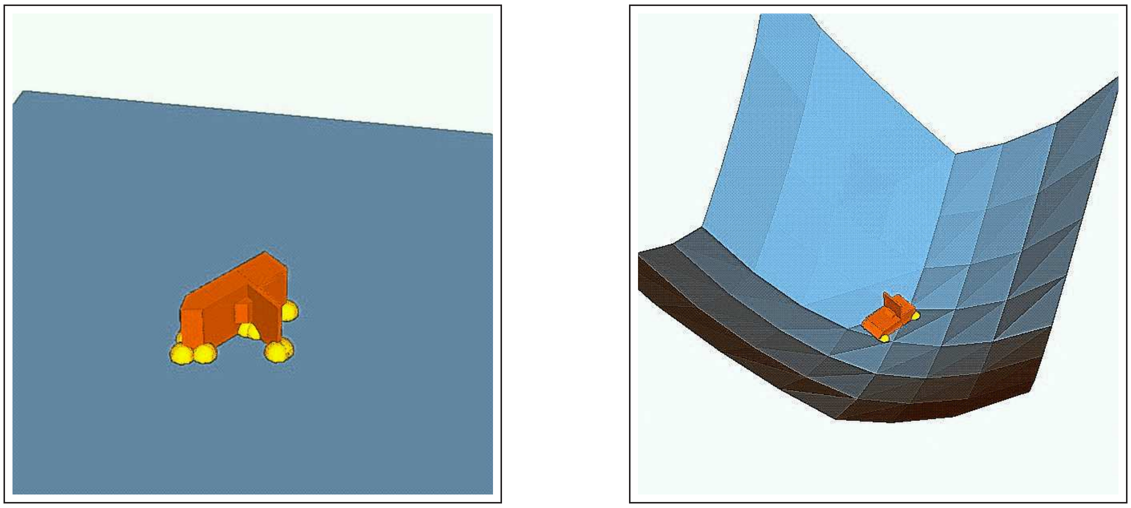

(a) Orange insulator cap on a flat surface (b) . . . and on an optimized bowl

Figure 5.2: Result of optimizing a surface for the orange insulator cap.

| | Stable Poses | | | | | | |
|---|---|---|---|---|---|---|---|
| | $(-1,0,0)$ | $(0,-1,0)$ | $(0,1,0)$ | $(.8,0,.6)$ | $(.7,0,-.7)$ | $(0,0,-1)$ | **Entropy** |
| Experimental, flat (1036 trials) | 0.271 | 0.460 | | 0.197 | 0.050 | 0.022 | 1.58 |
| Dynamic simulation, flat surface | 0.355 | 0.207 | 0.221 | 0.185 | 0.019 | 0.014 | 1.48 |
| Dynamic simulation, optimal bowl | 0.622 | 0.125 | 0.154 | 0.096 | 0.003 | 0.000 | 1.09 |

Table 5.2: Probability distribution function of stable poses for two surfaces. The initial velocity is zero and the initial rotation is uniformly random. The experimental data is taken from (Goldberg et al., 1999). There, $(0,-1,0)$ and $(0,1,0)$ are counted as one pose.

The format is the same as in table 1.1, except that the stable poses are now written as vectors. These vectors are the outward pointing normals (w.r.t. the center of mass) of the planes passing through the contact points. That way, a face with many vertices in contact with the surface will always be represented by the same vector, no matter which subset of the vertices is actually in contact. In the experimental setup of (Goldberg et al., 1999) the part was dropped from one conveyor belt on another. The initial drop height was 12.0 cm. In the experiments the part had an initial horizontal velocity of 5.0 cm/s. The second row correspond to computing the pose distribution when the part is dropped from 12.0 cm (but with initial velocity set to 0). The third row corresponds to a local minimum returned by the simulated annealing algorithm.

