

TAMPAR: Visual Tampering Detection for Parcel Logistics in Postal Supply Chains

Alexander Naumann
FZI and KIT
Karlsruhe, Germany
anaumann@fzi.de

Felix Hertlein
FZI and KIT
Karlsruhe, Germany
hertlein@fzi.de

Laura Dörr
FZI and KIT
Karlsruhe, Germany
doerr@fzi.de

Kai Furmans
FZI and KIT
Karlsruhe, Germany
furmans@kit.edu

Abstract

Due to the steadily rising amount of valuable goods in supply chains, tampering detection for parcels is becoming increasingly important. In this work, we focus on the use-case last-mile delivery, where only a single RGB image is taken and compared against a reference from an existing database to detect potential appearance changes that indicate tampering. We propose a tampering detection pipeline that utilizes keypoint detection to identify the eight corner points of a parcel. This permits applying a perspective transformation to create normalized fronto-parallel views for each visible parcel side surface. These viewpoint-invariant parcel side surface representations facilitate the identification of signs of tampering on parcels within the supply chain, since they reduce the problem to parcel side surface matching with pair-wise appearance change detection. Experiments with multiple classical and deep learning-based change detection approaches are performed on our newly collected TAMpering detection dataset for PARcels, called TAMPAR. We evaluate keypoint and change detection separately, as well as in a unified system for tampering detection. Our evaluation shows promising results for keypoint (Keypoint AP 75.76) and tampering detection (81% accuracy, F1-Score 0.83) on real images. Furthermore, a sensitivity analysis for tampering types, lens distortion and viewing angles is presented. Code and dataset are available at <https://a-nau.github.io/tampar>.

1. Introduction

The amount of valuable goods within postal supply chains is increasing steadily [1]. This trend implies the rising importance of safety and security considerations in transportation networks. One crucial aspect to improve safety and security along the supply chain is checking

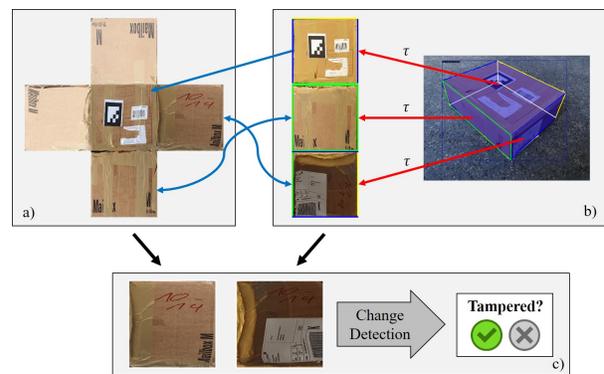


Figure 1. We detect tampering by comparing the full parcel texture from a database (a) with the viewpoint-invariant parcel side surfaces of a single image by exploiting parcel corner point predictions (b). Appearance change detection is performed for each pair of matching parcel side surfaces to identify tampering (c).

parcels for damages or signs of tampering [2]. Tampering detection, on which we focus in this work, tries to verify and guarantee the integrity of a parcel within the supply chain. Common cases of potential tampering are applying new tape after opening a parcel or attaching labels to or removing labels from a parcel. Of course, not all cases where such changes occur are relevant for safety and security considerations, e.g. new tape might be applied to prevent items from falling out of a damaged parcel.

In general, tampering detection for parcels requires a two-step pipeline: (1) We need to reliably identify a parcel by either its shipping label or unique visual cues on the packaging. Especially in scenarios with numerous visually similar parcels, using the latter can be challenging, while at the same time, the shipping label might not always be visible. (2) We need to compare the appearance of the packaging across time to detect changes that might stem from tampering. This is a challenging task since the photos used

for the comparison can show the objects of interest from different viewing angles and under varying lighting conditions.

In this work, we focus on step (2), since parcel re-identification has been studied by Clausen *et al.* [3], and present an approach for tampering detection for already identified parcels (cf. Fig. 1). Since last-mile delivery is considered as use-case, we assume that only a single RGB image is available, which should be compared against a reference from an existing database. Similar to Noceti *et al.* [1], appearance change detection is performed separately per parcel side surface, and if at least one of the parcel side surfaces has been tampered with, the whole parcel is considered to have undergone tampering. To tackle this problem, we suggest the usage of keypoint detection to identify the parcel corners as a first step toward change detection. Knowledge of the eight corner points of cuboid-shaped parcels enables computing normalized fronto-parallel views of all visible parcel side surfaces by applying a perspective transformation τ (cf. Fig. 1 (b)). These views eliminate the viewpoint variance and thus, facilitate change detection and potentially also re-identification of parcel side surfaces. We use the Parcel3D dataset [2] to demonstrate the capabilities of keypoint detectors for generating viewpoint-invariant parcel side surface representations from single RGB images. Additionally, we collected a novel dataset for change detection in postal supply chains, and present a detailed analysis of change detection approaches for tampering detection. The main contributions of our work are as follows

- we suggest an effective keypoint ordering for parcel detection and side surface segmentation,
- we present TAMPAR, a novel dataset for TAMpering detection of PARcels, and
- we propose and evaluate an approach for tampering detection, which exploits keypoint and change detection.

2. Related Work

We review related literature in logistics applications, 3D bounding box detection, keypoint estimation and change detection in the following.

Applications in Logistics. Karaca *et al.* [4] present an early approach using a stereo camera and feature matching techniques to track parcels along a conveyor belt. Clausen *et al.* [3] present an approach for tracking parcels inside a logistics facility. A Mask R-CNN-based [5] Siamese network [6] complemented with their so-called feature improver head is used to re-identify parcels. They create a manually labeled dataset of 3,306 images taken by 37 different cameras with a total of 14,248 parcels. The evaluation shows that currently around 81 % of parcels are tracked correctly.

