

Supervised Pose and Motion Estimation over Weakly Textured Industrial Objects

Valentin Borsu and Pierre Payeur

School of Information Technology and Engineering

University of Ottawa

Ottawa, ON, Canada

[vbors100, ppayeur]@site.uOttawa.ca

Abstract—Visually estimating the motion of automotive body parts over an assembly line represents a major challenge for classical feature detection, matching and tracking algorithms due to the lack of a rich surface texture. But as feature extraction and matching remain vital for accurate object pose and motion estimation, this paper presents a thorough investigation on the actual reliability of popular feature extraction and matching tools in terms of stability and robustness for industrial applications. Severe tracking errors that result from brightness variations and occlusions are corrected with the integration of an original supervisory approach that relies on the encoding of a minimum amount of a priori information about the general appearance of the objects. The proposed solution is experimentally validated on an application for quality control in the automotive industry.

Keywords— *feature extraction; feature matching; feature tracking; pose and motion estimation.*

I. INTRODUCTION

The integration of robotic systems in industrial manufacturing resulted in more efficient, accurate, cost effective and safer solutions for assembly and quality control operations. The way to conduct robotic procedures on static bodies has been addressed for long. For moving objects, current solutions often rely on preset motion parameters or use track encoders. However, for a robot to execute tasks on a freely moving object, accurate and real-time data about the pose and motion of the object is required. The surface appearance of the object then plays an important role.

Automotive body panels [1-3] are characterized by a very low number of distinctive visual features over their surface during the car assembly process. To account for the lack of features over such objects, current approaches for pose and motion estimation rely on structured lighting [4], edge tracking [5] or particle filtering [6]. However, these approaches limit the possibility of sustaining high production rates while handling several types of body parts, under the complexity of typical industrial environments, which are often characterized by occlusions caused by manipulator robots and sporadic appearances of factory associates in the view of the vision systems. Under these settings, the pose and motion estimator (PME) needs to run in real-time and not rely on exact 3D CAD models of the panels while embedding a high level of robustness not to deviate from its target. In order to overcome some of these difficulties, Kak *et al.* [1-3] used three different

vision systems, along with controlled backgrounds and the knowledge that the moving object contains circular features over its surface. These conditions hardly transpose in automotive assembly environments.

Sparse optical flow computation is a well established technique and is often based on Lucas-Kanade's (LK) [7, 8] approach. However, the LK tracker is extremely sensitive to cases where the brightness constancy assumption is not validated, such as when occlusions and photometric variations occur [9]. In order to overcome the sensitivity of the LK tracker to photometric variations, Jin *et al.* [10] introduced a hybrid model based on photometry and geometry for the characterization of variations exhibited by the feature patches during tracking. However, their formulation still remains sensitive to occlusions, and to the appearance of other objects.

The objective of this research is to develop an alternative solution that relies on passive vision only, and surmounts the constraints of the application from a software perspective rather than with the addition of extra hardware, which ensures greater flexibility. The proposed approach extends previous solutions for robotic tracking and marking of surface deformation defects [11, 12]. An original supervisory layer is added to the pose and motion estimation stage that allows to reliably detect and track a low number of available features throughout the displacement of the object. This paper concentrates on the analysis of the feature extraction and matching processes and details the supervisory layer operation, whose objective is to provide time-efficient, accurate and fault-tolerant visual servoing data to the robotic station.

In the following sections, the experimental setup is presented and the framework for the supervised PME is introduced. Subsequently, an experimental feature extraction analysis is proposed, which evaluates the suitability of five leading feature detectors [13-17]. The core of this analysis is supported by an original metric which empirically evaluates the correlated stability-robustness property of the feature extraction process. Based on the results provided by the analysis of the correlated stability-robustness metric, and in spite of the lack of rich texture over the automotive panels, the proposed PME builds upon classical computer vision algorithms for feature extraction, tracking and matching [7, 8, 17]. Finally, the improvements brought by the proposed supervised PME solution are demonstrated.

II. EXPERIMENTAL PLATFORM

A vision-robotic platform, shown in Fig. 1, is used for experimentation. It integrates modules for automated surface deformations detection over automotive panels [12], for part motion estimation and for robotic marking. Various automotive panels, such as unfinished doors and fenders, which are representative of the early manufacturing stage of automobiles, are tested on the system. A sled provides a linear motion at variable speeds which mimics the assembly line displacement over a length of 54 cm.