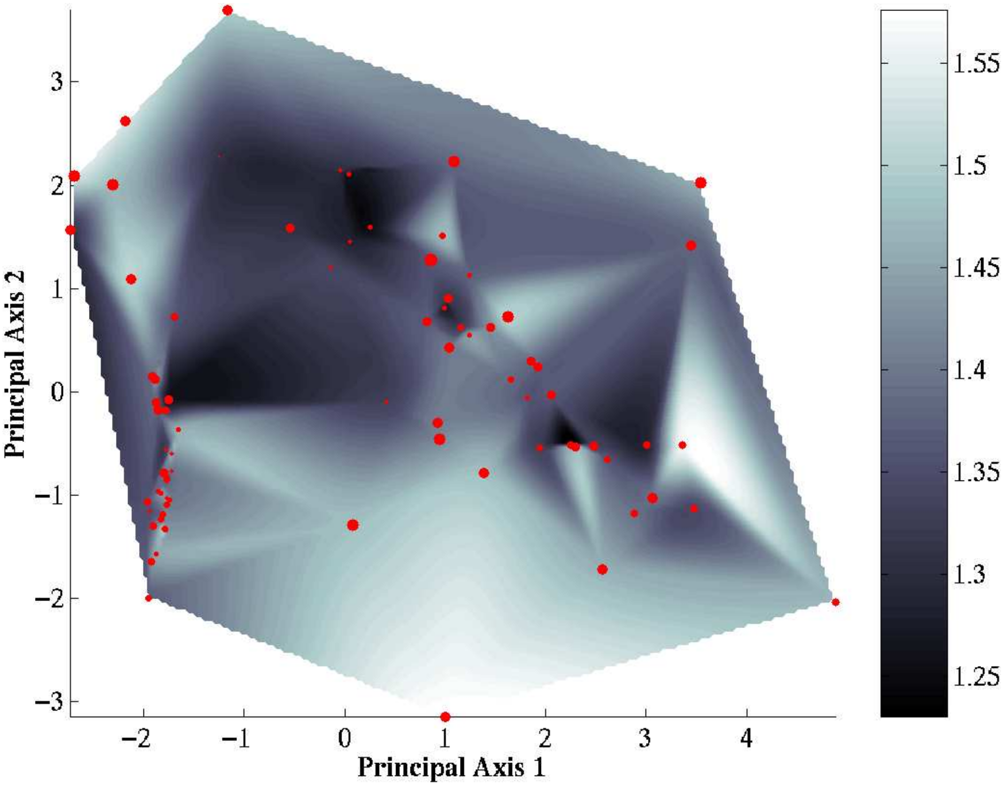


Figure 5.3: Entropy as a function of the two principal axes of the five-dimensional parameter space that was searched.

# 6 Discussion

We have shown a sufficient condition on the position and velocity of the simplest possible 'interesting' shape (i.e., the rod) that guarantees convergence to the nearest stable pose under some assumptions. This condition gives rise to regions in configuration phase-space, where each point within such a region will converge to the same stable pose. We have coined the term quasi-capture regions for these regions, since they are very similar to Kriegman's notion of capture regions.

The quasi-capture regions also apply to general polygonal shapes. However, we can no longer use the symmetry of the rod. So the quasi-capture expressions for general polygonal shapes become more complex. On the other hand, we might be able to orient planar parts by using a setup similar to the one described in section 5 and attaching two pins to the top of the part. Generalizing the quasi-capture regions to three dimensions is non-trivial and is an interesting direction for future research.

The simulation and experimental results show that the simulator is not 100% accurate, but that it is a useful tool for determining the most promising initial conditions for uncertainty reduction. In other words, the optimum predicted by the simulator will probably be near-optimal in the experiments. We can then experimentally search for the true optimum.

Another area where quasi-capture regions may be applied is in computer animation. Before a part comes to rest, there are many interactions between the part and the support surface. It turns out that these interactions are computationally very expensive. With our capture regions we can eliminate the last 'clattering' motions of the part, since we can predict what the final pose will be. For applications where fast animation is more important than physical accuracy, a pre-computed motion can be substituted for the actual motion.

With future research we hope to improve the constraints on the quasi-capture velocity by taking into account more information, such as the direction of the velocity vector. If improving the quasi-capture bounds is impossible, it might be possible to get better approximations for pose distributions. As noted in section 5.1 it is possible to get a good estimate of the maximal uncertainty reduction after only a small number of bounces of the rod. So another interesting line of research would be to find out how accurate these approximations are in general. We are also planning to do more experiments to verify our current and future analytic results.

# References


Akella, S., Huang, W. H., Lynch, K. M., and Mason, M. T. (1997). Sensorless parts orienting with a one-joint manipulator. In *Proc. 1997 IEEE Intl. Conf. on Robotics and Automation*.

Baraff, D. (1991). Coping with friction for non-penetrating rigid body simulation. *Computer Graphics*, 25(4):31–40.

Baraff, D. (1993). Issues in computing contact forces for non-penetrating rigid bodies. *Algorithmica*, pages 292–352.

Berkowitz, D. R. and Canny, J. (1996). Designing parts feeders using dynamic simulation. In *Proc. 1996 IEEE Intl. Conf. on Robotics and Automation*, pages 1127–1132.

Berkowitz, D. R. and Canny, J. (1997). A comparison of real and simulated designs for vibratory parts feeding. In *Proc. 1997 IEEE Intl. Conf. on Robotics and Automation*, pages 2377–2382, Albuquerque, New Mexico.

Bhatt, V. and Koechling, J. (1995a). Partitioning the parameter space according to different behavior during 3d impacts. *ASME Journal of Applied Mechanics*, 62(3):740–746.

