

TrackSort: Predictive Tracking for Sorting Uncooperative Bulk Materials

Florian Pfaff*, Marcus Baum*, Benjamin Noack*, Uwe D. Hanebeck*,
Robin Gruna†, Thomas Längle†, and Jürgen Beyerer†

Abstract—Optical belt sorters are a versatile, state-of-the-art technology to sort bulk materials that are hard to sort based on only nonvisual properties. In this paper, we propose an extension to current optical belt sorters that involves replacing the line camera with an area camera to observe a wider field of view, allowing us to observe each particle over multiple time steps. By performing multitarget tracking, we are able to improve the prediction of each particle’s movement and thus enhance the performance of the utilized separation mechanism. We show that our approach will allow belt sorters to handle new classes of bulk materials while improving cost efficiency. Furthermore, we lay out additional extensions that are made possible by our new paradigm.

I. INTRODUCTION

The annual mass of bulk materials handled worldwide is on the rise. The transport and handling of bulk materials is a 10 billion dollar/year industry, consuming 10% of all energy produced globally [1, Ch. 1.2]. To make the best use of bulk materials and avoid unnecessary transportation costs, early and efficient sorting of bulk materials is essential. One example is separating bulk materials from unnecessary or low quality components. Bulk materials can be sorted in various ways. While a lot of bulk materials can be sorted using methods that make use of their shape (e.g., trommel screens [2]) or density, some of them are either not suited for these types of sorting schemes or even indistinguishable regarding these properties.

A very versatile scheme that can be used to even sort particles that are hard to distinguish physically are optical belt sorters. A belt sorter [3], [4], shown schematically in Fig. 1, consists of three essential parts: a conveyor belt on which the bulk material is transported, suitable sensors such as a camera, and a mechanism to separate particles of different classes. The key advantage of these sorters, which we will explain in more detail in the next section, is the variety of bulk materials that can be sorted. Most bulk materials are in some way visually distinguishable using camera technologies such as RGB or near-infrared cameras and thus allow for reliable sorting decisions based on these

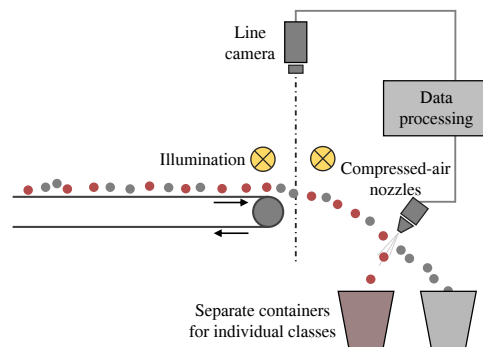


Fig. 1. Schematic drawing of an optical belt sorter.

properties. Currently, belt sorters are successfully being used for sorting (or removing foreign particles from) dried and frozen foods, industrial materials, and glass fragments for recycling, just to name a few examples.

While the current system successfully captures visual properties and can sort a variety of bulk materials appropriately, assumptions regarding the movement of the particles limit its applicability and cause conflicts in design decisions. We will go into more detail on the state of the art in Sec. II. A key component of our novel approach introduced in Sec. III is to replace the commonly used line camera with an area camera. This allows us to perform *predictive tracking*—we first track individual particles of the bulk materials for a while and then use the obtained knowledge about the particles’ movement for an accurate prediction of the particles’ future positions. *TrackSort*, the name of our project to improve belt sorters using tracking, also includes additional ideas for further improvements. Some of these ideas are also laid out in Sec. III. In Sec. IV, we show that using predictive tracking, we can predict the particles’ future positions more reliably than with the line camera approach, especially for particles whose motion strongly deviates from the simple model used in the line camera system.