For more details on literature regarding re-identification we refer to Ye *et al.* [7] and Khan *et al.* [8]. Naumann *et al.* [9] work towards parcel side surface segmentation. By combining plane segmentation [10] and contour detection [11], [12], they present an approach to refine parcel side surface segmentation masks without relying on any task-specific training data. Naumann *et al.* [2] tackle the problem of estimating the 3D shape of potentially damaged parcels from a single RGB input. They extend Cube R-CNN [13] by an iterative mesh refinement [14] and present Parcel3D, a dataset comprising over 13,000 images of cuboid-shaped and damaged parcels with full 3D annotations. Noceti *et al.* [1] present a multi-camera system for damage and tampering detection in postal supply chains. Damages are detected by finding the parallelepiped which best aligns with the captured images. For tampering detection a Histogram of Oriented Gradients (HOG) [15] for the parcel side surfaces is used. Rotation invariance is accomplished by considering all possible rotations with histogram intersection as similarity measure. Tampering is reported when the similarity of two feature vectors is below a certain threshold. Other works focusing on parcels consider synthetic training data generation [16], tracking inside a moving truck [17] and depalletization [18], [19]. Finally, Naumann *et al.* [20] present a detailed overview of computer vision applications in transportation logistics and warehousing.

3D Bounding Box Detection. Dwibedi *et al.* [21] present an early deep learning-based approach to estimate the 3D bounding box of cuboid-shaped objects. Generally, 3D bounding box detection is a common task for autonomous driving [22]. Approaches often rely on only estimating yaw, since they can exploit the fact that vehicles are driving on the road. Li *et al.* [23] exploit 2D/3D correspondences by estimating keypoints of cars to improve 3D bounding box detection. Rui *et al.* [24] introduce a framework for vehicle recognition from a single RGB image. They estimate a 3D bounding box which is used to compute normalized views for the front, side and roof view of a car by applying a perspective transformation. This information is fused with region-aligned features of the respective region of interest to estimate the vehicle model.

Keypoint Detection. Lots of research tackling keypoint estimation considers monocular human pose estimation, which is reviewed by Chen *et al.* [25] and Chen *et al.* [26]. Dörr *et al.* [27] treat the problem of packaging structure recognition. The goal is to identify the number, type and arrangement of small load carriers on a uniformly packed transport unit from a single RGB image. They extend CornerNet [28] to leverage keypoint estimation to detect objects based on four arbitrary corner points.

Change Detection. To detect signs of tampering, after re-identification, change detection is necessary. Change detection is most commonly applied for remote-sensing and street views and reviewed by Shi *et al.* [29]. A dataset for change detection in industrial environments has been presented by Park *et al.* [30]. Furthermore, Park *et al.* [31] propose the novel change detection approach SimSaC which is targeted towards industrial use-cases. SimSaC relies on dual task learning and exploits both, dense correspondence and mis-correspondence to increase robustness when encountering imperfect matches.

While Noceti *et al.* [1] also tackle the problem of tampering detection, they focus on a constrained environment with calibrated background, constant illumination and a multi-sensory setup. In contrast to that, our approach does not have any such constraints and relies just on a single RGB image as input. Consequently, ours is the only approach suitable for scenarios such as last-mile delivery and cannot be fairly compared to the work by Noceti *et al.* [1]. Furthermore, we rely on existing keypoint and change detection approaches and strive to combine them efficiently, however, we do not aim to develop novel approaches in these areas.

3. Approach

We present our approach for parcel keypoint detection in Sec. 3.1 and for change detection in Sec. 3.2. Details on our novel dataset TAMPAR are given in Sec. 3.3.

3.1. Parcel Keypoint Detection

We use a Mask R-CNN [5] with keypoint head and a ResNet-50-FPN [32], [33] backbone for our experiments. This choice is motivated by the fact, that we do not focus on improving keypoint detection techniques, but rather want to demonstrate the usefulness of well-established baselines for the use-case of parcel corner detection.

One key challenge for this use-case is to identify an unambiguous keypoint ordering which works well with Artificial Neural Networks since there are several options for ordering keypoints of a parcel. In contrast to the common application of 3D bounding box detection for autonomous driving, where vehicles have a well-defined front and back side, there is no such notion for parcels. To have a consistent, unambiguous keypoint ordering with explicit visual cues, we proceed as follows. We assume, that three parcel side surfaces are visible in each image and define the front of a parcel by choosing the visible parcel side surface whose normal aligns best with a left- and front-facing vector, i.e. $(x, y, z) = (1, 0, -0.5)$. From this, we derive our keypoint ordering definition, which is visualized in Fig. 2 and described in the following. We denote the number of visible α and invisible β parcel side surfaces that intersect

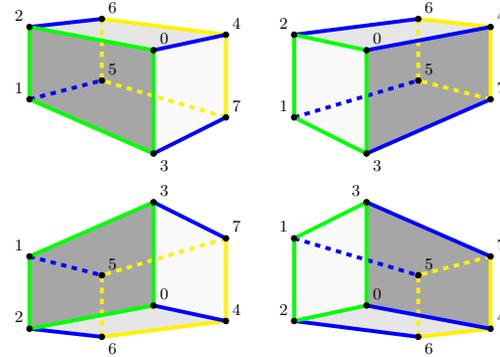


Figure 2. Visualization of the consistent and unambiguous keypoint ordering for a cuboid without well-defined front and back across different viewing angles. We highlight the front side in green and the back side in yellow.

in keypoint k_i as $k_i=(\alpha, \beta)$. On the front side (highlighted in green in Fig. 2), we define the keypoints:

- $k_0 = (3, 0)$: point of intersection of the three visible parcel side surfaces, which is located inside the convex hull of the parcel.
- $k_1=(1, 2)$: joint point of the two invisible parcel side surfaces, where only two visible parcel edges intersect.
- $k_2=(2, 1)$, leftmost: leftmost point of the remaining points, where three visible parcel edges intersect.
- $k_3=(2, 1)$, rightmost: remaining point, which is the rightmost point that belongs to two visible parcel side surfaces and one invisible one.

