Validating rigid body simulation of real particle shapes using pose estimation from high-speed video

Mathew Price[1] Garry Morrison[2]

[1]Department of Electrical Engineering
University of Cape Town
South Africa
mathew@dip.ee.uct.ac.za

[2]De Beers Group Services
Johannesburg
South Africa
Garry.Morrison@debeersgroup.com

Abstract

We present an image based method for computing 3D trajectories of rigid particles from experiments, where sequences are captured using a single high-speed video camera (HSV). The computed trajectories, which are representative of real world particle interaction, can then be directly compared with trajectories obtained from DEM simulations, hence aiding validation. Experiments consist of a diamond impacting a plate and are captured at 4000 frames per second. Simple image analysis is used to track the particle in each frame and to extract its 2D silhouette boundary. Using an approximate 3D model of the particle, generated from a multi-camera setup, a pose estimation scheme based on silhouette consistency is used in conjunction with a rigid body simulator to produce a registered 3D trajectory. We demonstrate how this information can be used to validate typical simulation characteristics, such as contact models, when applied to particles of arbitrary shape. Under reasonable conditions, the method can reliably estimate linear and angular particle velocities to within 1%.

1 Introduction

Understanding mass particle interaction in large scale processes is a complex problem. Simulation based techniques such as the discrete element method (DEM) have brought new insight to the design and optimisation of such processes. Emergent behaviour of the system is sought by simulating a large number of individual particles which interact according to specified mechanisms. These mechanisms operate at a single particle level. Hence incorporation of appropriate interaction models is important for successful simulation.

An important component of simulation based design and optimisation is the validation process. The validation needs to occur at both a global and a mechanism level. In multi-particle processing systems detailed global validation can be difficult, because tracking and measurement of individual particles is intractable. Hence one form of global validation takes the form of comparing gross trends such as particle throughput, density and power draw. Validation of interaction mechanisms can be no less problematic. Our goal is to develop a tool that can be used to assist in these validation processes.

Particle process simulation generally results in a vast simplification of the real world process in order to extract a computationally viable model. Various proposals have been made to circumvent problems caused by these simplifications. In the case of shape representation these include sphere clumping and superquadric representations (Williams, 1991). However, such methods, while possibly improving accuracy, come at the price of an increased computational overhead. The issue of suitable contact models and shape representation are interconnected because the material response of the contacting bodies may differ depending on the contact geometry.

Material response to particle contacts can range from elastic to elasto-plastic to plastic. In many cases the simple Hertz contact model is employed even though this is only applicable to purely elastic contacts. Whilst this may be appropriate in relatively static environments, e.g. soil mechanics, this is rarely the case in most dynamic applications. Empirical spring and dashpot contact models are often employed to capture a degree of the inelastic material response (Walton and Braun, 1986). While such models have the advantage of speed, they do not relate to material properties, making them difficult to generalise. Although some progress has been made in the development of phenomenological contact models that are parameterised by relevant material properties (Li et al., 2002), these are often computationally expensive. Moreover, in the case of real particle shapes, the rebound behaviour is difficult to validate.

Towards facilitating the validation process, we present an image based method capable of analysing the 3D motion of real particles within a controlled experiment, captured with a single high-speed video camera (HSV). Since these experiments involve actual particle interactions, the analysed 3D motion can be used, for instance, to assess the validity of corresponding DEM simulations. We limit our analysis to single impacts involving a single particle and a plate. Eventually, we hope to extend this to include particle-on-particle impacts.

2 Problem definition

Given a HSV sequence of a particle-plate impact experiment and a previously generated 3D model of the particle, we wish to determine the 3D trajectory of the particle. This requires estimating the particle's *pose*. Once the trajectory is known, a simulation of the experiment can be constructed by initialising the 3D model of the particle to the initial state of the estimated trajectory. Since the estimated trajectory represents the particle's motion in the real world, it can be used to gauge the accuracy of the simulated experiment if suitable material properties for the contacting bodies are known.

An approximate 3D model of the particle is generated using a six-camera setup as illustrated in Figure 1. The cameras are calibrated, i.e. the relative poses are known. The silhouette boundaries of the particle are extracted from each image and combined with the corresponding calibration parameters, forming a calibrated *silhouette set*. Volumetric carving followed by polygonisation (Bloomenthal, 1994) is used to generate a 3D mesh of the maximum bounded volume, known as the *visual hull* (Forbes et al., 2003). By repeating this process several times and merging the results, more accurate 3D models can be constructed. In our application, we merge ten runs for each particle, resulting in a 60-view calibrated silhouette set (i.e. the particle is effectively seen from 60 different viewpoints).