II. STATE OF THE ART

In this chapter, we describe the details of current belt sorters. While we use the current state-of-the-art belt sorter of the Fraunhofer IOSB as a reference, similar optical belt sorters are also developed and distributed by other companies. Conceptually similar slide sorters also exist in which a longer slide takes the role of the belt. In our explanation in the first subsection, we regard all hardware components

*Florian Pfaff, Marcus Baum, Benjamin Noack, and Uwe D. Hanebeck are with the Intelligent Sensor-Actuator-Systems Laboratory (ISAS), Institute for Anthropomatics and Robotics, Karlsruhe Institute of Technology (KIT), Germany. florian.pfaff@kit.edu, marcus.baum@kit.edu, benjamin.noack@kit.edu, uwe.hanebeck@kit.edu

†Robin Gruna, Thomas Längle, and Jürgen Beyerer are with the Fraunhofer Institute of Optronics, System Technologies and Image Exploitation (IOSB), Germany. robin.gruna@iosb.fraunhofer.de, thomas.laengle@iosb.fraunhofer.de, juergen.beyerer@iosb.fraunhofer.de

in the order they are passed by the particles. In the second subsection, we explain how motion models can be used for sorting bulk materials and describe the simple line camera model in the third subsection. Finally, we lay out why improvements to the current state of the art are desirable in the fourth subsection.

A. Hardware Components

The bulk material starts in a container that it is filled into. Using an agitator, the bulk materials are slowly and evenly applied to a slide that leads to the conveyor belt. As we will highlight later, the conveyor belt plays a central role to modify the motion of the individual particles. At the end of the belt, the particles begin a parabolic flight. In flight or on the belt, the particles are observed by a line camera that is oriented orthogonal to the belt direction. This high speed line camera observes visual features that are used to classify the particles. Slightly above the flight path of the particles, there is an array of compressed-air nozzles lined up orthogonal to the direction of the belt and parallel to the line camera. Each compressed-air nozzle can be activated independently. This allows for specifically targeting certain particles and changing their flight path using bursts of compressed air. By changing the flight path of one category of particles, these particles will land in a container closer to the end of the belt than those that fly along their parabolic flight path unobstructed.

B. General Model-Based Sorting

The key to reliable sorting in the above described fashion is to accurately hit particles of one category. Aside from the complexity and effort needed to control the compressed-air nozzles in real-time, it is necessary to accurately specify which nozzle to activate and precisely when to activate it. This is particularly difficult as particles cannot be targeted at the exact same time they are captured by the camera. The high speed of the belt, the delays of image processing, and the delays of the nozzles necessitate that the compressed-air nozzles may not be closer than at a certain distance to the line camera.

Consequently, the decision which nozzle to activate and when to activate it has to be merely based on when and where the particle was observed by the line camera at an earlier point in time. At this point, using proper calibration, we also know the relative position to the array of nozzles. However, to precisely specify when and where the particle will pass the array of nozzles, we require a model for the motion of the particles. If the distance between the camera and the array of nozzles is short or if the particle stays on the belt, a constant velocity model is reasonable for the prediction. For a given velocity, the velocity component along the direction of the belt can be used to calculate the time at which the particle will arrive under the array of nozzles. Using this temporal offset and the velocity component orthogonal to the belt direction, the spatial offset necessary for deciding which nozzle to activate can be calculated. However, this

model-based approach relies on estimates of each particle's velocity that may only be known inaccurately.

C. Line Camera Model

The line camera model is very simple. When using the line camera as the only sensor, we cannot—or at least not without sophisticated tricks—make statements about each particle's velocity. Therefore, we need to use strong, static assumptions. Since we know the direction and speed of the belt, we can deduce the velocity of particles that are lying on the belt calmly. For the line camera model, we assume that all particles behave like this. Note that under this assumption, no orthogonal movement is accounted for and thus the compressed-air nozzle at the corresponding location can simply be activated after a fixed delay after the measurement.

D. Limitations of Line Camera Systems

The assumptions that the simple model used in line camera systems is based on are very hard to actually ensure. The belt plays a key role in approximately realizing the model's assumptions. The belt's purpose is to calm the particles, bring them all to the same known speed, and minimize orthogonal movement. If this is achieved well enough, the simple model can be used successfully. However, there are bulk materials which do not adapt sufficiently for belts with reasonable lengths. These bulk materials are called *uncooperative* and limit the applicability of the state-of-the-art systems.

Since the belt is costly and takes up space, a short belt would be preferable and having no belt at all would be the best case. However, using the line camera system, costly long belts are often necessary to ensure that the particles adapt their motion to the belt velocity accurately enough. While additional considerations—such as the shape of the particle and the air resistance—also play a role, a key component to how fast a particle adapts its speed to the belt velocity is the friction between the particle and the belt. To improve friction, the belt is often fluted although flat belts would be preferable for easier cleaning.