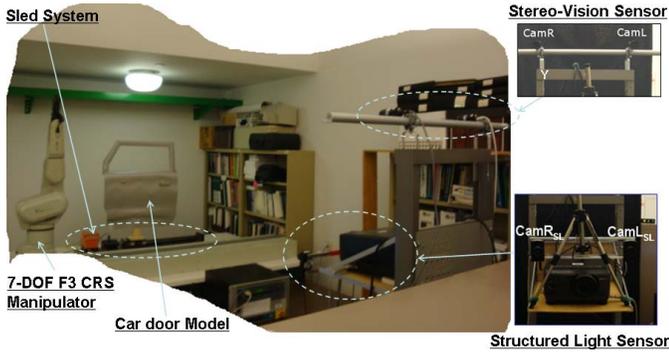


Figure 1. Experimental setup for vision tracking and robotic interaction with moving panels.

A calibrated stereo-vision sensor (SS) is used for estimating the pose and motion of the automotive panel. It is composed of two Point Grey Flea2 IEEE-1394b CCD cameras (CamL and CamR) with 8.5mm lenses and 640x480 pixels of resolution. A 44.5 cm baseline is used between the cameras, as it provides improved accuracy in reconstructing the sparse structure of the panel. The automotive part is located at approximately 310 cm from the acquisition system. The interaction with the panel is performed by an F3 7DOF CRS serial manipulator, and the inter-calibration between the SS and the robot's base [12] is available.

III. FEATURE EXTRACTION AND MATCHING ANALYSIS

The reliability of the feature extraction and matching processes is critical for accurate visual estimation of position and motion on weakly textured objects in an industrial environment. This section presents an experimental evaluation of several classical feature point detectors in order to derive guidelines for the development of motion estimators in real world applications. An analysis of Torr's "Structure and Motion toolkit" [18] for solving the feature correspondence problem is also performed in the last part of this section.

Most popular computer vision tools for keypoint extraction [13-17] can be very successful when dealing with richly textured objects. However, their performance severely deteriorates when facing constraints such as those found in the industrial application considered here. For conducting the study of these feature detectors, the properties of stability and robustness [19] are analyzed as they represent essential characteristics for reliable feature extraction. The robustness property is linked to the insensitivity to noise, and the proportion between false positives and true localized feature points. The stability property, which also embeds the

"repeatability" of feature extraction, is linked to the capability of a corner detector to identify the same points even though the images suffer from perspective distortion, zoom or illumination changes. This study also aims at intrinsically testing the variability of feature detection, since the input images contain a very low amount of details resulting from weak surface textures over the panels. The variability characteristic represents the ability of the feature extractor to still detect several feature points despite the nature of the image content.

The first part of this section is dedicated to the SIFT (Scale Invariant Feature Transform) keypoint detector [13] and its suitability to heavily constrained applications. Secondly, an original measure is introduced with the purpose of quantizing the correlated stability-robustness property of the feature extraction. Subsequently, the correlated stability-robustness measure is applied to four popular feature detectors in order to select the most appropriate feature extraction methodology. Finally, the performance of Torr's feature correspondence technique [18] is analyzed under the industrial context of the considered application.

A. SIFT Keypoint Detector

The SIFT feature detector [13] has received a lot of attention over the last few years, because of its capability to extract features which are invariant to scaling and rotation. But the repeatability [19] of the SIFT feature extractor has not been rigorously evaluated for real-life industrial applications where the lifespan of the detected features directly impacts the tracking accuracy. In order to empirically assess the repeatability of the SIFT feature detector, a scenario was considered in which the linear sled system was positioned approximately parallel to the baseline of the SS, whereas the velocity of the sled system was set to $v_{ss}=1.4\text{cm/s}$.

Under these settings, Fig. 2 illustrates the keypoints detected by SIFT on the first and last processed frames, respectively. It can be noticed that the descriptors associated to the features detected on the panel surface substantially differ between the two frames, and not only in the area occluded by the robot in the last frame. This degradation exemplifies the lifespan reduction of the extracted features, which negatively impacts the tracking process, as those features represent its main source of information.

