Impact-Aware Manipulation by Dexterous Robot Control and Learning in Dynamic Semi-Structured Logistic Environments

I.Model Report

Dissemination level	Public (PU)
Work package	WP1 - Modeling
Deliverable number	D1.3
Version	F-1.0
Submission date	25/07/2023
Due date	30/06/2023

www.i-am-project.eu

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 871899

Authors

Authors in alphabetical order		
Name	**Organisation**	**Email**
Maarten JONGENEEL	TUe	m.j.jongeneel@tue.nl
Jos DEN OUDEN	TUe	j.h.v.d.ouden@tue.nl
Alexander OLIVA	TUe	a.a.oliva@tue.nl
Alessandro SACCON	TUe	a.saccon@tue.nl
Jari VAN STEEN	TUe	j.j.v.steen@tue.nl

Control sheet

Version history			
Version	**Date**	**Modified by**	**Summary of changes**
0.1	13/05/2020	Jos DEN OUDEN	TOC & first contents
0.11	21/06/2023	Maarten JONGENEEL	General structure and first content
0.12	22/06/2023	Maarten JONGENEEL	Finished Sections 2.1, 2.2, 2.5, and 2.6
0.13	26/06/2023	Maarten JONGENEEL	Finished Sections 2.3, update Section 2.4
0.14	27/06/2023	Maarten JONGENEEL	Finished Section 1 and conclusions, update Section 2.4
0.15	29/06/2023	Jari VAN STEEN	Update Section 2.4.2.4
0.16	29/06/2023	Alexander OLIVA	Update Section 2.4.2.1 & Section 2.4.2.3
0.17	07/07/2023	Maarten JONGENEEL	Update Section 2.4.2.2
0.18	12/07/2023	Jos DEN OUDEN	Pre-final version (reviewed)
1.0	24/07/2023	Maarten JONGENEEL	Final version ready for submission

Peer reviewers		
	Reviewer name	**Date**
Reviewer 1	Jos DEN OUDEN	12/07/2023
Reviewer 2	Alessandro SACCON	24/07/2023

Legal disclaimer

TABLE OF CONTENTS

ABBREVIATIONS

Abbreviation	Definition
AGX	Algoryx
API	Application programming interface
CAD	Computer Aided Design
DOF	Degrees of Freedom
EC	European Commission
FPS	Frames per second
LWPR	Locally Weighted Projection Regression
PU	Public
RGB	Red Green Blue (color spectrum)
SE(3)	Special Euclidian Group in 3 dimensions
SO(3)	Special Orthogonal Group in 3 dimensions
TOC	Table of content
TUE	Technical University of Eindhoven
UPF	Unscented Particle Filter
WP	Work Package

EXECUTIVE SUMMARY

This deliverable D1.3 I.Model report, aims at providing an overview of the publications related to the I.Model framework. More specifically, for each of the tasks related to I.Model (T1.1, T1.2, T1.3, T1.4, T1.5, T2.1), this report summarizes the main achievements, limitations, and future directions. As the I.Model framework focuses mainly on developing parameter identification methods to obtain parametric impact laws from experimental data, enabling to predict the dynamic effect of impacts (relative contact speed higher than 0.1m/s, prediction accuracy in post-impact speed of 10%), this report will reflect on the developments related to these tasks. The modeling efforts of I.Model are then used for learning, sensing, and control of dynamic manipulation tasks, that are related to the remaining three objectives.

1 INTRODUCTION

1.1 I.AM. project background

Europe is leading the market of torque-controlled robots. These robots can withstand physical interaction with the environment, including impacts, while providing accurate sensing and actuation capabilities. I.AM leverages this technology and strengthens European leadership by endowing robots to exploit intentional impacts for manipulation. I.AM focuses on impact aware manipulation in logistics, a new area of application for robotics which will grow exponentially in the coming years, due to socio-economical drivers such as booming of e-commerce and scarcity of labor. I.AM relies on four scientific and technological research lines that will lead to breakthroughs in modeling, sensing, learning, and control of fast impacts:

1. I.Model offers experimentally validated accurate impact models, embedded in a highly realistic simulator to predict post-impact robot states based on pre-impact conditions;
2. I.Learn provides advances in planning and learning for generating desired control parameters based on models of uncertainties inherent to impacts;
3. I.Sense develops an impact-aware sensing technology to robustly assess velocity, force, and robot contact state in close proximity of impact times, allowing to distinguish between expected and unexpected events;
4. I.Control generates a framework that, in conjunction with the realistic models, advanced planning, and sensing components, allows for robust execution of dynamic manipulation tasks.

This integrated paradigm, I.AM, brings robots to an unprecedented level of manipulation abilities. By incorporating this new technology in existing robots, I.AM enables shorter cycle time (10%) for applications requiring dynamic manipulation in logistics. I.AM will speed up the take-up and deployment in this domain by validating its progress in three realistic scenarios: a bin-to-belt application demonstrating object tossing, a bin-to-bin application object fast boxing, and a case depalletizing scenario demonstrating object grabbing.

1.2 I.Model background

I.Model is the name for the impact prediction framework, that comprises parameter identification methods to tune parametric impact laws from experimental data, enabling to predict the dynamic effect of collisions (relative contact speed higher than 0.1m/s, prediction accuracy in post-impact speed of 10%). The open distribution of experimental data is also part of this objective, to promote the creation of standard practice to model and estimate the effect of physical impacts on robots.

The impact models will be validated on the basis of their ability to estimate and predict relevant features such as post-impact velocities and post-impact modes (contact establishment/detachment, stick/slip...), and impact propagation in articulated structures. I.Model is then used for learning, sensing, and control of dynamic manipulation tasks, that are related to the remaining three objectives.

1.3 Purpose of the deliverable

This deliverable D1.3 I.Model Report, aims at providing an overview of papers published related to the I.Model framework. It provides a summary of the main results and proposed algorithms of these papers, and summarizes the main achievements, limitations, and future directions related to I.Model (T1.1, T1.2, T1.3, T1.4, and T2.1).

1.4 Intended audience

The dissemination level of D1.3 is 'public' (PU) – meant for members of the Consortium (including Commission Services) and the general public. This document is intended to serve as an internal communication for the entire I.AM. consortium on the developments and publications related to I.Model.

2 SUMMARY AND RESULTS OF PUBLICATIONS RELATED TO I.MODEL

In this chapter, we will detail for each task associated to I.Model (T1.1, T1.2, T1.3, T1.4, and T2.1) the publications related to these tasks. Each of the following subsections contains a description of the task, an overview of the publications related to that task and a summary of the achievements, limitations, and future directions of this task.

2.1 Task T1.1 - Specifications for the impact motion database, including Taxonomy of Contact Transitions

In task T1.1, the focus lies on specifying a database to share open-access high-quality data related to robot-object-environment impact data in controlled environments. Datasets stored in the database will, in particular, include joint encoder and torque sensors data synchronized with high speed camera data and motion capture data recorded at high frequencies (360 FPS), including detailed information about the objects and robots employed (e.g., Franka Emika Panda arm) and detailed descriptions of the motion capture market set. Task 1.1 focuses on specifying what data is to be collected, how the data is stored (dataset structure), but more importantly, how the data is classified. This classification is done via a taxonomy of dynamic contact transitions. The collected data will be used to validate parametric impact models in T1.4 and to learn sensory-based impact models in T2.3. The metrics for evaluation of this task involve the creation of the data collection and data storage structure, as well as the possibility to upload and download datasets from the database via a web interface. Furthermore, detailed descriptions of classification (taxonomy) of data will be provided in each dataset, together with metadata files that allow for human and machine interpretability of the data.

2.1.1 Overview of publications related to the task

Details about the specification of the impact motion database are given in Deliverable D1.1 [1] and Deliverable D1.4 [2]. Although various publications use data that is stored in the database, the work presented in [3] is the single publication where the database is introduced and detailed, including the taxonomy of contact transitions. We will detail the main results results in the paragraph below.

2.1.1.1 Publication: "The Impact-Aware Robotics Database: Supporting Research Targeting the Exploitation of Dynamic Contact Transitions"

This publication should be seen as the main reference point to the database. In [3], we observe that, although some datasets related to impact-aware robotics do exist, our general observation is that most datasets are tailored to support a specific publication, with data often stored without considering its potential reuse for different purposes and by different users. In some cases, datasets also lack the support of metadata for human or machine interpretability, making it difficult to understand and use the data. Furthermore, we observe that some datasets have become unavailable unless an institution was chosen for data storage whose mission is to preserve the data for several years in the future. This is an important conclusion and motivation for the development of the Impact-Aware Robotics Database. In order to support data collection and data storage, all software has been made publicly available via GitLab[1]. When data is collected, it is important to store the data in a structured way to ensure compliance with

[1] See `https://gitlab.tue.nl/robotics-lab-public/impact-aware-robotics-database`

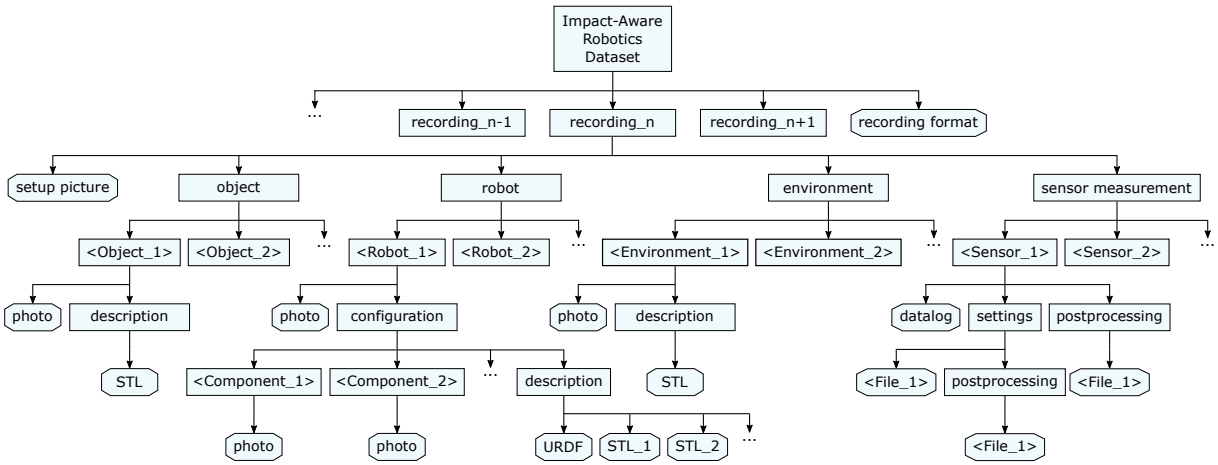

Figure 1: Structure of a dataset. Each dataset contains a set of recordings and a description of the recording format. Each recording is subdivided into object, robot, environment, and sensor measurement subfolders and contains a picture of the setup used during for the experiments. Image taken from [3].

FAIR principles. This structure is specified in [3] and for ease of readability also shown in Figure 1.

A web-interface has been developed [4] that provides a platform to search for and download datasets from the database. On this web interface, users can search for datasets via specific keywords such as keywords related to the objects, robots, environments, and sensors used in the dataset, as well as the contact transition taxonomy and recording purpose. The web interface contains all metadata of all objects, robots, and environments used in the datasets that are published in the database. Dedicated pages are created for objects, robots, and environments. On each of these pages, users can search for specific items (e.g., Box005 for object, or Conveyor002 for environment), view the detailed item information, and view all datasets where this item was used after which they can be downloaded.