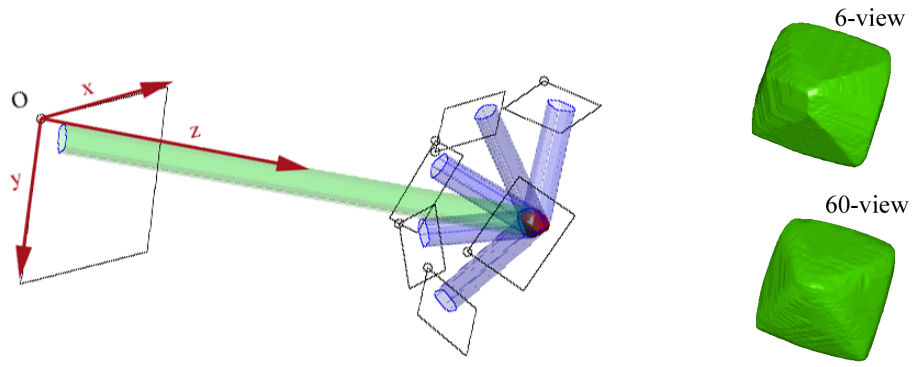

Figure 1: *Left: A typical 6-view calibrated silhouette set associated with the six-camera setup. The 3D model of the particle (visual hull) is the result of the intersection of all the visual cones. The unknown viewpoint associated with the a particle's pose in a HSV sequence is shown with respect to the calibrated silhouette set on the left. Right: Example of a 6- and 60-view visual hull.*

Assuming the particle undergoes unhindered rigid body motion both before and after impact, the full trajectory can be modelled as a train of piecewise trajectories separated by collision events. This simplifies estimating the particle's trajectory, since each piecewise component corresponds to a constant linear and angular momentum term. Constant forces, such as gravity, can also be included in the model with minimal additional complexity. Currently, we do not consider cases in which the particle rolls along the plate.

3 Pose estimation

A particle is moving with respect to a fixed camera whose relative pose is unknown. Equivalently, if we consider the object to be fixed, the problem can be visualised as someone attempting to move the camera around the object so that a projection of the object into the camera view corresponds to the image seen by the camera.

The relative pose between the world and the reference frame of a camera can be specified by a rigid transform $\mathtt{P}_{\mathtt{wc}}$ with a homogenous coordinate system:

$$\mathtt{P}_{\mathtt{wc}} = \begin{bmatrix} \mathtt{R} & \mathbf{t} \\ 0 & 1 \end{bmatrix}, \tag{1}$$

where $\mathtt{R}$ is a 3×3 rotation matrix and $\mathbf{t}$ is a 3×1 vector corresponding to a translation. Therefore, the pose estimation problem corresponds to finding appropriate values for $\mathtt{R}$ and $\mathbf{t}$. Figure 1 illustrates this scenario.

$\mathtt{P}_{\mathtt{wc}}$ can be estimated from silhouettes by maximising the silhouette consistency between the unknown camera and the calibrated silhouette set. A convenient measure of silhouette consistency is the epipolar tangency error, which measures the extent to which the visual cones, formed by rays passing through a camera centre and a corresponding silhouette boundary, correctly intersect. The primary advantage of this formulation is its efficiency, since it can be computed directly from silhouettes and does not require an explicit 3D model (e.g. triangular mesh). In addition, since the measure is based on a geometric constraint, it can be formulated as an objective function that can be minimised using standard nonlinear least-squares optimisation techniques.

3.1 The epipolar tangency error

Consider two views of an object as shown in the left hand image of Figure 2. Each camera ($C_\mathtt{A}$ and $C_\mathtt{B}$) images an object from different views resulting in the corresponding silhouettes $S_\mathtt{A}$ and $S_\mathtt{B}$. We assume a pinhole camera model, thus the visual cones of the silhouettes converge to the camera centres: c_a and c_b.

The projection of each camera centre into the opposite camera's image plane is called an *epipole*. Therefore, epipoles are points of intersection of the line created by joining the two camera centres (the baseline) and the respective image planes. In the figure, epipoles $e_{a\mathtt{B}}$ and $e_{b\mathtt{A}}$ corresponding to the projections of c_a and c_b into cameras $C_\mathtt{B}$ and $C_\mathtt{A}$ are shown. Epipolar tangency lines are constructed for each view by drawing lines from the epipoles to their corresponding silhouette boundary in each image, such that they touch the extents of their silhouettes. The epipolar tangency lines, which exist in 2D image space, touch their silhouettes at the epipolar tangency points (t_{a_1}, t_{a_2}) for $S_\mathtt{A}$ and (t_{b_1}, t_{b_2}) for $S_\mathtt{B}$. Since these points are projections of the same 3D points, called *frontier* points, they must be consistent across both views. That is, the ray passing through $c_\mathtt{A}$ and t_{a_1} must project exactly onto the epipolar tangency line $e_{a\mathtt{B}}t_{b1}$. This is known as the epipolar tangency constraint (Wong, 2001). In practice, there is inherent noise in the imaging process and the camera parameters are not without error, which leads to inexact projection of the epipolar tangency lines as illustrated in the right hand image of Figure 2.