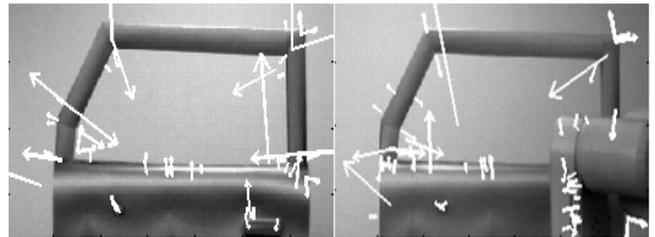


Figure 2. Extracted SIFT keypoints over the tracking sequence.

A limited lifespan has a direct effect on the repeatability property of the feature extraction, which is essential for accurate and reliable pose and motion estimation. Due to its limited repeatability, the SIFT detector was discarded from the assessment of the correlated stability-robustness, which will be performed in the next section.

B. Correlated Stability-Robustness Empirical Measure

For assessing the combined stability-robustness performance of different feature extraction processes, the data acquired with the scenario described in Section III.A was analyzed while assuming that a limited number of macro-features (MF) are manually pre-selected over the surface of a car door panel at the beginning of the motion cycle. Ten MFs are shown and numbered in Fig. 6a.

A quality measure called “success percentage” is introduced:

$$P_S[\%] = \frac{1}{N_{\text{frames}}} \sum_{i=1}^{N_{\text{frames}}} \frac{(n_{\text{MF}}^{\text{cd}})_i}{(n_{\text{MF}}^{\text{total}})_i} \quad (1)$$

where $(n_{\text{MF}}^{\text{cd}})_i$ is the number of correctly detected MFs (in the i -th frame) within a neighborhood of 5×5 pixels around the keypoints shown in Fig. 6a, and $(n_{\text{MF}}^{\text{total}})_i$ is the total number of detectable MFs in the i -th frame. It should be noticed that $(n_{\text{MF}}^{\text{total}})_i$ is not necessarily equal to the total number of MFs extracted on the car door panel, because of the occlusions occasionally caused by the manipulator robot, also present in the scene. Without loss of generality, the “success percentage” was computed only with the results obtained by processing the frames grabbed by the left camera of the stereo pair, CamL.

C. Classical Feature Extractors

Other than SIFT, the feature extractors investigated in this experimentation include the Harris and Stephens corner detector [14], followed by the same feature detector but with Noble’s validation gate [15], the SUSAN corner extractor [16], and finally, the Shi and Tomasi feature detector [17]. For the Harris and Stephens corner detector with Noble’s validation gate, and SUSAN feature extractor, the MATLAB framework developed by Garcia [20] was used in the analysis, whereas for the Harris and Stephens extractor [14] and the Shi and Tomasi’s approach [17], their OpenCV [21] implementations were selected. All the feature extractors were applied on image segments containing the car door’s window frame, which were used for the SIFT-processing as well. For the extraction of the regions of interest (ROI), the robust optical flow calculations, which will be introduced in Section IV, had a direct impact.

Since the constrained motion pattern of the car door was a translation along the X axis of CamR, as in Fig. 1, with no changes in orientation, the positions of the MFs within the ROI were not affected by the movement of the object. As a consequence, a 5×5 pixels patch was set for all MFs, shown in Fig. 6a, within the initial ROI belonging to the first frame grabbed by CamL that contained a full view of the car door. Then, the process of computing $(n_{\text{MF}}^{\text{cd}})_i$ consisted in verifying if the detected MFs lie within the pre-assigned patches. Due to the constrained motion of the car door, the positions of these patches in the extracted ROI remained constant, throughout the entire duration of the motion experiment. Figure 3 illustrates the feature extraction results obtained with the four selected approaches applied on the 10th frame segment during the motion of the rigid body. In this case, $(n_{\text{MF}}^{\text{total}})_{10} = 8$ is the total

number of detectable MFs, which were all correctly extracted with the Shi and Tomasi corner detector [17] and Harris detector with Noble’s validation gate [15].