Beside these two disadvantages, there is a third considerable issue that more directly translates into a monetary disadvantage. Customers using belt sorters are usually trying to achieve a certain sorting quality. If the results of a single sorting run are not satisfactory, clients may need to perform multiple runs to achieve their desired sorting quality. Since those additional runs result in more energy consumption and longer occupation of the belt sorter by one batch, cost efficiency is reduced if multiple runs are needed. Therefore, it is desirable to improve the model and thus the hit ratio and consequently also the sorting quality.

One additional disadvantage of the line camera system is that no knowledge about future particles arriving at the nozzles is available. If particles of different categories are coming directly one after another, a deliberate decision not to activate a nozzle in order to spare particles of the other category cannot be done.

III. OUR NEW PARADIGM

For our new paradigm, only one minimal hardware change is necessary: we replace the line camera with an area camera. Using this little change, we can achieve significant improvements using only algorithmic innovations. In the first subsection of this section, we will lay out the predictive tracking approach that is one of the most important ideas of TrackSort. For details on how we have implemented this approach, refer to Sec. IV-B. In the second subsection of this section, we describe further possible innovations of TrackSort to improve the performance of the system.

A. The Predictive Tracking Approach

Unlike a line camera, an area camera with a sufficiently high frame rate allows us to observe particles at multiple time steps and at multiple distances from the array of compressed-air nozzles. Observing particles at multiple points in time was already proposed by Kattentidt et al. [3] for their system that uses multiple line cameras. However, they used these multiple observations only for classification and had to put a lot of effort into calibration. In our basic predictive tracking approach, we only use the multiple observations to predict the motion of the particles. Since the aim of our predictive tracking approach is to target the particles more accurately, we will disregard different particle classes from now on. However, the ability to target individual particles more accurately obviously also improves the final sorting quality.

Using appropriate image processing, we locate all particles in one frame and determine their individual centers. This yields a set of noisy, unlabeled measurements in each time step. Tracking the position and velocity of particles without any labels is a classical multitarget tracking problem. While it is a common problem and a lot of research has been done in this area [5], [6], [7], it still remains hard to deal with. The key challenge in multitarget tracking is that due to the unknown labels, the measurements cannot be directly assigned to the current tracks. However, several approaches exist that are candidates for the use in our system. Since we have to keep real-time performance in mind, we have to limit ourselves to fast algorithms.

By the use of multitarget tracking, we obtain a model that is clearly superior to the line camera model. By observing the particles at multiple points in time and performing tracking, we can reliably estimate the velocity components that we can use for the constant velocity prediction model. As we were able to validate on real data, while particles of common uncooperative bulk materials do not move straight along the direction of the belt, they still follow an approximately linear path of motion. Therefore, the predictive tracking approach is promising and we were able to confirm its good performance for uncooperative bulk materials in our evaluation in Sec. IV.

Aside from the important aspect of new bulk materials becoming feasible, there are more inherent advantages. Using stochastic modeling, we can obtain uncertainties for each particle's position at every point in time. The uncertainties can be used to give estimates about the expected hit ratio

depending on parameters of the hardware and thus facilitate hardware optimizations.

B. Ideas for Improved Classification

Beside the inherent benefits of this approach, there is a lot of potential for further improvements enabled by tracking the particles over time. One of our ideas that we have already tested with promising results is motion-based classification. How the particles behave, e.g., how fast and how much they adapt to the belt speed, varies significantly for different classes of particles. As an easy thought experiment, imagine a semisphere and a sphere. While they are visually indistinguishable on a camera image from above, their motion behavior strongly differs. However, beside this artificial example, there are a lot of potential applications. For example, blemished fruits could be detected by their motion behavior. In a proof-of-concept experiment, we were able to tell peppercorns apart from similar looking airsoft gun ammunition with high accuracy. The concept could also be made more powerful by actively provoking motion of the particles, e.g., by using low amplitude vibrations.

