

Learning Data-Efficient Rigid-Body Contact Models: Case Study of Planar Impact

Nima Fazeli

Department of Mechanical Engineering
Massachusetts Institute of Technology MA, United States
nfazeli@mit.edu

Samuel Zapolsky

Toyota Research Institute
CA, United States
sam.zapolsky@tri.global

Evan Drumwright

Toyota Research Institute
CA, United States
evan.drumwright@tri.global

Alberto Rodriguez

Department of Mechanical Engineering
Massachusetts Institute of Technology MA, United States
albertor@mit.edu

Abstract: In this paper we demonstrate the limitations of common rigid-body contact models used in the robotics community by comparing them to a collection of data-driven and data-reinforced models that exploit underlying structure inspired by the rigid contact paradigm. We evaluate and compare the analytical and data-driven contact models on an empirical planar impact data-set, and show that the learned models are able to outperform their analytical counterparts with a small training set.

Keywords: Contact Modeling, Rigid-Body Dynamics, Data-Efficient Learning

1 Introduction

Rigid-body contact models are ubiquitous in robotics with applications, for example, in planning (Posa et al. [1]), control (Hogan and Rodriguez [2]), system identification (Fazeli et al. [3, 4], Zhou et al. [5]), and state-estimation (Koval et al. [6, 7]). In this paper we briefly review six established rigid contact models from the literature, and evaluate the predictive performance of these models on an empirical data-set. Next we highlight the limitations of these models as motivation for the need to involve data in the modeling process; and propose two classes of data-efficient learning paradigms inspired by rigid contact models and show that they are able to outperform their counterparts.

Prevalent contact models used in robotics assume rigid body dynamics and can generally be categorized as: i) pseudo-rigid models (i.e., a thin and inertialess compliant layer over a rigid core) and ii) rigid contact models (resulting from constraint equations). These models have the desired properties of computational efficiency and simplicity, as opposed to finite element models that are expensive to compute and not suitable for realtime applications. Penalty methods assume that the bodies in contact do not deform but allow for penetration as a proxy to deformation, and compute the imparted contact forces as a function of the amount of penetration. Adams multi-body simulator [8] (commercialized) and Dance [9] (open source) are penalty method simulation packages. Akin to these penalty methods are the relaxations via regularization approaches, as deployed in Todorov et al. [10].

Rigid-body dynamic simulation using the complementarity formulation as popularized by Anitescu and Potra [11] and Stewart and Trinkle [12] resolves contact without penetration, assumes perfectly rigid bodies, and treats contact as a instantaneous and discontinuous event. These approaches make use of contact models (such as Stronge [13], Wang and Mason [14], Whittaker [15]) that are amenable to the complementarity formulation. Regarding both categories, only a handful of studies (Fazeli et al. [16, 17]) have tried to provide formal empirical verification of contact models and model selection is treated mostly as a matter of convenience, highlighting the need for experimental work and data-driven approaches.

Motivated by the modicum of empirical work in contact model verification and selection, as well as the limited usage of data in improving and constructing effective contact models (Yu et al. [18]), we first review several common rigid body contact models and show their limitations in making predictions, then construct data-driven models using the underlying rigid-body contact assumption and demonstrate that they are able to out-perform the best possible predictions made by the models.

2 Background - Rigid-Body Contact Models

Consider Fig. 1a, where a planar object makes contact with a horizontal surface, the purpose of a contact model is to predict the imparted impulses and the resultant states of the object. We use m and I to represent the mass and second moment of inertia of the object, and (r_x, r_y) denote the position of the contact point with respect to the center of mass of the object in an inertial frame. We represent the configuration of the object with $\mathbf{q} = (x, y, \theta)^T$, the velocity with $\mathbf{v} = \dot{\mathbf{q}}$, the imparted impulse during impact with \mathbf{P} , and use superscripts i and f to denote pre- and post-contact variables. A contact model performs the mapping:

$$\mathbf{f}_{\text{contact}} : \{m, I, r_x, r_y, \mathbf{q}, \mathbf{v}\}_{10 \times 1} \xrightarrow{\mathbf{P}} \{\mathbf{q}, \mathbf{v}\}_{6 \times 1}^f \quad (1)$$

This mapping does not have a functional form yet. Rigid-body contact models assume that deformations are negligible during impact, and resolve contact instantaneously. This assumption means that the configuration of the object pre- and post-contact does not change and the impact happens at a single point on the body. Point contacts permit only linear impulses to be transmitted between the bodies. We can write the dynamic equations of motion for the instant before and after the impact for a rigid-body as:

$$\mathbf{v}^f = \mathbf{v}^i + \mathbf{M}^{-1} \mathbf{J}^T \mathbf{P} \quad (2)$$

where $\mathbf{M}_{3 \times 3}$ denotes the inertia matrix, $\mathbf{J}_{2 \times 3}$ is the contact Jacobian, and $\mathbf{P}_{2 \times 1}$ is the linear imparted impulse. The rigid-body assumption gives Eq. 1 a functional form in Eq. 2 but is still unresolvable since \mathbf{P} and \mathbf{v}^f are both unknowns.