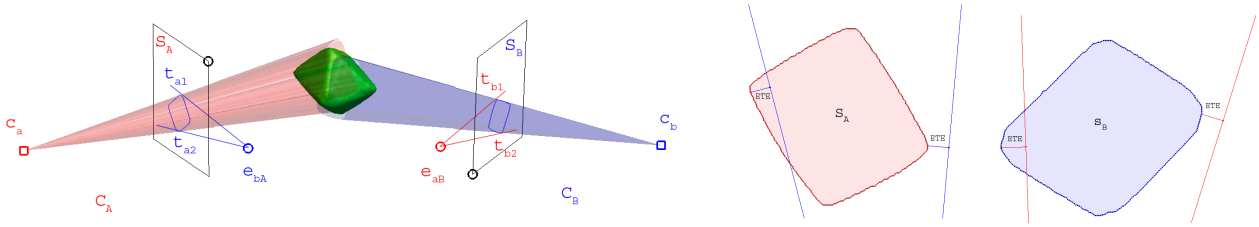

Figure 2: *Left: An object viewed from two cameras with corresponding epipoles and epipolar tangents shown. Right: Example of Epipolar Tangency Error (ETE).*

The perpendicular distance between an epipolar tangency point and a projected epipolar tangency line from another view is called the epipolar tangency error (ETE). The ETE is measured in pixels and provides a convenient measure of consistency between two silhouette projections. For each pair of silhouettes there are two epipolar tangency points for each image, resulting in four ETE's. A detailed explanation on computing the ETE can be found in (Forbes et al., 2003).

4 3D Trajectory estimation

Estimating the pose of a particle for an entire sequence would typically require estimating a rigid transform (Equation 1) for each frame. Our method simplifies this by using a rigid body motion model to predict the pose of the particle during frames in which the particle is moving unhindered, i.e. not colliding with any other body. After estimating the pose of the particle for the first frame, the camera parameters for the remaining frames can then be generated by simulating the particle according to the motion model parameters. The ETE is retained as a measure of how consistent the resulting trajectory is with respect to the calibrated silhouette set of the particle. This results in a pose estimation framework, which makes use of the silhouette consistency constraint while ensuring that only physically possible solutions are found. We assume that the particle is considered to be a rigid body, and particle motion between impact events follows unhindered rigid body motion.

4.1 Particle initialisation

The first task in determining the trajectory of a particle in a HSV sequence is to estimate its initial pose. The formulation discussed in Section 3 is used. Given the particle's silhouette at the beginning of the HSV sequence (unknown camera) and a calibrated silhouette set of the particle (six-camera setup), we wish to determine the relative pose of the unknown camera.

A perspective camera model with square pixels and no lens distortion is assumed for the unknown camera, which is specified by a 4×4 projection matrix $\mathtt{P}$:

$$\mathtt{K} = \begin{pmatrix} f & 0 & u_x \\ 0 & f & u_y \\ 0 & 0 & 1 \end{pmatrix} \qquad (2) \qquad\qquad \mathtt{P} = \begin{bmatrix} \mathtt{K} & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \mathtt{R} & \mathbf{t} \\ 0 & 1 \end{bmatrix}. \qquad (3)$$

$\mathtt{K}$ contains the camera's intrinsic parameters, focal length f and principal point (u_x, u_y), and is used to define the perspective projection. $\mathtt{R}$ and $\mathbf{t}$ form a rigid body transform, from Equation 1, that orients the world to the camera's reference frame. A convenient parameterisation for rotations is to use quaternions. This allows the nine parameter

rotation matrix $\mathtt{R}$ to be specified by a four parameter unit quaternion $\mathbf{Q}$. Another reasonable assumption is that the principal point (the projection of the camera centre onto the image plane) is approximately at the centre of the image. The unknown camera can therefore be specified by eight parameters:

$$\mathbf{X} = [\mathbf{f}\ \mathbf{Q}^{\mathtt{T}}\ \mathbf{t}^{\mathtt{T}}], \tag{4}$$

where $\mathbf{X}$ represents a concatenated vector of parameters.