Another potential advance enabled by our predictive tracking approach is a substantial improvement to the visual classification capabilities. Observing the particles at multiple time steps allows using multiple measurements for the classification. This not only enables us to average out outliers, e.g., regarding RGB values, but also allows us to make deliberate use of the measurements at different positions and points in time. One way to take advantage of this is to use the multiple perspectives to obtain 3D information about the regarded particles. Furthermore, the angle or the color of the lighting could be changed as particles travel along the belt, allowing us to obtain much more visual information than when we only use a single measurement.

IV. TESTS AND EXPERIMENTAL VALIDATION

We have developed a prototype of the predictive tracking approach for offline evaluation of real data. In the first subsection, we will describe the experimental hardware setup and the recorded data. We will provide some relevant details about our sample implementation in the second subsection. In the third subsection, we will go into detail on how we evaluated the prediction performance of the line camera model and the predictive tracking model. Finally, we will present the evaluation results in the fourth subsection.

A. Experimental System and Recorded Data

To record data for our evaluation, an experimental belt sorter of the Fraunhofer IOSB was modified. It only consisted of an agitator, a slide, a belt, and an area camera. Compressed-air nozzles were not included in the experimental system as our goal was only to evaluate the prediction performance offline. The belt was configured to run at a speed of 2.7 m s^{-1} . The area camera installed was an Allied Vision Bonito CL-400 configured to record at 220 frames per second and a resolution of 2320×1500 pixels. The exposure time was 0.8 ms and an area of approximately

68 cm \times 45 cm was observed. For illumination, an array of halogen floodlights was used. For camera calibration, we took a picture of a printed chessboard pattern before starting to record.

B. Sample Implementation

In our sample implementation, we put our focus on the multitarget tracking aspect. To allow us to rapidly develop a prototype for offline processing, Matlab and Mathematica were used. For the image processing part, we first subtracted the mean over all images from each individual image to remove effects of the lighting. In the second step, we used thresholding [8, Ch. 3] to perform a basic segmentation. Afterwards, connected components labeling [8, Ch. 3] was used to determine the full shape of the individual particles. In the last step, the centroid of each particle was determined and passed on to the tracking as a noisy position measurement. Albeit simple, these easy steps serve well to convert the image data into a multitarget tracking problem.

For the multitarget tracking, we used a global association likelihood (GAL) approach [5, Ch. 10.3] for measurement-to-track association and a simple Kalman filter for tracking. For the GAL approach, we filled a matrix with the association probabilities of each measurement to each current estimate. By taking the logarithm, it is possible to solve the association problem using bipartite graph matching algorithms. We included extra columns and rows to account for appearing and disappearing targets. For the appear and disappear probabilities, we took knowledge about the scenario into account. The probability for target appearance is high at the beginning of the image and falls off to close to zero in regions that cannot be reached in one time step by new particles. Likewise, the disappear probability is high at the end and lower at the beginning.

With optimized lighting conditions, all particles could be detected reliably and with a clean belt, no erroneous measurements were observed. However, using a bad lighting configuration, we observed missed detections (caused by reflections) and using a dirty belt, we observed erroneous measurements. For additional robustness against missed detections, we added a track score approach [6, Ch. 6]. Combining the sophisticated appear and disappear probabilities with the track score approach, we were able to achieve reliable tracking even with a lot of clutter (measurements that do not originate from an actual particle). The intense, temporarily Poisson distributed, spatially uniformly distributed clutter shown in Fig. 2 was artificially added and tested successfully with our implementation. The clutter in practical applications is far less severe, but usually not uniformly distributed.

Using the obtained associations, we used a Kalman filter for every individual track. We decided to use the coordinates x , y along and orthogonal to the belt direction, the speed $|v|$, and the direction of the motion α as our state variables. Although the direction as an angle is a quantity on a periodic domain, a simple Kalman filter was used. This is suboptimal but a good approximation when the angles are far from the border of periodicity [9], which was ensured




Fig. 2. Tracks are successfully maintained despite clutter (shown as crosses that are not part of a track).

by an appropriate choice of the zero angle in our scenario. Concerning track initialization, all new tracks started with their angle pointing in the belt direction and the speed was initialized with the mean speed of the last 100 speed estimates.