The impulse \mathbf{P} is constrained by the law of conservation of energy and that no penetration can occur. To derive a mathematical representation for these two constraints, we pre-multiply Eq. 2 by \mathbf{J} :

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

where \mathbf{v}_c denotes the velocity of the contact point, and \mathbf{M}_c represents the effective inertia at the point of contact. We may now express the conservation of kinetic energy for the contact point as:

$$(\mathbf{v}_c^f)^T \mathbf{M}_c \mathbf{v}_c = \alpha (\mathbf{v}_c^i)^T \mathbf{M}_c \mathbf{v}_c^i, \quad \alpha \in [0, 1] \quad (4)$$

Plugging in Eq. 3 we arrive at:

$$(\mathbf{P} + \mathbf{M}_c \mathbf{v}_c^i)^T \mathbf{M}_c^{-1} (\mathbf{P} + \mathbf{M}_c \mathbf{v}_c^i) = \alpha (\mathbf{v}_c^i)^T \mathbf{M}_c \mathbf{v}_c^i \quad (5)$$

Eq. 5 defines an ellipse in the impulse space (\mathbf{P}), with an offset of $-\mathbf{M}_c \mathbf{v}_c^i$ which is the incident momentum of the contact point. This ellipse is called the Energy Ellipse and is explained in more detail in (Chatterjee and Ruina [19], Fazeli et al. [17]). Fig. 1b shows an example of an Energy Ellipse, the line of maximum compression denotes impulses that yield zero separating velocities, and the line of sticking denotes the impulses that lead to a zero relative tangential velocity component. The lightly shaded region is the space of admissible impulses subject to the energy and non-penetration constraints. These constraints are still not sufficient to resolve the contact, contact models introduce additional constraints and parameters to select a unique imparted impulse from this admissible set.

2.1 Analytical Contact Models Studied

Analytical contact models proposed in the literature parameterize some region within the admissible impulse space, and select a unique impulse given pre-contact states and parameters. The selected impulse \mathbf{P} is then plugged into the equations of motion (Eq. 2) to predict the velocity vector for the next time instant.

In this study, we consider 6 commonly used analytical models: 1) **Anitescu-Potra Newton** [11], 2) **Anitescu-Potra Poisson** [11], 3) **Drumwright-Shell** [20], 4) **Mirtich** [21], 5) **Wang-Mason**

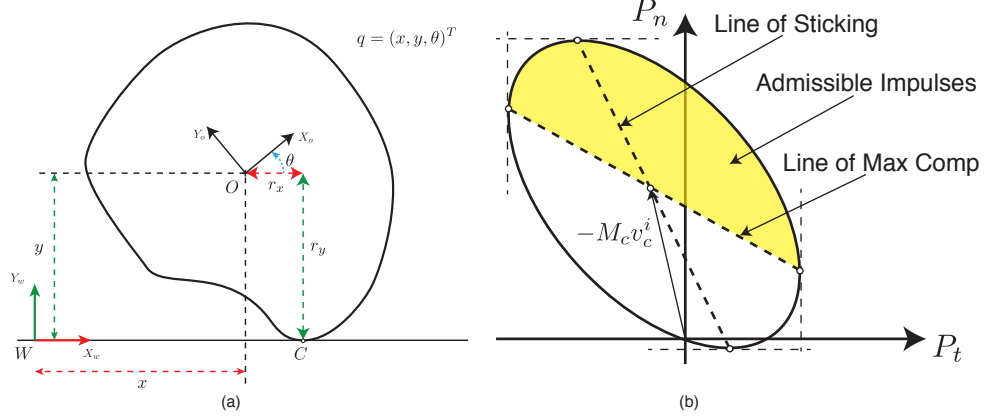


Figure 1: a) planar object making contact with a fixed horizontal surface, b) Energy Ellipse.