Suitable parameter values can be efficiently determined by iterative nonlinear least-squares optimisation — here the Levenberg-Marquardt method (Moré, 1977) is used (LM optimisation). The objective function for the optimisation is specified by a concatenated vector of ETE's, which are formed by considering all forward and reverse pairings between the unknown silhouette and each silhouette view in the calibrated set. This leads to $4n$ errors (from $2n$ pairs), where n is the number of silhouettes in the calibrated silhouette set.

4.2 Motion parameters

After determining the initial pose of the particle, the trajectory can be estimated by finding rigid body motion parameters that best reproduce the particle's motion in the HSV sequence. Each sequence is associated with an uncalibrated silhouette set, i.e. there is an unknown view of the particle for each frame in the sequence. We define the particle state by the vector:

$$\mathbf{Y}(t) = \left[\mathbf{x}(t)^{\mathtt{T}}\ \ \mathbf{Q}(t)^{\mathtt{T}}\ \ \mathbf{P}(t)^{\mathtt{T}}\ \ \mathbf{L}(t)^{\mathtt{T}}\right], \tag{5}$$

relating to position, orientation (quaternion representation), linear momentum and angular momentum respectively.

Under the assumption of unhindered rigid body motion, if we know $\mathbf{Y}(t_0)$ then all future states (until a collision) can be determined. From the initial pose estimation (Section 4.1), $\mathbf{x}(t_0)$ and $\mathbf{Q}(t_0)$ are already known. Therefore, we need to determine suitable values for the linear and angular momentum terms. As before, the iterative least-squares approach (LM optimisation) is used. The parameter vector for optimisation is:

$$\mathbf{Z} = [\mathbf{P}(\mathbf{t_0})^{\mathtt{T}}\ \mathbf{L}(\mathbf{t_0})^{\mathtt{T}}]. \tag{6}$$

Next, an objective error function must be defined. With any given parameter vector $\mathbf{Z}$, we associate a set of simulated states $\beta = \{\mathbf{Y}(t_0) \ldots \mathbf{Y}(t_{n-1})\}$, where n is the number of frames in the sequence. Since the derivative of the state vector is:

$$\frac{d}{dt}\mathbf{Y}(t) \quad = \quad \left[\mathbf{v}(t)^{\mathtt{T}}\ \ \frac{1}{2}\boldsymbol{\omega}(t)\mathbf{Q}(t)^{\mathtt{T}}\ \ \mathbf{F}(t)^{\mathtt{T}}\ \ \boldsymbol{\tau}(t)^{\mathtt{T}}\right], \tag{7}$$

(with $\mathbf{v}$, $\boldsymbol{\omega}$, $\mathbf{F}$ and $\boldsymbol{\tau}$ corresponding to linear velocity, angular velocity, force and torque respectively) β can be determined using a standard ODE solver. In our implementation we have found that a 4th order Runge-Kutta solver produces good results (Eberly, 1999).

In order to retain the ETE as the measure of silhouette consistency for the optimiser's objective function, the set of state vectors β must be converted to an equivalent set of camera projection matrices $\Psi = \{\mathsf{P}_1 \ldots \mathsf{P}_n\}$. Similarly, we define Ω to be the set of projection matrices corresponding to the calibrated silhouette set of the particle (i.e. from the six-camera setup). This representation may require a transformation if the simulation's reference frame differs from that of the camera. For simplicity, we use the conventional camera reference frame, depicted in Figure 1, for both systems. Having identical reference frames leads to the following direct relation between a state vector $\mathbf{Y}(t_i)$ and a corresponding camera projection matrix P_{i+1}:

$$\mathsf{P}_{i+1} \quad = \quad \begin{bmatrix} \mathtt{K} & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \mathsf{R}(t_i) & \mathbf{x}(t_i) \\ 0 & 1 \end{bmatrix}, \tag{8}$$

where $\mathsf{R}(t_i)$ is the rotation matrix corresponding to $\mathbf{Q}(t_i)$, $0 \leq i < n-1$ and K is the constant matrix of intrinsic camera parameters from Equation 2. If $S = \{S_1 \ldots S_n\}$ is the set of silhouettes associated with the set of projection matrices Ψ, then let $\{S, \Psi\}$ denote the silhouette set corresponding to the uncalibrated HSV sequence. Similarly, let $\{C, \Omega\}$, with $C = \{C_1 \ldots C_m\}$, be the calibrated silhouette set of the particle acquired from the six-camera setup. The silhouette consistency across both silhouette sets can therefore be computed by pairing each silhouette from $\{S, \Psi\}$ with each silhouette in $\{C, \Omega\}$ and vice versa, resulting in nm ETE's. The concatenation of these errors forms the objective function used by the LM optimisation. Final refinement of the solution can be conducted by adding inter-silhouette pairings within $\{S, \Psi\}$ itself. This enables the uncalibrated silhouette set to be optimised for internal consistency as well as with the calibrated silhouette set.