The backside (highlighted in yellow in Fig. 2) of the parcel is the one across from the front side, and we define the keypoint order as follows:

- $k_4=(2, 1)$: point that is part of two visible parcel side surfaces and thus, at this point three visible parcel edges intersect.
- $k_5=(0, 3)$: self-occluded keypoint, which is the point of intersection of the three invisible parcel side surfaces.
- $k_6=(1, 2)$, leftmost: leftmost point of the remaining points, where two visible parcel edges intersect.
- $k_7=(1, 2)$, rightmost: remaining point, which is the rightmost point where two visible edges intersect.

This keypoint ordering is used for training and evaluating corner point detection in the following. Note, that it is not invariant to horizontal, but only to vertical flipping of the image. Furthermore, technically, estimating seven keypoints would be sufficient to infer all eight, however, we want to show that the estimation even works for the self-occluded keypoint k_5 . The information on the seven visible keypoints can be utilized to compute viewpoint-invariant

parcel side surface representations by applying a perspective transformation. This, in turn, enables the composition of parcel texture mappings as visualized in Fig. 1 (a).

3.2. Change Detection

In our use-case, we assume that a postman takes a single image of a parcel which seems suspicious of potential tampering. First, the parcel keypoints are extracted and the viewpoint-invariant parcel side surfaces of size 400×400 pixels are computed as described in Sec. 3.1 and visualized in Fig. 1 (b). By exploiting this information, we can reduce the task of tampering detection of parcels to comparing fronto-parallel parcel side surface representations. If one parcel side surface has been tampered with, the parcel is considered tampered.

While the usage of viewpoint-invariant representations alleviates the problem of perspective distortion, change detection remains challenging since image alignment issues cannot fully be resolved, and additionally, the lighting might vary significantly (cf. Fig. 1 (c)). To cope with these issues, we use SimSaC [31]. SimSaC is a recent approach for robust change detection with imperfect matches. It estimates scene flow using correspondence maps at the same time as change masks by exploiting mis-correspondences. This enables robustness against geometric transformations and differences in lighting.

We benchmark SimSaC against several baselines, each combining an image homogenization approach and a similarity metric. For image homogenization, we utilize (cf. Fig. 3)

- *DexiNed*: Dense EXtreme Inception Network for Edge Detection [12]
- *Canny*: Adaptive Canny edge detection [11], [34]
- *Laplacian*: Laplacian filter
- *Mean Channel*: Per-channel mean alignment

As image similarity metrics, we consider

- Learned Perceptual Image Patch Similarity (LPIPS) [35],
- Structural Similarity (SSIM) [36],
- Multiscale Structural Similarity (MS-SSIM) [37],
- Complex Wavelet Structural Similarity (CW-SSIM) [38],
- HOG [15] feature similarity¹, and
- Mean Absolute Error (MAE).

A change is detected when the input and reference parcel side surface image after applying the image homogenization to both, have a low image similarity. Suitable thresholds for image similarity will be determined in Sec. 4.2.

¹We use 9 orientation bins, 8 pixels per cell and 2 cells per block.

3.3. Dataset

Our dataset resembles a use-case with multisensory setups within logistics facilities and a simple cell phone camera during the last-mile delivery. More precisely, we assume that multiple cameras are used to capture and segment all five visible parcel side surfaces in logistics facilities. Note that we also suppose that the side surface with the unique identifier is always visible, which means that the opposing side surface is never visible. Subsequently, the parcel ID and texture map, as visualized in Fig. 1 (a), are saved to a database. Finally, a single RGB image of a parcel with suspected tampering is taken during last-mile delivery and compared against its high-quality reference texture.

To generate a suitable dataset for this use-case, we proceed as follows. We use ArUco markers to uniquely identify parcels and the spatial relationships between their side surfaces. The parcel textures for the database are generated by taking several images of the parcel in its original state, i.e. without tampering. By manually labeling the parcel corner points, we automatically generate the full parcel texture by applying perspective transformations. Subsequently, we apply different types of tampering to three out of the five relevant parcel side surfaces. While real-world tampering attempts focus on a single parcel side surface, our dataset design enables a more diverse analysis of tampering detection by considering a larger number of tampering examples. As mentioned before, transferring side surface tampering to the parcel level is straightforward. We consider three different types of tampering flags, each with an *easy* and a *hard* to detect variant:

- *Label*: Adding a new shipping label (*easy*) or transportation hints (*hard*)
- *Tape*: Adding new tape, which covers more than 50% of the longer side (*easy*), or less than 25% of the shorter side (*hard*)
- *Writing*: Adding manually written text, using a pen with 5-15 mm (*easy*) or 1.5 – 3 mm of width (*hard*)

Note that adding written text usually would not be considered tampering. However, we strive to detect diverse appearance changes to reliably flag parcels for manual inspection. In total, we collect and annotate 296 images of 10 parcels for the training/validation and 614 images of 20 parcels for the test set. Since each image contains three visible parcel side surfaces, TAMPAR comprises 888 images for training/validation and 1842 images for testing change detection. The main difference to existing datasets such as Parcel3D [2] and Parcel2D Real [16] is that we have paired images of the same parcel across different points in time, i.e. before and after tampering.