[14], and 6) **Whittaker** [15]. The details of these models are summarized in Fazeli et al. [16]. In addition to these models, we define two post hoc models “**Best Post Hoc**” and “**Ideal Rigid Body Bound (IRB Bound)**”. The Best Post Hoc model chooses the analytical model that predicted the outcome most accurately for each trial (the best possible performance for the 6 models), and the IRB Bound chooses an impulse in the Energy Ellipse that best explains the trial given the outcome (establishing an absolute upper bound on the performance of any rigid contact model satisfying the standard assumptions made). We present the computation of the IRB Bound in sec. 3.2.

The key common feature of the six analytical models is that they are all characterized by two parameters, a normal restitution coefficient (ϵ) and a tangential friction coefficient (μ), though each model uses and interprets these parameters differently. The normal restitution regulates the magnitude of the imparted normal impulse (a larger value indicates a larger impulse and a larger “bounce”) and the tangential friction coefficient regulates the tangential impulse, and discriminates between sticking and sliding during the impact event. Fig. 2 shows the regions of predictions of the models within the Energy Ellipse for all possible values of the parameters, grouped based on graphical similarity. Though these regions appear similar graphically, they are constructed differently computationally.

Fig. 2 shows that no model is able to cover a large portion of the admissible space, and that the regions of predictions are quite similar. Further, we note that none of the six studied models is able to predict points to the opposite side of the line of sticking with respect to their own subspace. This phenomenon is referred to as “back-spin” [22] and occurs during our experiments. We can also infer the definitions of the parameters used in the models from the shape of the admissible regions, for example the Poisson restitution is a ratio of the normal momenta at contact, therefore the models that use it have vertical lateral boundaries.

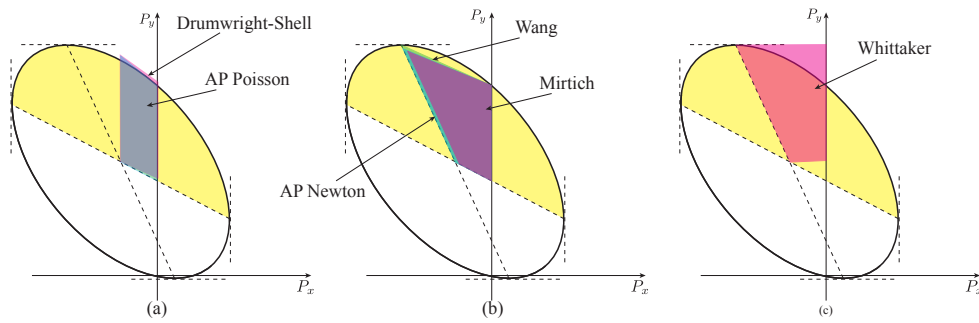


Figure 2: Space of admissible impulses demarcated by the contact models for all possible choices of parameters.

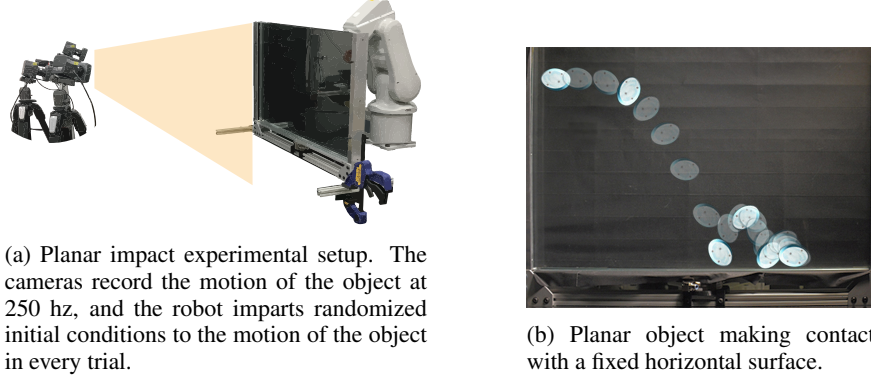


Figure 3: Experimental setup: robotic autonomous impact data collection.

3 Experiments and the Performance of the Analytical Models

3.1 Experimental Setup:

We devised the experimental setup in Fig. 3a for the evaluation and learning of the contact models. The experiment consists of planar impact of objects subject to gravity, and is designed for autonomous data collection with the robot. Fig. 3b depicts an example trajectory of an ellipse.

The setup consists of a dropping arena, where two flat sheets of glass constrain the motion of planar objects to the vertical plane, the motion tracking cameras to measure the states of the object, and the robot with a directional magnetic latch mechanism to manipulate the object (equipped with similar directional magnets) between the glass sheets. In each run of the experiment, the robot generates a random set of initial conditions for the drop, the object is released with these conditions, and the subsequent motion the object is recorded. In total 2000 drops were generated and 1718 of them were selected after filtering for data integrity (missing frames, reflection, unexpected releases etc.). Next, contact events were detected, pre- and post-contact states were extracted, and the physical properties of the object were carefully measured. For more details please see Fazeli et al. [17].

