On Approximate Nearest Neighbour Field Algorithms in Template Matching for Surface Quality Inspection

Gernot Stübl[1], Bernhard Moser[1], and Josef Scharinger[2]

[1] Software Competence Center Hagenberg (SCCH), Hagenberg
{Gernot.Stuebl,Bernhard.Moser}@scch.at
[2] Department of Computational Perception, Johannes Kepler University (JKU), Linz
Josef.Scharinger@jku.at

Abstract. Surface quality inspection is applied in the process of manufacturing products where the appearance is crucial for the product quality and customer acceptance, like for woven fabrics. The predominating approaches to detect defects are feature-based. Recently we investigated an alternative approach utilizing template matching in the context of regular or near-regular textured surface inspection. This paper reveals that the template matching approach belongs to the class of approximate nearest neighbour field (ANNF) algorithms which are common in a different field of image processing, namely structural image editing. By modifying a state-of-the-art ANNF algorithm the advantage of template matching algorithms for defect detection can be shown. Furthermore the importance of the chosen distance function is demonstrated in an explorative study and a concept to determine if the template matching approach is suitable for a given texture and defect type is demonstrated on a set of defect classes and texture types.

1 Introduction

Surface quality inspection is an important issue for producing industrial woven textiles. Typically the predominating approaches in this field are either hard to configure or rely on a high amount of training data. Recently a new approach based on template matching was proposed which needs only one defect-free reference image and has very low configuration effort [13]. However it had not been stated for which defect types the approach is suitable.

After an overview of the state-of-the-art, this paper explains the relation between the proposed approach and approximate nearest neighbour field (ANNF) algorithms in Section 2. Furthermore a modified state-of-the-art ANNF algorithm is utilized to show a lower bound of what is reachable with the template matching approach for defect detection. Thereafter a concept to determine whether the approach is suitable for a given texture/defect combination is presented. Since there is a need for a distance function in the algorithm the question which one can be used will be answered in Section 3 in an study. A conclusion as well as an outlook completes the work.

R. Moreno-Díaz et al. (Eds.): EUROCAST 2013, Part II, LNCS 8112, pp. 79–86, 2013.

1.1 Approximate Nearest Neighbor Fields

According to Barnes et al. [1] a nearest-neighbour field (NNF) is a function $f : A \mapsto \mathbb{R}^2$ of offsets, defined over all possible patch coordinates in an image A. For a patch coordinate $\mathbf{a}$ in A it gives the corresponding nearest neighbor $\mathbf{b}$ in image B, for some distance function d. Usually f is implemented as look-up table in two arrays (one for each coordinate) of the same size as the image A. There are various applications of NNFs ranging from structural image editing [1] to de-noising of images [3].

Typically an exact solution of a NNF is not needed which favours the use of an approximate nearest neighbour field (ANNF). However even the computation of an approximation is time consuming, see Kumar et al. [8] for a survey on tree based algorithms before 2009. In 2009 the foundation of a new generation of ANNF algorithms was given through the introduction of Patchmatch by Barnes et al. [1]. This algorithm utilizes image coherency by propagating matching results to nearby queries, which enables the ANNF generation to happen in interactive speed. A generalization [2] was introduced later by the same author.

Unfortunately the original definition is prune to get trapped in local minima and uses only short-distance propagation, which results in imprecise results. So Korman and Avidan [6] picked up the idea of neighbourhood propagation and presented a hashing based algorithm named Coherence Sensitive Hashing which overcomes these limitations. The neighbourhood propagation idea has also been combined with tree data structures in the work of Olonetsky and Avidan [11] as well as He and Sun [5]. In the work of the latter a new fast solution for the exact NNF problem is given too. Recently Sureka et al. [14] showed that the original Patchmatch enhanced with a mixed resolution approach cures its former drawbacks and brings the same or better results as the newer algorithms.

The original Patchmatch definition is based on the L_2 distance but also other distances are applicable. However most of the successor algorithms gain their speed through fast approximations to the L_2 distance.