Bhatt, V. and Koechling, J. (1995b). Three dimensional frictional rigid body impact. *ASME Journal of Applied Mechanics*, 62(4):893–898.

Böhringer, K. F., Bhatt, V., Donald, B. R., and Goldberg, K. Y. (1997). Algorithms for sensorless manipulation using a vibrating surface. *Algorithmica*. Accepted for publication.

Böhringer, K. F., Donald, B. R., and MacDonald, N. C. (1999). Programmable vector fields for distributed manipulation, with applications to MEMS actuator arrays and vibratory parts feeders. *Intl. J. of Robotics Research*, 18(2):168–200.

Boothroyd, C., Redford, A. H., Poli, C., and Murch, L. E. (1972). Statistical distributions of natural resting aspects of parts for automatic handling. *Manufacturing Engineering Transactions, Society of Manufacturing Automation*, 1:93–105.

Boothroyd, G., Poli, C., and Murch, L. E. (1982). *Automatic Assembly*. Marcel Dekker, Inc., New York; Basel.

Brost, R. C. and Mason, M. T. (1989). Graphical analysis of planar rigid-body dynamics with multiple frictional contacts. In *Fifth International Symposium on Robotics Research*, pages 293–300.

Bühler, M. and Koditschek, D. E. (1990). From stable to chaotic juggling: Theory, simulation, and experiments. In *Proc. 1990 IEEE Intl. Conf. on Robotics and Automation*, pages 1976–1981.

Caine, M. E. (1993). *The Design of Shape from Motion Constraints*. PhD thesis, MIT Artificial Intelligence Laboratory, Cambridge, MA. Technical Report 1425.

Canny, J. (1986). Collision detection for moving polyhedra. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 8(2):200–209.

Chatterjee, A. and Ruina, A. L. (1998). A new algebraic rigid body collision law based on impulse space considerations. *ASME Journal of Applied Mechanics*. Accepted for publication.

Chen, Y.-B. and Ierardi, D. (1995). The complexity of oblivious plans for orienting and distinguishing polygonal parts. *Algorithmica*, 14.

Christiansen, A. D., Edwards, A. D., and Coello Coello, C. A. (1996). Automated design of part feeders using a genetic algorithm. In *Proc. 1996 IEEE Intl. Conf. on Robotics and Automation*, volume 1, pages 846–851.

Daubechies, I., editor (1993). *Different Perspectives on Wavelets*, volume 47 of *Proceedings of Symposia in Applied Mathematics*. AMS.

Erdmann, M. A. (1994). On a representation of friction in configuration space. *Intl. J. of Robotics Research*, 13(3):240–271.

Erdmann, M. A. and Mason, M. T. (1988). An exploration of sensorless manipulation. *IEEE J. of Robotics and Automation*, 4(4):369–379.

Erdmann, M. A., Mason, M. T., and Vaněček, Jr., G. (1993). Mechanical parts orienting: The case of a polyhedron on a table. *Algorithmica*, 10:226–247.

Feldberg, R., Szymkat, M., Knudsen, C., and Mosekilde, E. (1990). Iterated-map approach to die tossing. *Physical Review A*, 42(8):4493–4502.

Goldberg, K., Mirtich, B., Zhuang, Y., Craig, J., Carlisle, B., and Canny, J. (1999). Part pose statistics: Estimators and experiments. *IEEE Trans. on Robotics and Automation*, 15(5).

Goldberg, K. Y. (1993). Orienting polygonal parts without sensors. *Algorithmica*, 10(3):201–225.

Goyal, S., Papadopoulos, J. M., and Sullivan, P. A. (1998a). The dynamics of clattering I: Equation of motion and examples. *J. of Dynamic Systems, Measurement, and Control*, 120:83–93.

Goyal, S., Papadopoulos, J. M., and Sullivan, P. A. (1998b). The dynamics of clattering II: Global results and shock protection. *J. of Dynamic Systems, Measurement, and Control*, 120:94–102.

Gudmundsson, D. and Goldberg, K. (1997). Tuning robotic part feeder parameters to maximize throughput. In *Proc. 1997 IEEE Intl. Conf. on Robotics and Automation*, pages 2440–2445, Albuquerque, New Mexico.

Hitakawa, H. (1988). Advanced parts orientation system has wide application. *Assembly Automation*, 8(3):147–150.

Jolliffe, I. T. (1986). *Principal Components Analysis*. Springer-Verlag, New York.

Kavraki, L. E. (1997). Part orientation with programmable vector fields: Two stable equilibria for most parts. In *Proc. 1997 IEEE Intl. Conf. on Robotics and Automation*, pages 2446–2451, Albuquerque, New Mexico.