4.3 Piecewise trajectories

Extending the above procedure so that motion across impact frames is handled transparently, we have introduced the notion of piecewise rigid body motion. In simple interactions, sets of frames relating to unconstrained rigid body motion

can be separated by a set of collision events. Therefore, the sequence can be described as a piecewise combination of rigid body trajectories, which are separated by a set of discontinuous impulses. The advantage of this formulation is that the complexities of mechanical interaction during contact can be bypassed by treating their combination as a unified event, which results in the instantaneous change of momentum. From the point of view of the particle there are two separate rigid body trajectories, which overlap at some time and hence share a common pose as illustrated in Figure 3.

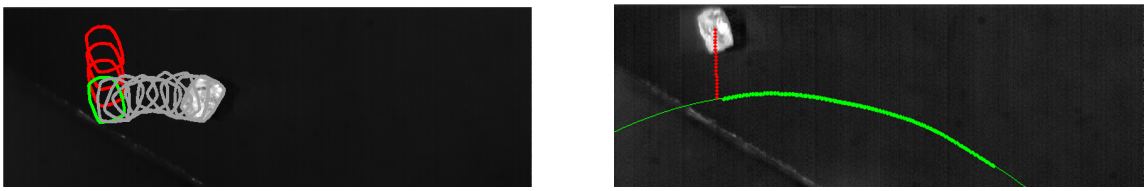

Figure 3: *Piecewise trajectory splitting for a particle impacting a plate (left) with automatic collision frame detection (right).*

Since the contact duration is substantially less than the inter-frame time, the exact collision time does not occur at any particular frame, but rather between frames. As we are estimating the motion of the particle so that it replicates an imaged experiment, discrete time steps are enforced. Given a set of silhouettes as input, the angles between successive centroids is used to automatically divide the set into piecewise rigid body trajectories. After estimating the parameters for the first piecewise trajectory, the impact frame is automatically determined by intersecting a simulation of the first piecewise trajectory with an approximation of the second piecewise trajectory. The approximate trajectory is formed by fitting a second order polynomial curve to the 2D centroids of the corresponding silhouettes (shown in Figure 3). If $S1 = \{S_1 \dots S_\alpha\}$ and $S2 = \{S_\alpha \dots S_n\}$ are sets of silhouettes corresponding to the first and second piecewise trajectories respectively, then the generalised 3D trajectory estimation method is as follows:

1. Estimate Ψ_1 for $S1$ using the trajectory estimation method (from Sections 4.1 and 4.2).
2. Approximate collision frame α using newly estimated trajectory (Figure 3).
3. Estimate Ψ_2 for $S2$ using projection matrix P_α from Ψ_1 as the initial pose.
4. Refine global solution $\{S1, \Psi_1\} \bigcup \{S2, \Psi_2\}$ for maximum silhouette consistency.

4.4 Initial parameter selection

A common issue with optimisation frameworks is their sensitivity to the initial parameters. This is especially notable when optimising using ETE as an error measure since the epipolar tangency constraint only constrains the true object to lie within the intersected volume of the viewing cones. It is therefore possible for the optimisation to become trapped in local minima, which results in the camera parameters being unable to move towards the correct solution without increasing the residual error term. This seems to occur often with more regularly shaped particles since several possible poses can be equally consistent. One workaround is to apply multiple optimisations with different starting points. In particular, the primary cause of locking in the initial pose estimation process is the choice of the camera's initial orientation. This problem can be alleviated by repeatedly sampling random orientations (quaternions) from a normal distribution. The search can be terminated by setting a suitable threshold on the RMS value of the ETE vector (e.g. 1 pixel).

In the case of estimating the motion parameters, the equivalent problem is encountered with the initial choice of angular momentum $\mathbf{L}(t)$. Here random angular momentum directions (3×1 vectors) are sampled from a normal distribution and the magnitude of the vector is limited so that exceptionally large momentum values, which cause numerical instability, are avoided.

5 Results

The results are split into several groups. First, we evaluate the primary components of the method: particle initialisation and trajectory estimation. In the latter case, synthetic sequences generated from real particles are used. The method is then applied to real data in the form of particle-on-plate impact tests, which are captured with a HSV camera. Since actual ground truth is not available for the HSV experiments, we use the overall RMS ETE as a performance metric. A low overall ETE implies that the silhouettes in a HSV sequence correspond closely with those obtained from the 3D model of the particle, simulated according to the estimated trajectory, i.e. there is a good fit between the estimated trajectory and the physically permissable trajectory. Finally, we show comparative results for particle velocity under different contact models.