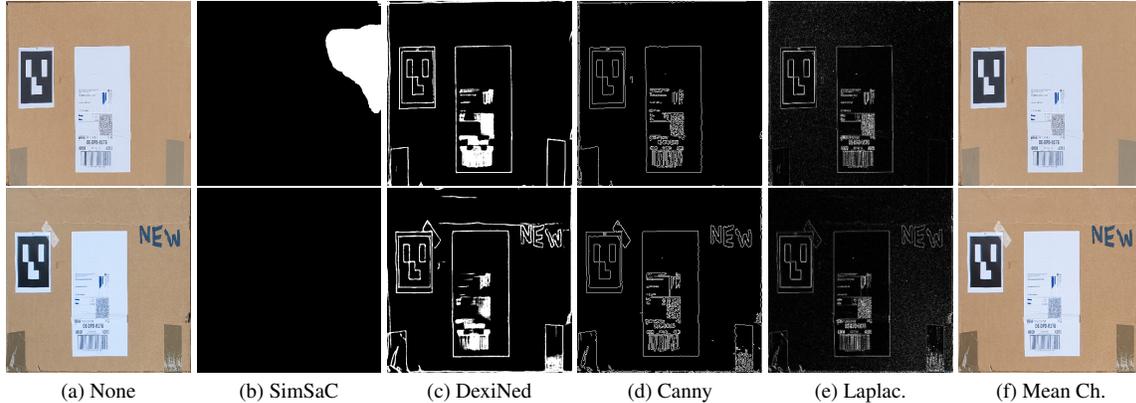


Figure 3. Examples of the different image homogenization methods before (top) and after (bottom) tampering. Note that SimSaC [31] is the only approach that localizes potential tampering and directly outputs change maps.

4. Evaluation

We first evaluate keypoint detection for parcel corners separately in Sec. 4.1. Subsequently, we evaluate the considered change detection approaches isolated and in combination with keypoint estimation in Sec. 4.2. Furthermore, we present a sensitivity analysis on the influence of the tampering type, lens distortion and viewing angles.

4.1. Parcel Corner Point Estimation

For all experiments, we use a ResNet-50-FPN [32], [33] that was pre-trained on MS COCO [39] as backbone and freeze its weights at stage four. We use Stochastic Gradient Descent with Momentum (SGD+M) with a batch size of 16 and a cosine learning rate schedule [40]. The initial learning rate is set to 0.001 and the final learning rate to 0 after 10 000 iterations. Moreover, a linear warm-up during the first 1000 iterations is applied.

Training is always performed on the synthetic dataset Parcel3D [2] which contains cuboid-shaped and damaged parcel images. For the evaluation, we consider synthetic and real-world data in the following. We evaluate bounding box detection, instance segmentation and keypoint detection, and summarize the quantitative results in Tab. 1.

For the evaluation of keypoint detection using Keypoint AP², it is necessary to define κ_i for each keypoint. This value is usually obtained by comparing redundantly annotated images to infer each keypoints' annotation precision. Since no redundantly annotated images are available, we select $\kappa_5 = 0.1$ for the self-occluded and $\kappa_i = 0.05$, $i \in \{0, 1, 2, 3, 4, 6, 7\}$ for the visible keypoints, which is close to the κ for human hips (0.107) and human wrists (0.062), respectively [39]. We argue that human wrists are a suitable approximation because the keypoints for Parcel3D and Parcel2D Real are computed from 3D

bounding boxes, frequently leading to a misalignment between the parcel corners in the image and the actual annotated keypoints. This misalignment is also present for damaged parcels, where the keypoints correspond to the ones of the pristine version of the parcel.

4.1.1 Synthetic Data

The quantitative results from Tab. 1 indicate excellent performance for bounding box detection and instance segmentation, with a Box AP of 93.62 and a Mask AP of 97.54. Likewise, keypoint detection achieves strong results with a Keypoint AP of 88.80.

Qualitative examples are presented in Fig. 4. For intact, i.e. cuboid-shaped, parcels keypoint detection enables computing high-quality fronto-parallel views of the parcel side surfaces as can be seen in Fig. 4a. Strong distortions of parcel side surface views, however, cannot be recovered and lead to low-quality representations, which are challenging to use for tampering detection. In the case of damaged parcels, the representations' quality strongly depends on the degree of deformation (cf. Fig. 4b). Strong deformations also impede tampering detection. Problematic cases can include imprecise or missing detections (cf. Fig. 4c).

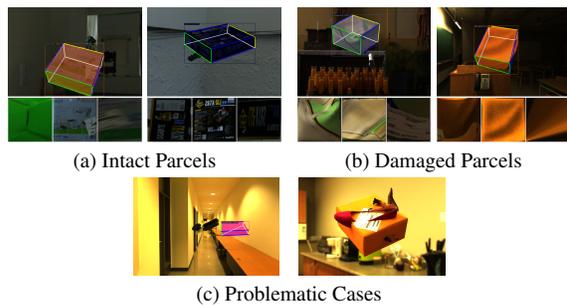


Figure 4. Exemplary qualitative results for synthetic parcels.

²See <https://cocodataset.org/#keypoints-eval> for details.

Dataset	Box		Mask		Keypoint	
	AP	AP ₇₅	AP	AP ₇₅	AP	AP ₇₅
Parcel3D	93.62 (0.1)	98.46 (0.2)	97.54 (0.2)	98.58 (0.3)	88.80 (0.2)	94.06 (0.2)
Parcel2D Real	84.88 (0.2)	97.28 (0.1)	85.02 (0.2)	96.92 (0.6)	75.76 (0.5)	85.36 (1.2)
TAMPAR (ours)	96.38 (0.2)	99.72 (0.5)	98.94 (0.2)	99.70 (0.5)	97.18 (0.5)	99.12 (0.4)

Table 1. Quantitative performance analysis of the ResNet-50-FPN for bounding box detection, instance segmentation and keypoint detection. We repeated all trainings five times and report *mean (standard deviation)*.

