

# Fundamental Limitations in Performance and Interpretability of Common Planar Rigid-Body Contact Models

Nima Fazeli, Samuel Zapsolsky, Evan Drumwright, and Alberto Rodriguez

**Abstract** The ability to reason about and predict the outcome of contacts is paramount to the successful execution of many robot tasks. Analytical rigid-body contact models are used extensively in planning and control due to their computational efficiency and simplicity, yet despite their prevalence, little if any empirical comparison of these models has been made and it is unclear how well they approximate contact outcomes. In this paper, we first formulate a system identification approach for six commonly used contact models in the literature, and use the proposed method to find parameters for an experimental data-set of impacts. Next, we compare the models empirically, and establish a task specific upper bound on the performance of the models and the rigid-body contact model paradigm. We highlight the limitations of these models, salient failure modes, and the care that should be taken in parameter selection, which are ultimately difficult to give a physical interpretation.

## 1 Introduction

Real-time planning, control, and state estimation play a central role in modern robotics. Many contemporary algorithms rely on accurate models of the dynamics of a robot as well as models of interactions of the robot with its environment ([3, 11, 22]). Dynamic frictional interaction between bodies is important to robotic systems undergoing contact, yet has proven to be a challenging phenomenon to model. Challenges include, but are not limited to, non-smooth dynamics (i.e., transitions from absence of contact to sticking/sliding contact), near-impulsive forces with large magnitudes, local deformations (e.g., difficult to predict elastic/plastic inden-


---

Nima Fazeli, Alberto Rodriguez  
Mechanical Engineering Department, MIT, USA e-mail: `<nfazeli, albertor>@mit.edu`

Samuel Zapsolsky, Evan Drumwright  
Toyota Research Institute, CA, USA e-mail: `<sam.zapsolsky, evan.drumwright>@tri.global`

tations), wave propagation through the bodies (e.g., vibrations along rods from impact), sensitivity to initial conditions, and acting friction—a complex phenomenon that emerges from the interactions between microscopic surface asperities—through necessarily coarse models.

Despite these difficulties, various contact models have been proposed and generally follow one of three paradigms: i) fully deformable, elasto-dynamic models; ii) pseudo-rigid, compliant models; and iii) rigid contact. Fig. 1 shows, conceptually, how each paradigm models contact.



**Fig. 1** Contact paradigms: a) Disk motion prior to contact, b) Elasto-dynamic (colors depict degree of stress), c) Compliant (pseudo-rigid), d) Rigid

The elasto-dynamic paradigm models deformations for contacting bodies over a finite contact time (i.e., impacts are correctly characterized as non-instantaneous phenomena), can approximately predict elastic/plastic deformations, and are able to model vibrations along the object. These modeling capabilities come at a high computational cost and require identifying greater numbers of system parameters; their greater predictive power is compromised by uncertainty in initial conditions of robotics applications. As such, these models have not been prevalent in robotics, where speed and simplicity play an important role and system identification cannot yet be conducted accurately in real time.

The class of pseudo-rigid, compliant contact models (often designated “penalty methods”) is the oldest, beginning with the seminal work of Hertz [10] followed by numerous authors [15, 17, 19, 25]. These models assume a rigid core and a thin, deformable layer and apply contact forces that are proportional to the magnitude of deformation. Simulation software such as DANCE [23] (open-source) and ADAMS [18] (proprietary) have been developed under this paradigm. These models are numerically challenging to use in simulation and control applications due to discontinuity in the Coulomb friction model [26] and nearly rigid contacts. Modeling pseudo-rigid contact from the opposite direction—softening otherwise rigid contacts—was pioneered by Lacoursière [16] in MathEngine and has been used in numerous rigid body dynamics simulations libraries since (e.g., ODE, Bullet, MuJoCo, DART).

Like full elastodynamic models, pseudo-rigid body models require identifying system parameters (Young’s Modulus and Poisson’s Ratio); when bodies are known to be nearly rigid (the case of the objects in the present study), the extra flexibility afforded by the pseudo-rigid body models can thus be undesirable.

Rigid contact models resolve contact without permitting geometric intersections between bodies. These methods assume that deformations are nonexistent and contacts “appear” discontinuously as a function of configuration—meaning that impacts must be resolved instantaneously. The rigid contact model distinguishes between non-impacting and impacting contacts; while the combination of complementarity conditions, Coulomb friction, and stiction constraints often yields a deterministic model for non-impacting contact, while impacting contact is underdetermined. Rigid impact models effectively add constraints (like restitution) to select points in the feasible impact space (concept described by Chatterjee and Ruina [2]). Rigid-body impact models span centuries of research, starting with Newton, to contributions from Poisson, and more “modern” treatments [30, 28, 29]. These models typically require 2–3 parameters (often referred to as normal restitution and friction coefficient) and are commonly used in robotic manipulation [3, 11] and locomotion [22].