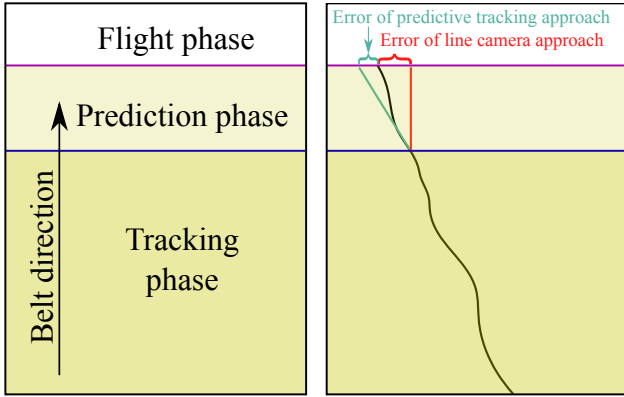
While a variety of other multitarget tracking algorithms exist, the described algorithm is easy to implement and showed promising performance. Furthermore, as use in a real-time system is intended, ease and speed of the algorithm are also key components for the choice of the algorithm.

For the planned real-time experimental platform, an FPGA will be used for image processing at high frame rates. The multitarget tracking algorithm is to be implemented using CPU and GPU computing. As there is an upper bound to the reasonable amount of particles that can be in the field of view simultaneously, optimizing the asymptotic run time is not necessary. Instead, the aim has to be to achieve real-time performance with hardware that does not significantly increase the price of the final machinery. Although solving the bipartite graph matching problem optimally is in $O(n^3)$, algorithms can be sped up by the use of GPUs [10] to allow for efficient solving for reasonable numbers of particles. Furthermore, the faster but suboptimal auction algorithm could be used. Alternatively, more sophisticated approaches such as using a gating method using priority kd -trees [7, Ch. 4] could be employed.

C. Evaluation Methodology

Our basic predictive tracking approach described in this paper only aims to improve the prediction accuracy. Therefore, we only evaluated the prediction accuracy and not the classification performance. First, we get into how we subdivided the images obtained, which is sketched in Fig. 3a. Afterwards, we describe how we derived an approximation of the ground truth for the evaluation.

For our evaluation, the flight phase starting at the end of the belt was disregarded as the drop off along the z -axis makes it harder to obtain an accurate ground truth. Please note that this only implies that we put our virtual array of compressed-air nozzles directly behind the end of the



(a) Partitioning of the images.

(b) Error calculation.

Fig. 3. Illustrations how the image is partitioned and how the errors are calculated. To allow for simpler illustration, we assume that the last measurement occurs at the border between the prediction and tracking phase.

belt. This results in a clear borderline between the flight and prediction phase which is a valid and reasonable design decision even for the real system. The subdivision of the part on the belt into a tracking and a prediction phase was to emulate the challenges in the real system. The prediction phase is the phase during which no new measurements of a particle can be accounted for anymore. Meaning, if we want to alter the flight path of a particular particle, we have to decide on which nozzle to activate based on the last estimate before the particle enters the prediction phase.

In our evaluation, we focused on the spatial prediction error. To calculate this error, we first tracked all particles over the whole field of view. Then, as sketched in Fig. 3b, we calculated the intersection of the track and the simulated array of compressed-air nozzles as an approximation of the ground truth position that we want to correctly predict. Then, we searched for the last measurement of a track before the start of the prediction phase. Using the estimated position at this time step as the starting point, the two models were compared. Shown in red is the prediction that is obtained when predicting straight along the belt direction, corresponding to the result of the line camera model. Shown in green is the prediction based on the predictive tracking approach. For this approach, the last estimate of the angle is used for the constant velocity model. The distances between the ground truth and the two predictions by the models are then seen as the prediction errors. Please note that we exaggerated the prediction error of the predictive tracking approach in Fig. 3b for visualization purposes.

D. Evaluation Results

One important benchmark for the prediction accuracy is the performance of the line camera model for bulk materials that can still be handled with sufficient accuracy using this simple model. For this, we chose the airsoft gun ammunition as one of the bulk materials to evaluate. To test our approach for bulk materials that the current system cannot properly handle, we also evaluated peppercorns, a bulk material that

we know to be uncooperative.

