2024 IEEE International Conference on Robotics and Automation (ICRA 2024)
May 13-17, 2024. Yokohama, Japan

Measuring Ball Joint Faults in Parabolic-Trough Solar Plants with Data Augmentation and Deep Learning

M.A. Pérez-Cutiño[1,2], J. Capitán[3], J.M. Díaz-Báñez[1] and J. Valverde[1,2]

***Abstract*— Automatic inspection of parabolic-trough solar plants is key to preventing failures that can harm the environment and the production of green energy. In this work, we propose a novel methodology to inspect ball joints in parabolic trough collectors, which is a relevant problem that is not adequately covered in the literature. Images collected by an Unmanned Aerial Vehicle are segmented using deep learning to extract ball joint components. In order to generate rich training datasets, we develop a novel data augmentation technique by rotating joints and adding synthetic image background, and demonstrate its impact on the object detection accuracy. Then two types of faults are analyzed: fluid leaks, by means of image color filtering; and geometric shape anomalies, by measuring joint angles of the robotic arms. We propose metrics to quantify these faults and evaluate the damage of the inspected components. Our experimental results with images from operating commercial plants show that we can automatically detect leaks and anomalous angular geometry with a low failure rate compared to human labeling.**

I. Introduction

Due to the growing development of Concentrated Solar Power (CSP) plants, automatic fault detection has become of interest in recent years [1], [2], [3], [4]. Manually inspecting a CSP plant is a tedious and time-consuming task, as hand-held devices have to be used to survey multiple kilometers of components, which could take days [5]. Instead, the use of Unmanned Aerial Vehicles (UAVs) for data collection and automatic fault detection is becoming commonplace to reduce inspection time in large plants [4].

Parabolic-Trough (PT) CSP plants consist of parabolic-shaped mirrors, the so-called Solar Collector Elements (SCEs), that concentrate the solar radiation on tubes that act as Heat Collector Elements (HCEs). Each solar collector (see Fig. 1) is a structure containing an SCE with three HCEs that rotates at discrete time intervals while tracking the Sun. The Ball Joint Assembly (BJA) is another component of special relevance, as it applies the necessary forces to rotate the whole system. Although there exist works in the literature for automatic fault detection in PT collectors [2], [6], [7], [8], they usually focus on SCEs and HCEs, being the inspection of a key element such as the BJA an under-researched area.

*This work is supported by grants PID2020-114154RB-I00, TED2021-129182B-I00 and DIN2020-011317, funded by MCIN/AEI/10.13039/501100011033 and the NextGenerationEU/PRTR.

[1]M.A. Pérez-Cutiño, J.M. Díaz-Báñez (dbanez@us.es) and J. Valverde (jvalverde@us.es) are with the Department of Applied Mathematics II, University of Seville, Seville, Spain.

[2]M.A. Pérez-Cutiño (m.perez@virtualmech.com) and J. Valverde are with Virtualmechanics S.L., Seville, Spain.

[3]J. Capitán is with the Multi-Robot Lab, University of Seville, Seville, Spain, jcapitan@us.es.

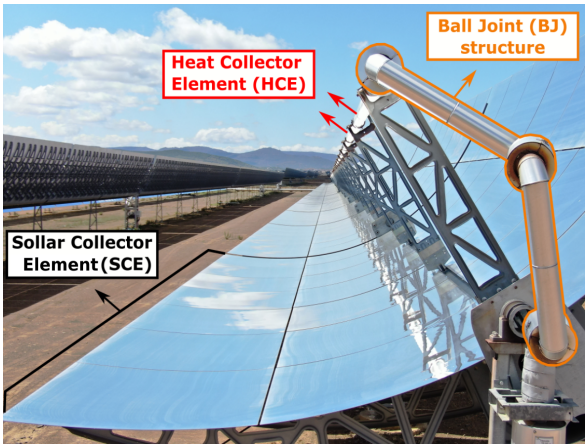

Fig. 1: Main components of a parabolic trough collector.

Different strategies have been proposed for fault detection in commercial CSP plants with the use of UAVs. In [9], slope mirror deviations were measured with the help of a remote controlled UAV. They use SCE corners as well as coded markers for accurate detection of the mirrors. Subsequently, this system was improved [10] to detect the HCE displacement with respect to its ideal position. Another approach for a similar task was implemented in [4], where they used the Detectron2 [11] library for image segmentation and automatic SCE identification. Coupling deep learning algorithms with other computer vision techniques allows them to achieve a robust method for large-scale optical analysis in PT solar plants. The work in [12] extracted thermal information from aerial images, which can be used to assess the state of the glass envelope of HCEs in a solar field. In [13], the data collected by a UAV and a car capture system were combined to detect broken thermal insulation in the ball joints and soiling levels in the SCEs. The authors report that thermal losses were measured with the use of in-house software, but the algorithms are not formally described.