Despite the variety and wide usage of existing rigid impact models, relatively little empirical evaluation and comparison among models has been conducted. The performance of such models has not been quantified; consequently, the limitations of these models are not well understood. Additionally, systematic parameter estimation methodologies have not been developed for the class of rigid impact models. In this paper we empirically evaluate the performance of six common rigid impact models used in robotics on a planar impact experiment, demonstrate that a fundamental upper-limit for performance of these models exists, and provide a systematic methodology for optimal parameter identification to maximize the models’ predictive performance. We further show that regions for good/poor predictions can be identified and demonstrate the wide variability in model parameters depending on the choice of data used to identify the models.

## 2 Rigid contact models

Rigid contact models couple the Newton-Euler equations of motion to contact equations (non-interpenetration constraints, Coulomb friction, stiction, *etc.*) to predict forces imparted to the bodies, and, subsequent motion by extension. The class of rigid contact models makes three simplifying assumptions: i) bodies undergoing contact sustain negligible deformation (coupling this assumption with convex geometry implies that a single point of contact need be considered); ii) all impacts and their resolutions are instantaneous, and therefore velocity changes discontinuously; and iii) the configuration of impacting bodies does not change during impact (follows from the instantaneity assumption).

Consider Fig. 2, where a planar object makes contact with a fixed horizontal surface. The configuration of the object is  $q_o \equiv [x \ y \ \theta]^T$ , the velocity  $v_o \equiv \dot{q}_o$ , the point of contact  $r \equiv [r_x \ r_y]^T$ , and the mass and second moment of inertia  $m$  and  $I$  respectively. Unless stated otherwise, all points and vectors defined in a “world” (stationary) reference frame. We can represent a rigid model for contact of a planar moving body with a fixed surface as a mapping:


$$f_c : \{\mathbf{q}_o, \mathbf{v}_o^i, \mathbf{r}, m, I\} \rightarrow \{\mathbf{v}_o^f\} \quad (1)$$

where the superscripts  $i$  and  $f$  denote the pre- and post-contact values (recall that the rigid body assumption, which implies instantaneity, entails that  $\mathbf{q}_o^i = \mathbf{q}_o^f$ , so we omit such subscripts from  $\mathbf{q}_o$ ). Under the rigid-body assumption, Newtonian mechanics relates the pre- and post-contact states by:

$$\mathbf{M} \frac{\mathbf{v}_o^f - \mathbf{v}_o^i}{dt} = \mathbf{J}(\mathbf{q}, \mathbf{r})^\top \mathbf{F}_c \rightarrow \mathbf{v}_o^f = \mathbf{v}_o^i + \mathbf{M}^{-1} \mathbf{J}^\top \mathbf{P} \quad (2)$$

where  $\mathbf{M} \in \mathbb{R}^{3 \times 3}$  is the generalized inertia matrix of the rigid body,  $\mathbf{P} = \mathbf{F}_c dt$  denotes the impulse imparted to the bodies due to contact ( $\mathbf{F}_c$  is the non-impulsive version of this force), and  $\mathbf{J} \in \mathbb{R}^{2 \times 3}$  denotes the Jacobian matrix that maps generalized velocities to contact-frame velocities and contact-frame impulses to generalized impulses (via the transpose of the Jacobian). Note that this relation holds for the moment before and after contact, but is insufficient to determine the contact impulse (the vector operations above yield three equations and six unknowns:  $\mathbf{v}_o^f$  and  $\mathbf{P}$ ).

Given the unknown imparted impulse we can solve for the post-contact velocity; choosing the imparted impulse is precisely the job of the contact models. To gain an understanding of how contact models choose  $\mathbf{P}$ , in sec. 2.1 we will briefly discuss the fundamental constraints of contact. This understanding will lay the foundation to the system identification formulation as well as the data-driven methods.



**Fig. 2** Planar object making contact with a horizontal surface. The contact point is denoted by  $C$ .

## 2.1 Fundamental constraints of contact

Under the rigid-body assumption, all contact must obey three fundamental constraints: i) only compressive impulses may be applied along the surface normal (contacts are not permitted to pull); ii) impacting bodies must not penetrate; and

iii) the total system momentum must be conserved. To obtain a visual representation of the feasible impulse space for 2D contact we adopt the treatment discussed in [2]. We write the change in momentum of the object due to contact as:

$$\mathbf{v}_o^f - \mathbf{v}_o^i = \mathbf{M}^{-1} \mathbf{J}^T \mathbf{P} \quad (3)$$

We write the change in velocity of the contact point by pre-multiplying both sides of the expression by  $\mathbf{J}$ :

$$\mathbf{J}(\mathbf{v}_o^f - \mathbf{v}_o^i) = \mathbf{J} \mathbf{M}^{-1} \mathbf{J}^T \mathbf{P} \rightarrow (\mathbf{v}_c^f - \mathbf{v}_c^i) = \mathbf{M}_c^{-1} \mathbf{P} \quad (4)$$