3.2 Rigid-body Contact Model Predictive Performances

An effective way of quantifying the predictive performance of the contact models is to compare the measured post contact velocities of the object with the values predicted by the models. To this end, first we find the optimal model parameters (system identification) for each model. For each model, we chose k trials of the experiment without replacement m times, and for each of m iterations we find the optimal parameters by solving:

$$\begin{aligned} \arg \min_{\mu, \epsilon} \quad & \sum_{n=1}^k \|\mathbf{P}_n - \hat{\mathbf{P}}_n\|_2 \\ \text{subject to} \quad & \dot{\mathbf{P}}_n = \mathbf{f}_c(\mu, \epsilon, \mathbf{q}_n^i, \mathbf{v}_n^i) \end{aligned} \quad (6)$$

Here \mathbf{f}_c denotes the contact model of interest, \mathbf{P}_n denotes the measured impulse for trial n , and $\hat{\mathbf{P}}_n$ is the estimated impulse from the contact model. Next we compute the average and standard deviation over the m iterations. Fig. 4a shows the surface of the cost function (Eq. 6) plotted as a function of the parameters for the Whittaker model for $k = 120$. Fig. 4b shows the convergence behaviour of the model parameters for $k = 1, \dots, 120$. For more details on the identification procedure see Fazeli et al. [17]. Tab. 1 shows the computed optimal model parameter values for each model.

With the identified optimal parameters we are ready to evaluate the predictive performance of the models, here we chose the l_2 norm of the error in predicted post-contact center of mass velocity of the object as the metric. Fig. 5 shows the performance of the six analytical models and the two post hoc models. The IRB Bound is the post hoc model that selects the impulse in the admissible space of the Energy Ellipse that best explains the trial and is computed by solving:

$$\arg \min_{\mathbf{P}} \quad \|\mathbf{v}_n^f - \mathbf{v}_n^i - \mathbf{M}^{-1} \mathbf{J}^T \mathbf{P}\|_2 \quad (7)$$

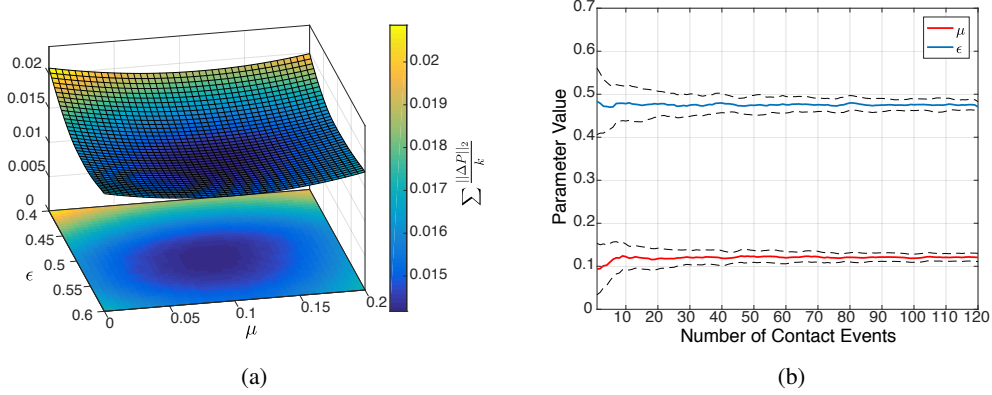


Figure 4: a) Whittaker model objective function for $k = 120$. b) Convergence of Whittaker model parameters as a function of number of trials. Solid colors show mean parameter value, and dashed black show 1 standard deviation.

Table 1: Optimal parameters for the models.

Model	μ	ϵ
DrumShell	0.081 ± 0.008	0.516 ± 0.009
AP Poisson	0.101 ± 0.007	0.526 ± 0.008
AP Newton	0.084 ± 0.007	0.547 ± 0.009
Mirtich	0.062 ± 0.009	0.558 ± 0.010
Wang-Mason	0.120 ± 0.008	0.537 ± 0.010
Whittaker	0.111 ± 0.008	0.484 ± 0.011

where $v_{3 \times 1}$ denotes the velocity of the center of mass of the object, and $P_{2 \times 1}$ is the estimated imparted impulse under the rigid-body contact model. All the contact models studied so far assume point contact and only allow a linear impulse to be imparted on the object and the IRB Bound gives an absolute bound on the performance of any model constructed from the rigid body contact model assumptions. In sec. 4 we learn models that relax some of these assumptions.