As stated above, the aforementioned works focus on inspecting SCEs and HCEs. However, BJAs also require special attention, as severe incidents due to faulty BJA systems have been reported [14], [15]. When left unattended, leakages of the Heat Transfer Fluid (HTF) can result in fires with toxic smoke, which is noxious to the environment. Despite research interest in CSP plants, there is still a clear need for automatic techniques and metrics to assess the state of BJAs in PT systems. In this work, we rely on images collected by UAVs on the solar field and propose the combination of deep learning and computer vision methods

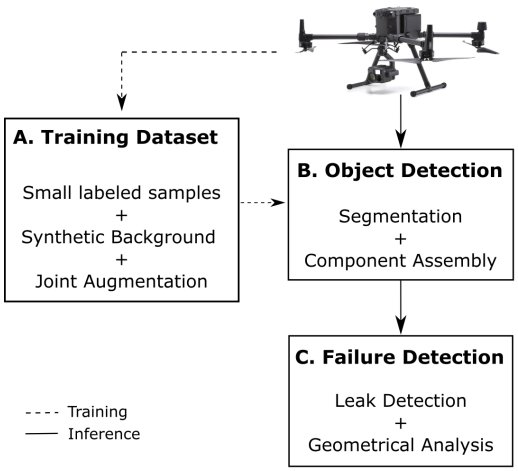

Fig. 2: Pipeline for detecting failures in BJAs.

for robust fault detection in BJA elements. First, in order to address the scarcity of available images of PT collectors (due to confidential issues), we develop a data augmentation technique that allows us to create rich datasets from smaller sample sets. Then we train a deep neural network to segment BJA components on the images, and assess the state of each BJA by analyzing its geometrical shape and the stains produced by the HTF when it leaks. Our main contributions can be summarized as follows:

- We tackle a relevant problem in CSP plants and propose a **novel methodology to inspect the BJA elements in PT collectors**, including new metrics and algorithms for leak detection and geometrical analysis.
- We propose **a new data augmentation technique based on what we call *Synthetic Background* and *Joint Augmentation*** in order to enrich BJA image datasets. The augmented dataset allows us to ease the training process for object segmentation.
- We assess the **capabilities of our methodology with extensive experiments** using data collected in **operating commercial CSP plants**, and demonstrate that faulty components are detected with high precision.

II. Methodology

We take as input images captured by a UAV and apply a processing pipeline with three main modules (see Fig. 2), which are described in the following sections.

A. The dataset

We created a dataset with images of BJAs captured with a UAV flying at 5 meters of altitude in 2 different PT plants. We manually labeled 37 images from one of the plants for training and 30 from the other for testing. Although an initial dataset with 37 images is not large enough for training machine learning algorithms, note that: 1) we are in a constrained environment with only one BJA per image, which reduces the data variability; 2) our final augmented dataset contains around 3,000 images; and 3) we apply transfer learning to improve the accuracy of our object detection results.

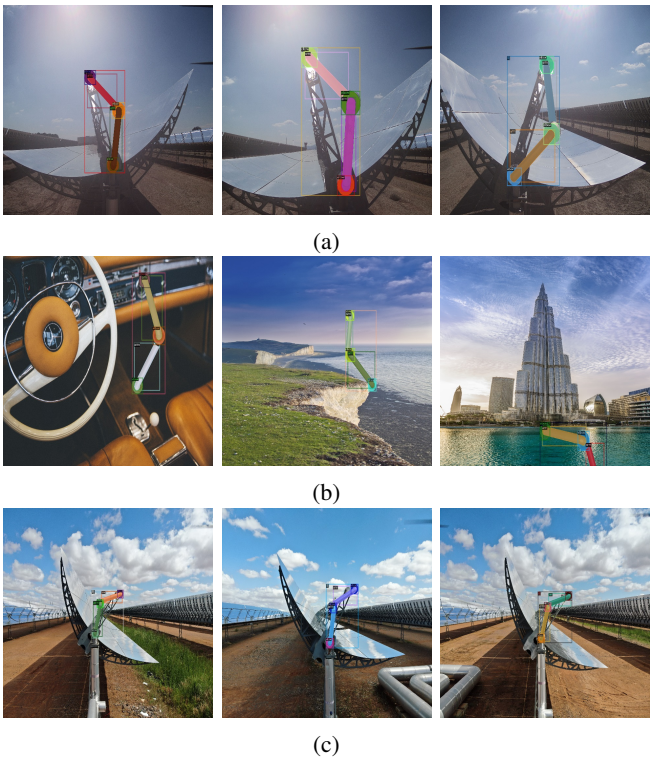

(a)

(b)

(c)

Fig. 3: (a) Training dataset samples without augmentation, (b) joint and background augmentation, (c) Test dataset samples. Each `arm` and `elbow` is filled with a random color, according to the ground truth segmentation.