4.1.2 Real Data

Due to the fact that training was only performed on the synthetic training dataset Parcel3D [2], a domain gap occurs when evaluating on the two real-world datasets Parcel2D Real [16] and TAMPAR. This domain gap manifests itself in the generally lower performance on Parcel2D Real compared to the evaluation on synthetic data, as seen in Tab. 1. At the same time, performance on TAMPAR is higher, presumably due to the simpler nature of the dataset - all images are high-quality and show only a single parcel in the center. Performance for bounding box detection and instance segmentation remains high with a Box AP of 84.88/96.38 and a Mask AP of 85.02/98.94, on Parcel2D Real and TAMPAR, respectively. The same holds true for the performance of keypoint detection, which reaches 75.76 and 97.18 Keypoint AP.

Quantitative inspection of the prediction results confirms the suitability of Parcel3D and our proposed keypoint ordering. Especially for cuboid-shaped parcels, as visualized in Fig. 5a, results look very promising for applications in tampering detection. Furthermore, we evaluate our approach on images of damaged parcels without ground truth annotations (cf. Fig. 5b). These qualitative impressions also underpin the suitability of our approach, however, detecting keypoints accurately for damaged parcels seems to be a more difficult task. Examples of failed detections include missing and imprecise keypoint localizations, as visualized in Fig. 5c.

4.1.3 Sensitivity Analysis: Lens Distortion

We investigate the influence of barrel distortion according to

$$r_{\text{src}} = r_{\text{dist}} \cdot (A \cdot r_{\text{dist}}^3 + B \cdot r_{\text{dist}}^2 + C \cdot r_{\text{dist}} + D)$$

with r_{src} being the radial distance from the image center in the input image, and r_{dist} the one, in the distorted output. We analyze six different settings, which are visualized in Fig. 6 by creating distorted dataset versions with parameter $A \in [-0.08, -0.04, -0.02, 0.04, 0.08, 0.16]$, $B = 0$, $C = 0$, and $D = 1.0$. Note that these datasets can be smaller in size, since we discard instances if the distortion

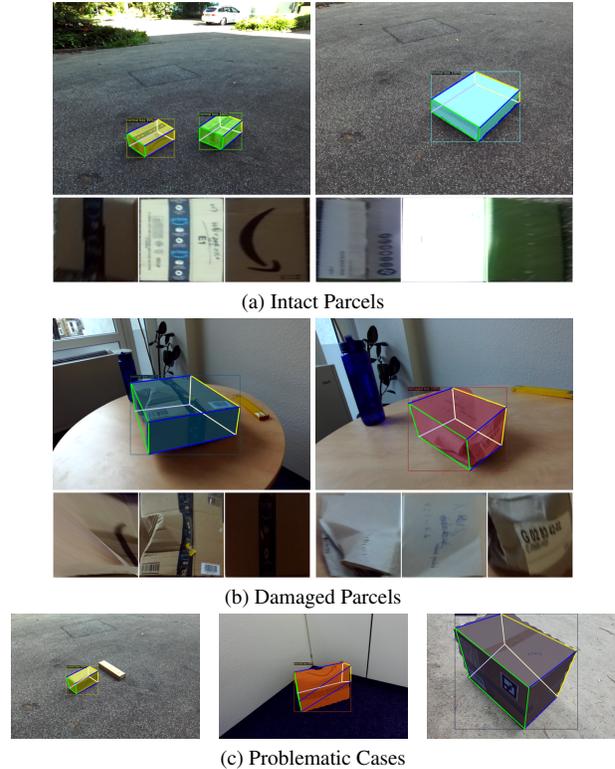


Figure 5. Exemplary qualitative results for real parcels.

corrupted the annotations (e.g. keypoints lie outside the image) or the ArUco detection. Results in Fig. 7 indicate that instance segmentation performance is robust w.r.t. distortion effects. While keypoint detection only degrades for pincushion distortions ($A < 0$), bounding box detection results are also affected by strong barrel distortions ($A > 0$).

4.2. Tampering Detection

We evaluate change detection in isolation by using the ground truth keypoint annotations, as well as in a combined system on our novel dataset TAMPAR by using keypoint predictions from Sec. 4.1. Furthermore, we analyze the influences of tampering types, lens distortion, and viewing angles.

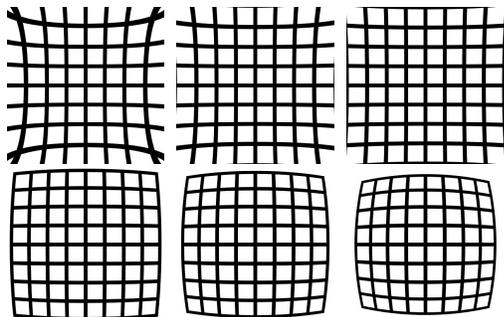


Figure 6. Visualization of the investigated distortion effects with parameter $A \in [-0.08, -0.04, -0.02, 0.04, 0.08, 0.16]$, $B = 0$, $C = 0$, and $D = 1.0$.

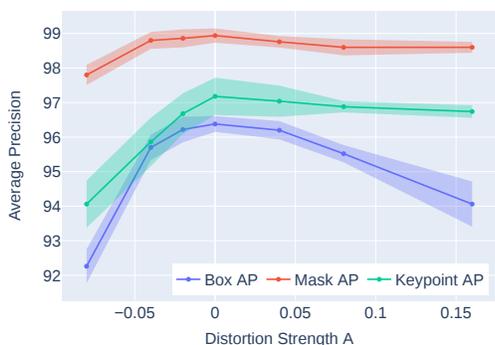


Figure 7. Quantitative performance analysis of the ResNet-50-FPN under different types of lens distortion. We repeated all trainings five times and report mean values with standard deviations as error boundaries.