2.1.2 Summary achievements, limitations, and future directions

With the work presented in [3], the tasks specified in T1.1 have been completed. However, the current version of the database only contains a limited amount of datasets (24 in total) recorded at institutions within the consortium. In the process of creating these datasets, we have continuously been updating the software tools for data collection and storage to be more modular and robust to handle new robots, objects, and environments and we have encountered situations were our software tools could not process data coming from new sensors. At the time of writing, we have not yet been contacted by institutions outside the consortium that want to upload datasets to the database. Although the current version of the database allows for this, it is hard to foresee possible shortcomings related to this. As future work, we plan to expand the database with new datasets, and we will continue the development of the web interface and the software used for data collection and storage. We are open to contributions of datasets and future improvements and developments.

2.2 Task T1.2 - Data Collection of Robot-Object-Environment Contact Transitions for Robotic Manipulation

In task T1.2, all partners collect datasets from experiments where robots are involved in manipulation tasks that involve dynamic contact transitions such as impacts. The impact motion data will be stored

in accordance with the policy and format defined in T1.1. The metric for evaluation of this task is the successful storage of impact data in the Impact-Aware Robotics Database, in particular for the tossing, boxing and grabbing validation scenarios as described in workpackage 5.

At the moment of writing this deliverable, the database contains 24 datasets with a total storage of 37GB. Most of these datasets are collected as part of the I.AM. project, coming from various partners within the consortium. Detailed information of the collected datasets can be found on the web-interface [4] and Deliverable D1.4 [2]. In particular, we highlight the datasets related to the scenarios of TOSS [5], GRAB [6], and BOX [7].

2.2.1 Overview of publications related to the task

Although no specific publication is connected to the pure task of collecting data, various publications rely on datasets that are created to support these publication. In the following paragraphs, we will highlight for various publications the associated datasets that are stored in the database.

2.2.1.1 Publication: "Learning Suction Cup Dynamics from Motion Capture: Accurate Prediction of an Object's Vertical Motion during Release"

This publication [8] is mainly related to T2.1, and we will describe its content in Section 2.6. However, we want to stress here that in this publication experiments are performed where objects are released from a suction gripper to learn the release dynamics of an object via motion capture. These experiments are all stored in a dataset [9] that is part of the Impact-Aware Robotics Database.

2.2.1.2 Publication: "Predicting the Post-Impact Velocity of a Robotic Arm via Rigid Multibody Models: an Experimental Study"

This publication [10] is mainly related to T1.4 and we will describe its content in Section 2.4. In this publication, experiments are performed with a KUKA LWR IV+ robotic arm that impacts a surface, with the purpose of validating the prediction of the post-impact velocity. This publication has an associate dataset [11]. However, this dataset is not part of the Impact-Aware Robotics database, as the data was collected before the structure of the database in tasks T1.1 was properly defined. Therefore, this dataset is not conform the database structure of T1.1.

2.2.1.3 Publication: "Experimental Validation of Nonsmooth Dynamics Simulations for Robotic Tossing involving Friction and Impacts"

This publication [12] is, at the time of writing this document, submitted to Transactions on Robotics, and focuses on the validation of nonsmooth dynamics simulations in the context of robotic tossing applications. The content is mainly related to Task T1.4 and details on this publications will be given in Section 2.4. In this publication, many experiments are performed with the purpose of parameter identification and validation of nonsmooth dynamics simulations. A total of 14 datasets are published related to this publication, which for ease of accessibility are grouped in a collection supporting the publication [13].

2.2.2 Summary achievements, limitations, and future directions

With a total of 24 datasets created by various partners of the consortium, the task of T1.2 to collect data of robot-object-environment contact transitions has been completed. The datasets are successfully stored on the Impact-Aware Robotics Database and can be viewed and downloaded via a web-interface [4]. We will continue the collection of impact-aware robotics data in the remaining part of the I.AM. project, and encourage the continuation of data collection and data storage beyond the I.AM. project, with the scope of making the database grow. The collaboration with 4TU.ResearchData [14] ensures the datasets will be stored and accessible for at least 15 years.

2.3 Task T1.3 - Physics Engine Interface and Impact Laws implementation for learning, planning, sensing, and control

Task T1.3 focuses on the development of an application suited for analyzing the I.AM. dataset and perform simulations to calibrate various learning, sensing, and control algorithms. The goal is to improve and refine the impact and contact models as well as the frictional multi-impact AGX solver focusing for a solution for real-time robot control. Details on the Algoryx Physics Engine are documented in Deliverable D1.2 Physics Engine API. Additionally, one publication that focuses on the correlation between measured and simulated pre- and post-impact data is discussed in Section 2.4 and 2.6 in this deliverable.

2.4 Task T1.4 - Validation and Identification of model-based Impact Laws

Task T1.4 focuses on reviewing the state-of-the-art on algebraic impact laws and parameter identification in the nonsmooth dynamics literature, with the focus on inelastic impacts with friction. It aims to develop parameter estimation routines for model-based impact laws (e.g, impact and friction coefficients) from experimental data.

2.4.1 Overview of publications related to the task

In the following paragraphs, we will summarize the most important findings of several publications that are linked to T1.4.

2.4.1.1 Publication: "Experimental Validation of Nonsmooth Dynamics Simulations for Robotic Tossing Involving Friction and Impacts"

An important objective of I.Model is to develop parameter identification techniques to estimate the contact parameters (friction, restitution) from recorded impact data. In this work [12], we evaluate the prediction performance of two nonsmooth rigid-body dynamics simulators on real-world data with spatial impacts in the context of robotic tossing. We perform a parameter identification procedure to find the coefficient of friction and restitution of different objects via a velocity-based and trajectory-based cost function. In both cases, we establish a cost function based on the difference between the experimental data and the prediction using a certain parameter set μ, e_N. The parameter values minimizing this cost function therefore result in a prediction that is closest to the experimental data, which we consider ground-truth. Therefore, we define the optimization problem as

$$(\mu^*, e_N^*) = \arg\min_{\mu, e_N} \frac{1}{M} \sum_{i=1}^{M} L(i; \mu, e_N), \tag{1}$$

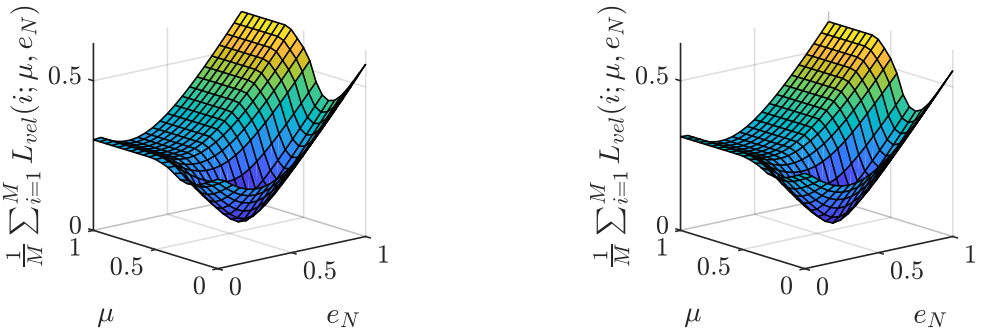

(a) (b)

Figure 2: Computed costs for the velocity-based metric as result of simulations with different parameters. Results of Algoryx Dynamics (a) and MATLAB (b) for impact experiments of Box006.

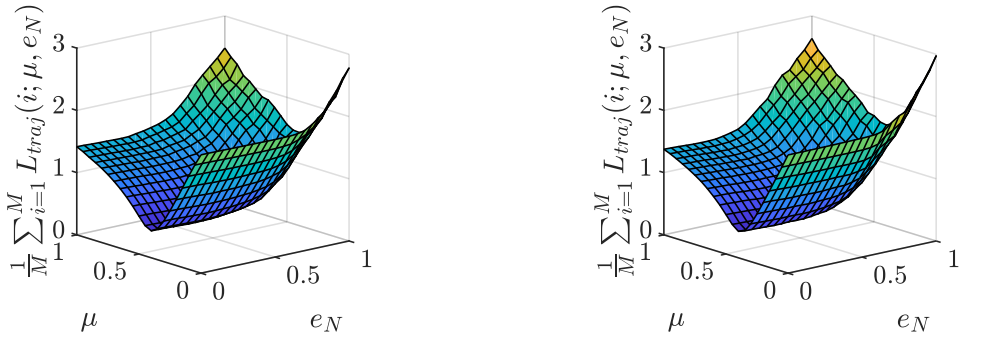

(a) (b)

Figure 3: Computed costs for the trajectory-based metric as result of simulations with different parameters. Results of Algoryx Dynamics (a) and MATLAB (b) for impact experiments of Box006.

as in [15], subject to $0 \leq \mu$ and $0 \leq e_N \leq 1$, where $L(i;\mu,e_N)$ is the loss function of experiment i depending on the parameter set $\{\mu, e_N\}$, and M is the size of the dataset. Hence, M is the total number of impact events for the velocity-based metric and the total number of trajectories for the trajectory-based metric, which my vary per dataset.

In the velocity-based approach, we consider the pre-impact state of the object as obtained from measurement data, denoted as $\mathbf{x}^- = \{\mathbf{H}^-, \mathbf{v}^-\}$. A trajectory of the box is then simulated with a varying set of contact parameters $\{\mu, e_N\}$ from the initial condition $\mathbf{x}^-$. The loss function computes the error between the post-impact velocity as a result of the simulations (denoted as $\tilde{\mathbf{v}}^+(\mu, e_N)$) with the post-impact velocity as obtained from experimental data (denoted as $\mathbf{v}^+$) for that specific measurement i. Hence, we write the velocity-based loss function as

$$L_{vel}(i;\mu,e_N) = \left\| \mathbf{W}\left(\mathbf{v}^+ - \tilde{\mathbf{v}}^+(e_N,\mu)\right) \right\|_2 \tag{2}$$

In the trajectory-based approach, we initialize the dynamical model with the state at release $\mathbf{x}(k_{rel})$ and simulate until the box has reached its rest-pose at $\mathbf{x}(k_{rest})$. The moment of release k_{rel} is defined at the

moment when the box is released from the robotic arm and is in free flight and the moment of rest k_{rest} is when the relative velocity between the box and the conveyor is zero. The loss function is based on the approach of [15], where for the varying set of parameter values $\{\mu, e_N\}$, it quantifies the difference between the simulated trajectory $\tilde{\mathbf{H}}(k)$ and the measured trajectory $\mathbf{H}(k)$ at each discrete time index $k_{rel} < k < k_{rest}$. As the position $\mathbf{o}$ and orientation $\mathbf{R}$ can be obtained from the transformation matrix $\mathbf{H}$, we define the trajectory-based loss function as

$$L_{traj}(\mathbf{x}(k_{rel} : k_{rest})_i, \mu, e_N) = \frac{1}{N} \sum_{k=k_{rel}}^{k_{rest}} \left(\frac{1}{l}\|\mathbf{o}_i(k) - \tilde{\mathbf{o}}_i(k)\|_2 + \| \log\left(\mathbf{R}_i^{-1}(k)\tilde{\mathbf{R}}_i(k)\right)\|_2\right) \quad (3)$$

In Figure 2 and 3 the velocity based and trajectory based cost functions are shown for simulations in Matlab and Algoryx. Our results show that these two identification criteria lead to different parameter values, and these criteria should be chosen in consideration of the application at hand. We perform long-range tossing experiments where the predictability performance of the simulators is measured in terms of the rest-pose error. These errors show, considering the parameters obtained via the trajectory-based cost function, a position error in the order of 5-10 cm, and a orientation error in the order of 8-16 degrees over a tossing trajectory of about 1.2 meters, with boxes whose maximum dimension is about 20cm. These errors are, considering the application at hand, rather small, and suggest that nonsmooth dynamics models can indeed be used to predict the rest-pose of box-tosses in logistics with sufficient accuracy, at least for objects with constant inertia properties. Further details are provided in [12].