We also study each impact individually and identify the contact parameters for each impact that most optimally explain the event. Fig. 6a shows the (μ_k, ϵ_k) for each drop, colored by the amount of error in the resulting prediction. We see that for a subset of the initial conditions, two parameters can be found such that the predicted and measured impulses are equal, but such alignment does not hold globally, indicating that a significant number of impacts exist for which the measured impulse does not lie within the polygonal predictive range of the models. Fig. 6b shows the distribution of the model parameters found for each drop, and while the mean of these distributions is very close to

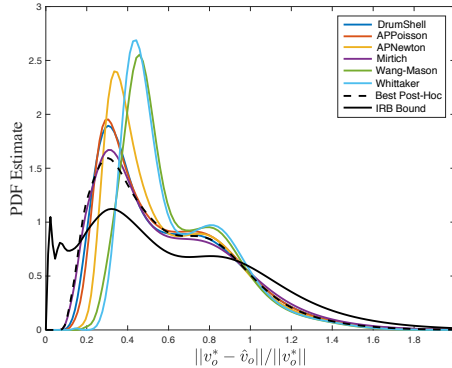


Figure 5: Normalized l_2 norm error in predicted post contact center of mass velocity.

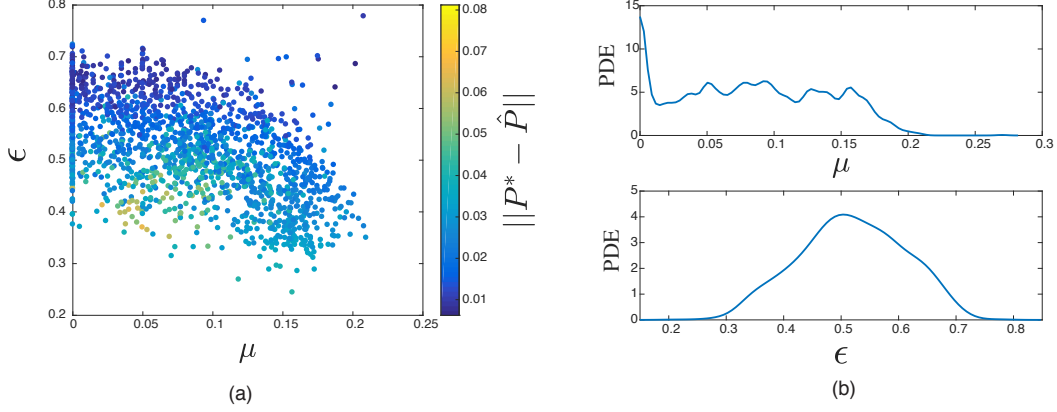


Figure 6: a) Anitescu-Potra Newton model parameters and corresponding error in predicted impulse, blue to yellow indicates increase in error, b) AP Newton model parameters (μ, ϵ) distributions.

the values in Tab. 1, we see a significant amount of variance. The data suggests that: i) the analytical models are not expressive enough to explain all impacts (the predictive ranges do not cover a large enough region in the feasible impulse space), and ii) the choice of contact parameters should depend on the initial conditions of the impact, which is rarely done in practice.

4 Data Driven Contact Models

Sec. 3.2 highlights the variability in optimal contact parameters across experiments, and indicates that a significant portion of the imparted impulses do not lie in the predictive range of the commonly used contact models. This is in part due to the fact that the analytical contact models rely on parametrization of the impulse space and introduce additional constraints to limit the set of admissible imparted impulses. As Chatterjee and Ruina [19] discusses, the parameters used are only very coarse approximations to very complicated and difficult to model phenomena, and there is no reason to expect that these or the constraints imposed on the interactions are physically meaningful.

In light of this, we propose two data-driven contact model formulations. The first class is purely data-driven and uses Gaussian Processes (GPs) trained on a feature space inspired by the analytical rigid contact models. The second class uses data to reinforce analytical models, either by informing parameter choice or correcting the predictions made by the models. The first class has weak priors, whereas the second class uses the analytical models which provide a much stronger prior.