In Fig. 3 the correctly localized MFs are marked with the symbol “●”, whereas for the erroneously detected MFs the symbol “○” is used. In order to compute $(n_{\text{MF}}^{\text{cd}})_i$, a 20×20 pixels patch was centered on the MFs in the first extracted ROI. The normalized cross-correlation, as a measure for quantizing how much the intensities of the selected regions have changed during tracking, when compared to the distribution of intensities in the first ROI, was used to compute the total number of detectable MFs.

The results obtained for the “success percentage” with each selected feature point extractor are summarized in Table I. In the case of the Harris and Stephens corner detector [14], a threshold $T_{\text{Harris_Stephens}}=0.01$ was used for the corner strength measure. As it can be noticed from Table I, a considerable improvement in the success percentage is obtained by applying Noble’s validation gate [15] with the same threshold, $T_{\text{Noble}}=0.01$. This improvement is a result of Noble’s corner strength measure, which diminishes the sensitivity of the Harris and Stephens corner extractor to image patches having contrast variations.

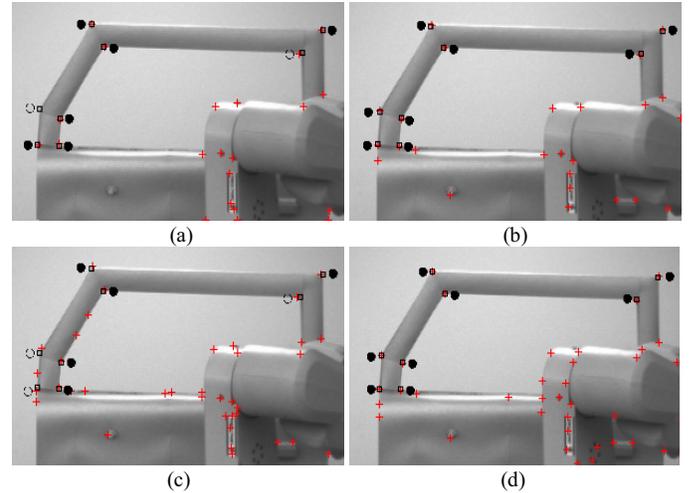


Figure 3. Feature extraction results: (a) Harris and Stephens corner detector, (b) Harris and Stephens corner detector with Noble’s corner strength measure, (c) SUSAN corner detector, (d) Shi and Tomasi corner detector.

The SUSAN corner detector [16] was used with its default thresholds [20] and gave the lowest “success percentage” from all the studied feature extractors. The Shi and Tomasi corner detector [17], used with the threshold $T_{\text{Shi_Tomasi}}=0.02$, gave the highest “success percentage” for the proposed stability-robustness correlated measure.

Several tests performed on similar objects led to the same conclusions. Therefore, the supervised pose and motion estimator, that will be described in section IV, builds upon the classical Shi and Tomasi corner detector [17], which is selectively applied on the ROIs containing the area of the MFs.

TABLE I. “SUCCESS PERCENTAGE” FOR THE FOUR FEATURE EXTRACTORS

Feature Extractor	Success Percentage $P_s(\%)$
Harris and Stephens	57.14
Harris and Stephens with Noble’s validation gate	75.25
SUSAN	47.52
Shi and Tomasi	77.08

D. Torr’s Structure and Motion Toolkit

In order to evaluate the feature matching process, the “structure and motion toolkit” developed by Torr [18] was also tested under the same experimental setup. Torr’s approach relies on a Maximum *A Posteriori* Sampling Consensus (MAPSAC) technique for matching features in the presence of outliers. It was applied on two ROIs extracted from the first two acquired frames, before the manipulator robot occluded the view of the SS. As shown in Fig. 4, the set of correspondences, obtained after judiciously tuning the two thresholds related to the maximum numbers of samples to be drawn and the proportion of inliers, has a large proportion of outliers, as multiple corners, extracted from CamL’s ROI, point to the same corner in the adjacent view. As a result, the unicity property of the feature matching [19] is highly compromised. This experiment shows that Torr’s toolkit requires a substantial amount of adaptation, when transferred to industrial parts that exhibit weak texture surfaces.