Here  $\mathbf{v}_c$  is the  $2 \times 1$  velocity vector of the point of contact  $c$ ,  $\mathbf{M}_c^{-1} \equiv \mathbf{J} \mathbf{M}^{-1} \mathbf{J}^T$  is the  $2 \times 2$  *contact space compliance matrix* (making  $\mathbf{M}_c$  the *contact space inertia matrix*). We note that the kinetic energy of the system post-contact may not exceed that of the energy prior to contact; we may write this constraint as:


$$\mathbf{v}_c^{fT} \mathbf{M}_c \mathbf{v}_c^f = \alpha \mathbf{v}_c^{iT} \mathbf{M}_c \mathbf{v}_c^i, \quad 0 \leq \alpha \leq 1 \quad (5)$$

Recall that the position of the object does not change during an instantaneous impact, so potential energy remains the same. We re-write this expression using (4), and after expansion and substitution:

$$(\mathbf{P} + \mathbf{M}_c \mathbf{v}_c^i)^T \mathbf{M}_c^{-1} (\mathbf{P} + \mathbf{M}_c \mathbf{v}_c^i) = \alpha \mathbf{v}_c^{iT} \mathbf{M}_c^{-1} \mathbf{v}_c^i \quad (6)$$

The representation of the expression above is that of an ellipse in 2D due to the positive definite nature of  $\mathbf{M}_c$ ; the term  $\mathbf{M}_c \mathbf{v}_c^i$  simply acts to translate the ellipse  $\mathbf{P}^T \mathbf{M}_c^{-1} \mathbf{P}$  in  $\mathbf{P}$ . The union of all ellipses for all  $0 \leq \alpha \leq 1$  gives what we call the Energy Ellipse [2].

Fig. 3 depicts an Energy Ellipse, where the shaded region is the admissible region based on the fundamental constraints of rigid contact. Two important features of the Energy Ellipse are the lines of sticking and maximum compression.



**Fig. 3** The Energy Ellipse: This ellipse is constructed from application of the conservation of energy law before and after the contact event. Points within the ellipse satisfy this law, and points in the shaded region satisfy the non-negative normal velocity component of the contact point.

The line of maximum compression delimits the region for which the separation velocity of the contact point changes sign; regions on and above this line obey the

fundamental contact constraints. Points along the line of sticking signify contacts for which the relative tangential velocity of the contact point is zeroed due to friction. The intersection of these two lines is the impulse imparted that eliminates all relative velocity between the bodies at the contact point. The diagram of the Energy Ellipse graphically shows that the three fundamental constraints governing the impact are not sufficient to select an impulse to resolve the impact. Contact models introduce additional assumptions and constraints to choose a unique point within the feasible space.

In the next section we briefly introduce the models considered in this paper, and demonstrate how they select points within the Energy Ellipse.

## 2.2 *Rigid contact models studied in this paper*


The purpose of a rigid contact model is to select a point in the feasible impulse space demarcated with the Energy Ellipse. In this section we very briefly introduce the six models we used in this study; and refer the interested reader to [6] for further details. The six models studied are all parameterized by two variables, denoted as  $(\mu, \epsilon)$  and often referred to as the coefficient of friction and restitution, respectively. When the contact-space compliance matrix is diagonal, the former is used to regulate the magnitude of tangential momentum of the contact point and the latter is used to regulate the normal component.

The first two models are **Anitescu-Potra Newton** [1] and **Anitescu-Potra Poisson** [1], which are closely related to Stewart-Trinkle [27], differing primarily in the restitution used for contact resolution as their names suggest. The **Drumwright-Shell** [4] also uses a Poisson-type contact model that computes contact forces that maximize kinetic energy dissipation. **Mirtich** [20] is an incremental collision model and a computational implementation of Stronge’s work [28] relating physical work done in the compression and restitution phases of an impact. **Wang-Mason** [29] relates the magnitude of the impulses in the compression and restitution phases of an impact using a Poisson-type restitution. **Whittaker** [30, 12] is an algebraic model that requires solution to a nonlinear system of equations and both 2D and 3D versions have been studied.

We now define two post hoc models, the **Best Post Hoc** and the **Ideal Rigid-Body Bound (IRB Bound)**. The Best Post Hoc model selects the best predicting rigid-body contact model, of the six considered in this paper, after each impact event, and the IRB Bound selects the best predicting impulse from the Energy Ellipse. We emphasize that the IRB Bound is able to select any possible feasible impulse, but the Best Post Hoc only selects a model (spanning a sub-space of the Energy Ellipse). Consequently the IRB Bound quantifies the absolute best performance feasible for any possible rigid-body contact model, and the Best Post Hoc quantifies the best possible performance for the models studied. Neither model is suitable for making predictions, since each makes their choice post hoc.