The dataset is labeled for image segmentation considering two classes: `arm` for the steel pipes in the BJAs; and `elbow` for the ball joints, i.e., the circular steel portions. Instead of less demanding labeling techniques like creating bounding boxes or assigning one label to an image from a set of predefined classes, we annotate the ground truth pixel-wise. This allows us to extract the segmented object from the image more precisely, which is later used to augment the dataset. A sample of our resulting set is shown in Fig. 3.

1) Synthetic Background: This is a well-known data augmentation strategy to improve image segmentation training [16], [17], [18]. This methodology is especially important in the context of monotone backgrounds, where a similar background embedded in the training data becomes the main feature for object detection [17]. Increasing feature activation for foreground objects can be achieved by adding richer backgrounds to the training data. We created a Synthetic Background (SB) dataset (370 images). First, BJAs are extracted from each image and individually augmented (10 times) with traditional methods such as random scaling and rotation. Then we insert the transformed objects in random positions in images selected from a background dataset [19].

2) Joint Augmentation: Alternatively, we propose another data augmentation technique. Our novel idea can be applied to objects rotating through a joint, such as arms and elbows.

Fig. 4: Joint Augmentation for 4 example angles. The image surrounded by a blue bounding box is the input image.

To simulate different positions in the BJA, we fix the elbow and rotate the upper portion of the object as depicted in Fig. 4. We rotate each object eight times from -90° to 150°, with 30° steps. Then we apply background augmentation as indicated in the previous section, resulting in 2,960 images.

B. Object detection

Our first objective is to obtain a pixel-wise classification of the objects for each image. There exist multiple algorithms for image segmentation in the literature [20], [21], [22]. We selected a deep learning detector called Mask-RCNN [20], from the Detectron2 library [11], to extract the `arms` and `elbows` on the image; then we assemble these components with a heuristic procedure.

1) Deep learning: Detectron2 has been used in the literature to segment SCEs in CSP plants [4], and specifically Mask-RCNN is a well-established method for segmenting objects in a wide variety of tasks [23], [24], [25]. We selected this method for its robustness and its availability for fast prototyping and testing, although our methodology is not limited to a particular detector, and other alternatives for image segmentation could be used. We pre-trained the network with the COCO dataset for instance segmentation [26]. When training our models, we used Stochastic Gradient Descent as optimizer with a base learning rate of 10^{-3}, a momentum of 0.9, and a weight decay of 10^{-4}. The output of Mask-RCNN is a list with the detected objects on the image, indicating for each of them its class and the set of pixels on the image (mask) corresponding to the object.

2) Component assembly: After image segmentation, we combine the objects detected with Mask-RCNN to obtain a single BJA, which allows us to improve the robustness of our method for object detection. We apply a heuristic that takes advantage of the shape of a BJA, which is known a priori. A BJA is made up of two `arms` attached to a central `elbow`, and two additional `elbows` attached to the `arms` (5 components in total). We take the set of masks corresponding to `arms` and `elbows` on the image, according to the Mask-RCNN detection results. If more than 5 objects are detected, then we apply the following procedure: 1) we compute the distance from each `elbow` to each `arm`; 2) we select as the central elbow of the BJA the one with the closest distance to two of the detected `arms`; 3) the arms associated with the central elbow are marked as the arms of the BJA; 4) from the remaining objects, the closest `elbow` to each BJA arm is selected as the other elbow of the arm. After discarding spurious detections, the algorithm returns the filtered list of detected objects, corresponding to the valid `arms` and `elbows`.

C. Failure detection

Providing relevant metrics to evaluate the state of the BJA is one of the main contributions of this work. Our proposal is the result of the experience acquired over the years in the analysis of failures in components of the CSP plants by the company Virtualmech [1].

1) Leak detection: When the HTF leaks over the BJA, it leaves permanent stains that can only be removed by cleaning or changing the component. Leakages do not occur at any time of the day, and they depend on several conditions, such as the position of the BJA and the thermal expansion of the system. As solar plant inspection is conducted on specific time intervals, it is usually the case that the HTF does not leak during inspection. Therefore, robust methods are needed that can detect this failure using the data captured during the inspection of the plant.

Our leak detection algorithm is based on computer vision techniques for color-based filtering applied to the BJA detected with the Object Detection module. First, we mark the pixels on the image with a color in the cooper palette[2] as possible locations of HTF leaks, as the color of the HTF is similar to the brown color. This filter itself is not enough, as brown is a common color in CSP plants (e.g., soil and soil reflections in the mirrors). Therefore, we obtain a precise location of the HTF stains by applying the object masks resulting from our object detection method. As we know the number of pixels corresponding to an HTF leak, and the area of the `arms` and `elbows` on the image, we can compute the percentage of stains covering each component of the BJA.