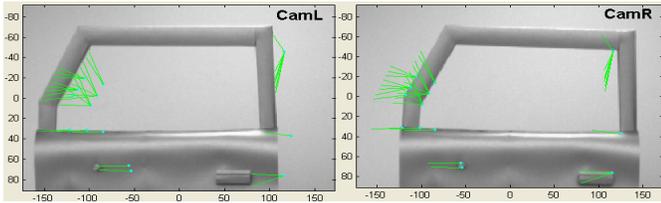


Figure 4. Feature matches computed with Torr’s toolkit.

IV. SUPERVISED POSE AND MOTION ESTIMATOR

The proposed supervised pose and motion estimation framework, illustrated in Fig. 5, builds upon a pre-selection of a minimum set (typically 6 to 10) of macro-features (MFs) over the structure of the object. In a factory, the selection is manually performed by an installation engineer over the left view (CamL) of the stereo images, and only once, when configuring the vision-robotic station for a specific type of panels. These MFs are represented by a few corner points which are distinctive over the panel.

In the case of a car door panel, the selected MFs can belong to the inner and outer frame of the door window. Figure 6a shows the location of the MFs as they were pre-selected, along with the associated 9x9 pixels patches, which are used to refine their initial locations on the first image. The automated refinement of the MFs uses the Shi and Tomasi corner detector [17] within the extracted patches. Once the MFs are refined in the left view, their correspondences with the right view (CamR) are computed. Given that the stereoscopic cameras are approximately parallel, the pyramidal implementation of the LK tracker [7, 8] is used to guide the correspondences [11].

Figure 6b illustrates the disparity vectors extracted from the right initialization frame.

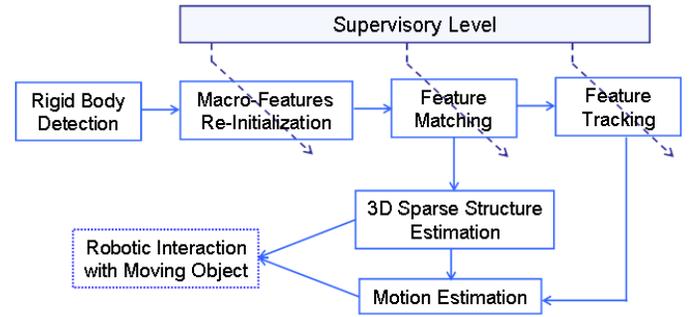


Figure 5. Supervised pose and motion estimation framework.

Following the initialization procedure, the PME is triggered. First, a rigid body detection module informs the PME about the full appearance of the panel in the field of view of the SS. The supervised MFs re-initialization block ensures the automatic detection of the MFs in subsequently acquired stereo frames, whenever a new panel enters the field of view. It builds upon the Shi and Tomasi corner detector and the pyramidal LK tracker but is fully automated.

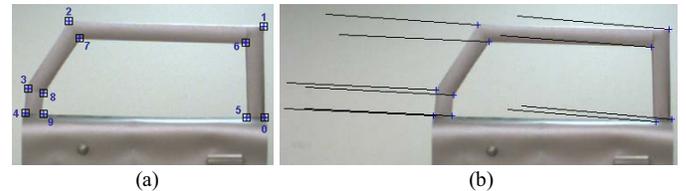


Figure 6. (a) Car door’s MFs, (b) MFs disparity vectors over the right initialization frame.

Following re-initialization, the MFs position information is stored in two buffers that will be used by the supervisory system. The first buffer contains the Euclidean distances between each 2D MF and all of the other MFs in the extracted set. The second buffer stores the relative x and y displacements, expressed with respect to the image plane, between each MF and all of the other MFs. Subsequently, the MFs are being tracked frame-by-frame through the pyramidal LK tracker, which provides the necessary data for estimating the motion undergone by the panel between two successive frames. The motion estimations are used to guide the robotic interaction with the moving panel, given that the inter-calibration between the robot’s base and the SS is available [12].

The objective of the proposed supervisory layer is to overcome the limitations observed with classical computer vision methods for feature extraction, matching and tracking, in the case of objects with insufficient surface texture.

V. SUPERVISORY LAYER OPERATION AND VALIDATION

In the proposed framework, the tracking of the MFs is performed in both stereo views. Feature matching is straightforward given that the MFs are stored in an indexed data structure. However, the LK tracker is known to provide poor performance when variations in the brightness occur [9], such as when the robot or people temporarily occlude a certain

part of the MFs' area. This has a direct impact on the feature matching process. To address this problem, a validation gate is embedded in the supervisory layer to correct erroneous motion vectors and to recover the MFs lost in the tracking, assuming a rigid structure for the object.