4.2.1 Pipeline Evaluation

Considering the input and reference parcel we first perform marker-based side surface matching. Subsequently, we apply the image homogenization to both (input and reference side surface view) and compute their image similarity using all metrics mentioned in Sec. 3.2. We denote the combination of image homogenization *Method A* and similarity *Metric B* as (*Method A*, *Metric B*), and seek to determine the best image similarity metrics and corresponding thresholds. SimSaC [31] poses a special case since it uses the input and reference image to output change maps. This enables localization of tampering, which is advantageous in practice, however, not evaluated in this work. Instead, we compare the binary change map to a black image (i.e. the change map corresponding to no changes) to compute image similarity. We summarize the evaluation results using simple thresholding by training a decision tree of depth one per method using all similarity metrics as input in Tab. 2.

Results in Tab. 2 using predicted keypoints show that (*SimSaC*, *LPIPS*) yields the best performance and reaches 0.81 accuracy and an F1-Score of 0.83. The by far highest precision is also achieved by (*SimSaC*, *LPIPS*), which indicates cautious change detection for our use-case. The highest recall is reached by (*DexiNed*, *HOG*) and (*Mean Ch.*, *LPIPS*), however, at the cost of precision. Performance differences between using predicted and ground truth keypoint positions are comparatively small due to the high accuracy of the keypoint detection (cf. Tab. 1).

4.2.2 Sensitivity Analysis: Tampering Types

The analysis of performance differences across tampering types in Tab. 3, shows that *labels* can be detected most reliably, while *tape* and especially *writing (hard)* are more difficult to recognize. Surprisingly, when detecting *writing* performance deteriorates when using ground truth keypoint annotations. One potential reason for this might be, that inaccurate keypoints enlarge the region of interest unproportionally.

4.2.3 Sensitivity Analysis: Lens Distortion

We analyze the influence of six different degrees of distortion (cf. Fig. 6) on the tampering detection quality using predicted keypoints and (*SimSaC*, *LPIPS*). These distortions imply that our simple perspective transformation cannot accurately create normalized side surface views and the change detection approach needs to handle these inaccuracies. The results in Fig. 8 suggest robustness w.r.t. distortions, with a slight negative trend for distortions with distortion strength $A > 0$. This is in line with the fact, that our approach can cope with lens distortion effects across the two real-world dataset TAMPAR and Parcel2D Real, while being trained on different, synthetic data.

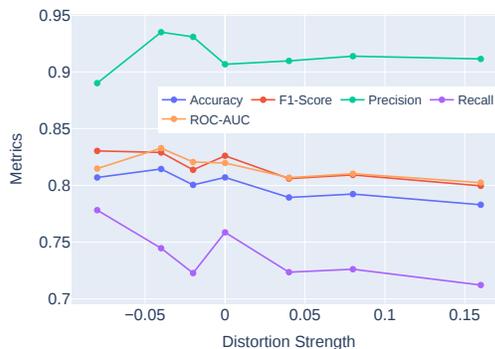


Figure 8. Sensitivity analysis for tampering detection w.r.t. to the distortion strength A using pred. keypoints and (*SimSaC*, *LPIPS*).

Method	Metric	Accuracy	Precision	Recall	F1-Score	ROC-AUC
None	LPIPS/MS-SSIM	0.66/0.65	0.66/0.65	0.91/0.93	0.76/0.76	0.60/0.58
SimSaC	LPIPS/MAE	0.81/0.80	0.91/0.93	0.76/0.72	0.83/0.81	0.82/0.82
DexiNed	HOG/SSIM	0.60/0.62	0.60/0.63	1.00/0.91	0.75/0.74	0.48/0.54
Canny	MS-SSIM/SSIM	0.60/0.60	0.62/0.61	0.91/0.91	0.74/0.73	0.52/0.52
Laplacian	LPIPS/LPIPS	0.65/0.68	0.71/0.71	0.72/0.80	0.71/0.75	0.64/0.65
Mean Ch.	LPIPS/MS-SSIM	0.63/0.65	0.62/0.65	0.99/ 0.94	0.76/0.77	0.53/0.58

Table 2. Quantitative performance analysis of the tampering detection using a decision tree with depth one. The metric indicates the selection for thresholding during the training of the decision tree. We report metric names and scores for *predicted* / *ground truth* keypoints.

Tampering Type	Label		Tape		Writing	
	easy	hard	easy	hard	easy	hard
Number of Samples	606	570	462	546	624	498
Recall (Pred. Keypoints)	1.00	1.00	0.58	0.48	0.87	0.52
Recall (GT Keypoints)	1.00	0.99	0.59	0.49	0.80	0.36

Table 3. Sensitivity analysis on the performance differences across tampering types using (*SimSaC*, *LPIPS*).

4.2.4 Sensitivity Analysis: Viewing Angles

We approximate the viewing angle per parcel side surface by considering the angle between the x- and y-axis, and the polygon spanned by the four side surface corner points. No clear trend emerges from this analysis in Fig. 9, which suggests that our approach is robust w.r.t. a reasonable spectrum of viewing angles. Note, however, that TAMPAR does not feature extreme viewing angles. Due to the strong distortions under such viewing angles, we expect the performance of tampering detection to degrade heavily.