2) Geometrical analysis: Apart from HTF leaks, another important failure in BJAs is related to *anomalous* geometrical shapes adopted by its components during solar tracking. The right angular geometry of the BJA could be obtained theoretically with computational models, depending on the specific solar conditions in the inspection instant. Then this expected BJA shape could be compared with the one detected during the inspection. However, developing these computational models is time consuming, and their evaluation requires extensive computational resources. We acknowledge that our method can be used along with computational models, but we propose a different strategy based on field measurements to ease the failure detection procedure.

The geometry of the BJA is defined as a triangle with vertices in the center of the *elbows* detected with the Object Detection module (there are three elbows per BJA). We propose the angle corresponding to the central elbow as the evaluation metric, because this value is scale invariant. With this information, two different methods are used to assess the geometry of the BJA:

- Compare the geometry of assemblies from consecutive locations in the plant. As they are inspected almost at the same time of the day, their shape should be similar.
- Select a BJA and record it during solar tracking. The resulting angles can be used to compile the interval of

[1] https://virtualmech.com/

[2] https://matplotlib.org/stable/users/explain/colors/colormaps.html

valid operational values, which is then used in future inspections to evaluate the geometric state of the BJA.

The second method is intended to be used during the installation of a new component. Alternatively, it can also be used for a more detailed inspection of BJAs with possible failures.

III. Experimental Results

In this section, we present several experiments to evaluate the performance of our methodology, more specifically the modules for object detection, leak detection, and geometrical analysis. All images were captured with an H20T camera (with a resolution of 4056×3040 pixels) on board a DJI M3T UAV, except for the tracking experiment for geometrical analysis, where a static GoPro camera was used. The neural networks were trained and evaluated on a computer with 40GB of GPU and 32 cores. All evaluations are carried out over the images from the test set.

A. Object detection

Table I depicts the results obtained with our object detection algorithm without the use of the Component Assembly (CA) module. We consider four neural networks, depending on the data used for training: (1) refers to the original dataset, without using any data augmentation; (2) refers to the Synthetic Background dataset; (3) adds the data resulting from Joint Augmentation (JA); and (4) is the same as (3) but including an additional post-training phase in the original dataset for transfer learning. Notice we do not evaluate JA as a single component because the objects generated with JA must be added into a custom background. We measure the standard metric AP (Average Precision) defined in the COCO dataset [26], where AP-`arm` and AP-`elbow` only consider the class `arm` and `elbow`, respectively.

The results show the importance of using synthetic data in conjunction with real data. (3) and (4) get the best results, the latter being more accurate for image segmentation. Compared to using only real-world data, the AP increases around 15 units, while AP-`elbow` improves in more than 20. As elbows are essential for component assembly as well as geometrical analysis, this is a key reason to select (4) as the network to evaluate the methods for failure detection presented in the following sections.

The Component Assembly module improves the detection of BJAs by removing spurious detected elements. In our application, the number of removed detections is not that relevant in comparison to the total number of legitimate detections. Therefore, the CA module does not significantly vary the accuracy results in terms of typical detection metrics. In our case, only the SB option reached a considerable increase in the proposed metrics (i.e., ≥ 0.01): 0.06 and 0.07 in AP-`arm` for bounding box detections and image segmentation, respectively, also improving the metrics considering both classes. However, CA is still key for our methodology as it allows us to compile the full BJA, which is later used for leak detection and geometry analysis. We show a visual confirmation about the advantages and disadvantages of the CA module in Fig. 5. For networks with low detection capabilities, CA is not appropriate as it would not be capable of obtaining the full ball joint structure. See, for instance, Fig. 5 (a), where the network fails to detect the central `elbow`. However, networks with better performance will benefit from using the CA module to remove spurious objects, as in Fig. 5 (b), (c), (d), where the object detected in the left bottom corner would be filtered. Interestingly, the confidence assigned by the network to this spurious detection differs according to the training data. Networks trained with real CSP images result in lower confidence (better) as these are typical objects in a PT plant, while networks trained on synthetic data are more confident in their wrong predictions because these objects were not present in the background images used to build the datasets. Finally, confidence for the spurious object in Fig. 5 (a) is lower than confidence in Fig. 5 (d) because the network (a) was trained on the real images 16 times more than the network (d).

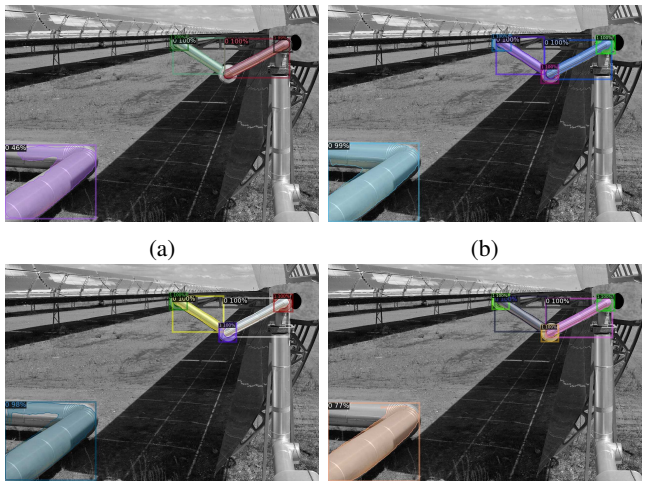

(a) (b) (c) (d)

Fig. 5: Output of the 4 trained networks before applying the CA module. (a) Trained in the original dataset, (b) trained in the SB dataset, (c) trained in SB+JA, and (d) trained in SB+JA and post-trained in the original dataset. The confidence in class assignment obtained by Mask-RCNN for the spurious detected object (bottom left of each image) is: (a) 46%, (b) 99%, (c) 98%, (d) 77%. Background in grayscale and zoom into the region of interest to improve visualization.

B. Leak detection

We determine the stains produced by HTF leaks through a Color-based Filtering (CF). When combined with the output of the CA module, it is possible to assess the state of `arms` and `elbows` individually. However, to provide a robust strategy to evaluate our methods, we propose the following classification task:

- Label every image in the test set according to leak impact: `low`, `high`. Images with low leak impact are those with small or non-visible stains in the BJA

Training data	Task									
	Bounding Box Detection					Image Segmentation				
	AP-arm	AP-elbow	AP	AP50	AP75	AP-arm	AP-elbow	AP	AP50	AP75
(1) Original	70.87	64.71	67.72	94.95	80.93	58.58	30.51	44.55	85.77	40.07
(2) SB	78.69	65.30	72.00	98.96	87.43	70.90	36.08	53.49	94.77	54.10
(3) SB + JA	**84.22**	71.19	77.70	**99.00**	93.26	**71.72**	41.01	56.36	94.91	58.98
(4) SB + JA + post-train original	81.70	**75.67**	**78.69**	97.52	**95.97**	66.80	**51.35**	**59.08**	**97.01**	**67.82**

TABLE I: Object detection results considering different datasets for training. The Mask-RCNN algorithm used outputs both bounding boxes and a segmented image; we report both accuracy results separately.

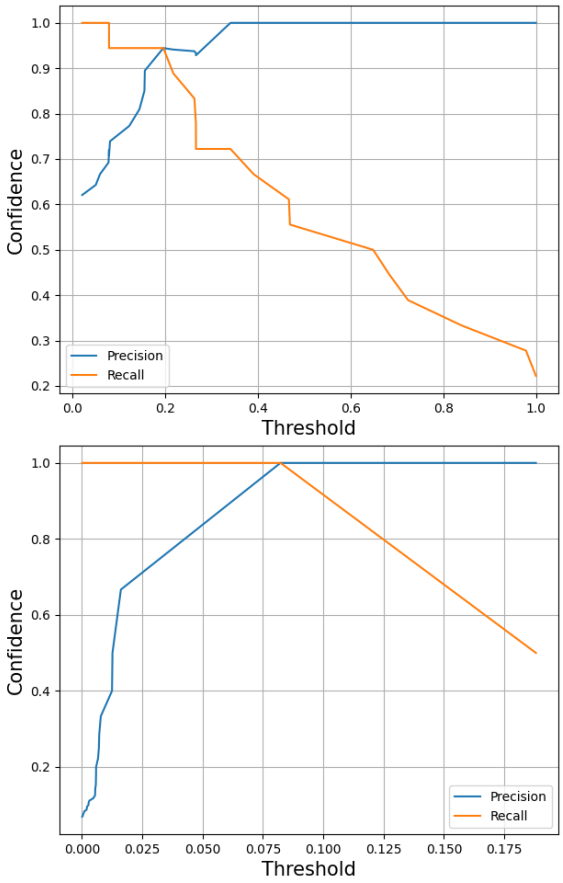

Fig. 6: Precision and recall confidence depending on the classification threshold. The leak impact classification is based on the stained percentage in the elbows (upper), and the arms (lower).

surface.

- Using CF + CA, obtain the stained percentage covering the `elbows`, and the stained percentage covering the `arms` on every image.
- Evaluate the output of our algorithms in the labeled dataset.

We created the labels for the classification task manually, by human visual inspection, obtaining a total of 18 images associated with class `high` and 11 with class `low`. An image of the test set where one of the elbows was not detected by our deep neural network was not used, as the CA module cannot provide an accurate output with missing detections.

The percentage of stain in each individual component is evaluated according to the algorithm described in Section II-C.1. We assume this percentage as a confidence value for our classification task. However, we do not know beforehand which threshold is appropriate to separate each class. The typical value 0.5 is not adequate in our case, as a lower number might also be an indicator of a relevant HTF leak. Therefore, we plot the precision and recall curves using different threshold values. As the elbows and arms have different dimensions, the same amount of HTF leak results in different stained percentages in terms of the total area of the individual component. Hence, we make two individual plots (see Fig. 6) that allow us to define the proper threshold values to evaluate our algorithm in the proposed classification task.

		Predicted Label	
		low	**high**
True Label	**low**	10	1
	high	1	17

TABLE II: Confusion matrix for the task of predicting the leak impact in the whole BJA.

We determined that an image has a `high` leak impact if the percentage of stains on the elbows is greater than or equal to 0.19 (19%), or the percentage of stains on the arms is greater than or equal to 0.8 (8%); otherwise we consider it to be of class `low`. Table II depicts the confusion matrix obtained for the classification task, with only two errors. In conjunction with visual validation, this result confirms the capabilities of our HTF leak detection method. Our algorithm is not only accurate, but also allows us to obtain a numerical assessment of the inspected component. This value can be used throughout different inspection periods to evaluate the evolution of the leaks in a BJA.

C. Geometrical analysis

We evaluate the performance of our algorithm for geometrical analysis with two tests. First, we select a set of consecutively located BJAs in the solar plant, and the images associated with them. In the set, we ensure that at least one of the BJAs has an anomalous geometrical position. Then we evaluate our method by computing the central angle in all the images to compare their differences. As the images were captured at similar times of the day (it can be observed by the position of the mirrors), the central angle in all of the images should be similar. Figure 7 depicts the output of our method for a set of 6 different consecutive BJAs in a solar plant. The BJA associated with the image with the red number is evidently at a different angle, which is correctly

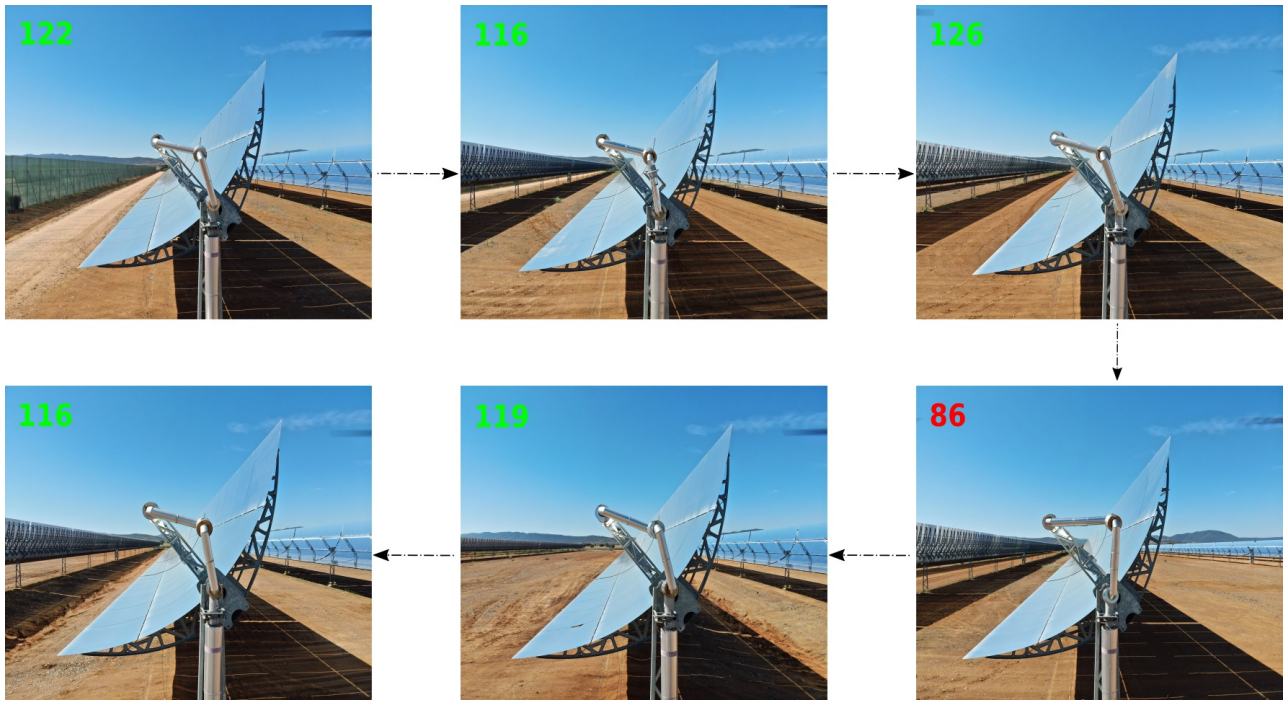

Fig. 7: Consecutive BJAs in a solar plant. Images of neighboring elements are indicated with a dashed arrow. The top left numbers indicate the central angle detected by our system: in green, the angles within a similar range, and in red, the anomalous value.

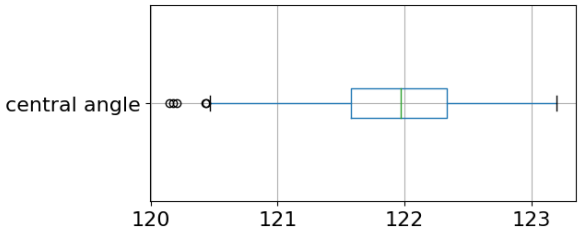

Fig. 8: Measured values for the central angle of a BJA recorded while tracking the sun.

identified by our method (more than 30° difference from the rest of the images).

We designed a second test to evaluate our method with a short video for a fixed BJA. As the SCE tracks the Sun, we compute the angle of the central elbow during the whole tracking process. Since the video has a short duration, angle values are expected to be within a small range. The goal of the test is twofold: 1) assess the robustness of our method, and 2) provide a valid operational range for the recorded period. Figure 8 shows the results obtained. Our method provides stable values for the angle during the tracking, with a range of less than 4° between the smallest and the largest values. This indicates that we can provide valid operational ranges for the BJA angle values, which can be used in future inspections to analyze the geometry of the system. The full video of this experiment is available online [3].

[3] https://youtu.be/7dXW_Pkthjw

IV. Conclusions

In this work, we have presented a novel methodology for automatic inspection of Ball Joint Assemblies in Parabolic-Trough solar plants, using images collected by a UAV in real plants. To increase our training data, we have developed a new data augmentation technique based on Synthetic Background and Joint Augmentation, and we have demonstrated its positive impact improving the accuracy of our deep learning models for image segmentation. Moreover, we have combined our algorithms for object detection with computer vision techniques to determine faulty ball joints with fluid leaks or anomalous geometrical shapes. We designed two metrics to quantify the impact of these leaks and geometrical anomalies, respectively. Our experimental results show that we can detect leaks automatically with a low failure rate when compare with human labeling. We also demonstrate that our algorithm is able to detect ball joint components with faulty angular geometry, compared to neighboring collectors or during solar tracking. Our angular metric for geometry analysis allows us not only to detect the malfunctioning components but also to establish a valid range of operation for future inspections.

We believe that this work represents a significant step forward in the automation of the inspection of ball joint assemblies in Parabolic-Trough plants, which is a relevant industrial problem not adequately covered in the literature. As a future work, we will study other key related problems, like the geometrical analysis in 3D, detecting ball joints with loss of insulation, or detecting fumes.

REFERENCES

[1] F. Rodríguez, W. D. Chicaiza, A. Sánchez, and J. M. Escaño, "Updating digital twins: Methodology for data accuracy quality control using machine learning techniques," *Computers in Industry*, vol. 151, p. 103958, 2023.

[2] S. Ruiz-Moreno, A. J. Gallego, and E. F. Camacho, "Artificial neural network-based fault detection and isolation in a parabolic-trough solar plant with defocusing strategy," *Solar Energy*, vol. 262, p. 111909, 2023.

[3] M. El Ydrissi, H. Ghennioui, A. Alae, M. Abraim, I. Taabane, A. Farid, *et al.*, "Dust InSMS: Intelligent soiling measurement system for dust detection on solar mirrors using computer vision methods," *Expert Systems with Applications*, vol. 211, p. 118646, 2023.

[4] D. Kesseli, V. Chidurala, R. Gooch, and G. Zhu, "A combined computer vision and deep learning approach for rapid drone-based optical characterization of parabolic troughs," *Journal of Solar Energy Engineering*, vol. 145, no. 2, p. 021008, 2023.

[5] G. Espinosa-Rueda, J. L. N. Hermoso, N. Martinez-Sanz, and M. Gallas-Torreira, "Vacuum evaluation of parabolic trough receiver tubes in a 50 mw concentrated solar power plant," *Solar Energy*, vol. 139, pp. 36–46, 2016.

[6] C. Q. Gómez Muñoz, A. Arcos Jimenez, F. P. García Marquez, M. Kogia, L. Cheng, A. Mohimi, and M. Papaelias, "Cracks and welds detection approach in solar receiver tubes employing electromagnetic acoustic transducers," *Structural Health Monitoring*, vol. 17, no. 5, pp. 1046–1055, 2018.

[7] M. El Ydrissi, H. Ghennioui, E. G. Bennoua, and A. Farid, "Geometric, optical and thermal analysis for solar parabolic trough concentrator efficiency improvement using the photogrammetry technique under semi-arid climate," *Energy Procedia*, vol. 157, pp. 1050–1060, 2019.

[8] S. Ruiz-Moreno, A. J. Gallego, A. J. Sanchez, and E. F. Camacho, "Deep learning-based fault detection and isolation in solar plants for highly dynamic days," in *International Conference on Control, Automation and Diagnosis (ICCAD)*. IEEE, 2022, pp. 1–6.

[9] C. Prahl, B. Stanicki, C. Hilgert, S. Ulmer, and M. Röger, "Airborne shape measurement of parabolic trough collector fields," *Solar energy*, vol. 91, pp. 68–78, 2013.

[10] C. Prahl, M. Röger, B. Stanicki, and C. Hilgert, "Absorber tube displacement in parabolic trough collectors–a review and presentation of an airborne measurement approach," *Solar Energy*, vol. 157, pp. 692–706, 2017.

[11] Y. Wu, A. Kirillov, F. Massa, W.-Y. Lo, and R. Girshick, "Detectron2," https://github.com/facebookresearch/detectron2, 2019.

[12] M. Pérez-Cutiño, J. Valverde, and J. Díaz-Báñez, "Detecting broken receiver tubes in CSP plants using intelligent sampling and dual loss," *Applied Intelligence*, vol. 53, no. 24, pp. 29 902–29 917, 2023.

[13] M. Lamghari, A. Kriouile, H. Bouzekri, and S. Rachidi, "Innovative approach to perform optimized and highly efficient operation & maintenance activities for 100mw-scale CSP plants," in *SolarPACES: International Conference on Concentrating Solar Power and Chemical Energy Systems*, vol. 2126, no. 1, 2019, p. 120009.

[14] "Fire at Casablanca CSP plant," https://www.hoy.es/prov-badajoz/incendio-casablanca-valdecaballeros-20220702213951-nt.html, accessed: 2023-09-01.

[15] "Fire at Manchasol II CSP plant," https://www.eldiario.es/castilla-la-mancha/extinguen-horas-incendio-declarado-plantas-termosolares-alcazar-san-juan_1_9233030.html, accessed: 2023-09-01.

[16] H. Jo, Y.-H. Na, and J.-B. Song, "Data augmentation using synthesized images for object detection," in *International Conference on Control, Automation and Systems (ICCAS)*. IEEE, 2017, pp. 1035–1038.

[17] P.-Y. Chen, J.-W. Hsieh, M. Gochoo, and Y.-S. Chen, "Mixed stage partial network and background data augmentation for surveillance object detection," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 12, pp. 23 533–23 547, 2022.

[18] N. Li, F. Song, Y. Zhang, P. Liang, and E. Cheng, "Traffic context aware data augmentation for rare object detection in autonomous driving," in *International Conference on Robotics and Automation (ICRA)*, 2022, pp. 4548–4554.

[19] J. Li, J. Zhang, S. J. Maybank, and D. Tao, "Bridging composite and real: towards end-to-end deep image matting," *International Journal of Computer Vision*, vol. 130, no. 2, pp. 246–266, 2022.

[20] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask R-CNN," in *IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 2961–2969.

[21] Z. Zhang, Q. Liu, and Y. Wang, "Road extraction by deep residual U-Net," *IEEE Geoscience and Remote Sensing Letters*, vol. 15, no. 5, pp. 749–753, 2018.

[22] J. Wang, K. Sun, T. Cheng, B. Jiang, C. Deng, Y. Zhao, D. Liu, Y. Mu, M. Tan, X. Wang, *et al.*, "Deep high-resolution representation learning for visual recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 43, no. 10, pp. 3349–3364, 2020.

[23] N. Lüling, D. Reiser, A. Stana, and H. W. Griepentrog, "Using depth information and colour space variations for improving outdoor robustness for instance segmentation of cabbage," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 2331–2336.

[24] X. Xu, M. Zhao, P. Shi, R. Ren, X. He, X. Wei, and H. Yang, "Crack detection and comparison study based on faster R-CNN and mask R-CNN," *Sensors*, vol. 22, no. 3, p. 1215, 2022.

[25] B. Serhan, H. Pandya, A. Kucukyilmaz, and G. Neumann, "Push-to-see: learning non-prehensile manipulation to enhance instance segmentation via deep q-learning," in *International Conference on Robotics and Automation (ICRA)*, 2022, pp. 1513–1519.

[26] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft COCO: common objects in context," in *European Conference on Computer Vision (ECCV)*, 2014, pp. 740–755.