Fig. 4 depicts the admissible contact regions for each model, grouped based on graphical similarity. This similarity is not necessarily indicative of the way the re-

regions are constructed. These regions are the union of all predicted impulses for all admissible values of the model parameters.



**Fig. 4** Regions of the admissible contact space covered by each model. Note the similarity between the models and the limited coverage of the admissible space.

Fig. 4 (a) shows that the Drumwright-Shell and AP Poisson models cover very similar regions of the space. Fig. 4 also provides intuition for the definitions used in contact parameters. For example, the coefficient of restitution regulates the magnitude of the normal impulse; therefore, models with vertical side boundaries use the ratio of the post to pre-contact normal momenta. Fig. 4 (b) depicts the regions from the AP Newton, Wang-Mason, and the Mirtich models. We note that the definition of the coefficient of restitution used in these three models relates the relative normal velocities pre- and post-contact, resulting in fairly similar regions. Fig. 4 (c) depicts the Whittaker model, for which the horizontal boundaries of the feasible region imply that the coefficient of friction relates the tangential momenta pre- and post-contact, differing significantly in interpretation from the previous models.

Two key observations are: i) no single model is able to span the full admissible region of contact (we further note that the regions covered by the various models are quite similar in shape); ii) any given model will make predictions only up to the line of sticking, but not beyond it; therefore the models are not able to predict “back-spin” (i.e., the tangential velocity of the contact point post-contact changes direction). This issue is sometimes alleviated by contact models that utilize an additional tangential restitution, but the physical interpretation and utility of such restitution is debatable [2].

### 3 System identification formulation and model evaluation

System identification applied to inertial parameter estimation of robotic systems has received considerable attention within the robotics research community [7, 5, 9, 13, 14] and several practical implementations have been developed and applied to industrial robots [21]. Control architectures have also been developed to estimate and adapt to uncertainties over these dynamic parameters [11, 24], and have been shown to exhibit asymptotic convergence to desired trajectories in the case of free space motion of manipulators.

System identification is a necessary step to maximize the predictive performance of these models, yet to our knowledge no systematic formulation of this step has been proposed. To this end, we first propose a system identification formulation in impulse space, leveraging the Energy Ellipse. Next, we briefly discuss the experimental setup and data collection protocols used in this study; we then apply the system identification procedure to find the optimal parameters for each of the studied models, allowing us to empirically compare predictions and use the IRB Bound and Best Post Hoc models to establish upper limits in predictive performance. Finally, we study individual experimental trials and demonstrate the large variability in model parameters across experimental trials (indicating challenges in interpreting parameters). We also quantify the predictive range of each model by evaluating the number of outcomes across individual instances of the experiment that can be explained by *any* choice of parameter set.

To perform system identification, we first compute the impulse imparted on the object from pre- and post-contact states (Fazeli et al [7]) under the ideal rigid body impact model by solving:

$$\mathbf{P}_{IRB} = \arg \min_{\mathbf{P}} \|\mathbf{M}(\mathbf{v}_o^f - \mathbf{v}_o^i) - \mathbf{J}^T \mathbf{P}\|_2 \quad (7)$$

Then we can identify the parameters for a single impact by solving the following optimization program:

$$(\mu^*, \epsilon^*) = \arg \min_{(\mu, \epsilon)} \|\mathbf{P}_{IRB} - \text{Model}(\mu, \epsilon)\|_2 \quad (8)$$

$$\text{s.t.} \quad 0 \leq \mu \leq \mu_s \quad 0 \leq \epsilon \leq 1 \quad (9)$$

The optimization program is convex within the constraints (as will be explained in more detail in Sec. 3.2.1) and takes place in the space of imparted impulses and the Energy Ellipse. The upper bound on the coefficient of friction  $\mu_s$  is the lowest value of the coefficient of friction for which the model predicts sticking contact and can be found with a line search using a bisection method. For values of  $\mu > \mu_s$  and a fixed  $\epsilon$ , the predicted impulse from the model will remain the same and the cost function will not have a gradient, making the optimization ill-posed numerically. This implies that pre and post contact conditions exist which do not provide sufficient information to estimate  $\mu$  correctly and render this parameter unobservable, more on this in Sec. 3.2.2.


To perform system identification for a batch of impacts, we simply sum up the cost function over the set of impacts (Sec. 3.2.1).

### 3.1 Experimental data collection

In order to evaluate and compare the contact models we conceived the experimental setup depicted in Fig. 5 [6, 7]. The purpose of the apparatus is to autonomously collect measurements of the trajectory of falling planar objects initialized from a specifiable set of initial conditions. We chose this experiment because it captures



key features of a dynamic frictional interaction while being amenable to careful monitoring and instrumentation resulting in high quality measurements; for further details see [6].



**Fig. 5** Experimental setup: Autonomous data collection for planar impact data. A planar object is constrained to move between two lubricated vertical glass planes. The robot uses a directional magnetic latch to pick up the object across the glass and provide control over position and velocity of the object. The Vicon motion capture system is used to track the states of the object.