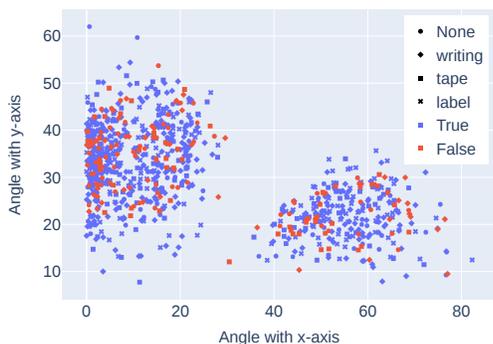


Figure 9. Sensitivity analysis for tampering detection w.r.t. to the viewing angle per side surface using predicted keypoints and (*SimSaC*, *LPIPS*). Tampering types are encoded with different geometries and the prediction correctness with color-coding.

5. Conclusion

In this work, we propose a tampering detection pipeline for parcels that leverages keypoint and change detection. We utilize the parcel keypoints to compute viewpoint-invariant parcel side surfaces, which reduces the problem of tampering detection to (1) identifying identical parcels and their side surfaces across time and (2) applying change detection on them. We propose an unambiguous keypoint ordering for parcels and evaluate well-established baseline algorithms for the task of parcel corner point detection on real-world data. Our approach reaches 75.76 and 97.18 Keypoint AP on two real-world datasets when trained only on synthetic images from Parcel3D [2]. Moreover, we introduce the first publicly available benchmark for tampering detection on parcels called TAMPAR. We propose a systematic approach for tampering detection which combines image homogenization approaches to alleviate lighting differences with several image similarity metrics. To make a prediction, the most suitable metric is chosen and simple thresholding is applied. Our approach performs best when combining SimSaC [31] with LPIPS [35] and reaches 0.81 accuracy and an F1-Score of 0.83. The additional sensitivity analysis w.r.t. tampering types, lens distortion and viewing angles demonstrates the robustness of our approach.

Future work can incorporate recent advances in keypoint detectors and also exploit shape priors by utilizing a vanishing point loss [24] or enforcing 2D/3D consistency [23]. Furthermore, fine-tuning SimSaC [31] for tampering detection and more complex decision rules such as random forests are expected to improve performance significantly provided suitable datasets are available.

References

- [1] N. Noceti, L. Zini, and F. Odone, "A multi-camera system for damage and tampering detection in a postal security framework," *EURASIP Journal on Image and Video Processing*, vol. 2018, no. 1, p. 11, 2018.
- [2] A. Naumann, F. Hertlein, L. Dörr, and K. Furmans, "Parcel3D: Shape Reconstruction from Single RGB Images for Applications in Transportation Logistics," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops, 2023*, pp. 4402–4412.
- [3] S. Clausen, C. Zelenka, T. Schwede, and R. Koch, "Parcel Tracking by Detection in Large Camera Networks," in *Pattern Recognition*, T. Brox, A. Bruhn, and M. Fritz, Eds., ser. Lecture Notes in Computer Science, Cham: Springer International Publishing, 2019, pp. 89–104.
- [4] H. N. Karaca and C. Akınlar, "A Multi-camera Vision System for Real-Time Tracking of Parcels Moving on a Conveyor Belt," in *Computer and Information Sciences - ISCIS 2005*, p. Yolum, T. Güngör, F. Gürgen, and C. Özturan, Eds., ser. Lecture Notes in Computer Science, Berlin, Heidelberg: Springer, 2005, pp. 708–717.
- [5] K. He, G. Gkioxari, P. Dollar, and R. Girshick, "Mask R-CNN," in *IEEE International Conference on Computer Vision (ICCV)*, Venice: IEEE, 2017, pp. 2980–2988.
- [6] J. Bromley, I. Guyon, Y. LeCun, E. Säckinger, and R. Shah, "Signature verification using a "Siamese" time delay neural network," in *Proceedings of the 6th International Conference on Neural Information Processing Systems*, ser. NIPS'93, San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1993, pp. 737–744.
- [7] M. Ye, J. Shen, G. Lin, T. Xiang, L. Shao, and S. C. H. Hoi, "Deep Learning for Person Re-Identification: A Survey and Outlook," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 6, pp. 2872–2893, 2022.
- [8] S. D. Khan and H. Ullah, "A survey of advances in vision-based vehicle re-identification," *Computer Vision and Image Understanding*, vol. 182, pp. 50–63, 2019.
- [9] A. Naumann, L. Dörr, N. Ole Salscheider, and K. Furmans, "Refined Plane Segmentation for Cuboid-Shaped Objects by Leveraging Edge Detection," in *2020 19th IEEE International Conference on Machine Learning and Applications (ICMLA)*, 2020, pp. 432–437.
- [10] C. Liu, K. Kim, J. Gu, Y. Furukawa, and J. Kautz, "PlaneRCNN: 3D Plane Detection and Reconstruction From a Single Image," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Long Beach, CA, USA: IEEE, 2019, pp. 4445–4454.
- [11] J. Canny, "A Computational Approach to Edge Detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-8, no. 6, pp. 679–698, 1986.
- [12] X. Soria, E. Riba, and A. Sappa, "Dense Extreme Inception Network: Towards a Robust CNN Model for Edge Detection," in *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)*, Snowmass Village, CO, USA: IEEE, 2020, pp. 1912–1921.
- [13] G. Brazil, A. Kumar, J. Straub, N. Ravi, J. Johnson, and G. Gkioxari, "Omni3D: A Large Benchmark and Model for 3D Object Detection in the Wild," in *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE, Vancouver, Canada, 2023.
- [14] G. Gkioxari, J. Malik, and J. Johnson, "Mesh RCNN," in *IEEE/CVF International Conference on Computer Vision*, 2019.
- [15] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, vol. 1, 2005, 886–893 vol. 1.
- [16] A. Naumann, F. Hertlein, B. Zhou, L. Dörr, and K. Furmans, "Scrape, Cut, Paste and Learn: Automated Dataset Generation Applied to Parcel Logistics," in *2022 21st IEEE International Conference on Machine Learning and Applications (ICMLA)*, 2022, pp. 1026–1031.
- [17] H. Hu, F. Immel, J. Janosovits, M. Lauer, and C. Stiller, "A Cuboid Detection and Tracking System using A Multi RGBD Camera Setup for Intelligent Manipulation and Logistics," in *2021 IEEE 17th International Conference on Automation Science and Engineering (CASE)*, 2021, pp. 1097–1103.
- [18] P. Arpentì, R. Caccavale, G. Paduano, G. Andrea Fontanelli, V. Lippiello, L. Villani, *et al.*, "RGB-D Recognition and Localization of Cases for Robotic Depalletizing in Supermarkets," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 6233–6238, 2020.