As the absolute delay for processing is not yet known and because the length of the prediction phase depends on the speed of the particles, we evaluated our approach for multiple sizes of the prediction phase, once for a 3 cm prediction phase and once for a 6 cm prediction phase. To show that our approach also performs well for shorter belt lengths for which the line camera model shows a significant decrease in performance, we applied our bulk materials once at 100 cm before the end of the belt and once at only 50 cm before the end of the belt. We then calculated the error between the ground truth for the position under the virtual array of compressed-air nozzles and the predictions of the line camera model and the predictive tracking approach. From the errors of approximately 100 particles in each dataset, we calculated the RMSEs shown in Tab. I. In the first row, the type of bulk material and at what distance to the end of the belt it was applied is stated. In the first column, the model used and the length of the prediction phase is denoted.

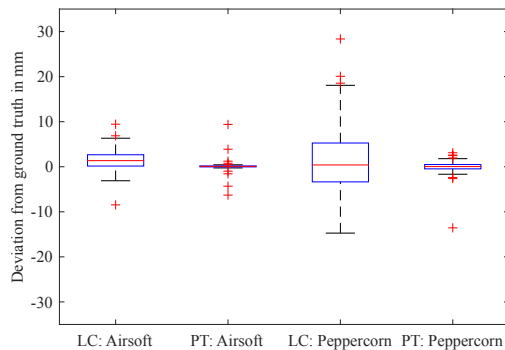
By comparing the first two entries of the column of the peppercorn dataset with 100 cm belt length to the last two entries of the column of the airsoft dataset with 100 cm belt length, we can see that using our predictive tracking approach, we obtain better prediction results for uncooperative bulk materials than are achieved by the line camera model for a bulk material that is much easier to handle. This suggests a massive improvement and that new types of bulk materials that were infeasible using the line camera model will be reliably hit using the predictive tracking approach. As can be seen by comparing the results of datasets with 100 cm belt length to the ones with 50 cm belt length, the performance of our proposed approach deteriorates only slightly (if at all) for shorter belt lengths.

For easier visual comparison of the occurring deviations, we have created boxplots for multiple datasets in Fig. 4. The box ranges from the lower quartile to the upper quartile and lengths of the whiskers are limited to the 1.5-fold of the interquartile range. Measurements with stronger deviation are regarded as outliers and are shown using red crosses. The sign of the values along the y -axis denotes the direction in which the prediction deviated from the ground truth.

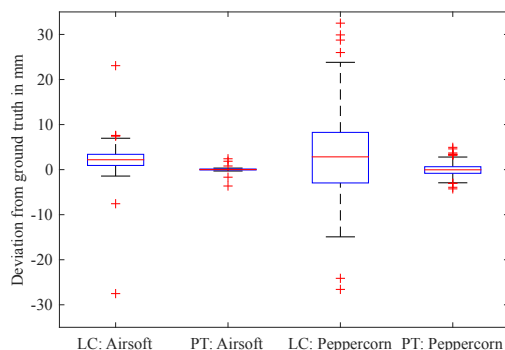
The two boxes on the left in Fig. 4a show the difference in performance for the airsoft dataset with 100 cm belt length. The interquartile range can be seen to be far smaller for the predictive tracking approach. On the right, the results for the more difficult peppercorn dataset are shown. While the performance of the line camera model is unacceptable, the performance of the predictive tracking approach on the very right is even superior to the performance of the line camera model in the easier airsoft dataset. This was a universal pattern that we observed throughout the evaluated datasets. In Fig. 4b, we show the results for the same bulk materials when a belt length of 50 cm is used. While the quality of both approaches is impeded, the predictive tracking approach handles the peppercorn dataset with a performance superior to that of the line camera model in the easiest of the shown scenarios.