The experimental setup is composed of a “dropping arena” designed to constrain the motion of planar objects to a plane using lubricated glass planes. The planar objects are equipped with markers to be tracked using the Vicon motion capture system at 250 Hz, and directional magnets allowing the robot to latch onto the object and prescribe desired initial state to the object for each drop.

We collected a total of 2,000 drops. In some instances the tracking system may lose frames or suffer spurious errors due to reflections; to detect and remove these effects each experimental run was subjected to a series of tests to validate the correctness of the measurements. The three most important tests were: i) the energy test: verifies constant energy during ballistic motion and energy loss during times of contact, ii) a frame drop test, and iii) a deviation test: checks deviations of trajectories during ballistic motion from parabolic paths. A total of 1,718 of the 2,000 drops passed all tests, and we used this set for the rest of this study. An example of the recorded configurations for one experiment is shown in Fig. 6.

To detect the time of contact we use the fact that the object undergoes ballistic motion, consequently the acceleration of the object is smooth. Any impact event induces a spikes in the second time derivative of the state vector, and we use a threshold on these spikes to detect events. We further use the ballistic motion of the object to fit parabola to the pre- and post-contact trajectories, thereby mitigating the effects of noise on the state measurements and yielding better approximations to velocities of the object; for more details see [6].

### 3.2 Results of system ID and model prediction performance

We conduct the identification in two scenarios, once in batch (Sec. 3.2.1 ensemble parameters), and once for individual impacts (Sec. 3.2.2 individual best parameters).



**Fig. 6** The configurations of the object as a function of time during an experimental run.  $\theta$  is in radians but has been scaled by the radius of gyration of the object.

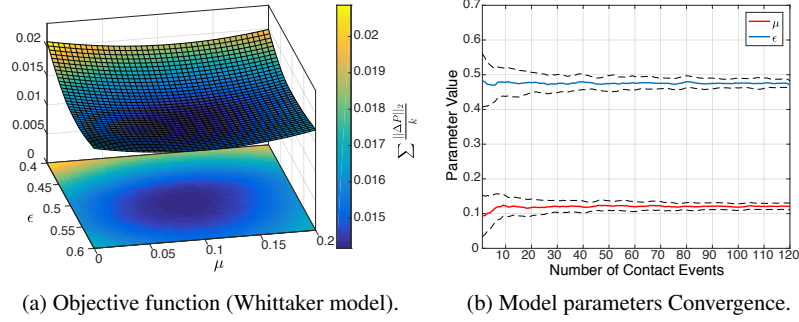
The former scenario will allow us to evaluate the performance of the models in predicting outcomes on previously unseen data, and we will use the latter to show the large variation in parameters across impacts and limitations in the predictions made by the models.

### 3.2.1 Ensemble parameter identification

Ensemble parameters were identified from a randomly selected subset of the data as per the procedure described in Section 3. Fig. 7 (a) shows the shape of the cost function. This function is convex, which is to be expected since the region of admissible impulses predicted by all models is convex and the values of the predicted impulses grow proportionally to the magnitudes of the parameters. Intuitively, if there exists a pair of parameters for a given impact that perfectly predicts the impulse, then the cost associated with that pair would be zero and all other points within the convex set would have error values that grow proportional to deviations from the optimal pair. For a point outside the set of admissible impulses—due to the imposed constraint  $\mu \leq \mu_s$ —then there is a point that lies on the boundary of the set for which the distance takes on a minimum.

Fig. 7 (b) shows the convergence of the parameters as a function of the number of data used for the Whittaker model; for every  $k$  contact events we selected 20 sets of size  $k$  at random without replacement from the experiments. Tab. 1 lists the results of the identification for all 6 models for  $k = 120$ .

Tab. 1 shows that the values for the coefficient of friction found for each model are much lower than the value (0.3) estimated from an inclined surface experiment between PLA (material of the object used) and Aluminium (material of the impact plate). We believe that the major contributing factor to the discrepancy is the complex interaction that the coefficient of friction tries to model in the highly dynamic dropping experiments. The object can undergo a number of rich and difficult to



**Fig. 7** a) Cost function plotted as a function of the parameters for the Whittaker model for 120 drops. b) Convergence behavior of the model parameters as a function of number of contacts used for the identification, the solid colored lines are the average parameter values over a random set of 20 samples and the dashed black lines depict 1 standard deviation.