For the supervisory operation to be reliable, the only condition is that the scaling effects exhibited by the object between two subsequent frames are minimal. The validation gate structure is illustrated in Fig. 7. It consists of finding a pair of tracked MFs that have a high level of confidence in both views. For that matter, the first test verifies that the difference between the Euclidean distance between the MFs (MF_i, MF_j) forming the pair, d_{ij} , and their corresponding distance saved in the buffer of Euclidean distances, \hat{d}_{ij} , is within ± 2 pixels. If the first test is passed, then a second test checks whether the difference between the relative displacement of the selected MFs, δ_{ij} , and the reference relative displacement stored in the buffer of relative displacements, $\hat{\delta}_{ij}$, is also within ± 2 pixels. Finally, if the second test is also validated, a third test is performed on the norms of the motion vectors, v_{OF}^i, v_{OF}^j associated with each MF. This last test verifies whether the difference between the motion vectors in the x and y directions of both MFs forming the pair is within ± 2 pixels. This mechanism ensures the consistency of the features distribution throughout the sequence, and validates the correlation between the motion vectors in both views, in order to eliminate divergence and false estimates. A margin of 2 pixels was experimentally selected based on the size of the panels in the image plane, when imaged from a distance of approximately 310cm.

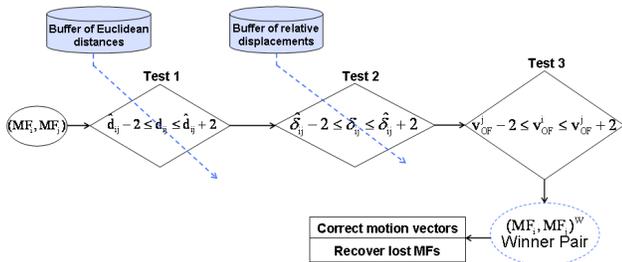


Figure 7. Validation gate to reinforce accurate MFs pairs tracking.

The three-step validation gate ends as soon as one pair of MFs, that passes all three tests, is found. This MFs pair $(MF_i, MF_j)^W$, is regarded as a “winning” pair and the corresponding motion vector is used to correct other erroneous motion vectors and to recover the MFs that are lost during the tracking.

Figures 8a and 8b present two processed frames, grabbed by CamL, over the tracking sequence, when running the supervisory layer. The motion vectors initially returned by the LK tracker are marked in red, the corrected motion vectors are represented in black, and the recovered motion vectors are drawn in white. Additionally, Fig. 8c was extracted from a scenario in which additional lighting was provided from a set of lamps mounted on top of the SS, in order to inspect the robustness of the PME to shadows and specular reflections on the panel surface. It can be noticed that all corrected and

recovered motion vectors point to the proper MF regions, as shown in Fig. 6a. The validation gate performance was tested on 60 running scenarios involving various positions and movements of the robot and people entering the field of view of the SS. The correction/recovery of the motion vectors demonstrated a 95.4% success rate. As a result, the proposed validation gate provides the feature tracker with a robustness and stability level that cannot be achieved with the LK tracker alone, due to its instability in non-constant brightness scenarios.