	Airsoft 100 cm belt	Peppercorn 100 cm belt	Airsoft 50 cm belt	Peppercorn 100 cm belt
Predictive tracking with 3 cm prediction phase	0.54 mm	0.80 mm	0.12 mm	0.86 mm
Predictive tracking with 6 cm prediction phase	1.24 mm	1.37 mm	0.62 mm	1.57 mm
Line camera with 3 cm prediction phase	1.26 mm	3.32 mm	2.31 mm	5.10 mm
Line camera with 6 cm prediction phase	2.73 mm	7.04 mm	5.32 mm	10.63 mm

TABLE I. RMSE of the line camera and predictive tracking approaches for different belt lengths, bulk materials, and lengths of the prediction phase.



(a) Comparison for datasets with 100 cm belt length.



(b) Comparison for datasets with 50 cm belt length.

Fig. 4. Boxplots comparing the performance of the line camera approach with the predictive tracking approach for multiple datasets with a 6 cm prediction phase and differing lengths of the belt. LC stands for the line camera model, PT for the predictive tracking approach.

All in all, our results strongly imply that predictive tracking will not only allow for successfully targeting classes of bulk materials that currently cannot be hit reliably but also allow for shorter belt lengths to be used. Please note that this is only the advantage of the increased accuracy due to the predictive tracking approach, additional ideas of TrackSort are expected to further broaden the applicability by improving the classification performance.

V. CONCLUSIONS

In this paper, we have introduced a cost-efficient extension to current belt sorters whose improvements of prediction accuracy stem from mostly algorithmic enhancements. For this, we have converted the bulk material sorting problem into a multitarget tracking problem using simple image processing techniques. Then, we tested a suitable multitarget tracking algorithm and adjusted it to the problem at hand. The promising results of our offline evaluation suggest that

higher hit rates will be achieved after integrating the predictive tracking paradigm into a fully functional belt sorter. Furthermore, sorting additional kinds of bulk materials will become feasible and shortening the belt will be possible, improving applicability and cost-efficiency.

Additionally, we described further potential enhancements that are enabled by performing tracking. The significant potential benefits for the classification performance enabled by multiple observations and motion-based features will help turn optical belt sorters into an even more versatile sorting scheme.

Exploring the potential benefits for the classification is one area of future work. There are also a variety of ways to improve the predictive tracking approach such as using IMM for more accurate modeling. Another area of future work on TrackSort is the real-time implementation on fully functional experimental hardware. While we know that a real-time implementation is feasible, the time constraints will make this a non-trivial and time consuming—but definitely rewarding—task.

ACKNOWLEDGMENTS

The authors would like to thank their colleague Dennis Heddendorp and all past and present students that have assisted them in this project, namely Frederick Doll, Juan Hus-sain, Christian Tesch, Sebastian Dingler, Robin Sandkühler, and Yassine Marrakchi.

REFERENCES

- [1] J. Duran, *Sands, Powders, and Grains: An Introduction to the Physics of Granular Materials*. Springer, New York, 2000.
- [2] Y.-S. Chen, S.-S. Hsiau, H.-Y. Lee, Y.-P. Chyou, and C.-J. Hsu, “Size separation of particulates in a trommel screen system,” *Chemical Engineering and Processing: Process Intensification*, vol. 49, no. 11, 2010.
- [3] H. Kattentidt, T. De Jong, and W. Dalmijn, “Multi-sensor identification and sorting of bulk solids,” *Control engineering practice*, vol. 11, no. 1, 2003.
- [4] T. Längle, “Quality control of bulk material in factories of the future,” in *23rd ISPE International Conference on CAD/CAM, Robotics and Factories of the Future - Cars & FOF’07*, 2007.
- [5] R. P. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Artech House, Inc., 2007.
- [6] S. Blackman and R. Popoli, *Design and Analysis of Modern Tracking Systems*, 1999.
- [7] M. Liggins II, D. Hall, and J. Llinas, *Handbook of Multisensor Data Fusion: Theory and Practice*, 2nd ed. CRC press, 2009.
- [8] L. Shapiro and G. Stockman, *Computer Vision*. Prentice Hall, 2001.
- [9] G. Kurz, “Directional Estimation for Robotic Beating Heart Surgery,” Ph.D. dissertation, Karlsruhe Institute of Technology, Intelligent Sensor-Actuator-Systems Laboratory, Karlsruhe, Germany, 2015.
- [10] C. N. Vasconcelos and B. Rosenhahn, “Bipartite graph matching computation on GPU,” 2009.