1.2 Template Matching for Defect Detection in Fabrics

Automated defect detection of fabrics is a widespread topic. It targets a lot of products, reaching from woven textiles over laces to air-bag hoses.

A huge amount of algorithms for visual defect detection can be found in literature, for surveys see Kumar [7] and the newer one by Ngan et al. [10]. The idea of template matching for defect detection was first reported by Chin [4] for the inspection of printed circuit boards using the Traditional Image Subtraction (TIS) method. With TIS defects are detected by subtracting a perfect aligned golden template from a test image. Sanby et al. [12] developed the idea further by adding a fine alignment through correlation. Other authors like Yazdi and King [15] proposed mechatronical solutions. Ngan was the first who used area-wise comparison instead of pixel-wise comparison. In his proposed golden image subtraction (GIS) method the diversities in a window around a pixel are accumulated as energy function and stored in a defectmap which has the size of the test image minus the window size. This approach does not use fine alignment. Template matching in general became an outsider in fabric defect detection because of its sensitivity to noise as reported by most of the authors.

Recently we investigated a new template matching approach suitable for near regular textures based on Hermann Weyl's discrepancy norm as noise robust dissimilarity measure [13]. For further information on the discrepancy norm and its properties, see Moser [9]. The approach is area based and utilizes a fine registration step with a main idea reverse to the previous publications: The test image is split into overlapping patches of at least periodicity size and each patch is registered on the defect-free reference image. The residual cost of registration is then stored as energy function in the defectmap. Therefore if a good registration for a patch beneath a mean registration error can be achieved then it is considered defect free. The main benefit is the low configuration afford.

An appropriate speed is gained through the use of a coarse-to-fine structure. However full pixel-wise evaluation is computationally still too expensive, so a lock-wise working principle is utilized. This is also the major drawback of the algorithm, since block processing results in a coarse defectmap, see the second column of Figure 1.

2 Patchmatch for Defect Detection

Reconsidering the finest level of the mentioned template matching approach, the registration of periodicity sized patches in the reference image generates an approximate nearest neighbour field. So both the template matching approach and Patchmatch do basically the same job. The main difference is that Patchmatch utilizes neighbourhood propagation and has a pixel-wise working principle.

Having this in mind a modified Patchmatch with adjustable window size and discrepancy norm (DN) as distance function can be used to get an impression of the template matching approach's advantage. It is of no importance that the original Patchmatch gives imprecise results, as long as the results are below a mean registration error.

Without giving an exact runtime complexity the execution time of Patchmatch is empirically below a pixel-wise but higher than a block-wise template matching approach. This is sound because the template matching approach does not utilize neighbourhood propagation.

Defectmaps generated with DN-Patchmatch are shown in the third column of Figure 1. The results are more precise than with the template matching approach. Nevertheless it raises the question why Patchmatch has to be run with the discrepancy norm as distance function. In the template matching approach this was clear since the algorithm gained speed by using an optimization step which utilized the monotonicity property of the discrepancy norm. However Patchmatch does not rely on any optimization step. Therefore the fourth column of Figure 1 depicts defectmaps for the same samples calculated with the L_2 norm. Any of Patchmatch's current successor algorithms would give similar results since they are approximating L_2. Examining the results the discrepancy