4.1 Class I - Data-Driven Contact Models

A contact model representation (Eq. 1) can be learned directly from the data using the feature space:

$$\mathcal{X} = \{m, I, r_x, r_y, \mathbf{q}, \mathbf{v}\}^i \rightarrow \mathcal{Y} = \{\mathbf{q}, \mathbf{v}\}^f \quad (8)$$

Using this feature space gives the learned models a large amount of flexibility at the cost of dimensionality of the space (10-D). For impacts that are near rigid we can exploit the structure and assumptions of rigid body contact to reduce the dimension of the feature space. The Energy Ellipse demarcates the region in space for all feasible imparted impulses under the rigid-body assumption that satisfy the law of conservation of energy. This space is a function of the effective inertia at the contact point (\mathbf{M}_c , square symmetric 2×2 matrix) which captures the position of the center of mass with respect to the contact point as well as the inertia of the object, and the velocity of the contact point \mathbf{v}_c . As such we propose two feature spaces that exploit this structure:

$$\mathcal{X}_1 = \{\mathbf{M}_{c,11}, \mathbf{M}_{c,12}, \mathbf{M}_{c,22}, \mathbf{v}_{c,t}, \mathbf{v}_{c,n}\} \quad (9)$$

$$\mathcal{X}_2 = \{(\mathbf{M}_c \mathbf{v}_c)_t, (\mathbf{M}_c \mathbf{v}_c)_n\} \quad (10)$$

$$\mathcal{Y}_1 = \{P_t, P_n\} \quad (11)$$

Here \mathcal{X}_1 and \mathcal{X}_2 are respectively 5- and 2-D, a significant reduction from the 10-D space. \mathcal{X}_1 attempts to capture the key features required to reconstruct the space of admissible impulses, while

\mathcal{X}_2 is a coarse approximation of this space, where we consider the aggregate effect of inertia and velocity as the linear momentum of the contact point. We refer to these two models as “**Data-driven Rigid Contact Models**” due to their similarity to the analytical models.

The data-driven rigid contact models inherit the same properties as the analytical models, they assume a point contact that occurs instantaneously which results in an instantaneous linear impulse. Real objects often undergo small deformations during contact over a finite time horizon where forces are distributed over a finite area. We can allow greater expressibility of our contact models by permitting an instantaneous wrench (force and torque) prediction, i.e. using $(\mathcal{X}_1, \mathcal{Y}_2 = \{P_t, P_n, \tau\})$. We refer to this model as “**Data-driven Contact Model**” since this model deviates from the rigidity assumption.

To learn these contact models we used Gaussian Processes (GPs) ([23]) with an Automatic Relevance Determination (ARD) Square Exponential Kernel. We use this kernel to compensate for the relative differences in magnitudes of the features.

4.2 Class II - Data Reinforced Contact Models

The models presented in sec. 4.1 are able to make predictions that span a larger subspace of the Energy Ellipse than their analytical counterparts, but can conceivably produce infeasible predictions until trained on a sufficient number of data.

In contrast analytical models, by construction, produce feasible predictions for reasonable choices of parameters. We can use data to reinforce these models and to this end we propose two formulations: i) using an analytical model and learning the residual errors in prediction using a GP, and ii) learning the optimal choice of model parameters for an analytical model given pre-impact states. Analogous to the naming convention in sec. 4.1 we will refer to these models as “**Data-reinforced Rigid Contact Models**”. The former type of model is amenable to using an instantaneous wrench (i.e. \mathcal{Y}_2) and we will refer to the resulting model as “**Data-reinforced Contact Model**”. In both cases we use the feature space \mathcal{X}_5 and the same kernels as sec. 4.1.

4.3 Data-driven Model Performance

Fig. 7 shows the ℓ_2 norm error in predicted post contact velocity of the data-driven models compared to the Best Post Hoc model (best performance of the analytical models) and the IRB Bound model (best performance of any rigid-body contact model). The rigid models of either class are able to outperform the Best Post Hoc model, but only approach the IRB Bound model. These models are able to outperform their analytical counterparts, but because they are built using the same features and assumptions as the analytical models, they will not be able to exceed the IRB bound. In Class I, the \mathcal{X}_1 shows a marginally better performance over \mathcal{X}_2 , as it has a richer description of the impact event. The learned models with the extra explanatory power of wrench are able to outperform the IRB bound significantly and make far better predictions as they are not restricted by the same assumptions.

Fig. 8 shows the expectation of the ℓ_2 -norm error in predicting the center of mass velocity as a function of the number of data used in the training, and here we show the best performing of each type of model for clarity. The data-driven and data-reinforced contact models are able to consistently and significantly outperform the IRB Bound due to their added expressibility, but require more data to converge (approximately 300 samples). The rigid learned models tend to converge to their optimal performance in approximately 100 samples, which is about double the number of samples required for the reliable identification of the parameters of the analytical models.