- [19] D. Chiaravalli, G. Palli, R. Monica, J. Aleotti, and D. L. Rizzini, "Integration of a Multi-Camera Vision System and Admittance Control for Robotic Industrial Depalletizing," in *2020 25th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA)*, vol. 1, 2020, pp. 667–674.
- [20] A. Naumann, F. Hertlein, L. Dörr, S. Thoma, and K. Furmans, *Literature Review: Computer Vision Applications in Transportation Logistics and Warehousing*, 2023. arXiv: [2304.06009](https://arxiv.org/abs/2304.06009) [cs].
- [21] D. Dwivedi, T. Malisiewicz, V. Badrinarayanan, and A. Rabinovich, *Deep Cuboid Detection: Beyond 2D Bounding Boxes*, 2016. arXiv: [1611.10010](https://arxiv.org/abs/1611.10010).
- [22] E. Arnold, O. Y. Al-Jarrah, M. Dianati, S. Fallah, D. Oxtoby, and A. Mouzakitis, "A Survey on 3D Object Detection Methods for Autonomous Driving Applications," *IEEE Transactions on Intelligent Transportation Systems*, vol. 20, no. 10, pp. 3782–3795, 2019.
- [23] P. Li, H. Zhao, P. Liu, and F. Cao, "RTM3D: Real-Time Monocular 3D Detection from Object Keypoints for Autonomous Driving," in *Computer Vision – ECCV 2020*, A. Vedaldi, H. Bischof, T. Brox, and J.-M. Frahm, Eds., vol. 12348, Cham: Springer International Publishing, 2020, pp. 644–660.
- [24] Z. Rui, G. Zongyuan, D. Simon, S. Sridha, and F. Clinton, "Geometry-Constrained Car Recognition Using a 3D Perspective Network," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 01, pp. 1161–1168, 2020.
- [25] Y. Chen, Y. Tian, and M. He, "Monocular human pose estimation: A survey of deep learning-based methods," *Computer Vision and Image Understanding*, vol. 192, p. 102 897, 2020.
- [26] H. Chen, R. Feng, S. Wu, H. Xu, F. Zhou, and Z. Liu, "2D Human pose estimation: A survey," *Multimedia Systems*, 2022.
- [27] L. Dörr, F. Brandt, A. Naumann, and M. Pouls, "TetraPackNet: Four-Corner-Based Object Detection in Logistics Use-Cases," in *DAGM German Conference on Pattern Recognition*, 2021.
- [28] H. Law and J. Deng, "CornerNet: Detecting Objects as Paired Keypoints," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 734–750.
- [29] W. Shi, M. Zhang, R. Zhang, S. Chen, and Z. Zhan, "Change Detection Based on Artificial Intelligence: State-of-the-Art and Challenges," *Remote Sensing*, vol. 12, no. 10, p. 1688, 2020.
- [30] J.-M. Park, J.-H. Jang, S.-M. Yoo, S.-K. Lee, U.-H. Kim, and J.-H. Kim, "ChangeSim: Towards End-to-End Online Scene Change Detection in Industrial Indoor Environments," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 8578–8585.
- [31] J.-M. Park, U.-H. Kim, S.-H. Lee, and J.-H. Kim, "Dual Task Learning by Leveraging Both Dense Correspondence and Mis-Correspondence for Robust Change Detection With Imperfect Matches," in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, New Orleans, LA, USA: IEEE, 2022, pp. 13 739–13 749.
- [32] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, NV, USA: IEEE, 2016, pp. 770–778.
- [33] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature Pyramid Networks for Object Detection," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 936–944.
- [34] G. Jie and L. Ning, "An Improved Adaptive Threshold Canny Edge Detection Algorithm," in *2012 International Conference on Computer Science and Electronics Engineering*, vol. 1, 2012, pp. 164–168.
- [35] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The Unreasonable Effectiveness of Deep Features as a Perceptual Metric," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 586–595.
- [36] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004.
- [37] Z. Wang, E. Simoncelli, and A. Bovik, "Multiscale structural similarity for image quality assessment," in *The Thirty-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, vol. 2, 2003, 1398–1402 Vol.2.
- [38] M. P. Sampat, Z. Wang, S. Gupta, A. C. Bovik, and M. K. Markey, "Complex Wavelet Structural Similarity: A New Image Similarity Index," *IEEE Transactions on Image Processing*, vol. 18, no. 11, pp. 2385–2401, 2009.

- [39] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, *et al.*, “Microsoft COCO: Common Objects in Context,” in *Computer Vision – ECCV 2014*, D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, Eds., ser. Lecture Notes in Computer Science, Cham: Springer International Publishing, 2014, pp. 740–755.
- [40] I. Loshchilov and F. Hutter, “SGDR: Stochastic gradient descent with warm restarts,” in *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*, OpenReview.net, 2017.