2.4.1.2 Publication: "Predicting the Post-Impact Velocity of a Robotic Arm via Rigid Multibody Models: an Experimental Study"

In this work, several impact experiments between a 7DOF torque controlled robotic arm and a sturdy table have been performed, at different velocities and impact angles, see also Figure 4. The main contribution is proposing a methodology that allows for a meaningful quantitative comparison between the recorded post-impact data, that exhibits a damped oscillatory response after the impact, and the post-impact velocity prediction derived via the readily available rigid-body robot model, that presents no oscillations and that is the one typically obtained via mainstream robot simulator software.
As known from the field of nonsmooth mechanics, the impact map is the result of the combination of an impact equation and an impact law. For a robot manipulator, the impact equation is derived from

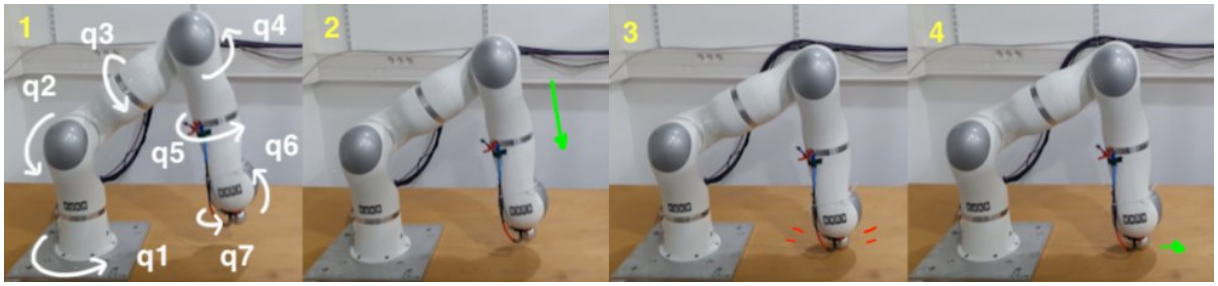

Figure 4: Example of a vertical impact test at 0.2 m/s conducted on a KUKA LWR IV+ robotic arm. The figure shows four snapshots from a recorded video. Time evolves from left to right and from top to bottom. Starting from a rest pose (1), also illustrating the robot 7 DoFs, the robot moves down (2). After impact (3), the end-effector slides on the table (4), keeping contact. Image taken from [10].

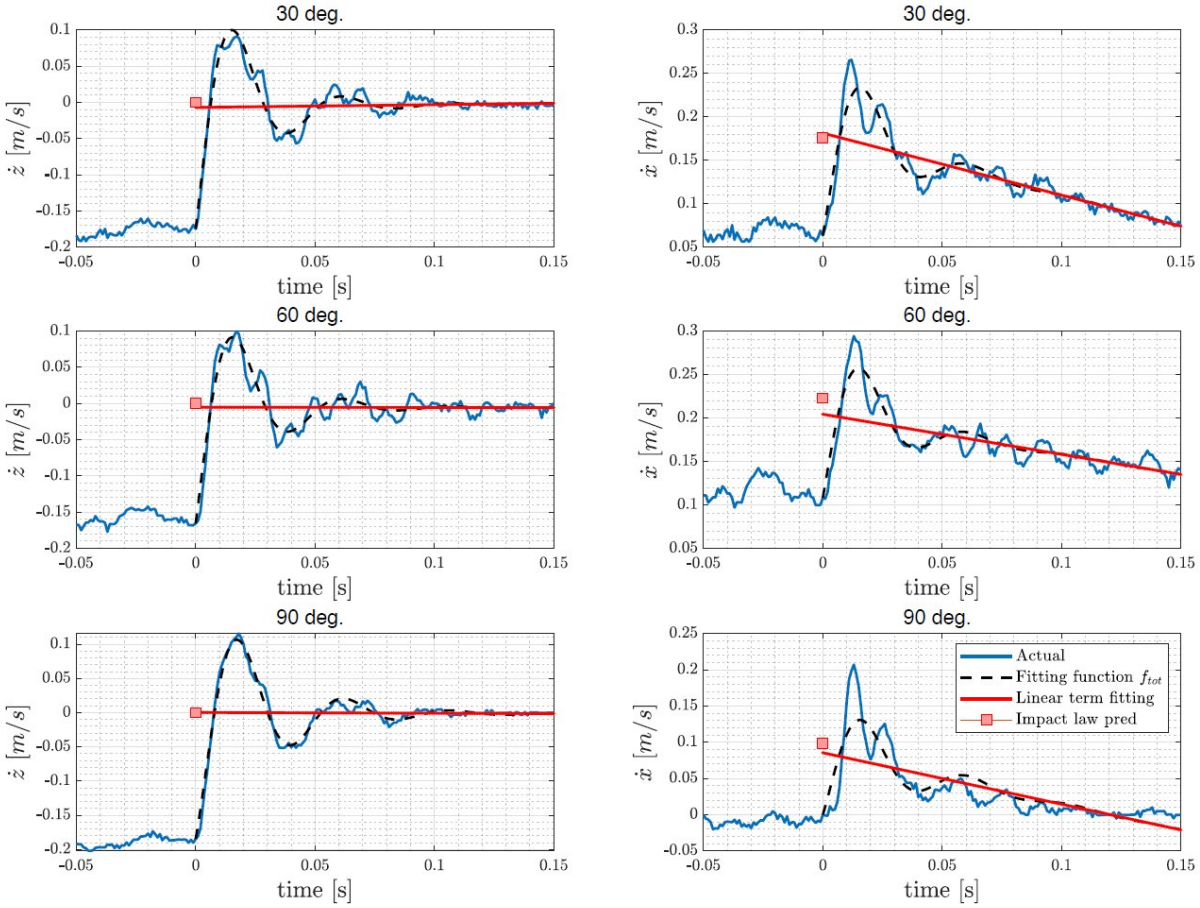

Figure 5: Cartesian velocity fitting results for impact at 0.2 m/s under angles ranging from 30 to 90deg w.r.t. to the horizontal table. The x-direction (right) is tangential to the table and the z-direction (left) is normal. Note how the fully rigid-model impact law prediction (red square) agrees with the prediction resulting from filtering out the post-impact transient (left end point of the red segment). Image taken from [10].

an identified rigid body model or readily derived from 3D CAD design. When performing impact experiments, however, clear oscillatory transients are visible: in our experiments, these last for about a hundred milliseconds, see also Figure 5.

The proposed post-impact velocity prediction procedure, based on the idea of removing post-impact oscillation components, shows to be in very good agreement with a fully-rigid robot impact model, see also Figure 5. This procedure, which is not tailored to a specific robot, can be used to assess if a fully rigid-body impact model can provide reliable post-impact velocity prediction that are of use in impact-aware robot planning, control, and perception. For the specific impact experiments performed with the KUKA LWR IV+, we obtained a 7.3% relative error and 8 mm/s absolute error difference between the measured and predicted sliding velocity, over a set of 18 experiments. This good accuracy is rather surprising giving the fact that we are comparing an ideal rigid-body impact model with real experiments that contain unmodeled post-impact (structural) vibrations, that are “filtered out” by the proposed procedure. Further investigation is required to assess what level of accuracy is required to achieve satisfying performance

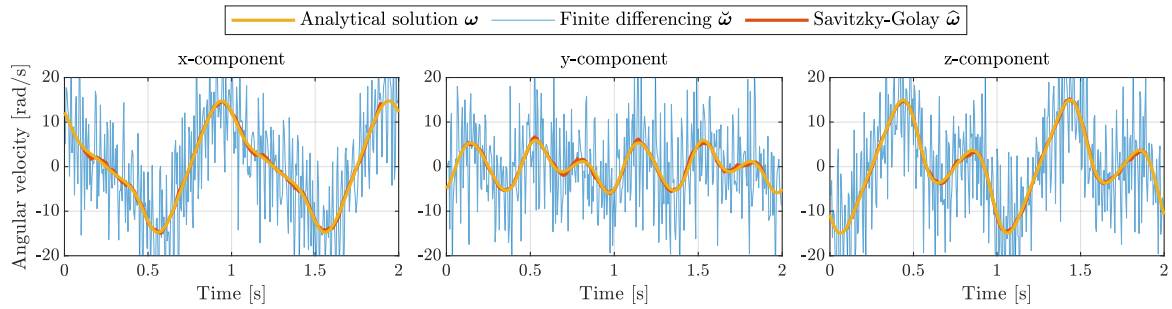

Figure 6: Simulation results showing the x-component (left), y-component (middle), and z-component (right) of the angular velocity as obtained by the analytical solution, finite differencing, and the Savitzky-Golay filter.

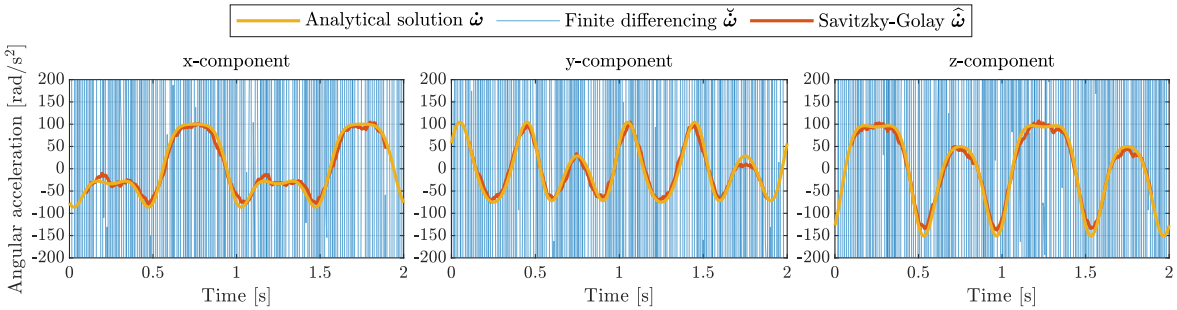

Figure 7: Simulation results showing the x-component (left), y-component (middle), and z-component (right) of the angular acceleration as obtained by the analytical solution, finite differencing, and the Savitzky-Golay filter.

in impact-aware manipulation for specific applications.

2.4.1.3 Publication: "Geometric Savitzky-Golay Filtering of Noisy Rotations on SO(3) with Simultaneous Angular Velocity and Acceleration Estimation"

Many validation studies make use of motion capture systems to obtain accurate pose measurements of various bodies at high frequency. However, these systems only provide information about position and orientation of the objects, they do not provide information of higher order derivatives such as velocities of accelerations of these rigid bodies. Naive approaches to obtain these velocities and accelerations from motion capture, such as finite differencing, drastically amplify zero-mean quantization or measurement noise to a level that makes the estimated velocity and in particular the acceleration useless. Therefore, some more advanced methods, typically referred to as data filtering (causal) or data smoothing (noncausal) are required.