5 Discussion & Conclusion

Contact models fulfill an important role in robotics and are used extensively in locomotion [24] and manipulation [25], yet the literature contains little work in empirical evaluation and comparison of the models prevalent in the field [10]. In this paper we show that empirical data plays an important role in highlighting the limitations of analytical models, and in allowing the construction of data-driven contact models that outperform their analytical counterparts. We also demonstrate that by exploiting known structure and assumptions about physical interactions, we can build models that

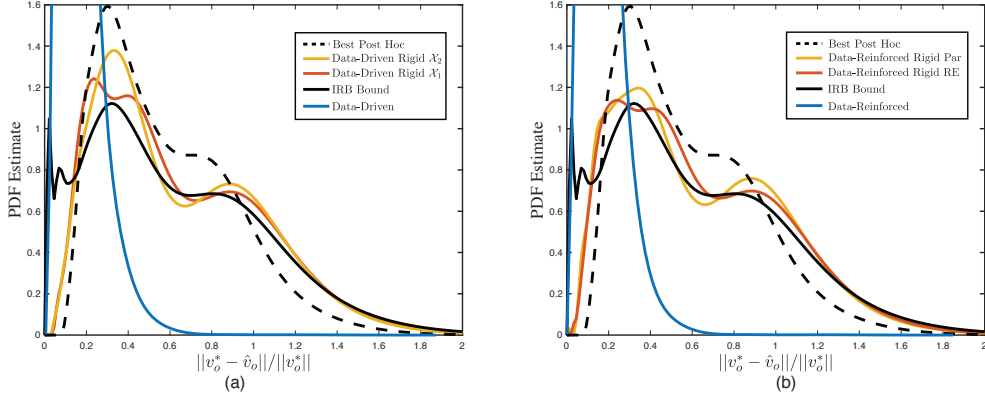


Figure 7: The estimated probability density function of the ℓ_2 norm error in predicted post contact velocity for 450 data samples used to train the models. a) Class I, b) Class II.

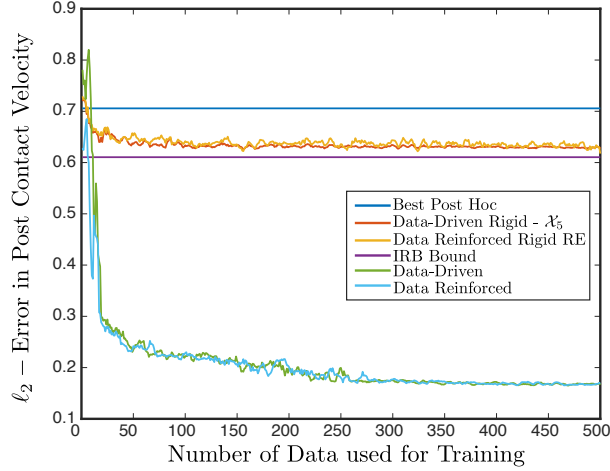


Figure 8: The ℓ_2 norm error in predicted post contact velocity for the models studied vs. the number of data samples used to train models.

are more data-efficient. In particular, here we were able to reduce a 10-D space to 5-D. We further show significant improvements in model prediction when we allow models to deviate slightly from the requirements of rigid contact, i.e. using instantaneous wrench instead of linear impulse.

In this study we considered two classes of models, the purely data-driven (weak priors), and the data reinforced models (with strong priors). The challenge with using a purely data-driven model is the potentially poor or infeasible predictions made by these models before being trained on enough data. In contrast, augmenting analytical contact models with data allows for feasible predictions with adjustments made as more data is collected. With no data, the predictions rely on the prior generated by the underlying analytical models with no adjustment, which will be feasible but may be inaccurate. Our results suggest that these two formulations saturate at similar performances for this experimental setup, and the data reinforced models may be preferred due to the strong priors they provide.

The present study focuses on the impact of a single object, and learns a contact model suitable for the predictions of outcomes. Two interesting questions to ask are: how well does the model learned for this object work in making predictions for other objects of different shapes and physical properties? Is it possible to use this model as a prior and learn a correcting model with far fewer samples?