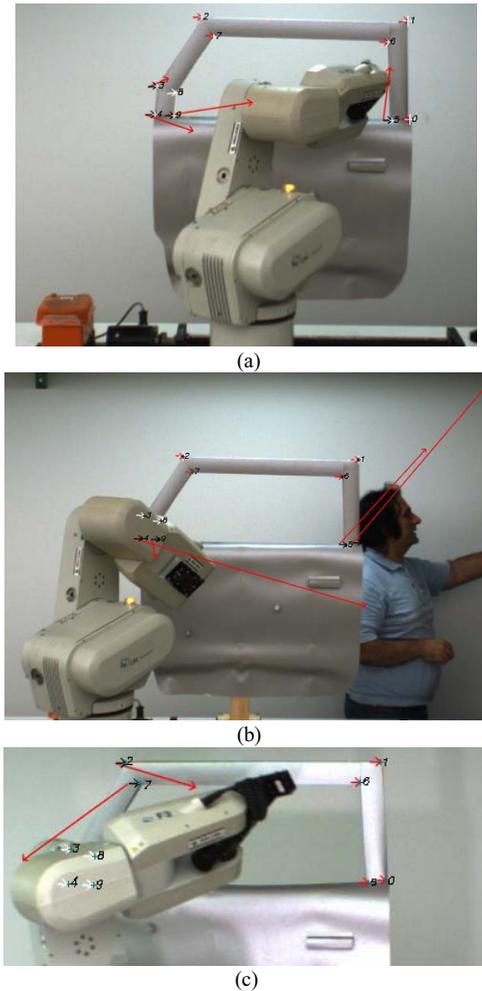


Figure 8. Motion vector results where: (a) the robot, (b) a person and the robot are present in the scene, (c) the scene contains shadows and specular reflections.

When it is impossible to find a winning pair in one of the frames, the results from the winning pair found in the adjacent view (left/right stereo-frame) are used as a best estimate, given that the cameras forming the SS are approximately parallel. In such a situation, as shown in Fig. 9a, the entire set of MFs can still be recovered based on the winning motion vectors extracted from the adjacent frame acquired by CamR.

Finally, in the situation where the search for a winning MFs pair is unsuccessful in both the left and right views, the motion vector related to the previous set of frames in the tracking sequence is used. Figure 9b and 9c show segments from two

frames grabbed by CamL, in which no winning pairs could be found. Thus, the motion vectors, in blue, were successfully recovered according to the motion vectors obtained from the processing of the previous pair of frames.

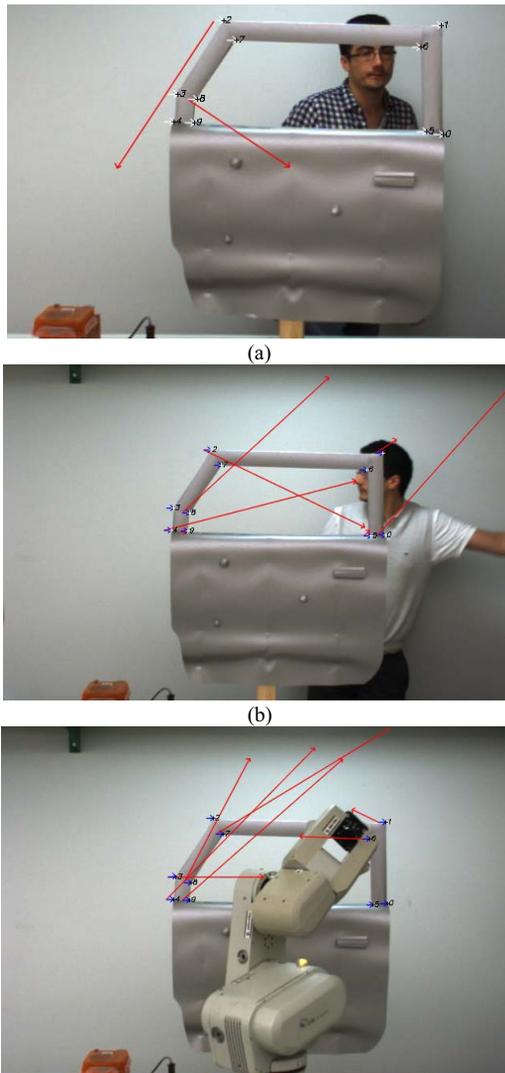


Figure 9. Motion vector results where: (a) the motion vectors are recovered from the adjacent view, (b), (c) the motion vectors are recovered from the previous frame.

VI. CONCLUSIONS

This paper addresses the problem of pose and motion estimation over weakly textured industrial objects and focuses on the stability-robustness analysis of the feature extraction process. It introduces a supervised pose and motion estimation system that provides accurate and reliable motion vectors in real-time while relying only on passive stereoscopy and classical feature point trackers. The proposed supervisory layer significantly increases the robustness of a pose and motion estimator to operate on weakly textured objects, even when surfaces exhibit specular reflections, and in industrial environments where occlusions or sporadic appearance of workers in the field of view of the vision-based pose and motion estimator occur.