This paper focuses on the problem of smoothing a rotation trajectory corrupted by noise, while simultaneously estimating its corresponding angular velocity and angular acceleration. To this end, we develop a geometric version of the Savitzky-Golay filter on SO(3) that avoids following the conventional practice of first converting the rotation trajectory into Euler-like angles, performing the filtering in this new set

of local coordinates, and finally converting the result back on SO(3).

In particular, the estimation of the angular acceleration requires the computation of the right-trivialized second covariant derivative of the exponential map on SO(3) with respect to the (+) Cartan-Schouten connection. We provide an explicit expression for this derivative, creating a link to seemingly unrelated existing results concerning the first derivative of the exponential map on SE(3). A numerical example is provided in which we demonstrate the effectiveness and straightforward applicability of the proposed approach. The angular velocity and angular acceleration as a result of the simulations are shown in Figures 6 and 7, respectively. Considering Figures 6 and 7, we clearly see, as expected, how finite differencing leads to poor estimations of the angular velocity and angular acceleration profiles. The noise on the orientation data is amplified to a level that makes the data useless for post-processing processes. Instead, our approach using the Savitzky-Golay filter clearly results in an accurate estimation of the velocity and acceleration profiles, closely following the analytical solution. An implementation of the developed geometric Savitzky-Golay filter is made available with this work[2] to promote its use and validation by the research community.
Future investigations will include: A performance comparison with other filtering approaches (e.g., Savitzky-Golay filtering using Euler-like angles), derivation of the filter for unit quaternions and left multiplicative noise (in this paper, we consider right multiplicative noise), investigation of filter parameters tuning depending on noise level, and application of the filter on experimental noisy rotation signals.

2.4.1.4 Publication: "Model Based 6D Visual Object Tracking with Impact Collision Models"

Estimating the position and orientation of objects in a 3D space from RGB images is an important problem in robotic applications. In robotic manipulation, for example, accurate pose estimation is advantageous for executing advanced grasping tasks. In this work, we focus on tracking rigid objects that collide with the environment. This can be of particular interest in real-world applications such as robot batting [16] or robot soccer [17]. For sake of brevity, we focus specifically on box-shaped objects as this is also relevant for logistics applications, e.g., for robot tossing. Our approach, however, can be extended to objects with other geometries.

To the best of the authors knowledge, there is no specific work in literature focusing on 6D object tracking in the presence of collisions. A substantial body of work has been recently published on 6D object pose estimation, which can be used as a measurement in tracking algorithms. Often, a Particle Filter (PF) is employed, as this filter maintains several hypotheses over time, which allows to approximate the full posterior distribution of systems with non-linear dynamics and non-Gaussian state distributions [18], [19], as is the case for our specific case of tossing a box impacting a surface. However, particle filters do not take the current observation into account to create a proposal distribution, and classical tracking algorithms based on particle filters are based on constant velocity models, that tent to lose track of the objects upon impact.

In our previous work [20] we have shown that, in a simulation environment, by using the Unscented Particle Filter (UPF) [21], which does take the current observation into account to create a proposal distribution, and by incorporating impact and friction on a nonsmooth dynamical motion model in the filter, the objects can be tracked significantly more accurate compared to classical methods based on particle

[2]The source code is publicly available at `https://gitlab.tue.nl/robotics-lab-public/savitzky-golay-filtering-on-so3`

filters using constant velocity models.

In this work, we validate our approach [20] on real-world data. We perform experiments on a setup where the objects and environments can be tracked with sub-millimeter accuracy using a motion capture system. The observations are taken from a single RGB camera and we make use of the known 3D model of the object and the objects' color characteristics to predict its appearance in the 2D image. We show that our approach leads to accurate tracking performance of the box in the presence of collisions and motion blur.

2.4.2 Overview of ongoing work

In this subsection, we are going to summarize the outcomes of ongoing projects, each of which may result in a future publication or a paper is currently in preparation.

2.4.2.1 In preparation: "On the Modeling and Identification of Active Bellow Suction-cups"

This work addresses the topic of modeling and stiffness parameter identification of active vacuum bellow suction-cups. An in depth analysis of the constructive symmetries of the suction-cup allows us to simplify the structure of the stiffness matrix with consequent reduction of the collected data for the identification procedure. The proposed spring model is validated through numerical simulations to assess its physical correctness by monitoring the total energy of the system. The predicted rest pose in static conditions for the identified parameters is compared with data from a real setup.

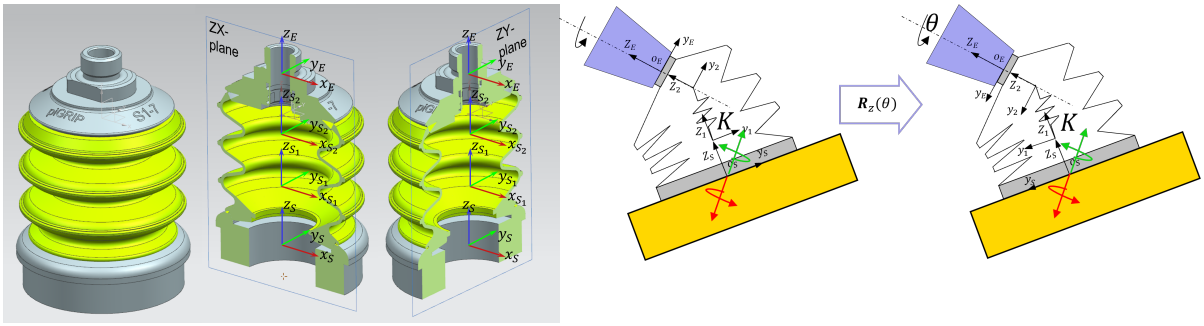

(a) (b)

Figure 8: (a) CAD model of the Suction-cup with its cross-sections on the ZX- and ZY-planes to highlight its axial symmetry. (b) The configuration dependent spring wrench in a certain pose of the tool-arm (left) and after a rotation of θ about the symmetry axis (right).

One of the main contributions of this work is the systematic analysis of the symmetry properties of the suction-cup that allows to find a structure in the stiffness matrix of the proposed model, significantly reducing the number of parameters to identify (from 21 to just 5). We proved that the equivariance property of the model for rotations about the axis of symmetry holds for the stiffness matrix structure found and we assessed this property on real data collected at the Vanderlande's Innovation Lab at the TU/e campus. This findings not only allows to reduce the number of parameters to identify, but also provides guidelines about the physical experiments to be conducted in order to identify such reduced

set of parameters. This translates into a simpler experimental procedure that do not requires to gather data spanning the whole configuration space but just a very reduced one.

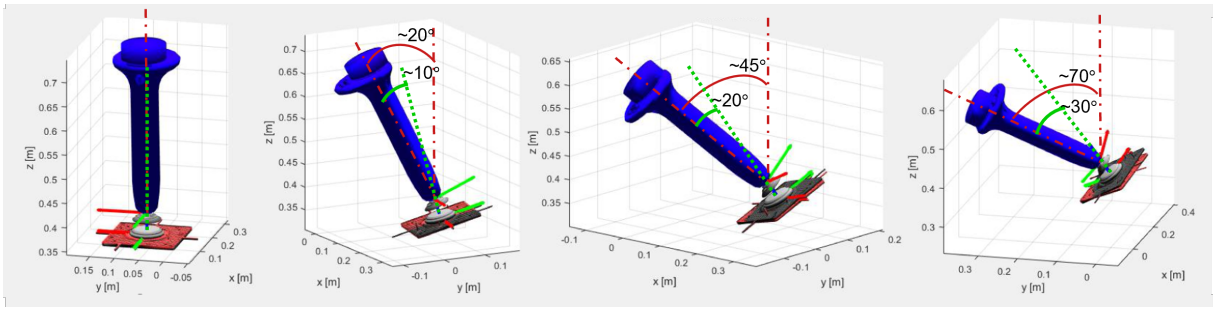

Figure 9: Predicted rest pose of the package (black plate) using the proposed model and the identified reduced set of parameters compared to measured data (red plate). From left to right, the relative inclination angle of the lip with respect to the attachment point of the suction-cup and the inclination of the Tool-Arm increases according to Table 1.

The results obtained for a payload of 1747 grams and CoM location $c = [0,\ 0,\ -0.0805]^\top [m]$, the estimated rest pose error of the suction-cup lip with respect to the measured one, for each of the validation experiments in Fig. 9, is reported in Table 1.

Quantity	Exp 1	Exp 2	Exp 3	Exp 4	dimension
Tool-Arm inclination	0	~ 20	~ 45	~ 70	[deg]
relative inclination	0	~ 10	~ 20	~ 30	[deg]
position error	1.8	1.7	4	6.5	[mm]
orientation error	TBD	TBD	TBD	TBD	[deg]

Table 1: Some of the results obtained for a payload of 1747 grams in the validation set.

2.4.2.2 In preparation: "Refined Post-Impact Velocity Prediction for Torque-Controlled Flexible-Joint Robots"

This work [22] is a continuation of the currently published paper [10] and focuses on the post-impact prediction for torque controlled flexible joint robots. Predicting the post-impact velocity for torque-controlled flexible-joint robots allows the implementation of impact-aware control schemes which exploit intentional collisions for robotic manipulation and locomotion. Starting from the existing rigid-robot post-impact velocity prediction approach [10], this paper shows how an improvement in the prediction quality can be obtained by taking into account impact surface friction, joint motor inertias, gear ratios, and low-level joint torque control gains. Compared to the previous rigid-robot approach, the paper also proposes a more robust method to estimate the low-frequency post-impact velocity trajectory from experimental data via a polynomial fit, to remove unmodeled and secondary vibratory effects. The approach is illustrated by means of simulation and validated on 50 experimental trials on a commercially available torque-controlled robot. The recorded impact data and prediction algorithms are shared openly for reproducibility [23].

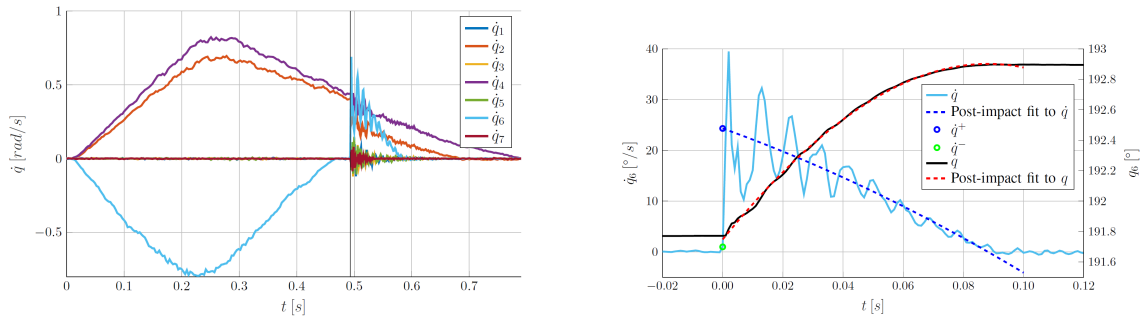

(a) (b)

Figure 10: Joint velocities during an impact experiment. Full experiment with the vertical line denotes the impact time (a) and close-up of joint 6 velocity (b).