5.1 Particle initialisation

Since the extent to which the trajectory can be estimated for a particle is largely related to the success of the pose estimation framework, we first show the effectiveness of the particle initialisation process (Section 4.1).

Using a sixty-view silhouette set of a known particle, a leave-one-out validation scheme is used in which each silhouette is removed from the set in turn and its pose is estimated using our method. The estimated pose is then compared with the silhouette's original calibrated pose, which is deemed to serve as ground truth. Figure 4 shows the results for three runs of six particles (using different starting points), where three particles are regular in shape and three are of an irregular nature. This results in three pose estimates for each silhouette from a sixty-view set and ultimately

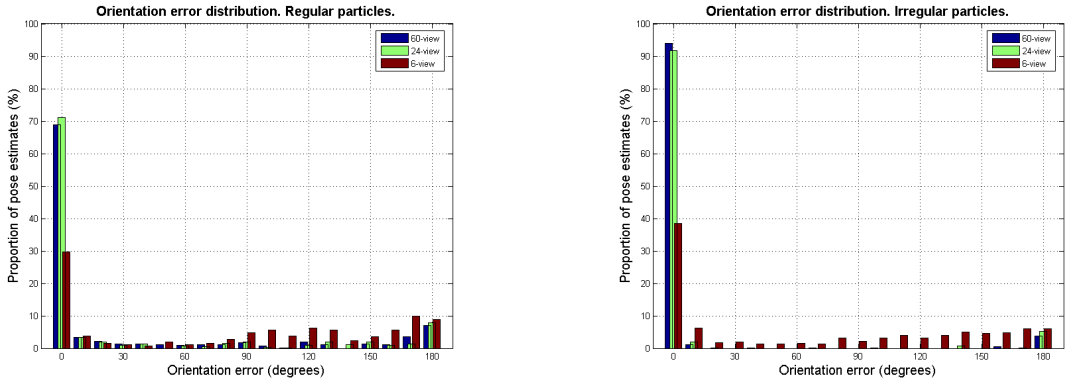

Figure 4: *Orientation error distributions. Left: regular particles (66% of the estimates are accurate to within 5 degrees). Right: irregular particles (88% of the estimates are accurate to within 5 degrees). Each bin represents 10 degrees.*

540 pose estimates for regular and irregular particles alike. Errors are binned according to the angular orientation error in degrees. From the distributions it is evident that single view poses from irregular particles are more consistently estimated than for regular particles. The small peaks located near the 180 degree bin show cases where the antipodal view has been found instead. This occurs because random orientations close to the antipodal view may appear similar and therefore result in a local minimum in the ETE function. Experiments conducted with smaller silhouette sets (6-and 24-view) show that the reliability of the solution degrades gracefully with decreasing number of views.

5.2 Trajectory estimation

The second group of tests evaluates the accuracy of the trajectory estimation method. This is accomplished by randomly generating non-contacting simulated trajectories of randomly selected particles and using our method to estimate the linear and angular components of the motion. Simulations were conducted in which the sequence length, linear momentum and angular momentum were randomly varied. In addition, particles were initialised to views from their calibrated silhouette sets, thus particle initialisation does not contribute any error in the results. Although the silhouettes extracted from the simulated trajectory are synthetic, the estimation uses the particle's original calibrated silhouette set, which corresponds to real data. Therefore, these tests provide a realistic means of assessing our method given that the segmentation in the HSV sequence is ideal. Figure 5 summarises the results, which are based upon the linear and angular velocity errors.

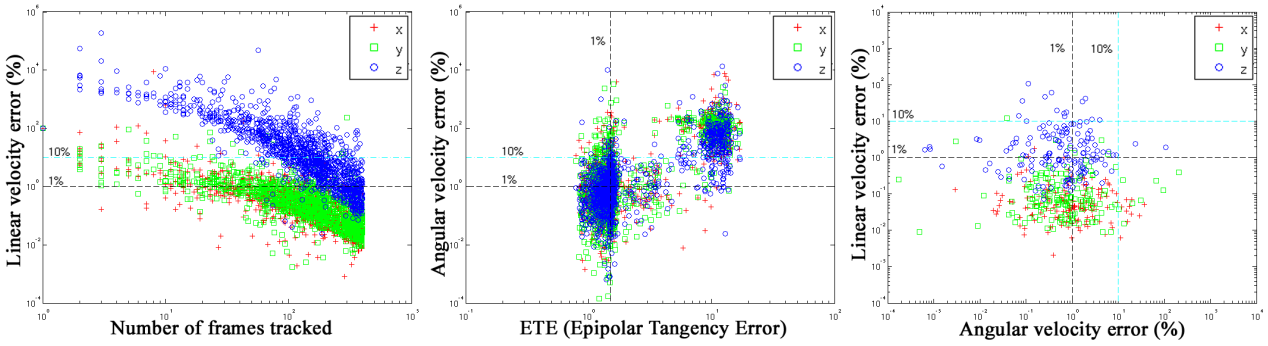

Figure 5: *Linear and angular velocity estimation error for 1200 simulated trajectories. Rightmost plot shows comparative linear and angular velocity errors for simulations where the sequence length is greater than 100 and the ETE is below 1.5 pixels.*