ACKNOWLEDGMENTS

The authors acknowledge the financial support from Precarn Inc., and the collaboration of Neptec Design Group Ltd and Honda Canada to this research.

REFERENCES

- [1] Y. Yoon, G. N. DeSouza and A. C. Kak, "Real-time tracking and pose estimation for industrial objects using geometric features", *Proc. of the IEEE Intl. Conf. on Robotics & Automation*, Taiwan, pp. 3473-3478, 2003.
- [2] G. N. DeSouza and A. C. Kak, "A subsumptive, hierarchical and distributed vision-based architecture for smart robotics", *IEEE Transactions on Systems, Man and Cybernetics*, part B, vol. 34(5), pp. 1988-2002, 2004.
- [3] Y. Yoon, J. Park and A. C. Kak, "A heterogeneous distributed visual servoing system for real-time robotic assembly applications", *Proc. of the Intl. Conf. on Robotics and Automation*, USA, pp. 4422-4424, 2006.
- [4] D. Scharstein and R. Szeliski, "High accuracy stereo-depth maps using structured light", *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*, USA, vol. 1, pp. 195-202, 2003.
- [5] G. Zhu, Q. Zeng and C. Wang, "Efficient edge-based object tracking", *Elsevier Journal of Pattern Recognition*, vol. 39, pp. 2223-2226, 2006.
- [6] X. Wu, D. Pi and X. Jiang, "Particle filter based pose and motion estimation with non-Gaussian noise", *Proc. of the IEEE Intl. Conf. on Industrial Technology*, India, pp. 979-982, 2006.
- [7] B. D. Lucas and T. Kanade, "An iterative image registration technique with an application to stereo-vision", *Proc. of the 1981 DARPA Image Understanding Workshop*, pp. 121-130, 1981.
- [8] C. Tomasi and T. Kanade, "Detection and tracking of point features", *Carnegie Mellon University Technical Report*, CMU-CS-91-132, USA, 1991.
- [9] J. Weber and J. Malik, "Robust computation of optical flow in a multi-scale differential framework", *Intl. Journal of Computer Vision*, vol. 14, pp. 67-81, 1995.
- [10] H. Jin, P. Favaro and S. Soatto, "Real-time feature tracking and outlier rejection with changes in illumination", *Proc. of the IEEE Intl. Conf. on Computer Vision*, Canada, pp. 684-689, 2001.
- [11] V. Borsu and P. Payeur, "Pose and motion estimation of a moving body with few features", *Proc. of the IEEE Intl. Workshop on Robotic and Sensor Environments*, Lecco, Italy, pp.161-121, 2009.
- [12] V. Borsu, A. Yogeswaran and P. Payeur, "Automated surface deformations detection and marking on automotive body panels", *Proc. of the IEEE Intl. Conf. on Automation Science and Engineering*, Canada, pp. 551-556, 2010.
- [13] D. Lowe, "Distinctive image features from scale-invariant keypoints", *Intl. Journal of Computer Vision*, vol. 60, no. 2, 2004, pp.91-110.
- [14] C. Harris and M. Stephens, "A combined corner and edge detector", *Proc. of the 4th Alvey Vision Conference*, UK, pp. 147-151, 1988.
- [15] A. Noble, "Description of image surfaces", *PhD Thesis*, Department of Engineering Science, Oxford University, 1989.
- [16] S. Smith and J. M. Brady, "SUSAN – a new approach to low level image processing", *Intl. Journal of Computer Vision*, vol. 23, no. 1, pp. 45-78, 1997.
- [17] J. Shi and C. Tomasi, "Good features to track", *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*, USA, pp. 593-600, 1994.
- [18] P. H. S. Torr, "A structure and motion toolkit in Matlab", *Microsoft Research Technical Report*, MSR-TR-2002-56, 2002.
- [19] É. Vincent and R. Laganière, "An empirical study of some feature matching strategies", *Proc. of the Intl. Conf. on Vision Interface*, Canada, pp. 139-145, 2002.
- [20] V. Garcia, "Keypoint extraction", [on-line], <http://www.mathworks.com/matlabcentral/fileexchange/17894-keypoint-extraction>, 2007.
- [21] G. Bradsky and A. Kaehler, *Learning OpenCV – Computer vision with the OpenCV library*, O'Reilly, USA, 2008.