References

- [1] M. Posa, C. Cantu, and R. Tedrake. A direct method for trajectory optimization of rigid bodies through contact. *The International Journal of Robotics Research*, 33(1):69–81, Oct. 2013.
- [2] F. Hogan and A. Rodriguez. Feedback Control of the Pusher-Slider System: A Story of Hybrid and Underactuated Contact Dynamics. In *Workshop on Algorithmic Foundation of Robotics (WAFR)*, 2016.
- [3] N. Fazeli, R. Kolbert, R. Tedrake, and A. Rodriguez. Parameter and contact force estimation of planar rigid-bodies undergoing frictional contact. *The International Journal of Robotics Research*, 2016.
- [4] N. Fazeli, R. Tedrake, and A. Rodriguez. Identifiability analysis of planar rigid-body frictional contact. In *Robotics Research*, pages 665–682. Springer, 2018.
- [5] J. Zhou, R. Paolini, J. A. Bagnell, and M. T. Mason. A convex polynomial force-motion model for planar sliding: Identification and application. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2016.
- [6] M. C. Koval, M. R. Dogar, N. S. Pollard, and S. S. Srinivasa. Pose estimation for contact manipulation with manifold particle filters. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4541–4548. IEEE, 2013.
- [7] M. Koval, N. Pollard, and S. Srinivasa. Pose estimation for planar contact manipulation with manifold particle filters. *The International Journal of Robotics Research*, 7(34):922–945, 2015.
- [8] Z. Li and S. Kota. Virtual prototyping and motion simulation with adams. *J. Comput. Inf. Sci. Eng.*, 1(3):276–279, 2001.
- [9] A. Shapiro, P. Faloutsos, and V. Ng-Thow-Hing. Dynamic animation and control environment. In *Proceedings of graphics interface 2005*, pages 61–70. Canadian Human-Computer Communications Society, 2005.
- [10] E. Todorov, T. Erez, and Y. Tassa. Mujoco: A physics engine for model-based control. In *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, pages 5026–5033. IEEE, 2012.
- [11] M. Anitescu and F. A. Potra. Formulating dynamic multi-rigid-body contact problems with friction as solvable linear complementarity problems. *Nonlinear Dynamics*, 14:231–247, 1997.
- [12] D. E. Stewart and J. C. Trinkle. An implicit time-stepping scheme for rigid body dynamics with inelastic collisions and coulomb friction. *Intl. J. Numerical Methods in Engineering*, 39(15):2673–2691, 1996.
- [13] W. J. Stronge. Rigid body collisions with friction. *Proc. of the Royal Society of London A*, 431(169–181), 1990.
- [14] Y. Wang and M. T. Mason. Two-dimensional rigid-body collisions with friction. *ASME J. Appl. Mech.*, 59:635–642, 1992.
- [15] E. T. Whittaker. *A Treatise on the Analytical Dynamics of Particles and Rigid Bodies*, 4th ed. Dover, 1944.
- [16] N. Fazeli, E. Donlon, E. Drumwright, and A. Rodriguez. Empirical evaluation of common contact models for planar impact. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3418–3425, May 2017.
- [17] N. Fazeli, S. Zapolsky, E. Drumwright, and A. Rodriguez. Fundamental limitations in performance and interpretability of common planar rigid-body contact models. *arXiv preprint arXiv:1710.04979*, 2017.

- [18] K.-T. Yu, M. Bauza, N. Fazeli, and A. Rodriguez. More than a Million Ways to be Pushed. A High-Fidelity Experimental Data Set of Planar Pushing. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2016.
- [19] A. Chatterjee and A. Ruina. *A New Algebraic Rigid Body Collision Law Based on Impulse Space Considerations*, volume 65. Dec 1998.
- [20] E. Drumwright and D. A. Shell. Modeling contact friction and joint friction in dynamic robotic simulation using the principle of maximum dissipation. In *Proc. of Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 2010.
- [21] B. Mirtich. *Impulse-based Dynamic Simulation of Rigid Body Systems*. PhD thesis, University of California, Berkeley, 1996.
- [22] A. Chatterjee and A. Ruina. A new algebraic rigid body collision law based on impulse space considerations. *ASME Journal of Applied Mechanics*, 65(4):939–951, Dec 1998.
- [23] C. E. Rasmussen and C. K. Williams. *Gaussian processes for machine learning*, volume 1. MIT press Cambridge, 2006.
- [24] S. Kuindersma, R. Deits, M. Fallon, A. Valenzuela, H. Dai, F. Permenter, T. Koolen, P. Marion, and R. Tedrake. Optimization-based locomotion planning, estimation, and control design for the atlas humanoid robot. *Autonomous Robots*, 40(3):429–455, 2016.
- [25] N. Chavan Dafle and A. Rodriguez. Prehensile Pushing: In-hand Manipulation with Push-Primitives. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6215 – 6222, 2015. URL <http://dx.doi.org/10.1109/IROS.2015.7354264>.