Kechen, Z. (1990). Uniform distribution of initial states: The physical basis of probability. *Physical Review A*, 41(4):1893–1900.

Kriegman, D. J. (1997). Let them fall where they may: Capture regions of curved objects and polyhedra. *Intl. J. of Robotics Research*, 16(4):448–472.

Krishnasamy, J. (1996). *Mechanics of Entrapment with Applications to Design of Industrial Part Feeders*. PhD thesis, Dept. of Mechanical Engineering, MIT.

Lin, M. C. and Canny, J. F. (1991). A fast algorithm for incremental distance calculation. In *Proc. 1991 IEEE Intl. Conf. on Robotics and Automation*, pages 1008–1014, Sacramento, CA.

Luntz, J. E., Messner, W., and Choset, H. (1997). Parcel manipulation and dynamics with a distributed actuator array: The virtual vehicle. In *Proc. 1997 IEEE Intl. Conf. on Robotics and Automation*, Albuquerque, New Mexico.

Lynch, K. M. (1999). Toppling manipulation. In *Proc. 1999 IEEE Intl. Conf. on Robotics and Automation*, pages 2551–2557, Detroit, MI.

Lynch, K. M., Shiroma, N., Arai, H., and Tanie, K. (1998). The roles of shape and motion in dynamic manipulation: The butterfly example. In *Proc. 1998 IEEE Intl. Conf. on Robotics and Automation*.

Marigo, A., Chitour, Y., and Bicchi, A. (1997). Manipulation of polyhedral parts by rolling. In *Proc. 1997 IEEE Intl. Conf. on Robotics and Automation*, pages 2992–2997.

Mason, R., Rimon, E., and Burdick, J. (1997). Stable poses of 3-dimensional objects. In *Proc. 1997 IEEE Intl. Conf. on Robotics and Automation*, pages 391–398.

Mattikalli, R., Baraff, D., and Khosla, P. (1994). Finding all gravitationally stable orientations of assemblies. In *Proc. 1994 IEEE Intl. Conf. on Robotics and Automation*.

Mirtich, B. and Canny, J. (1995). Impulse-based simulation of rigid bodies. In *Proc. 1995 Symposium on Interactive 3D Graphics*.

Mirtich, B., Zhuang, Y., Goldberg, K., Craig, J., Zanutta, R., Carlisle, B., and Canny, J. (1996). Estimating pose statistics for robotic part feeders. In *Proc. 1996 IEEE Intl. Conf. on Robotics and Automation*.

Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P. (1992). *Numerical Recipes in C: The art of Scientific Computing*. Cambridge University Press, second edition.

Rao, A., Kriegman, D., and Goldberg, K. (1995). Complete algorithms for reorienting polyhedral parts using a pivoting gripper. In *Proc. 1995 IEEE Intl. Conf. on Robotics and Automation*.

Routh, E. J. (1897). *Dynamics of a System of Rigid Bodies*. MacMillan and Co., London, sixth edition.

Sanderson, A. C. (1984). Parts entropy methods for robotic assembly system design. In *Proc. 1984 IEEE Intl. Conf. on Robotics and Automation*, pages 600–608.

Smith, C. E. (1991). Predicting rebounds using rigid-body dynamics. *ASME Journal of Applied Mechanics*, 58:754–758.

Strang, G. (1989). Wavelets and dilation equations: A brief introduction. *SIAM Review*, 31:613–627.

Symon, K. R. (1971). *Mechanics*. Addison-Wesley, Reading, MA.

Trinkle, J. C., Farahat, A. O., and Stiller, P. F. (1995). First-order stability cells of active multi-rigid-body systems. *IEEE Trans. on Robotics and Automation*, 11(4):545–557.

Trinkle, J. C. and Zeng, D. C. (1995). Prediction of the quasistatic planar motion of a contacted rigid body. *IEEE Trans. on Robotics and Automation*, 11(2):229–246.

Wang, Y. and Mason, M. T. (1987). Modelling impact dynamics for robotic operations. In *Proc. 1987 IEEE Intl. Conf. on Robotics and Automation*, pages 678–685.

Whittaker, E. T. (1944). *A Treatise on the Analytical Dynamics of Particles and Rigid Bodies*. Dover, New York, fourth edition.

Wiegley, J., Goldberg, K., Peshkin, M., and Brokowski, M. (1996). A complete algorithm for designing passive fences to orient parts. In *Proc. 1996 IEEE Intl. Conf. on Robotics and Automation*.

Wiegley, J., Rao, A., and Goldberg, K. (1992). Computing a statistical distribution of stable poses for a polyhedron. In *30th Annual Allerton Conf. on Communications, Control and Computing*.

Zumel, N. B. (1997). *A Nonprehensile Method for Reliable Parts Orienting*. PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA.