

Direct Matrix Factorization and Alignment Refinement: Application to Defect Detection

Zhen Qin
University of California, Riverside
zqin001@cs.ucr.edu

Peter van Beek Xu Chen
Sharp Laboratories of America
{pvanbeek, chenx}@sharplabs.com

Abstract—Defect detection approaches based on template differencing require precise alignment of the input and template image; however, such alignment is easily affected by the presence of defects. Often, non-trivial pre/post-processing steps and/or manual parameter tuning are needed to remove false alarms, complicating the system and hampering automation. In this work, we explicitly address alignment and defect extraction jointly, and provide a general iterative algorithm to improve both their performance to pixel-wise accuracy. We achieve this by utilizing and extending the robust rank minimization and alignment method of [12]. We propose an effective and efficient optimization algorithm to decompose a template-guided image matrix into a low-rank part relating to alignment-refined defect-free images and an explicit error component containing the defects of interest. Our algorithm is fully automatic, training-free, only needs trivial pre/post-processing procedures, and has few parameters. The rank minimization formulation only requires a linearly correlated template image, and a template-guided approach relieves the common assumption of small defects, making our system very general. We demonstrate the performance of our novel approach qualitatively and quantitatively on a real-world dataset with defects of varying appearance.

Keywords-Defect Detection; Industrial Inspection; Template Matching; Image Alignment; Matrix Factorization

I. INTRODUCTION

Defect detection on fabric, electronic panels, surface mounted devices (SMD), printed circuit board (PCB), etc. is a critical procedure in industry. Highly accurate, consistent and scalable automatic defect detection systems are needed to improve process yield and product quality. However, designing such a system is very challenging, due to the highly unpredictable characteristics of real-world defects. As hinted in Fig. 1, real-world defects possess extraordinary variability in terms of appearance, shape and size, and are often specific to the object class inspected. It is desirable to design a general defect detection method that works for various known or even unknown defects and inspected object types.

Motivated by the following reasons, we seek a defect detection solution based on the template differencing approach:

This work was done during the first author’s internship at Sharp Laboratories of America.



Figure 1: Examples of various real-world defects, indicated by the green ellipses. Defects may have arbitrary shapes, be specific to the object class examined, possess weak intensities, and have large sizes. The inspected objects themselves also vary strongly.

- This approach is very general and makes no assumptions about defect/object characteristics, by only seeking to detect “differences” between the input image and a defect-free template.
- Training data is generally not needed. In our case, we only need a single defect-free template image, which is usually not a problem. Acquiring large amounts of (labeled) training data can be problematic and time-consuming.

Template differencing approaches to defect detection highly rely on accurate alignment between the input and template images. To the best of our knowledge, all existing defect detection approaches apply alignment as a pre-processing step. However, alignment can easily be affected by the existence of defects, especially if the defects are large or have strong edges/texture. In order to tackle this problem, we propose to conduct alignment refinement while handling estimated defects in an explicit error measurement, jointly optimizing for both. We utilize and extend a method termed Robust Alignment by Sparse and Low-rank decomposition (RASL) [12], which was proposed for batch alignment of hundreds of images. Our formulation is aimed at optimal decomposition of an image matrix that contains both the input and template image into a low-rank part relating to aligned defect-free images and an error component that captures the defect signal. Extracting a defect mask from the error component becomes trivial and requires no post-processing.

Our optimization approach applies recent advances in robust rank minimization that have caught much attention

in the computer vision community. The work of Robust Principle Component Analysis (RPCA) [3] showed that it is possible to efficiently recover low-rank matrices despite significant corruptions. Further, it has been shown that the claim holds when entry-wise noise exists [21]. In these works, the initial formulation of optimizing a matrix rank and l_0 norm is relaxed to optimizing a nuclear norm and l_1 norm. Very recent work [19] [20] shows that this relaxation may not be necessary and direct factorization is possible.

Our contributions include:

1. We propose an effective and efficient optimization framework for defect detection termed Direct matrix Factorization and Alignment Refinement (DFAR). We iteratively improve both alignment and defect detection performance using an efficient algorithm. To the best of our knowledge, our application of robust rank minimization (jointly with alignment) to industrial defect detection is a first in literature.
2. Our derivation initially follows that of RASL ([12]); however, our algorithm is different from RASL in several aspects. Our algorithm does not rely on the relaxation of matrix rank to nuclear norm and l_0 norm to l_1 norm; instead we apply direct factorization as in [19] [20]. Our formulation considers entry-wise noise, not considered in RASL. We construct the image matrix from both input images and template images; this template-guided approach relaxes the sparsity assumption of the defects. We found that directly applying batch alignment did not work.
3. We demonstrate the precision-recall performance and efficiency of our algorithm with several experiments on a real-world industrial defect image dataset. We compare our approach to RASL as well as to several existing (template-based and template-free) defect detection methods.

The proposed algorithm (DFAR) is very general, training-free, only assumes a template image that is linearly correlated with respect to the input, and provides accurate alignment and defect localization, relieving the need for non-trivial application-specific pre/post-processing procedures. DFAR only has a few parameters that have clear meaning, making the algorithm easy to use.

Before presenting the proposed defect detection approach in detail in Sec. II and experimental results in Sec. III, we discuss related work in the following subsection.

A. Related Work

Several defect detection methods based on template differencing have been proposed. Often, the input image is firstly registered with a defect-free template using various registration algorithms, followed by localized direct differencing [13], optical flow [17], normalized cross correlation (NCC) [16], wavelet-based processing [6], or Hausdorff distance measurement [5]. Unfortunately, existing defect

detection methods are still largely application-specific. For example, [13] and [17] are aimed only at semiconductor wafers; [6] and [5] focus on SMD and PCB assembly respectively. Furthermore, most systems occasionally need training and manual intervention [17], mostly work only for small defects, and most are sensitive to parameter changes.

These undesirable features result from various assumptions made during algorithm design, for example: assumptions about the existence of some fixed landmark points for alignment [5], or assumptions about the nature of the inspected surface [16]. More importantly, almost all existing template differencing methods assume a superior registration/alignment performance; but even state-of-art alignment algorithms might cause misalignment and raise false alarms to some extent. Therefore, in practice, manual adjustment (then perfect alignment is assumed), the assumption of small defects (which shall hurt alignment performance less), and/or non-trivial pre/post-processing procedures are applied, introducing complexity and reducing generality.

Several methods other than template differencing are the following.

- Detection based on supervised classification is one major category of detection methods for targets such as human and vehicles [18][7]. In real-world industrial scenarios, it is difficult to gather a reasonable size of training samples with labeled defect masks. Given the high intra- and inter-class variance of possible defects, designing features is also difficult.
- Defect detection based on saliency detection [2][14][9][4] also poses challenges. State-of-art saliency detection methods typically estimate a coarse and subjective saliency support on natural images which is far from the industrial accuracy requirement. Also, these methods usually make a number of assumptions.
- Anomaly detection methods such as PHOT [1], which analyze images in the Fourier domain, may only locate small defects on uniformly textured or periodic patterned images, such as fabric surface. We shall show in experiments that our algorithm generalizes to such types of images.

II. PROPOSED DEFECT DETECTION APPROACH

We first give an overview of the detection system and introduce the problem formulation of decomposing vectorized image matrix to achieve simultaneous alignment refinement and defect estimation. Then we introduce our effective and efficient optimization algorithm with steps that can be solved by off-the-shelf algorithms.

A. System Overview

Our proposed framework is presented in Fig. 2. We first use a hierarchical framework that locates the input candidate on the whole template image. After evaluating a number of existing registration techniques, we find [8] provides best

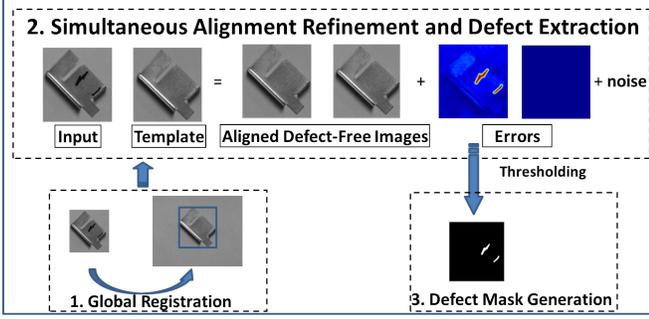


Figure 2: An overview of our system. After initial registration, alignment is refined while maintaining an explicit measurement of the errors, where defect mask can be easily generated from. Notice the misalignment and illumination difference between the input and template images. For the error images, color closer to red indicates higher intensity.

overall accuracy and scalability and we omit the details here. However, even with this state-of-art registration method, initial misalignments still generally exist. We next describe our proposed method for simultaneous alignment refinement and defect extraction, given the input image and the corresponding (but not perfectly aligned) template image.

B. Problem Formulation

Suppose we have the well-aligned, defect-free single-channel input image I_1^0 and template images $I_2^0 \dots I_n^0 \in R^{w \times h}$, we define $vec : R^{w \times h} \rightarrow R^m$ as the operator that stack the corresponding pixels as a vector. The single-channel images may be grayscale or luminance data, or simply contain the green channel of an RGB image. We use multiple template images as will be explained in Sec. II-D. Then the matrix $A \doteq [vec(I_1^0) | \dots | vec(I_n^0)] \in R^{m \times n}$ should be low-rank. Low-rank indicates that the input image should be linearly correlated with the template image. Another way to view it is that the columns of matrix A should be nearly constant regardless of global intensity changes due to illumination.

However, the images we observe are neither well-aligned nor defect-free, which can be represented as $I_i = (I_i^0 + e_i + \epsilon_i) \circ \tau_i^{-1}$, where e_i is an additive error component wherein we intend the defects are contained and is assumed to be sparse, $\tau_i^{-1} \in \mathbb{G}$ is a transformation that models the practical misalignment, and ϵ_i models the noise. In this work, \mathbb{G} corresponds to the similarity, affine or planar homography group of parametric transforms. Now our aim is to decompose the aligned observed image matrix $D \circ \tau \doteq [vec(I_1 \circ \tau_1) | \dots | vec(I_n \circ \tau_n)] \in R^{m \times n}$ as

$$D \circ \tau = A + E + \epsilon, \quad (1)$$

where $E \doteq [vec(e_1) | \dots | vec(e_n)] \in R^{m \times n}$ and ϵ models real-world entry-wise noise. In other words, we want to

decompose the aligned observed images $D \circ \tau$ into: a low-rank component A that should relate to the defect-free background, a sparse error component E that we expect to contain the defects, and a noise term. A direct formulation of the problem can be posed as a constrained optimization problem:

$$\begin{aligned} \min_{A, E, \tau} & \|D \circ \tau - E - A\|_F \\ \text{s.t.} & \text{rank}(A) \leq K \\ & \|E\|_0 \leq \gamma \end{aligned} \quad (2)$$

where $\|\cdot\|_F$ is the Frobenius norm, K is the rank constraint on the low-rank approximation A , and γ is the maximal number of non-zero entries in E . Intuitively, we can approximate the low-rank component reliably, since $D \circ \tau - E$ can be viewed as the aligned image matrix excluding the defects. We will directly solve this optimization problem in the primal form.

C. Optimization

We solve this problem effectively and efficiently without relaxing either the rank or the cardinality constraint, or referring to the Lagrangian. We term our algorithm as direct factorization and alignment refinement (DFAR).

First, since the dependence of $D \circ \tau$ is complicated on the transformations τ , when the change in τ is small, we can approximate the dependency by linearizing about the current estimate of τ , as commonly done in the alignment literature [15] [12]. Then the optimization problem becomes

$$\begin{aligned} \min_{A, E, \Delta\tau} & \|(D \circ \tau + \sum_i J_i \Delta\tau_i \mu_i^T - E) - A\|_F \\ \text{s.t.} & \text{rank}(A) \leq K \\ & \|E\|_0 \leq \gamma \end{aligned} \quad (3)$$

where $J_i \doteq \frac{\partial}{\partial \delta} vec(I_i \circ \delta)|_{\delta=\tau_i}$ is the Jacobian of the i -th image with respect to the transformation τ_i , and μ_i is the i -th standard basis for R^n (for a compact presentation). Since the linearization only holds locally, we repeatedly linearize about the current transformations and solve the problem of the form of Eq. 3.

Given the current estimate of transformations τ and error E (the very first transformation is identity transform and all elements of E are zeros), we take advantage of the nice decomposable structure of our formulation and apply block coordinate descent with respect to A , E and, $\Delta\tau$. The resulting algorithm is described in Alg. 1 and Alg. 2. Note that Alg. 1, the outer loop, is very similar to that of RASL [12], and is not discussed further.

The solution to step 1) in Alg. 2, a low-rank approximation problem, is directly and optimally given by standard truncated Singular Value Decomposition (SVD) approximation to the matrix $(D \circ \tau + \sum_i J_i \Delta\tau_i \mu_i^T - E)$, or we can accelerate this process by referring to partial SVD algorithms such as [10].

Algorithm 1: DFAR - Outer Loop

Data: Images I_1, \dots, I_n , initial transformations τ , the initial outliers E , Maximal rank K , Maximal number of errors γ

Result: Low-rank component A , Error E , Transformation τ

while not converged do

1) compute the Jacobian w.r.t. transformations

$$J_i \leftarrow \frac{\partial}{\partial \delta} \left(\frac{\text{vec}(I_i \circ \delta)}{\|\text{vec}(I_i \circ \delta)\|_2} \right) \Big|_{\delta=\tau_i}, \quad \text{for } i = 1, \dots, n;$$

2) warp and normalize the images

$$D \circ \tau \leftarrow \left[\frac{\text{vec}(I_1 \circ \tau_1)}{\|\text{vec}(I_1 \circ \tau_1)\|_2} \cdots \frac{\text{vec}(I_n \circ \tau_n)}{\|\text{vec}(I_n \circ \tau_n)\|_2} \right];$$

3) Inner Loop: see Algorithm 2

$$(A, E, \Delta\tau) \leftarrow$$

$$\arg \min_{A, E, \Delta\tau} \|(D \circ \tau + \sum_i J_i \Delta\tau_i \mu_i^T - E) - A\|_F$$

$$\text{s.t. } \begin{aligned} \text{rank}(A) &\leq K \\ \|E\|_0 &\leq \gamma \end{aligned}$$

4) Update transformations $\tau \leftarrow \tau + \Delta\tau$

end

Algorithm 2: DFAR - Inner Loop

Data: Warped Data Matrix $D \circ \tau$, the current E , Maximal rank K , Maximal error γ

Result: Updated low-rank component A , Error E , Transformation Refinement $\Delta\tau$

while not converged do

1)

$$A = \arg \min_A \|(D \circ \tau + \sum_i J_i \Delta\tau_i \mu_i^T - E) - A\|_F$$

$$\text{s.t. } \text{rank}(A) \leq K$$

2) $\Delta\tau =$

$$\arg \min_{\Delta\tau} \|\sum_i J_i \Delta\tau_i \mu_i^T - (E + A - D \circ \tau)\|_F$$

3)

$$E = \arg \min_E \|E - (D \circ \tau + \sum_i J_i \Delta\tau_i \mu_i^T - A)\|_F$$

$$\text{s.t. } \|E\|_0 \leq \gamma$$

end

To solve step 2) in Alg. 2, since the transformations are applied to each image individually, the unconstrained minimization problem is simply a least square problem with closed-form solution for each image i . We use the Moore-Penrose pseudoinverse here.

The solution to step 3) in Alg. 2, the error detection problem with l_0 -norm constraints, can also be solved effectively and efficiently by using [11]: compute a histogram of the entries of E , and then find a threshold value such that the histogram quantile (cumulative sum) above that threshold corresponds to γ out of mn pixels. Then we use that threshold to set other (small) entries (pixels) of E to zero.

This basically only requires one $\frac{\gamma}{mn}$ quantile computation.

The above three procedures are easy to implement. Each procedure provides the global optima for each step within the feasible region, and thus the algorithm is going to converge. Notice that since we do not apply relaxations to the constraints, the formulation is non-convex and may get trapped in a local minimum. As a remedy, we initialize DFAR by running RASL (a convex relaxation) for few iterations and are able to get better results than using DFAR or RASL alone. The running time of DFAR (with initialization time included) is also less than that of applying RASL alone under the same stopping criterion, as discussed in Sec. III-E.

D. Template-Guided Matrix Decomposition

We use multiple template images (actually, identical copies of a single template image) to construct the template-guided data matrix D . This strategy improves defect detection performance for two reasons:

First, when decomposing the image matrix containing only two images, the defects might be captured by either error image, as the objective values will be the same. We usually observe weaker errors in both error components. However, by using $(n - 1)$ template images, the low-rank component will be guided towards the template, since turning on one element of the template error image means $(n - 1)$ times the cost. This results in a more complete and stronger defect response in the input error image.

Second, robust rank minimization methods assume sparse errors in the entire data matrix (see Eq. 2) in general. When the defect is large (say occupying $m\%$ of the input, where $m > 80$), it is hardly sparse in terms of the whole matrix with two images (still possessing $\frac{m}{2}\%$ of the elements). However, when we use $(n - 1)$ template images, the defects only possess $\frac{m}{n}\%$ of the entries. In other words, even if the defects are not sparse in terms of the input, the template-guided trick can always make the defects sparse for the whole data matrix. We empirically find that $n = 4$ works very well in practice. Fig. 3 illustrates the difference between the two strategies. By applying this trick, our algorithm generalizes to large defects. Although we spend more running time as the data matrix becomes larger, the huge performance gain obtained for defect detection is worthwhile, as will be shown in the experiments.

E. Downsampling

Image down-sampling is a common technique to obtain acceleration, often at the expense of accuracy. We found that our algorithm achieves a slight accuracy improvement when applied on down-sampled images, while greatly reducing processing time. We identified two reasons for this fact.

First, linearization of the transformation works well when the initial misalignment is not too large [15]. After down-sampling, the initial misalignment is significantly reduced, leading to improved results.

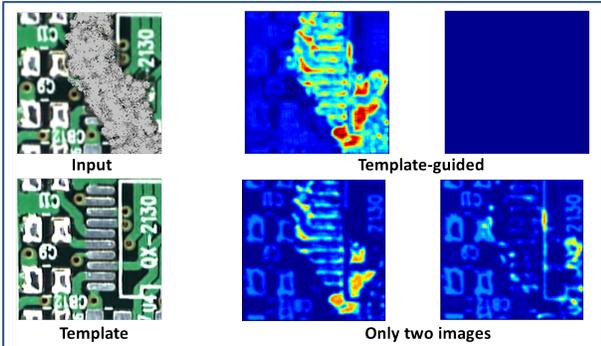


Figure 3: Different decomposition results from the two strategies. Images on each row on the right are the error images of the input and the template. Decomposing two images results in weaker and incomplete error components, especially for large defects.

Second, in practice there exist small local geometric deviations between the input and the template, which may cause false alarms. Such small distortions are less evident in the down-sampled images.

F. Parameter Setting

Our algorithm has two explicit parameters (K and γ in Eq. 2). Both parameters have clear meaning and thus are easy to set (K the rank constraint of the low-rank matrix, and γ the maximum percentage of pixels that contain defects). K is simply set to be 1, as we expect the defect-free input image to be linearly correlated with the low-rank component of the template. As for γ , the cardinality constraint on errors, is set to be $\frac{1}{n} \times (\# \text{ of elements of } D)$, as we only expect entries in the input image to be labeled as defects.

III. EXPERIMENTS

We demonstrate the superior accuracy and generality of our framework over existing template-based and template-free defect detection methods on real-world industrial datasets.

A. Qualitative Results

Our framework is general as we make few specific assumptions about the data. Here we first select several examples showing the strong generality of our method on different application scenarios in Fig. 4. Our method generalizes to textured images, metal objects, circuit boards, and other objects and materials.

B. Comparative Results on Real-world LCD TV Panel Dataset

We conducted extensive comparative experiments on a challenging real-world LCD TV panel dataset, which consists of 152 images (1024×768 pixels) almost evenly

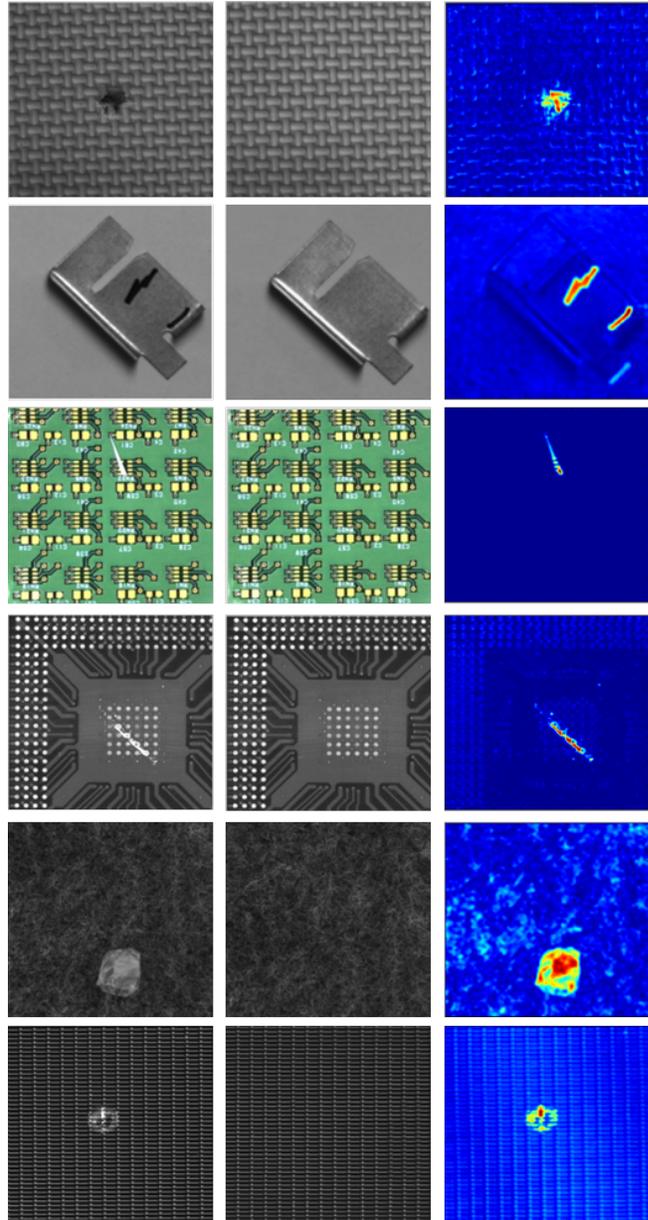


Figure 4: Sample images on different application scenarios using our method without any parameter tuning from Sec. III-B. For each row are the input, the template and the input error image (shown as a pseudo-color map overlaid on the input image). Notice the initial misalignment for several cases.

distributed into four types of defects and includes defect-free images. The complete template image used for registration is of the size 2160×1785 . Each image consists of circuit patterns, slowly varying background, and possibly defects. The circuit patterns are generally linear structures (intersecting lines, and rectangular structures), with strong edges, illustrated in Fig. 5. Such structures are sensitive to

misalignment during defect detection. We are unable to show complete input images (with defects) in this paper due to confidentiality reasons. We describe the four defect types and their characteristics in Tbl. I.

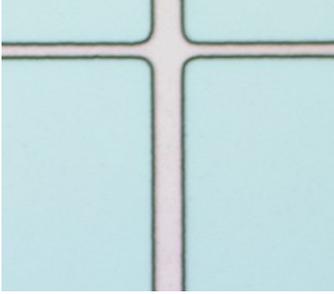


Figure 5: Small crop of input image illustrating a defect-free region with a circuit pattern.

Table I: Characteristics of defects in our LCD TV panel image dataset.

Defect	Description
Type 1	Very dark blobs, often small, often somewhat circular
Type 2	Superfluous material with same color as circuit lines
Type 3	Large transparent rainbow patterns with saturated colors
Type 4	Small dots patterns, with dark intensity

The abbreviations of the examined algorithms and their respective assumptions made are summarized in Tbl. II. DF is the proposed method without alignment refinement. RASL is the state-of-art on rank minimization and alignment refinement method but was never explored in defect detection. We used the same template-guided data matrix for DFAR, DF and RASL. PHOT is the state-of-art for template-free defect detection on textured images. We adapted PHOT to the template-based scenario, using PHOT-Refined. Since we are not using textured images, PHOT raises severe false alarms on edge structures. Thus after registration, we ran an edge detector on the resulting template and suppressed the corresponding edge pixels on the input, largely removing false alarms on the edges. NCC and Golden Matching are standard defect detection algorithms used in industry. Whenever global registration is needed, we use the state-of-art work of [8].

We show sample result images in Fig. 6, illustrating detection performance of DFAR, NCC, and PHOT qualitatively. The figure only shows cropped output images, again for confidentiality reasons. In the output images for the DFAR algorithm, we generally observe a very strong response in the defect region pixels, and a much weaker response in non-defect region pixels. This qualitatively indicates a very high precision and good recall.

For quantitative performance, we report precision-recall curves in Fig. 7 for each method. Precision and recall were determined for each image by varying thresholds on the

Table II: Methods tested and associated assumptions underlying each method: (1) Linearly correlated template image, (2) Superior initial alignment, (3) Uniformly textured input image, (4) small defects, (5) complex background without low-contrast regions, (6) identical template and (7) edge structure.

Method	Assumptions
DFAR (proposed)	(1)
DF (DFAR without alignment refinement)	(1)(2)
RASL [12]	(1)
PHOT [1]	(3)(4)
PHOT-Refined	(1)(2)(4)(7)
NCC [16]	(1)(2)(4)(5)
Golden Matching [13]	(2)(6)

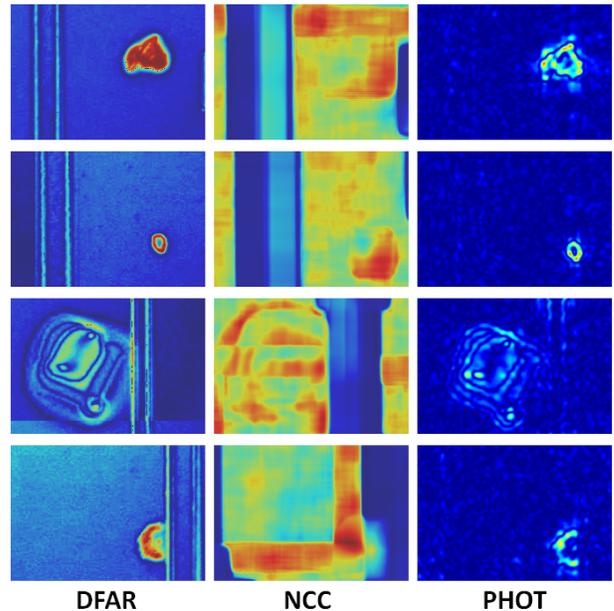


Figure 6: Sample (cropped) result images with defect strength shown with pseudo-color (before thresholding). Left column shows DFAR results, middle column shows NCC results, right column shows PHOT results. Rows correspond to type 1, 2, 3 and 4 defects, from top to bottom.

output error/difference/defect strength image, and comparing the resulting defect pixel masks with manually labeled ground truth images, and subsequent averaging over all images in a category.

We have the following observations and analysis:

- Robust rank minimization algorithm (DFAR, DF and RASL) consistently outperform other methods significantly due to the mild assumptions made, validating our idea of applying robust rank minimization to defect detection.
- DFAR generally performs much better than DF, showing the effectiveness of alignment refinement.
- Our complete algorithm (DFAR) always has comparative

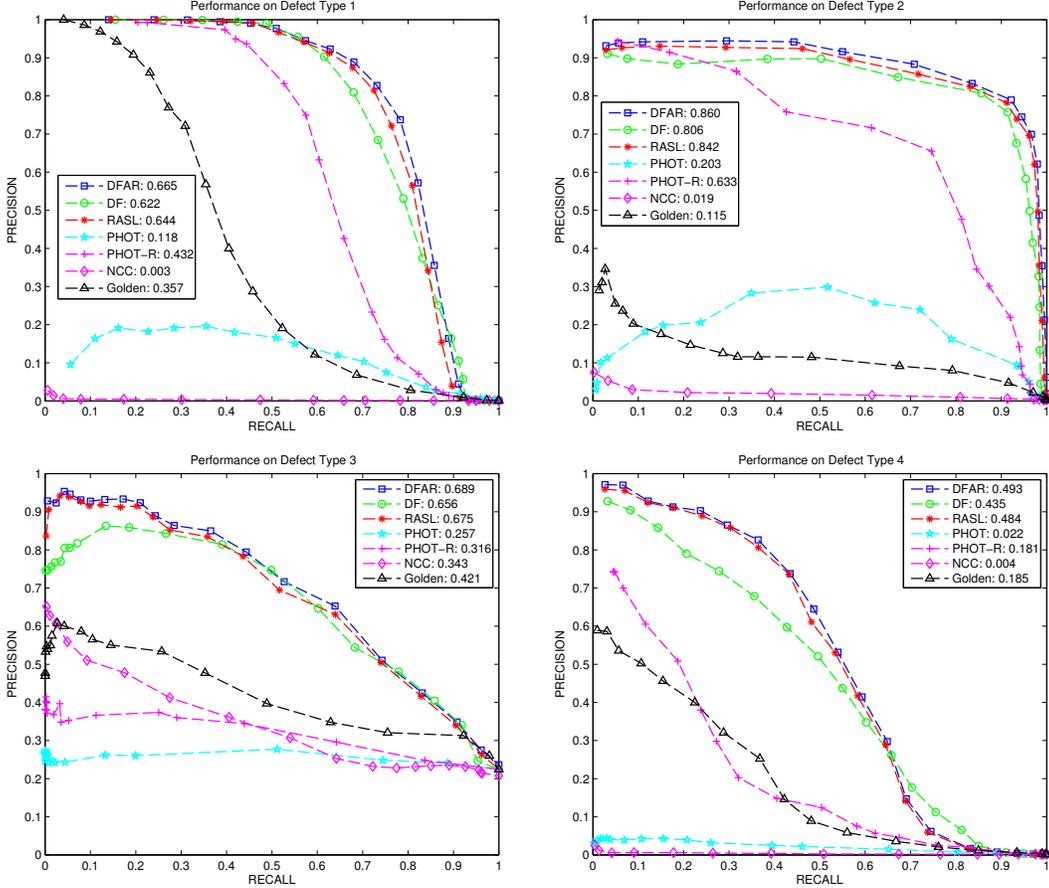


Figure 7: Precision-Recall curves for different defect types in the LCD TV panel dataset.

or slightly better performance than RASL and possesses faster convergence, thus achieving the best overall performance. The performance gain comes from our more faithful (non-relaxed) formulation and entry-wise noise handling.

- NCC performs worst, even though we carefully selected its parameters (such as window size). This is not surprising as NCC is even undefined for uniform regions. For real-world low-contrast regions, noise dominates the response and produces severe false alarms. NCC does not generalize to our inspected objects.
- PHOT does not perform well as it is designed for textured images. We observe a lot of false alarms along strong edges. Our modified PHOT-Refined gets good results especially for defect Type 1 where defects are mostly small. However, there still exists a large gap between its performance and that of our method, because when there is a large defect (e.g. for most Type 3 defects), PHOT would treat the large defect as texture itself resulting in severe miss detections.
- Standard approaches such as Golden Matching provide reasonable results. But is it still far from satisfactory

because of the misalignment, lack of an identical template image and the challenge of our dataset.

As our algorithm provides superior alignment result and avoids false alarms, we can easily set a single threshold to differentiate between images with or without defect. We empirically set a threshold and have not found any outliers in our dataset. For other methods, we find this task to be nontrivial, so we omit the comparison in this paper.

C. Direct Decomposition VS. Robust PCA

One major difference between our work and RASL is, RASL relaxes rank to nuclear norm and l_0 norm to l_1 norm. Here we compare their performance with the same fixed alignments so as to more faithfully evaluate the effect of relaxation. In other words, we remove the alignment refinement of our method and that of RASL to reduce it to Robust PCA [3]. We report the results in Fig. 8.

We can see that the original rank minimization can indeed achieve better performance than the relaxed version with reasons analogous to that analyzed in Sec. III-B.

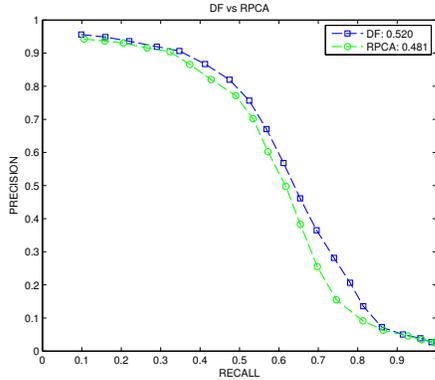


Figure 8: Precision-Recall curve for Direct Factorization and RPCA over all images.

D. Effect of Template-guided Decomposition

Here we evaluate the effectiveness of our idea of decomposing a template-guided data matrix introduced in Sec. II-D. We compare its performance with that of using only two images and report the results in Fig. 9. We can see that the template-guided method results in improved performance.

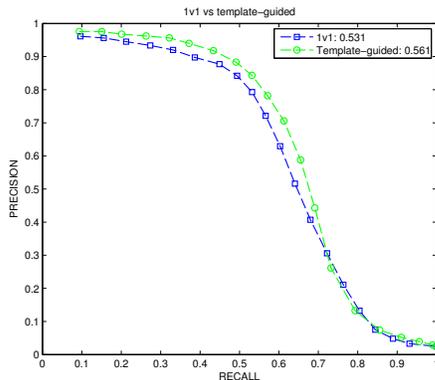


Figure 9: Precision-Recall curve over all images for decomposing template-guided data matrix and that consisting of only two images.

E. Running Time

On a regular PC, our unoptimized Matlab implementation of the proposed DFAR algorithm (including initialization time) takes ~ 40 seconds to process a 1024×768 image, while RASL takes ~ 50 seconds. As discussed in Sec. II-E, we found down-sampling to result in large efficiency gains with essentially no loss in precision-recall performance. When down-sampling by 2, the processing time of DFAR decreased to ~ 8 seconds per image, and when down-sampling by 4, the processing time of DFAR was only ~ 2 seconds.

IV. CONCLUSION

We introduced a fully automatic defect detection system which couples defect extraction and alignment refinement. We proposed an effective and efficient optimization algorithm, as well as several techniques to achieve high accuracy and efficiency. We compared with several classical and state-of-art defect detection methods on challenging real-world datasets and have shown convincing results. The generality, robustness and scalability make our system desirable for real-world industrial inspection tasks.

REFERENCES

- [1] D. Aiger and H. Talbot. The phase only transform for unsupervised surface defect detection. In *CVPR*, 2010. 2, 6
- [2] A. Borji and L. Itti. Exploiting local and global patch rarities for saliency detection. In *CVPR*, 2012. 2
- [3] E. Candes, X. Li, Y. Ma, and J. Wright. Robust principal component analysis? In *Journal of the ACM*, 2011. 2, 7
- [4] L. Chai, H. Zhang, Z. Qin, J. Yu, and Y. Qi. Multi-feature content-based product image retrieval based on region of main object. In *ICICS*, 2011. 2
- [5] C. Chen, S. Lai, S. Liu, T. Ku, and S. Yeh. Optical pcb inspection system based on hausdorff distance. In *SPIE-IS & T Electronic Imaging*, 2005. 2
- [6] H. Cho and T. Park. Wavelet transform based image template matching for automatic component inspection. In *J. Adv. Manuf. Technol.*, 2010. 2
- [7] P. Felzenszwalb, D. McAllester, and D. Ramanan. A discriminatively trained, multiscale, deformable part model. In *CVPR*, 2008. 2
- [8] S. Hinterstoisser, V. Lepetit, S. Ilic, P. Fua, and N. Navab. Dominant orientation templates for real-time detection of texture-less objects. In *CVPR*, 2010. 2, 6
- [9] Y. Jia and M. Han. Category-independent object-level saliency detection. In *ICCV*, 2013. 2
- [10] R. M. Larsen. PROPACK - software for large and sparse SVD calculations. <http://soi.stanford.edu/~rmunk/PROPACK>. 3
- [11] Z. Lu and Y. Zhang. Penalty decomposition methods for l_0 norm minimization. In *Department of Mathematics, Simon Fraser University, Tech. Rep.*, 2010. 4
- [12] Y. Peng, A. Ganesh, J. Wright, W. Xu, and Y. Ma. RASL: robust alignment by sparse and low-rank decomposition for linearly correlated images. In *PAMI*, 2012. 1, 2, 3, 6
- [13] N. Shankar and Z. Zhong. Defect detection on semiconductor wafer surfaces. In *Microelectronic Eng.*, 2005. 2, 6
- [14] X. Shen and Y. Wu. A unified approach to salient object detection via low rank matrix recovery. In *CVPR*, 2012. 2
- [15] R. Szeliski. Image alignment and stitching: a tutorial. In *MSR-TR-2004-92*, 2004. 3, 4
- [16] D. Tsai, C. Lin, and J. Chen. The evaluation of normalized cross correlation for defect detection. In *Pattern Recognition Letters*, 2003. 2, 6
- [17] D.-M. Tsai, I.-Y. Chiang, and Y.-H. Tsa. A shift-tolerant dissimilarity measure for surface defect detection. In *IEEE Trans. on Industrial Informatics*, Feb. 2012. 2
- [18] P. Viola and M. Jones. Robust real-time object detection. In *IJCV*, 2001. 2
- [19] L. Xiong, X. Chen, and J. Schneider. Direct robust matrix factorization for anomaly detection. In *ICDM*, 2011. 2
- [20] T. Zhou and D. Tao. Godec: randomized low-rank & sparse matrix decomposition in noisy cases. In *ICML*, 2011. 2
- [21] Z. Zhou, X. Li, J. Wright, E. Candes, and Y. Ma. Stable principal component pursuit. In *ISIT*, 2010. 2