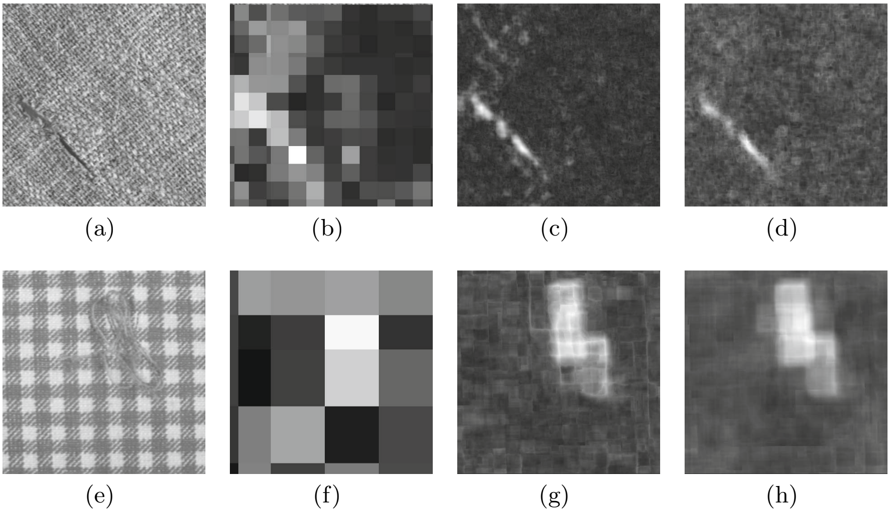

Fig. 1. Processed sample images (first column) with template matching approach (second column) and a modified Patchmatch algorithm with discrepancy norm (third column) as well as L_2 norm (fourth column) as distance functions. The localization of the template matching approach is imprecise due to the block processing principle. Patchmatch results are more precise in the localization and on these samples the discrepancy norm brings better contrast of defects than the $L2$ norm.

norm defectmaps seem to have a better contrast than the ones generated through L_2 norm. This is caused by two factors: the distribution of dissimilarity values and the relationship between dissimilarity values and defect patterns.

The effect of a better contrast using discrepancy norm has been already mentioned in previous publications, see [13]. One indication for this behaviour is the probability distribution (PDF) function of the discrepancy norm values. Figure 2 depicts a Monte Carlo estimation of the discrepancy and the L_2 norm PDF. The distributions have been calculated using 500000 random samples for 8×8px patches in a value range for difference patches of $[-255, 255]$. Please observe the scaling of the x-axis. While the L_2 norm produces a Gaussian curve around its mean, the discrepancy norm PDF is positively skewed and placed more to the left side. Therefore if defects produce high distance values, they can be better distinguished from the mean value.

Whether a defect produces high distance values is important, since this defines for which defects the detection algorithm is suitable. High values of the L_2 norm are intuitive for humans since they simply represent the sum of squared differences. On the other side the discrepancy norm gives high values if the difference between two patches contains a huge partial sum area. As this is not easy to interpret an explorative study with different defect types has been carried out which is described in the next section.

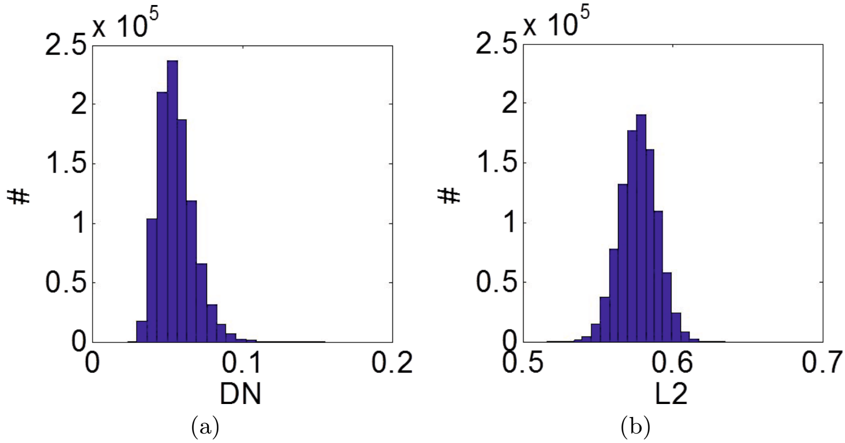

Fig. 2. Monte Carlo estimation of probability density functions for L_2 and discrepancy norm for 500000 random samples of 8×8px patches in a value range of $[-255, 255]$. Please observe the scaling of the x-axis. The distribution of L_2 norm is Gaussian around the mean, while the discrepancy one is positively skewed. Assuming that a defect has a value right away from the mean, there is a better chance to detect it with discrepancy norm than with L_2.