model interactions in a very short period of time, such as small local deformations, finite-time impact, vibrations, and impulsive torsional frictional forces (in conjunction with the modeled linear ones). The use of just two parameters, which in this case are state independent, is not sufficient to describe the resulting behavior, and the identified parameters lose meaning as they are compensating for all the unmodelled dynamics. In contrast, the inclined surface experiment is slow and controlled, and one parameter is more descriptive of the observations. This is a known phenomenon, [2] discusses the simplicity of impact models, in that they attempt to capture a complicated phenomenon in a minimalistic way and with intuitive parameters. This intuition is valuable: it gives a direct interpretation to the way these parameters affect the outcome of a prediction, but it is unreasonable to expect the values found for these parameters to match those found from other experiments. The coefficient of friction in this context is not necessarily the same as that measured from the inclined surface experiment, though they have the same name; this thought will be elaborated upon in Sec. 3.2.2. It is important to note that using the inclined surface experimental value for  $\mu$  leads to significantly poorer predictive accuracy. Fig. 8

**Table 1** Average-best parameter identified values and the percentage a model was chosen as the best and worst model


Model	$\mu$	$\epsilon$	Best%	Worst%
DrumShell	$0.081 \pm 0.008$	$0.516 \pm 0.009$	4	11
AP Poisson	$0.101 \pm 0.007$	$0.526 \pm 0.008$	7	6
AP Newton	$0.084 \pm 0.007$	$0.547 \pm 0.009$	25	0
Mirtich	$0.062 \pm 0.009$	$0.558 \pm 0.010$	17	43
Wang-Mason	$0.120 \pm 0.008$	$0.537 \pm 0.010$	16	10
Whittaker	$0.111 \pm 0.008$	$0.484 \pm 0.011$	31	30

depicts the estimated PDF of the  $\ell_2$ -norm error of predicted  $v_o$  (we normalize the angular velocity by the radius of gyration so that the linear and angular terms are

of the same order). The Best Post Hoc model is also plotted to show a performance upper bound on the analytical model; that is to say, if, for each trial, we knew *a priori* which model would perform best. Tab. 1 shows the number of times each model was chosen as the best and worst performing for the dataset. We can see that the Whittaker model is quite erratic, while the most consistent model is the AP Newton model, with a significant number of best predictions and no worst.


The IRB Bound model is the result of directly searching for the best impulse that explains the difference between pre- and post-contact velocities subject to the constraints of rigid-body point contact. This post hoc model quantifies the absolute best possible performance for any rigid-body contact model under the assumption of point contact.

It is important to note that the error introduced by the rigid body assumptions is not distinguishable from error resulting from uncertainty in sensor measurements, but our setup was designed to minimize the latter errors. The implication of this ambiguity is that it may be impossible to construct a zero error model without violating the rigid body or rigid contact models, meaning that the best we can hope to achieve in terms of performance is the “Ideal” line. We can further examine the errors made



**Fig. 8** Estimated probability density functions (PDFs) of the  $\ell_2$ -norm error in post-contact velocity of the center of mass.

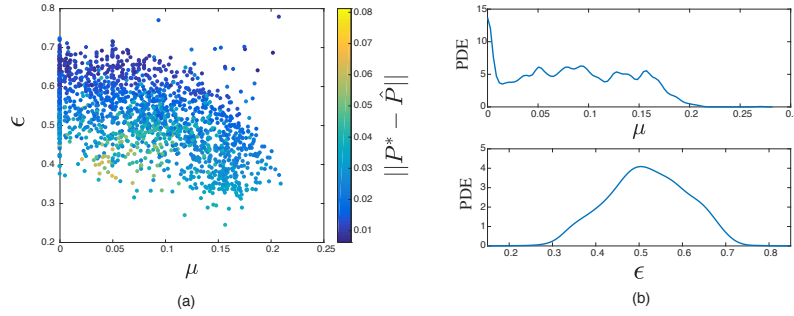
in the models’ predictions to quantify regions of the input space in which we expect models to perform comparatively better and worse. This information is valuable if we are interested in selecting initial conditions that are easier to predict or rejecting initial conditions that may be more difficult to predict (for purposes of control). Fig. 9 depicts the performance of the AP Poisson model in predicting the outcome of a contact (error in predicting the velocity of the COM) as a function of the normal and tangential components of the pre-contact system momentum measured at the contact point. The region in which the tangential momentum at the contact point is near zero or changes direction proves to be the most difficult to predict; we will discuss this phenomenon in Section 3.2.2.



**Fig. 9** Performance of AP Newton model depicted in the momentum space as a function of the normal and tangential components of the pre-contact momentum of the system measured at the contact point. Blue regions represent high accuracy, yellow regions predict high error.

### 3.2.2 Parameter identification for individual impacts

The optimization in Section 3.2.1 to find the ensemble parameters hides some of the details of individual trials. In this section we find the values of  $(\mu, \epsilon)$  that best explain each drop, for example Fig. 10 shows the identified parameters for the AP Newton model, and we make the following observations:




**Fig. 10** Identified values of  $(\mu, \epsilon)$  for each drop, where a) parameter values colored by the ability of the identified pair to predict the impulse, and b) distribution of the parameters.