A set of 50 experiments was recorded where a 7 DoF Franka Emika Panda robot with a spherical aluminum endeffector impacted a horizontal plate made of steel. During the recording process 10 different initial robot configurations along with the respective actuation joint torques required to impact the plate were set. For each of the 10 experiment types, 5 repetitions were run in order to prove the repeatability of the experiment type and little sensitivity of the prediction for very similar responses. These 10 groups of experiments can be subdivided in three main types depending on motion direction, each in turn subdivided according to different end effector orientations.

As in [10], we are mainly interested in predicting the gross velocity jump after an impact, discarding the impact-induced oscillations. Therefore, we need a procedure to filter out these oscillations. Differently from what is done in [10], a polynomial least squares fit is obtained from the post-impact joint position time-based signal, and the gross velocity at impact time is estimated by evaluating the derivative of the polynomial (in [10], the nonlinear fitting is performed on the velocity signal). The fitting is applied to samples over a span of 100 ms, starting at the impact time. We have found that a third order polynomial is sufficient in capturing the evolution of joint position trajectories. The interval length was selected as a result of the qualitative analysis of the velocity signals throughout the 50 recorded experiments, that showed that 100 ms corresponds in all cases to the approximate duration of the oscillations present in the post-impact response of the robot (cf. oscillations in Figure 10a). This approach turned out to be simpler and more robust than the nonlinear least-square approach proposed in [10], that can get stuck in local minima, providing essentially the same result when both methods converge.
The polynomial fit applied to the oscillatory post-impact response has good results (errors less than 10%), as illustrated by the representative post-impact response of joint 6 in Figure 10b. This fit is used for assessment of the prediction accuracy as explained later in the paper in detail, since the oscillations provoked by the impact are disregarded by the predictor. However, note that the observed impact time in Figure 10b is not completely accurate, as the ante-impact velocity is initially zero and there is a significant growth to about 1 deg/s when the impact is detected. In regard to the impact time detection, this study also shows explicitly that different joints exhibit a gradual lag in the time of impact, with the first joint at the base showing signs of impact at a later instance, when compared to more distal joints (e.g., joint 6).

2.4.2.3 Ongoing: "Analyzing the sim-to-real gap for dynamic bin packing operations involving sliding frictional contacts"

The goal of this study is the validation of the standard nonsmooth contact and impact models, implemented within the I.AM. project by the physics engine of Algoryx dynamics, against experiments in the real world in the case of the BOX scenario. There are multiple aspects that can be validated within the AGX Dynamics simulator such as, e.g., the suction-cup gripper and the contact friction model. These aspects are yet little explored in the multi-body and robotics simulation literature and the validation of the friction and suction-cup model are the main objectives of this work.

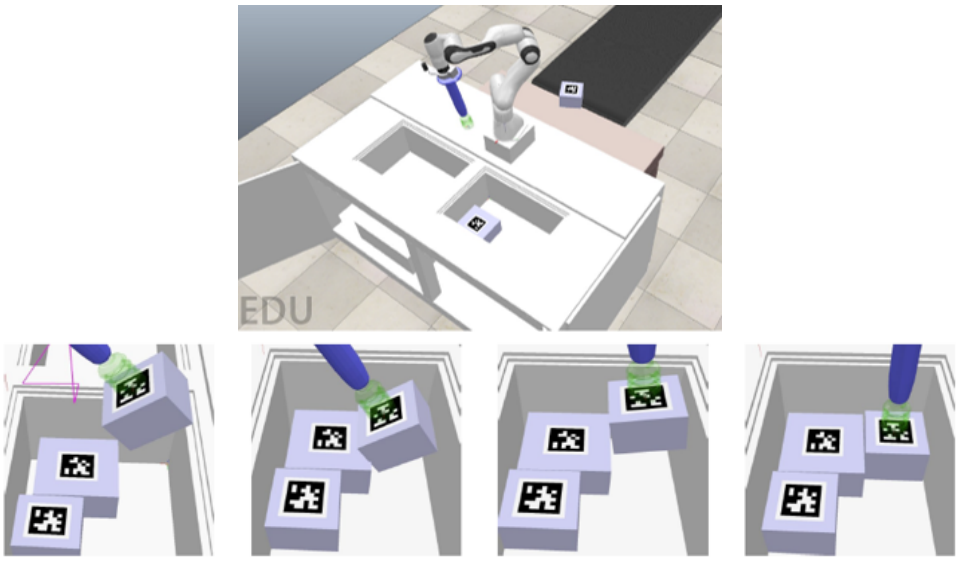

Figure 11: Physics simulation of a bin packing example that make use of the proposed motion primitives to stuff boxes in a tote and maximize the fill rate. On top, a screenshot of the CoppeliaSim scene of a replica of the panda setup at Vanderlande's Innovation Lab at the TU/e. On the bottom, a sequence of images showcasing the rotate-push motion.

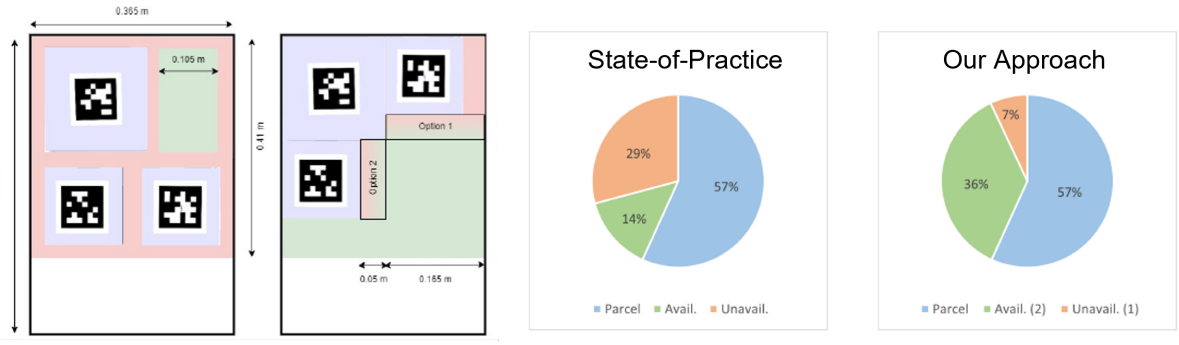

(a) (b) (c) (d)

Figure 12: Results of the bin packing simulation. The figure shows the final configuration of the boxes in the tote when using State-of-Practice (SoP) algorithms (a) and our approach (b) and the parcel occupancy area (blue), available area (green) and unavailable area (orange) for both SoP (c) and our (d) methods.

The successful validation of the models in this scenario allows the development of dynamic trajectory planners that autonomously generate advanced motion primitives, that we proposed in another ongoing project, and would be possible to test and validate them on such physics engine. The proposed motion primitives were tested on a research-grade physics simulator (CoppeliaSim with Newton physics engine) and a screenshot of the simulation environment is reported in Fig. 11 together with a sequence showing the rotate-push motion for stuffing a box in a tight space.
A quantitative analysis of the results (limited for now to the single experiment here reported) have shown and increment of the fill rate of the tote of about 32% compared to the state-of-practice approach, as illustrated in Fig. 12.

2.4.2.4 Ongoing: "Validation of nonsmooth simulation framework in dynamic manipulation of objects through pushing and rapid grasping."

Current ongoing work in the context of the GRAB scenario concerns the validation of the AGX Dynamics simulation framework in the context of impact tasks between robots and objects. In particular, we are focusing on rapid pushing motions, where contact between a silicon flat end effector of a robot and an object is established at nonzero velocity, as well as rapid grasping motions, where two robots with the same silicon end effectors simultaneously grasp objects with nonzero velocity. Both motions are highly relevant for the depaletization considered in the GRAB scenario. Having a reliable simulator is central in other objectives of the I.AM. project, such as impact classification considered in WP3 and impact-aware control considered in WP4. Hence, the objective of this work is to validate the simulation framework for a range of impact motions and objects and perform identification for the relevant simulation parameters. Preliminary results comparing a box push with 0.4 m/s between real world experiments and simulations with AGX Dynamics using the GLUE framework are shown in Figure 13. In both cases, the motion is

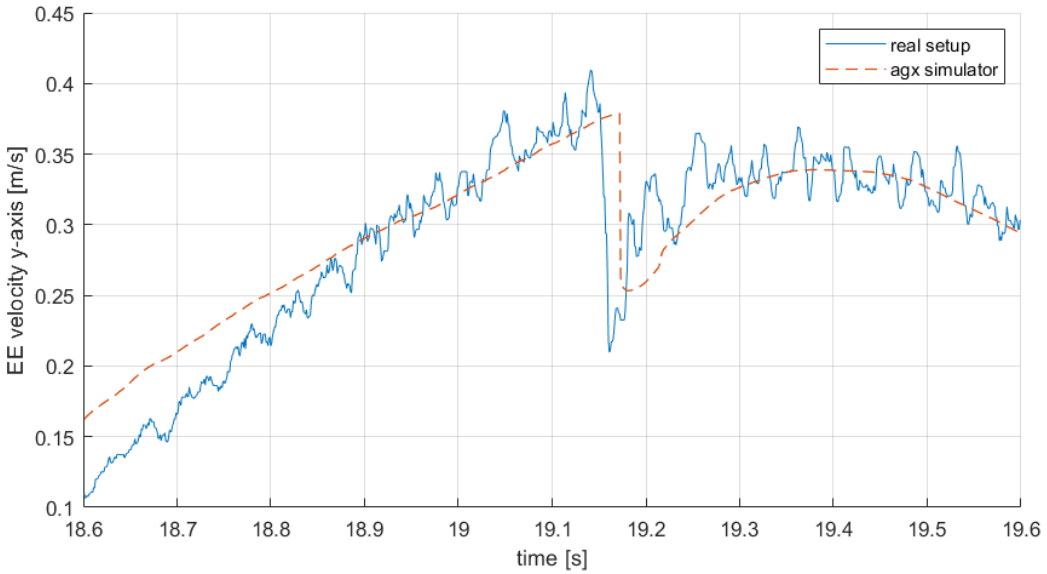

Figure 13: Cartesian velocity in normal impact direction around the time of impact for a task pushing a box with the Franka Emika robot, compared between simulation with AGX Dynamics and real experiments.

performed through the mc_rtc control framework using the same impedance-based QP controller with an identical reference. The rigid part of the impact, indicated by the velocity jump around 10.1/10.2 seconds, appears to be captured well by the simulations, as a similar jump with additional oscillations as a result of the unmodeled flexibility in the robot is observed in the real data.

2.4.3 Summary achievements, limitations, and future directions

Developing parameter estimation routines from experimental data and predicting post-impact behavior of a robot arm impacting the surface are the two main sub-tasks of T1.4. With two related publications (the work in [12] and [10]) and the work as discussed in Section 2.4.2.2 in preparation, we can conclude that the goals of Task T1.4 have been achieved. These works, in particular [12] and Section 2.4.2.2, provide clear expected accuracy metrics in terms of predicting the post-impact state (few centimeter error in terms of rest position and about 10 degrees error in the orientation for the rest pose) and post-impact velocities as a function of the pre-impact posture and speed (error smaller than 10%). See also Section 2.6.

As always in research, further modeling efforts are ongoing to improve and expand the results. For example, the work presented in Section 2.4.2.1 discusses the modeling efforts on the suction cup, that will improve robot motion planning, and will allow for payload estimation via suction cups. The suction-cup model discussed in Section 2.4.2.1 focuses on (quasi-)static conditions (model stiffness) and analyses the structural properties of symmetric suction-cups. The treatment of the dynamic case, which also involves the modeling and identification of the damping wrench, will be treated once the static part is completed. These models, which require the work presented in [24] for validation, will be further developed and related publications submitted in the remaining period of the project.