3 Experimental Evaluation for Quality Inspection of Woven Fabrics

As already mentioned, approximate nearest neighbor field algorithms like Patchmatch and the template matching approach can be used to detect a specific defect, if the defect versus the texture produces a high distance value and if this distance value sets itself off against the mean value. Both facts are determined through the distance function in use. In the following explorative study the discrepancy and the L_2 norm are tested regarding their capability to distinguish defective from non-defective regions in textures of defined defect types. The followed procedure can be seen as concept to check if a defect is detectable using the template matching approach.

The evaluation is carried out on a small set of defect classes. For industrial usage a more comprehensive study has to be performed. The defect types are based on the TILDA[1] C2 textures, namely *hole*, *color defect*, *open seam*, *thread on texture*, and *lint on texture*. A sixth type was added for illustration purpose which is not based on TILDA but on the DAGM 2007 Challenge C6[2]. Each class contains ten examples with hand labelled defects. To compare both distance functions, two classes of distance values are created for each distance function d: class C_0 between random patches in defect free regions and class C_1 between random patches of

[1] Available from Universität Freiburg, Institut für Informatik, Lehrstuhl für Mustererkennung und Bildverarbeitung (LMB).

[2] Available from `http://klimt.iwr.uni-heidelberg.de/dagm2007/prizes.php3`

defect free and defective regions. The number of class members is chosen empirically to keep the results in various random trials stable.

A distance function is assumed to have good separation ability if the accuracy of an one dimensional linear classifier trained on the generated classes is good. For the classifier LibSVM with a linear kernel and $c = 100$ is used. Since the classifier has only descriptive usage an extra test set was omitted.

Table 1. This table gives the mean accuracy in percent for defect classification on different defect types, evaluated with discrepancy norm (DN), L_1, L_2 and combinations DN/L_1,DN/L_2. As classifier LibSVM with linear kernel and $c = 100$ is used. Using DN gives good results, except for defect types *thread on texture* and *open seam*, where the template matching approach does not seem to be suitable. A combination of DN and L_2 would give slightly better results than using DN alone, while a combination of DN and L_1 brings no benefits.

	Accuracy				
Class Name	**DN**	L_2	L_1	**DN/L_2**	**DN/L_1**
Hole C2R2	77.55	66.09	58.19	78.93	75.75
Color defect C2R2	86.02	71.79	61.23	89.28	83.27
Open seam C2R2	57.32	51.97	50.75	58.77	56.88
Thread on texture C2R2	57.12	52.87	51.30	61.12	56.97
Lint on texture C2R2	88.05	85.47	82.42	89.42	86.47
DAGM 2007 C6	75.50	57.28	51.80	79.80	65.92

The first and second columns in Table 1 summarize the results. For the classes *hole*, *color defect*, and *lint on texture* DN is the better choice. Both distance functions perform bad on classes *open seam* and *thread on texture*. It seems that the template matching approach with the two distance functions is not applicable on these classes. The most surprising results are with class *DAGM 2007 C6*. While for L_2 norm the approach does not seem to be suitable, for DN the accuracy is quite good. However this cannot be generalized since there might be texture/defect combinations, where this is reverse.

When plotting the distance values of L_2 and DN against each other as in Figure 3(a) it looks promising to utilize a two dimensional classifier with both distance values as features. The forth row of Table 1 shows the result for a DN/L_2 combined classifier, which has slightly better accuracy than the DN classifier alone. Unfortunately the combination doubles the costs for distance computation. Nevertheless due to the construction principle of the discrepancy norm the L_1 norm can be computed simultaneously with very little overhead by adding one variable and one instruction. Since the L_1 and the L_2 norm behave very similar, see second and third column in Table 1 and Figure 3(b), a combination of L_1 and DN is suggested. The results of that combination can be seen in the fifth column of Table 1. It shows up that there is no improvement for the classification for the examples using L_1 and DN together, nevertheless there might occur a defect/texture combination where the classifier can benefit from the diverse nature of the two distance functions.