In Figure 5, the component velocity errors relative to the cameras' reference frames are shown. In-plane velocities

correspond to x,y values, while z represents out-of-plane velocities. Two dominant effects can be seen: The linear velocity estimation improves as the sequence length increases; and the angular velocity estimation improves as the ETE decreases.

Interestingly, even though we only make use of one high-speed camera, out-of-plane linear velocity can still be reasonably estimated when the sequence is long enough. Naturally, this will depend on the ratio between the actual particle velocity and the frame rate of the HSV camera. In our simulations a fixed frame rate of 4000fps was used and particle velocity varied between 0 and 4m/s.

While linear velocity depends on the translation of the particle in the camera, angular velocity is dependent on the difference in orientation. Therefore, a low ETE corresponding to consistent epipolar tangents tends to result in lower angular velocity error. Additionally, the existence of the second cluster on the right hand side of the middle plot illustrates cases in which insufficient angular momentum starting points were used. In our simulations, a maximum of 100 starting points was used. Therefore, greater overall consistency may be achieved by only accepting solutions once the ETE is below a specified threshold.

The previous results indicate that if the sequence is reasonably long, the particle shape is irregular and segmentation is accurate, then in most cases the linear and angular velocities of the particle can be estimated to within 1% and 10% respectively.

5.3 Particle-on-plate experiments

One application for our trajectory estimation method involves the validation of contact models used in DEM simulators in relation to real world processes. Typically, this type of validation has been accomplished using analysis of sphere impacts, since Hertzian type models are often used and accurate tracking of more complicated particle shapes is problematic. Here we show how this can be achieved for realistic particle shapes, thereby enabling a closer comparison between the predicted motion of a particle within a simulation and a real world interaction.

HSV experiments consisting of irregular diamond particles impacting an angled steel plate are used. A particle dropper is used to ensure that the particle is consistently dropped from a known height. This allows the pre-impact velocity to be verified. Sixty-view calibrated silhouette sets of the particle are used for the registration process. Trajectory estimation is applied and the solution is accepted when the overall ETE falls below 1.5 pixels. The estimated trajectory is then compared with DEM simulated trajectories of the same particle under two types of contact:

Impulse model: Based on the work of Mirtich (1996), this rigid body collision response model handles collisions of non-penetrating polyhedra by applying a series of impulses to the contacting bodies. This is sometimes referred to as the hard-sphere model, because material properties are not considered. Impulses are computed based on Stronge's hypothesis, which relates the normal work in the compression phase of the collision to the normal work in the restitution phase. This guarantees that the effects of the normal forces are purely dissipative, thereby ensuring that the addition of frictional components does not add energy to the system.

Hertz-Mindlin (no-slip) model: Building on Hertz theory, which provides solutions for the normal force-displacement under elastic spherical contact (soft-sphere model), Cattaneo, Mindlin, Deresiewicz theory (Mindlin and Deresiewicz, 1953) presents a solution in the case of varying oblique forces. Using several simple loading scenarios, tangential evolution of the contact point can be determined, resulting in improved contact modelling. In our simulations we use the simplified no-slip model as described in Di Renzo and Di Maio (2004), which typically appears in commercial DEM packages.

Figure 6 shows comparative results of the DEM simulations corresponding to each contact model versus the HSV registered trajectory. Particle shape is modelled by a sphere clump, which is generated using the method described