Despite having shown that the suggested suction-cup model properly predicts the rest pose of the lip for significant deformations ($> 20[deg]$), it only describes the suction-cup's behavior locally. As a result, it is evident that the pose prediction error grows with increasing deformation. A possible solution we are currently exploring, which is however beyond the scope of this task and the I.AM. project, is to sensorize the suction-cup. This solution will allows to directly measure the suction-cup's deformation and use the idenfitied model to predict interactioin forces and torques. This lead to develop a new technology that opens up multiple research possibilities e.g. the development of 6D data-driven models for both holding and release phases or the possibility to use the suction-cup as a virtual force/torque sensor, just to name a few.

In the ongoing work of Section 2.4.2.3, we are building the whole simulation scenario and updating the physical setup (a 6D force/torque sensor was added to the robot's flange) to systematically validate the contact, friction and suction-cup models implemented in AGX Dynamics. Then, we expect to narrow the gap between the real setup and the simulation by optimizing the model parameters using the collected data on the real setup. At this point we would be able to perform more realistic simulations for our motion primitives within AGX Dynamics. This constitutes an advantage with respect to the simulations performed in CoppeliaSim, since they were performed using a rigid suction-cup. Embedding the developed motion primitives into contact motion planners that allows to specify non-prehensile motions will ultimately allow to autonomously plan and execute complex bin packing scenarios in a realistic simulation before transferring it to the real setup.

In the ongoing work of Section 2.4.2.4, the continuation includes performing validation for rapid grasping motions as well as continue validation of different pushing motions like the experiments shown in Figure 13. On top of this, aside from visual comparison, these results are to be formalized for extensive quantitative comparison. Furthermore, a parameter identification is to be performed to match the

real scenario as well as possible to the simulation framework for a range of grasp/pushing motions with different objects.

2.5 Task T1.5 - Modeling Benchmarks and Progress Definition and Evaluation

Task T1.5 focuses on the benchmarks and performance evaluation of tasks T1.1, T1.2, T1.3, and T1.4. The benchmarks for T1.1 and T1.2 are related to the Impact-Aware Robotics database, and can be considered as the successful launch of the database. The benchmarks for T1.3 and T1.4 are more specifically related to the publications linked to these tasks, and are specified in Section 2.3 and 2.4.

2.6 Task T2.1 - Learning uncertainty models at impact

In Task T2.1, the focus lies on learning uncertainty models at impact and release. While some of the physics of impact can be modeled explicitly (under various hypotheses), other aspects such as deformations incurred at impact, fast transient dynamics, friction of two materials against each other, restitution coefficients, or release dynamics from a suction cup cannot be modeled so easily. Learning models of these effects can be more practical. This task will enrich the models developed in I.Model using data generated in WP1. Furthermore, the idea is to see how a (normal) uncertainty distribution centered about the ideal pre-impact configuration and velocity, transforms under the nonsmooth impact models developed in T1.4. The metrics for evaluation associated to this task involve predicting distributions of object and robot's states after impact on training and testing datasets using learned models. In the following paragraphs, we will detail the publications related to T2.1.

2.6.1 Overview of publications related to the task

2.6.1.1 Publication: "Learning Suction Cup Dynamics from Motion Capture: Accurate Prediction of an Object's Vertical Motion during Release"

When objects are released from a suction gripper, pressurized air is blown into the suction cup (green part in Figure 14) to speed up the release of the object. During this phase, called the *release phase*, pressure is building up inside the suction cup, which makes it expand. The suction cup can be considered as a spring-damper system, whose characteristics drastically change under varying pressure. This suggests that creating a model that describes these dynamics will likely be difficult task.

In this publication [8], we take a data-driven approach where we study the interaction dynamics of an active vacuum suction gripper during the vertical release of an object. Object and suction cup motions are recorded using a motion capture system, see also Figure 14 where snapshots of a slow motion recorded video of a release experiment are shown. As the object's mass is known and can be changed for each experiment, a study of the object's motion can lead to an estimate of the interaction force generated by the suction gripper. We show that, by learning this interaction force, it is possible to accurately predict the object's vertical motion as a function of time. This result is the first step toward 3D motion prediction when releasing an object from a suction gripper.

To learn the force model, we have employed the Locally Weighted Projection Regression (LWPR) incremental learning approach [25] and employed the freely available implementation [26]. LWPR is a machine learning approach for nonlinear function approximation that learns incrementally in (potentially)

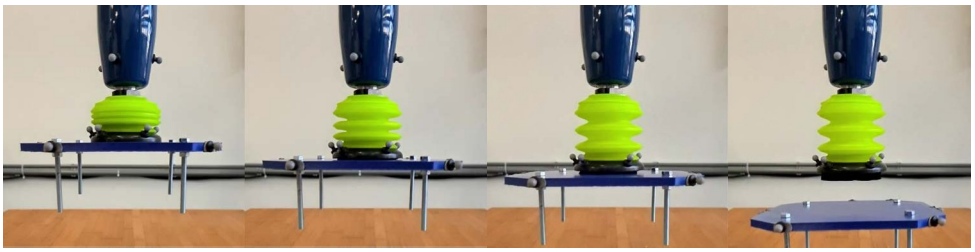

Figure 14: Snapshots of a slow motion recorded video during a typical vertical dropping experiment. The object is a plastic plate, that provides a sealant surface and four metal rods that are used to attach additional weights. The whole sequence takes less than 150 ms. Image taken from [8].

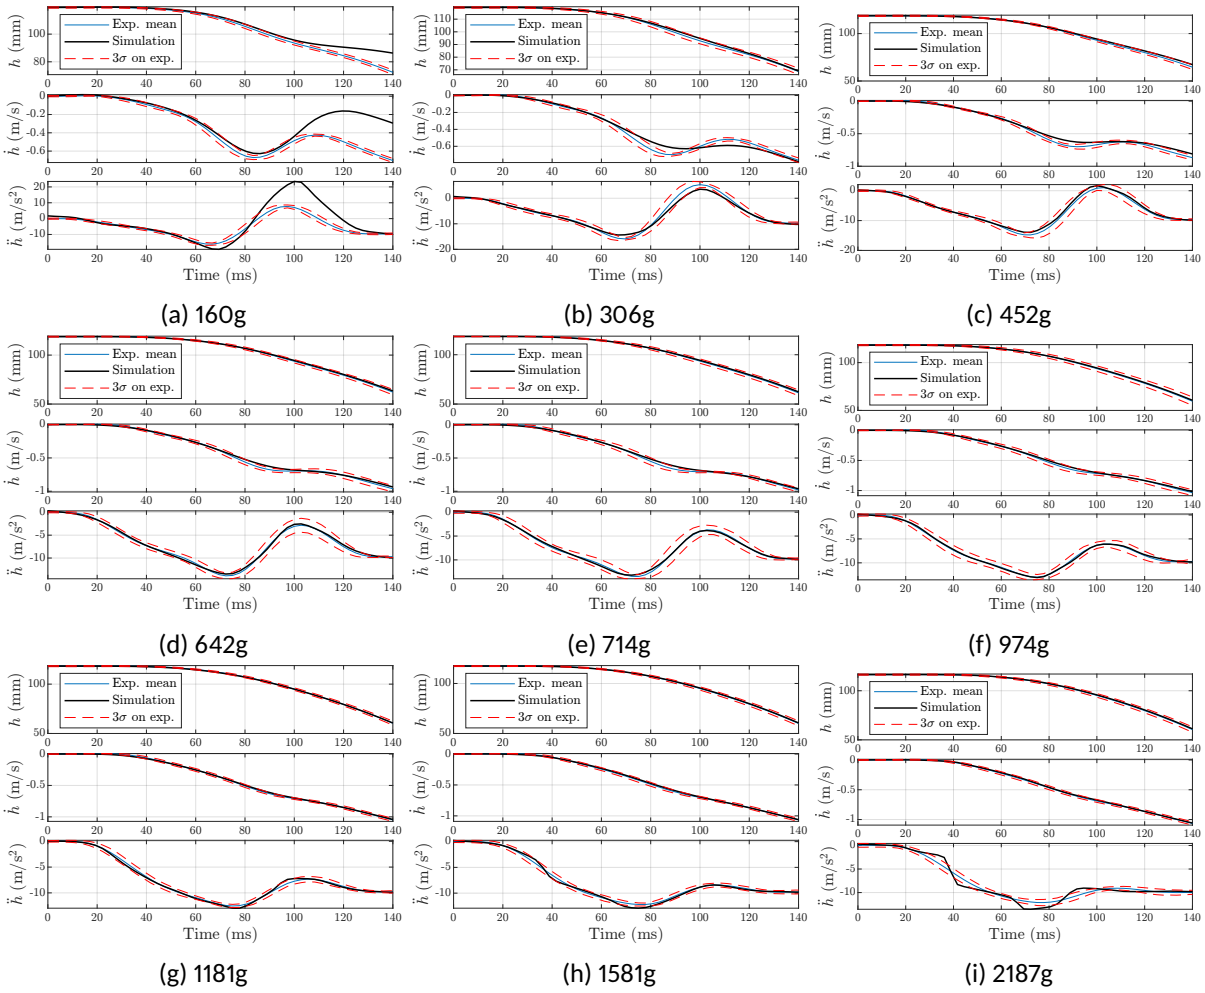

(a) 160g (b) 306g (c) 452g

(d) 642g (e) 714g (f) 974g

(g) 1181g (h) 1581g (i) 2187g

Figure 15: Predicted vs measured object height, velocity, and acceleration for a selection of object weights. Average measured object trajectories are plotted in blue, with 3σ intervals shown in red. The model prediction is shown in black. Edge cases (a) and (i) are not expected to be accurate, as they correspond to extrapolation of the experimental data.

high-dimensional spaces. This method was used because it is a well established method, and because in the future we will use it on a highly dimensional input space corresponding to the full 6D dynamics.

Consider Figure 15. The figure shows that by employing LWPR to learn the release dynamics from experimental data, we obtain quite accurate predictions of the motion of the object during release, for an experiment that was *not* included in the training set. More specifically, for each value of m, a force model is learned using the experiments with a different mass and then a simulation is performed by integrating the dynamics using m, the average acceleration $\ddot{a}(t)$, and the average initial condition $z(0)$ and $\dot{z}(0)$ from the experiments related to that value of m. Not surprising, the predictions are poorer for the edge cases ($m = 2187$g, but especially for $m = 160$g) as they correspond to extrapolation rather than interpolation. Excluding these edge cases, the *maximum absolute error* of the predicted trajectory from the average motion obtained from the experiments is of 2.14mm for the position, 0.09m/s for the velocity, and 2.65m/s^2 for the acceleration, respectively. The *root-mean-square error*, again excluding the edge cases, is 0.75mm, 0.023m/s, and 0.381m/s^2, respectively. [3].