1. We observe a sharp peak at  $\mu = 0$ .
2. The lower left region in Fig. 10a shows a loss in performance.
3. The distribution of  $\epsilon$  is uni-modal.
4. The means of the distributions of  $\mu$  and  $\epsilon$  are very close to the “optimal” values computed for the ensemble parameters.
5. Neither distribution is tightly localized about a single value. This is perhaps the most important feature to be noted since it is common practice to select a single value of the parameters.

The fact that such variation in parameter values exists, particularly for the value of  $\mu$ , begins to highlight issues of interpretation of the parameters. The intuitive expectation for a coefficient of friction is that it should be constant for all impacts, but these experiments demonstrate that this is not the case. The coefficient of friction in these models acts more similarly to a restitution variable that regulates the magnitude of the tangential impulse as a function of the pre-contact conditions rather than as the typical notion of the coefficient of friction (e.g., that used in the inclined surface experiments).

Part of the challenge in identifying  $\mu$  is that not all impacts are informative. Impacts that stick provide no information other than a lower bound for  $\mu$ . Effectively, any coefficient of friction above the value found by the identification is also able to explain the interaction equally well which partially explains the dispersion in identified  $\mu$ . To prevent parameter drift and an ill-posed optimization, we use  $\mu_s$  as a bound on the coefficient of friction.

We now look at the anomalous cases for which  $\mu = 0$  in isolation; Fig. 11a shows such an example. When the contact point lies directly below the center of mass of the object, the Energy ellipse’s major axis aligns with the normal impulse direction and the convex set of predictions for the outcome of the contact collapses to a very narrow band, resulting in a very small predictive range. A small prediction range coupled with sensor noise the predictive range of the models. In the case shown in Fig. 11, the space spans to the left, but the measured impulse lies to the right; consequently, the value of the coefficient of friction tends to zero to get as close as possible to the target value. This issue highlights the difficulties in interpreting  $\mu$  as the coefficient of friction within the context of these impact models. The second



**Fig. 11** a) An example of  $\mu = 0$ , this impact is characterized by the point of contact being located directly under the center of mass of the object. In this situation, the Energy Ellipse’s axis is aligned to the normal impulse direction. b) Example of a measured impulse outside the convex set of predictions made by the model due to “spin-reversal”.

set of anomalous cases is characterized by the inability to predict “spin-reversal”, an example of which can be seen in Fig. 11b). In these scenarios, the measured impulse does not fall within the predictive range of the model and a residual error will always be present, thereby explaining the region in Fig. 10 where the parameters take on values with non-zero errors. Tab. 2 quantifies the fraction of data that lies within the

convex set of each model. The relative ratios are quite low which indicates that the models suffer from potentially restrictive prediction spaces. Note that models with similar regions in the energy ellipse exhibit similar ratios, exemplified by the AP Poisson and the Drumwright-Shell models.

**Table 2** Fraction of drops that lie within the predictive range of each model

	DrumShell	AP Poisson	AP Newton	Mirtich	Wang-Mason	Whittaker
<b>Fraction</b>	0.355	0.357	0.458	0.413	0.455	0.456

## 4 Conclusion and Summary

Rigid contact models are popular due to their simplicity and computational efficiency yet they exhibit clear limitations in predictive power and accuracy. In this paper, we empirically demonstrated the inability to predict back-spin and challenges with both system identification and interpretation of model parameters.

The limitations of these analytical models can partially be attributed to the parameterization of a subset of the admissible impulse space. The independence of the parameters from the inputs to the models and the limited regions for prediction (Tab. 2) contribute to the errors in predictions. Two approaches can remedy these shortcomings, we may either learn the mapping directly from the impulses and do away with the parameterization all together (data-driven approach) or use the model predictions as a prior and correct the impulse predictions using a data-driven model (data-augmented approach). In Fazeli et al [8] we explore these approaches and compare performances across analytical and learned models.

The rigid models studied assume a single point contact and instantaneous impact. This assumption is only an approximation and the impact event occurs over a finite time and involves a small amount of local deformation leading to a distributed impact impulse over a patch. In Fazeli et al [8] we develop a data-driven contact model that accounts for this deformation using an instantaneous wrench and demonstrate that the resulting model can significantly outperform the rigid contact models.

This study demonstrates the importance of data in model validation and highlights the importance of experimentation in parameter selection, and its potential role in developing data-driven contact models [8, 31]. We hope that this study serves as motivation to experimentally validate other models frequently leveraged in robotics and leads to their better understanding.