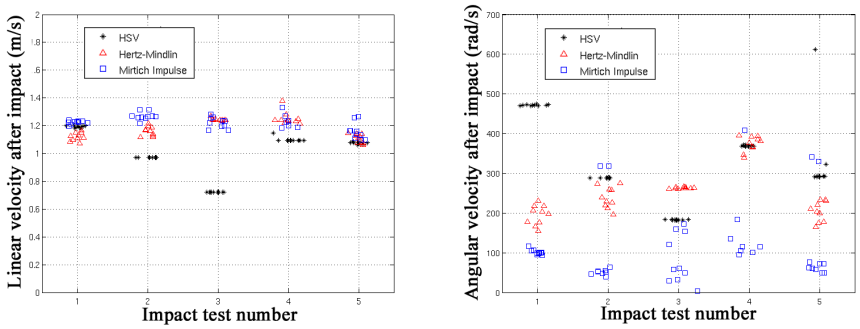

Figure 6: *Linear and angular velocities for simulations of 5 particle-plate HSV experiments. Ten runs of each experiment were processed in order to demonstrate repeatability. HSV asterisks correspond to velocities obtained from our registered trajectories. Other markers correspond to velocity from simulations using Hertz-Mindlin and Mirtich impulse contact models.*

in Price et al. (2007). For the purposes of these comparisons, we consider the generated sphere clumps to be an adequate representation of particle shape. A more complete assessment of simulation variation due to particle shape representation can be found in Price et al. (2007). Simulation parameters such as the Coefficient of Restitution and impact point are extracted from the registered trajectory and optimised in order to allow the simulation to reproduce the experiment as closely as possible. In terms of linear velocity, both models perform well. This is attributed to the material hardness of the contacting bodies and the low impact velocities, which minimises the effects of plasticity and justifies the use of a hard-sphere model. However, the Hertz-Mindlin model does appear to perform better in terms of angular velocity, and demonstrates more consistency.

6 Conclusions and future work

A method for estimating the 3D trajectory of rigid particles from high-speed video (HSV) experiments using a single camera is described. This enables quantitative feedback from real particle interactions, which can be used for process optimisation or simulation validation. In this paper a demonstrative example corresponding to comparative analysis of DEM contact models has been shown.

The use of the epipolar tangency error provides an appropriate framework for pose estimation because it incorporates all geometric information available from the images. Extensions such as texture analysis may aid estimation and boost consistency, but are unlikely to significantly improve results.

Synthetic testing has shown that the following factors affect the reliability of the method: the number of calibrated views available of particle; the amount of shape irregularity of the particle; the number of random starting points used for estimation; and the length of the HSV sequence.

Accurate estimation of out-of-plane velocities has proved difficult because only a single HSV camera is used and consistency depends on the extent to which silhouette scale changes with such motion. Distant cameras corresponding to long focal lengths create near parallel viewing rays that impair the differentiation of motion along the viewing axis. Future experimentation will explore dual camera setups.

Comparisons between our registered trajectories and identical simulations based on Hertz-Mindlin and Impulse contact models show strong correlations. Overall, Hertz-Mindlin provides more consistent results, since it accounts for the elastic nature and material properties of the contacting bodies.

References

Jules Bloomenthal. *Graphics Gems IV*, pages 324 – 349. Academic Press Professional, Inc., 1994.

Alberto Di Renzo and Francesco Paolo Di Maio. Comparison of contact-force models for the simulation of collisions in dem-based granular flow codes. *Chemical Engineering Science*, 59:525–541, 2004.

David Eberly. Numerical methods for ordinary differential equations. Technical report, Geometric Tools, Inc., March 1999. Available at `http://www.geometrictools.com`.

Keith Forbes, Anthon Voigt, and Ndimi Bodika. Using silhouette consistency constraints to build 3D models. In *Proceedings of the Fourteenth Annual South African Workshop on Pattern Recognition*, 2003.

L. Y. Li, C. Y. Wu, and C. Thornton. A theoretical model for the contact of elastoplastic bodies. *Part C: Journal of Mechanical Engineering Science*, 216:421–431, 2002.

R. D. Mindlin and H. Deresiewicz. Elastic spheres in contact under varying oblique forces. *ASME Journal of Applied Mechanics*, 20:327–344, 1953.

Brian Vincent Mirtich. *Impulse-based Dynamic Simulation of Rigid Body Systems*. PhD thesis, University of California at Berkeley, 1996.

J. J. Moré. The levenberg-marquardt algorithm: Implementation and theory. In G A Watson, editor, *Lecture Notes in Mathematics*, chapter Numerical Analysis, pages 105–116. Springer Verlag, 1977.

M. Price, V. Murariu, and G. R. Morrison. Sphere clump generation and trajectory comparison for real particles. *Submitted to DEM 2007*, 2007.

O. R. Walton and R. L. Braun. Viscosity, granular-temperature, and stress calculations for shearing assemblies of inelastic, friction discs. *Journal of Reology*, 30(5):949–980, 1986.

J. R. Williams. Superquadrics and model dynamics for discrete elements in concurrent design. Technical Report IESL 91-12, Intelligent Engineering Systems Laboratory, Massachusetts Institute of Technology, 1991.

K.-Y. K. Wong. *Structure and Motion from Silhouettes*. PhD thesis, University of Cambridge, October 2001.