2.6.1.2 Publication: "Experimental Validation of Nonsmooth Dynamics Simulations for Robotic Tossing involving Friction and Impacts"

Part of Task T2.1 is understanding how uncertainty distributions transform under the nonsmooth impact models developed in T1.4. In this publication [12], a sensitivity analysis is performed on the identified contact parameters (coefficient of friction and coefficient of restitution) in order to understand how uncertainty in the identified parameters reflects on the rest-pose prediction. Therefore, we vary the contact parameters around their optimal values to evaluate the predicted rest-pose sensitivity.
We perform a sensitivity analysis by changing the value of one of the parameters μ or e_N while keeping the other constant. If we vary a parameter, we consider a range of $[-0.05, 0.05]$ with steps of 0.01 around the optimal parameters listed in Table 2. In cases where the optimal value is equal to zero, we consider a range of $[0.00, 0.10]$. Figure 16 shows the results of this sensitivity study for three randomly selected experiments with Box005 and Box006, where the median of the simulation results are shown in full color, and the results of the varied parameters are shown with lower opacity.
Considering these results, we observe that on average, a variation in the coefficient of friction (μ) with ± 0.05 around the mean value, which corresponds to a variation of about $\pm 10\%$, leads to a variation in the rest-pose position of about ± 100mm, which corresponds to a variation of about $\pm 50\%$ of the maximum dimension of the box. At the same time, the variation in the coefficient of friction with ± 0.05 around the mean value leads to a variation in the rest-pose orientation of about ± 3deg. This shows the rest-pose position is rather sensitive the coefficient of friction, while the rest-pose orientation is not. This is also clearly visible in when we consider the corresponding results, shown in the first and third row of Figure 16.
A variation in the coefficient of restitution (e_N) with ± 0.05 around the mean value, corresponding to a variation of about $\pm 10\%$, leads to a variation in the rest-pose position of about ± 4mm, which corresponds to a variation of about $\pm 2\%$ of the maximum box dimension. At the same time, the variation in the coefficient of restitution with ± 0.05 around the mean value leads to a variation in the rest-pose orientation of about ± 3deg. This means the rest-pose position and orientation are rather insensitive to a variation in the coefficient of restitution. This is clearly represented by the second and fourth row of Figure 16. These results make sense when considering Figure 3, where both costs (for Algoryx Dynamics and MATLAB) show the insensitivity to a variation in e_N for values of $e_N < 0.6$.

[3] The code that is used to obtain the results, is made publicly available through `https://gitlab.tue.nl/robotics-lab-public/learning-1d-suction-cup-dynamics`

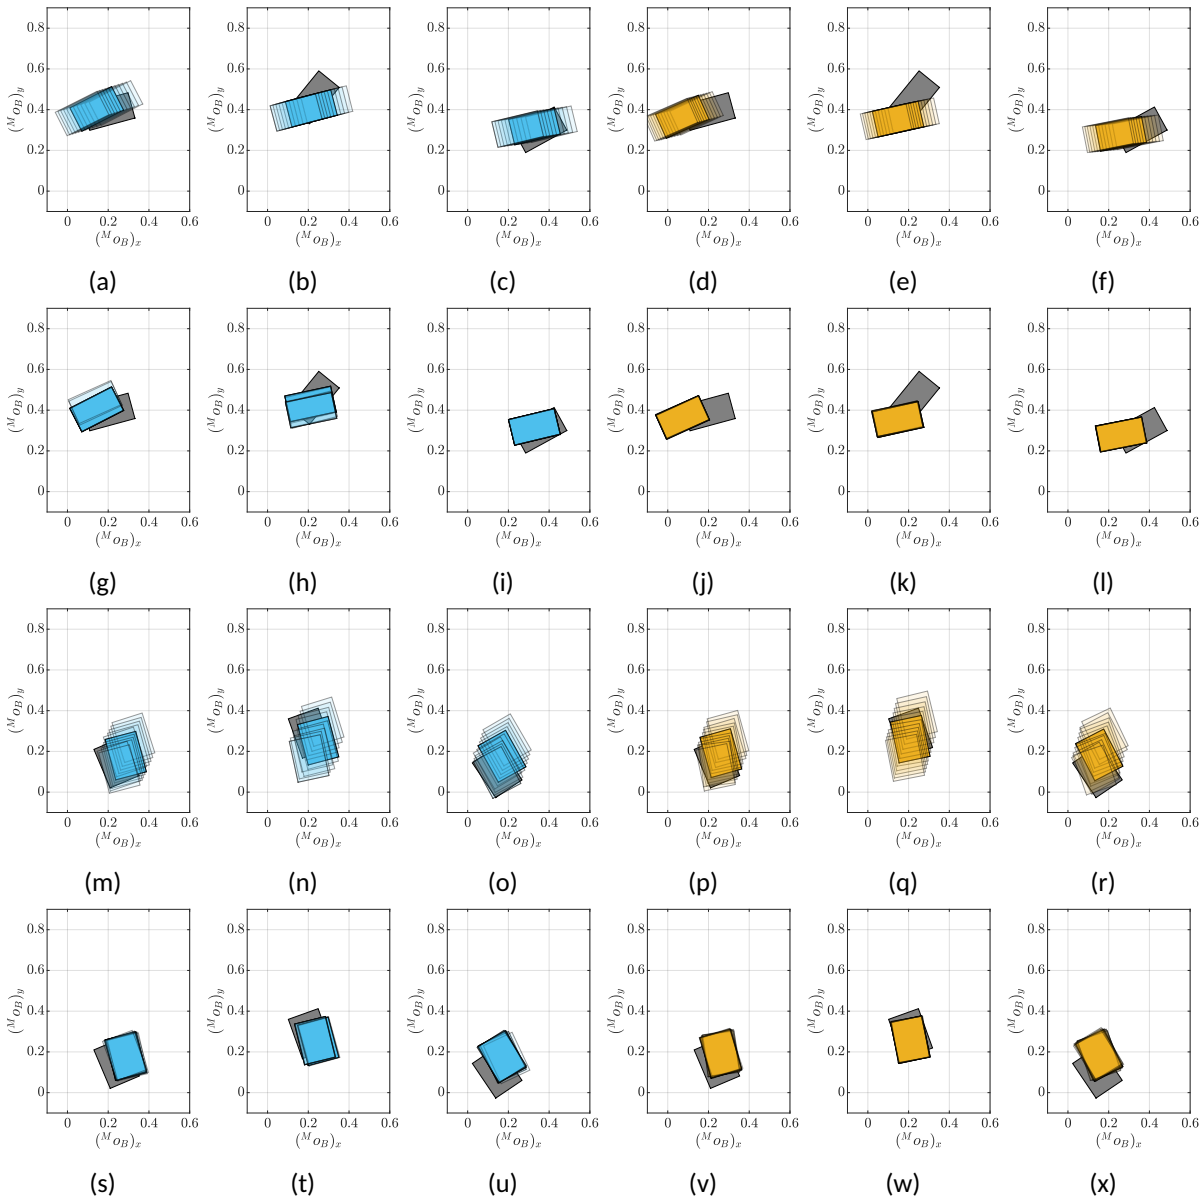

Figure 16: Resulting rest-poses for varying parameters. First row: result for varying μ for Box005. Second row: results for varying e_N for Box005. Third row: results for varying μ for Box006. Fourth row: results for varying e_N for Box006. The simulation results of Algoryx Dynamics and MATLAB are shown in blue (■) and orange (■), respectively, and the measured rest-pose is in gray (■).

2.6.2 Summary of ongoing work

In ongoing work, as a continuation of [27], we want to investigate how controlled perturbations in the pick pose of the suction gripper on the box, affect the outcome of a toss. In this work, we refer to this as the pick pose sensitivity on robotic tossing. We define the sensitivity as the variation between the outcomes of successive tossing demonstrations, for changed pick poses, under the same measurement conditions. By tossing outcomes, we refer to the rest position and orientation of the box after a toss.

	Velocity-based				Trajectory-based			
	MATLAB		Algoryx		MATLAB		Algoryx	
	μ	e_N	μ	e_N	μ	e_N	μ	e_N
Box004	0.60	0.40	0.65	0.35	0.45	0.05	0.45	0.10
Box005	0.45	0.35	0.40	0.30	0.45	0.10	0.40	0.00
Box006	0.25	0.40	0.25	0.40	0.40	0.25	0.40	0.25
Box007	0.40	0.60	0.35	0.55	0.35	0.45	0.35	0.40

Table 2: Resulting parameters found for velocity-based and trajectory-based parameter identification

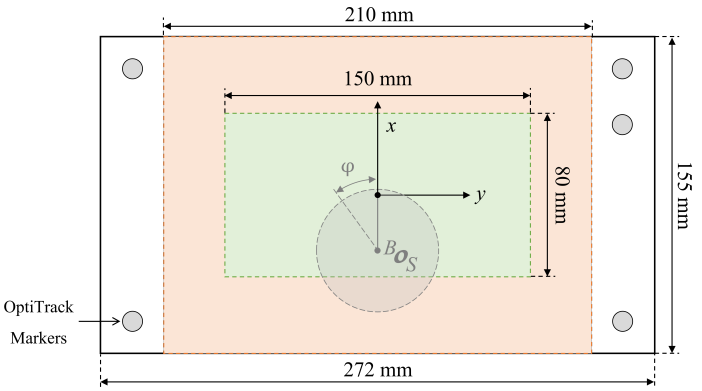

Figure 17: Schematic top view of Box005, where the orange marked area is the area where the suction gripper is allowed to pick the box, and the green area depicts the area where the center of the suction gripper can pick the box. Image taken from [27].

By exploring the effect of inaccurate box picking, on tossing results, it can be explored which picking accuracy is required, and what tossing precision can be achieved for a given accuracy. The goal of the pick pose sensitivity exploration is to determine if reliable and predictable tossing can be achieved, while disturbances in the pick pose of the suction gripper on a box are introduced. Moreover, this exploration might help to get an idea of what are acceptable limits on deviations in the pick pose, and how each independent pick pose parameter affects the outcome of a toss. Furthermore, by analyzing the deviation and precision of tossing results for changed pick poses on one object, it can be explored whether tossing can be applied in the crossorter infeed or traysorter infeed processes.

In the experiments carried out in this work, a box located on a fixed position on the conveyor is picked from different pick positions and orientations. Figure 17 shows schematically a top view of Box005, where the pick pose can be varied by changing the parameters for x, y, and φ. To vary the pick poses, we sample the parameters x, y, and φ from a multivariate Gaussian distribution.

The covariance for x, y, and φ, are chosen to be $3\sigma_{\mathrm{x_p}} = 75$ mm, $3\sigma_{\mathrm{y_p}} = 40$ mm, and $3\sigma_{\varphi_\mathrm{p}} = 30°$,

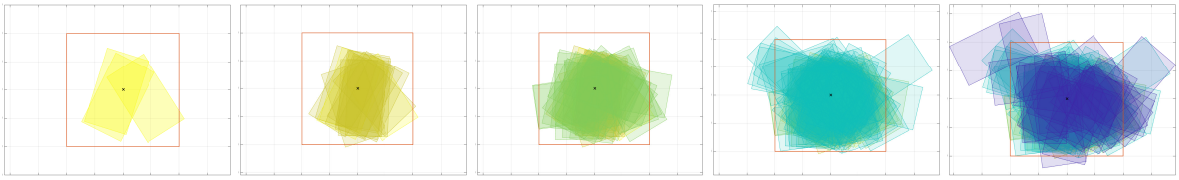

(a) $[\mathbf{0}, \frac{\mathbf{1}}{\mathbf{2}}\sigma]$ (b) $[\mathbf{0}, \frac{\mathbf{3}}{\mathbf{4}}\sigma]$ (c) $[\mathbf{0}, \sigma]$ (d) $[\mathbf{0}, \mathbf{2}\sigma]$ (e) $[\mathbf{0}, \mathbf{3}\sigma]$

Figure 18: Box' rest-poses for various ranges of deviations in the pick pose, for **short range** tossing. The orange square depicts the target area size of 400×400 mm in the traysorter process.