## References

- [1] Anitescu M, Potra FA (1997) Formulating dynamic multi-rigid-body contact problems with friction as solvable linear complementarity problems. *Nonlinear Dynamics* 14:231–247
- [2] Chatterjee A, Ruina A (1998) A new algebraic rigid body collision law based on impulse space considerations. *ASME J of Applied Mechanics* 65(4):939–951
- [3] Chavan Dafle N, Rodriguez A (2015) Prehensile Pushing: In-hand Manipulation with Push-Primitives. In: *IEEE/RSJ Intl. Conf. on Intell. Robots Systems (IROS)*, pp 6215 – 6222
- [4] Drumwright E, Shell DA (2010) Modeling contact friction and joint friction in dynamic robotic simulation using the principle of maximum dissipation. In: *Proc. Work. Alg. Found. Robotics (WAFR)*

- [5] Fazeli N, Tedrake R, Rodriguez A (2016) Identifiability analysis of planar rigid-body frictional contact. Intl Symp Robotics Research (ISRR)
- [6] Fazeli N, Donlon E, Drumwright E, Rodriguez A (2017) Empirical evaluation of common contact models for planar impact. In: 2017 IEEE International Conference on Robotics and Automation (ICRA), pp 3418–3425
- [7] Fazeli N, Kolbert R, Tedrake R, Rodriguez A (2017) Parameter and contact force estimation of planar rigid-bodies undergoing frictional contact. The International Journal of Robotics Research
- [8] Fazeli N, Zapolsky S, Drumwright E, Rodriguez A (2017) Learning data-efficient rigid-body contact models: Case study of planar impact. arXiv preprint arXiv:171005947
- [9] Gautier M, Khalil W (1988) On the identification of the inertial parameters of robots. In: IEEE Conf. on Decision and Control, pp 2264–2269
- [10] Hertz H (1881) On the contact of elastic solids. J Reine Angew Math 92:156–171
- [11] Hogan F, Rodriguez A (2016) Feedback Control of the Pusher-Slider System: A Story of Hybrid and Underactuated Contact Dynamics. In: Proc. Work. on Alg. Found. Robotics (WAFR)
- [12] Kane TR, Levinson DA (1985) Dynamics: Theory and Applications. McGraw-Hill, N.Y.
- [13] Khosla P, Kanade T (1985) Parameter identification of robot dynamics. In: IEEE Conf. on Decision and Control, pp 1754–1760
- [14] Koval M, Pollard N, Srinivasa S (2015) Pose estimation for planar contact manipulation with manifold particle filters. The Intl Journal of Robotics Research 7(34):922–945
- [15] Kraus PR, Kumar V (1997) Compliant contact models for rigid body collisions. In: Proceedings of Intl. Conf. on Robotics and Automation, vol 2, pp 1382–1387 vol.2
- [16] Lacoursière C (2007) Ghosts and machines: Regularized variational methods for interactive simulations of multibodies with dry frictional contacts. PhD thesis, Umeå University
- [17] Lankarani H, Nikravesh P (1990) A contact force model with hysteresis damping for impact analysis of multibody systems. J Mech Design 112(3):369–376
- [18] Li Z, Kota S (2001) Virtual prototyping and motion simulation with adams. J Comput Inf Sci Eng 1(3):276–279
- [19] Marhefka DW, Orin DE (1999) A compliant contact model with nonlinear damping for simulation of robotic systems. IEEE Trans Systems, Man, Cyber 29(6):566–572
- [20] Mirtich B (1996) Impulse-based dynamic simulation of rigid body systems. PhD thesis, University of California, Berkeley
- [21] Nubiola A, Bonev IA (2013) Absolute calibration of an abb irb 1600 robot using a laser tracker. Robotics and Computer-Integrated Manufacturing 29(1):236–245
- [22] Posa M, Cantu C, Tedrake R (2013) A direct method for trajectory optimization of rigid bodies through contact. Intl J Robotics Research 33(1):69–81
- [23] Shapiro A, Faloutsos P, Ng-Thow-Hing V (2005) Dynamic animation and control environment. In: Proceedings of graphics interface 2005, Canadian Human-Computer Communications Society, pp 61–70
- [24] Slotine JJE (1987) On the Adaptive Control of Robot Manipulators. Intl J Robotics Research 6(3):49–59
- [25] Song P, Kraus P, Kumar V, Dupont P (2000) Analysis of rigid body dynamic models for simulation of systems with frictional contacts. J Applied Mech 68
- [26] Stewart DE (2000) Rigid-body dynamics with friction and impact. SIAM, vol 42, pp 3–39
- [27] Stewart DE, Trinkle JC (1996) An implicit time-stepping scheme for rigid body dynamics with inelastic collisions and coulomb friction. Intl J Num Methods Engr 39(15):2673–2691
- [28] Stronge WJ (1990) Rigid body collisions with friction. Proc Royal Soc Lond A 431(169–181)
- [29] Wang Y, Mason MT (1992) Two-dimensional rigid-body collisions with friction. ASME J Appl Mech 59:635–642
- [30] Whittaker ET (1944) A Treatise on the Analytical Dynamics of Particles and Rigid Bodies, 4th ed. Dover
- [31] Yu KT, Bauza M, Fazeli N, Rodriguez A (2016) More than a Million Ways to be Pushed. A High-Fidelity Experimental Data Set of Planar Pushing. In: IEEE/RSJ Intl. Conf. on Intell. Robots and Systems (IROS)