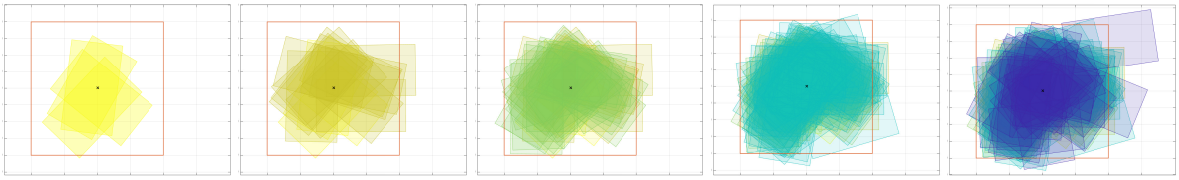

(a) $[\mathbf{0}, \frac{\mathbf{1}}{\mathbf{2}}\sigma]$ (b) $[\mathbf{0}, \frac{\mathbf{3}}{\mathbf{4}}\sigma]$ (c) $[\mathbf{0}, \sigma]$ (d) $[\mathbf{0}, \mathbf{2}\sigma]$ (e) $[\mathbf{0}, \mathbf{3}\sigma]$

Figure 19: Box' rest-poses for various ranges of deviations in the pick pose, for **long range** tossing. The orange square depicts the target area size of 400×400 mm in the traysorter process.

respectively. As such, the pick poses are defined within the bounds for successful picking, with a probability of 99.7%. In total, 125 different pick poses are generated.

The results in [27] show that, although a wide range of pick poses lead to unpredictable tossing, many tosses resulted in a box rest-pose that does lay close to the mean rest-pose of the repeatability experiments. This indicates that predictable tossing might be possible, even though the pick pose is not constant. This raises the question of how precise the robot must pick the box in order to perform predictable tossing operations. To analyze this, we divided the sampled pick pose distribution to explore what acceptable deviations in the pick pose are, for which robot tossing leads to predictable outcomes. To this end, we presented the sensitivity results, that the box' rest-poses with large deviations in the pick pose can be distinguished from tossing outcomes with smaller deviations in the pick pose. The tossing outcomes for various ranges of deviations in the pick pose are depicted in Figures 18 and 19, for short range (inside the kinematic range of the robot) and long range (outside the kinematic range of the robot) tossing, respectively.

Here, we use σ to denote each pick pose parameter covariance $\sigma_{\mathbf{x_p}}$, $\sigma_{\mathbf{y_p}}$, and $\sigma_{\varphi_\mathbf{p}}$. It is clear that tossing further leads to more unpredictable behavior, and deviations further from the ideal pick pose result in worse tossing outcomes. However, for the scenario of tossing on a crossorter (as considered here), picked objects with a deviation up to σ will be tossed successfully[4]

2.6.3 Summary achievements, limitations, and future directions

Various publications can be directly related to Task T2.1 and ongoing work shows potential of future publications focusing on understanding the uncertainties introduced by impacts. In particular, a data

[4]Although small parts of the boxes may appear outside the target window, this will not result in the box falling out the window, and therefore these tosses are considered as successful.

driven approach has been successfully employed to a create model for object release from a suction cup that shows remarkably accuracy and reproducibility. The works [10, 22] have focused on validation of rigid impact maps in the context of robot/environment impacts, and can now be used to understand how uncertainty distributions around the pre-impact configuration and velocity of a robotic arm transform under the nonsmooth impact map. Taking a statistical perspective, experimental work has been carried out to quantify how uncertainties in the picking position or uncertainty in the knowledge of impact model parameters [12] affect the final pose in robot tossing experiments, fully inline the task goal of quantifying post-impact state uncertainty.

3 CONCLUSION

This deliverable D1.3 I.Model report provides an overview of the publications related to the I.Model framework. More specifically, for each of the tasks related to I.Model (T1.1, T1.2, T1.3, T1.4, T1.5, T2.1), this report summarizes the main achievements, limitations, and future directions. The project will continue for more than half a year and the models discussed in this deliverable will be further developed until then.
Aside from this deliverable, Deliverable D1.2 Physics Engine API also covers a large part on the implementation of the developed models of I.Model into Algoryx Dynamics, which will be published in September 2023. Furthermore, the deliverables D1.1 and D1.4, already published September 2020 and December 2022, provide further details on the development of the open dataset, as part of I.Model.

4 REFERENCES

[1] M. J. Jongeneel, A. Saccon, and J. den Ouden, "Deliverable D1.1: Publication of I.AM. Dataset," Eindhoven University of Technology, Eindhoven, The Netherlands, Tech. Rep. D1.1, 2020.

[2] ——, "Deliverable D1.4: Publication of I.AM. Dataset," Eindhoven University of Technology, Eindhoven, The Netherlands, Tech. Rep. D1.4, 2022.

[3] M. J. Jongeneel, S. Dingemans, and A. Saccon, "The Impact-Aware Robotics Database," *Preprint on HAL*, Dec. 2022. [Online]. Available: https://hal.science/hal-03900923

[4] A. Saccon and M. J. Jongeneel, "TU/e Impact-Aware Robotics Database," Eindhoven, The Netherlands, 2022. [Online]. Available: https://impact-aware-robotics-database.tue.nl/

[5] M. J. Jongeneel and S. Dingemans, "I.AM. archive containing long range box-toss experiments with Box005 and Box006 for validation of impact models," 2023. [Online]. Available: https://doi.org/10.4121/21399657

[6] M. J. Jongeneel, S. Dingemans, and G. van den Brandt, "I.AM. Archive containing dual-arm grabbing experiments," 2022. [Online]. Available: https://doi.org/10.4121/21554478

[7] M. J. Jongeneel, S. Dingemans, and A. Oliva, "I.AM. Archive containing experiments of stacking boxes in totes, relevant for the BOX scenario." 2022. [Online]. Available: https://doi.org/10.4121/21554478

[8] M. L. S. Lubbers, J. van Voorst, M. J. Jongeneel, and A. Saccon, "Learning Suction Cup Dynamics from Motion Capture: Accurate Prediction of an Object's Vertical Motion during Release," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2022)*, October 2022, pp. 1541–1547. [Online]. Available: https://doi.org/10.1109/IROS47612.2022.9982211

[9] M. J. Jongeneel, A. Saccon, J. Voorst, van, and M. Lubbers, "Impact Aware Manipulation (I.AM.) archive containing suction cup release experiments," 2022. [Online]. Available: https://data.4tu.nl/articles/dataset/Impact_Aware_Manipulation_I_AM_archive_containing_suction_cup_release_experiments/20536569/2

[10] I. Aouaj, V. Padois, and A. Saccon, "Predicting the post-impact velocity of a robotic arm via rigid multibody models: an experimental study," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 2264–2271. [Online]. Available: https://doi.org/10.1109/ICRA48506.2021.9561768

[11] A. Saccon and V. Padois, "Dataset and related MATLAB scripts for the IEEE ICRA 2021 paper 'Predicting the Post-Impact Velocity of a Robotic Arm via Rigid Multibody Models: an Experimental Study'," 2021. [Online]. Available: https://doi.org/10.4121/14192429

[12] M. J. Jongeneel, L. Poort, N. van de Wouw, and A. Saccon, "Experimental Validation of Nonsmooth Dynamics Simulations for Robotic Tossing involving Friction and Impacts," *Preprint on HAL*, Feb. 2023. [Online]. Available: https://hal.science/hal-03974604

[13] M. J. Jongeneel, "Data underlying the publication: Validating rigid-body dynamics simulators on real-world data for robotic tossing applications," *4TU.ResearchData*, Nov 2022, Dataset Collection. [Online]. Available: https://doi.org/10.4121/c.6278310

[14] The 4TU.Federation, "4TU.ResearchData data and software repository," 2008-2022. [Online]. Available: https://data.4tu.nl/info/en/

[15] B. Acosta, W. Yang, and M. Posa, "Validating Robotics Simulators on Real-World Impacts," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 6471–6478, 2022.

[16] Y.-B. Jia, M. Gardner, and X. Mu, "Batting an in-flight object to the target," *The International Journal of Robotics Research*, vol. 38, no. 4, pp. 451–485, 2019. [Online]. Available: https://doi.org/10.1177/0278364918817116

[17] S. Olufs, F. Adolf, R. Hartanto, and P. Plöger, "Towards Probabilistic Shape Vision in RoboCup: A Practical Approach," in *RoboCup 2006: Robot Soccer World Cup X*, G. Lakemeyer, E. Sklar, D. G. Sorrenti, and T. Takahashi, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 171–182.

[18] A. Doucet, N. De Freitas, and N. Gordon, *Sequential Monte Carlo Methods in Practice*, 1st ed. Springer New York, 2001.

[19] O. Cappe, S. J. Godsill, and E. Moulines, "An Overview of Existing Methods and Recent Advances in Sequential Monte Carlo," *Proceedings of the IEEE*, vol. 95, no. 5, pp. 899–924, 2007.

[20] M. J. Jongeneel, A. Bernardino, N. van de Wouw, and A. Saccon, "Model Based 6D Visual Object Tracking with Impact Collision Models," in *2022 American Control Conference (ACC)*, June 2022, pp. 3850–3856. [Online]. Available: https://doi.org/10.23919/ACC53348.2022.9867622

[21] R. Van Der Merwe, A. Doucet, N. De Freitas, and E. Wan, "The Unscented Particle Filter," in *Proceedings of the 13th International Conference on Neural Information Processing Systems*, ser. NIPS '00, Cambridge, USA, Aug. 2000, pp. 563–569.

[22] C. A. R. Arias, W. Weekers, M. Morganti, V. Padois, and A. Saccon, "Refined Post-Impact Velocity Prediction for Torque-Controlled Flexible-Joint Robots," *Preprint on HAL*, Jul. 2023. [Online]. Available: https://hal.science/hal-04148817

[23] M. Jongeneel, C. Andrés Rey Arias, and S. Dingemans, "Iam archive containing impact experiments with a franka emika panda robot," 2023. [Online]. Available: https://doi.org/10.4121/21655631

[24] M. J. Jongeneel and A. Saccon, "Geometric Savitzky-Golay Filtering of Noisy Rotations on SO(3) with Simultaneous Angular Velocity and Acceleration Estimation," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2022)*, October 2022, pp. 2962–2968. [Online]. Available: https://doi.org/10.1109/IROS47612.2022.9981409

[25] S. Vijayakumar, A. D'Souza, and S. Schaal, "Incremental Online Learning in High Dimensions," *Neural computation*, vol. 17, pp. 2602–34, Jan. 2006.

[26] "Locally Weighted Projection Regression," https://web.inf.ed.ac.uk/slmc/research/software/lwpr, University of Edinburgh, 2019, [Online; accessed 29-Jun-2022].

[27] T. van Gorp, "Towards robot tossing by demonstration through intuitive teleoperation," Master's thesis, Eindhoven University of Technology, Department of Mechanical Engineering, 2022, DC 2022.038.