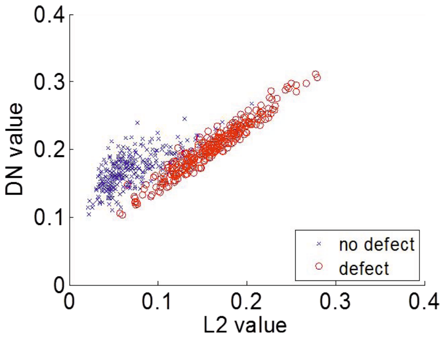

(a)

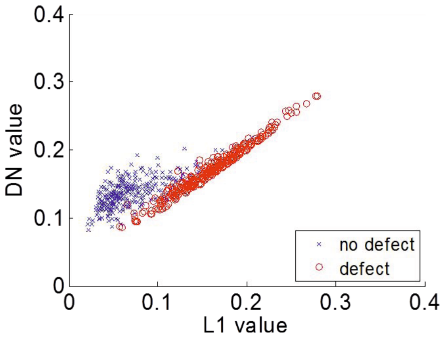

(b)

Fig. 3. In Subfigure (a) C_0 and C_1 for TILDA C2R2E1N12 are plotted with discrepancy and L_2 norm values on the axis, Subfigure (b) shows the same with L_1 values. The behaviour of L_1 and L_2 is similar and visually a two dimensional linear classifier would be able to separate the classes.

4 Conclusion and Outlook

In this paper it was revealed that the recently published template matching approach for defect detection can be seen as an approximate nearest neighbour field (ANNF) algorithm like the state-of-the-art in this field Patchmatch. The original Patchmatch was modified to illustrate the advantage of ANNF algorithms in the domain of defect detection. Since the proposed template matching approach suffers from imprecise results a combination of both algorithms is advised: template matching for pre-detection of suspicious blocks and Patchmatch for the exact location of the defect. This is an interesting topic for future research.

Furthermore it was shown that the separation between defective and non-defective regions depends strongly on the used distance function. In an explorative study a way to check whether ANNF algorithms for a given texture and defect type makes sense is demonstrated on a small set of defect classes and example images. The evaluated distance functions are the L_1, L_2 and DN norm. Since the construction principle of the discrepancy norm is very different to the L_p norm family also the sensitiveness to various defect classes differs. However a way to combine the discrepancy with the L_1 norm is possible with very little computational overhead, which might bring benefit for some types of textures/defects.

Summing up it can be said that ANNF algorithms are useful for defect detection on near regular textures with little configuration afford. For industrial usage the explorative study has to be extended to cover more defect classes.

Acknowledgement. This work was supported in part by the Austrian Science Fund (FWF) under grant no. P21496 N23, and by the European Fund for Regional Development under Regionale Wettbewerbsfähigkeit OÖ 2007-2013. We also want to thank the reviewers for their comments and suggestions.

References

1. Barnes, C., Shechtman, E., Finkelstein, A., Goldman, D.: PatchMatch: A randomized correspondence algorithm for structural image editing. ACM Transactions on Graphics (Proc. SIGGRAPH) 28(3) (August 2009)
2. Barnes, C., Shechtman, E., Goldman, D.B., Finkelstein, A.: The generalized PatchMatch correspondence algorithm. In: Daniilidis, K., Maragos, P., Paragios, N. (eds.) ECCV 2010, Part III. LNCS, vol. 6313, pp. 29–43. Springer, Heidelberg (2010)
3. Buades, A., Coll, B.: A non-local algorithm for image denoising. In: Conference on Computer Vision and Pattern Recognition (CVPR), pp. 60–65 (2005)
4. Chin, R.T.: Automated visual inspection: 1981 to 1987. Computer Vision, Graphics, and Image Processing 41(3), 346–381 (1988)
5. He, K., Sun, J.: Computing nearest-neighbor fields via propagation-assisted KD-trees. In: Conference on Computer Vision and Pattern Recognition (CVPR), pp. 111–118. IEEE (2012)
6. Korman, S., Avidan, S.: Coherency sensitive hashing. In: International Conference on Computer Vision (ICCV), pp. 1607–1614. IEEE (2011)
7. Kumar, A.: Computer-vision-based fabric defect detection: A survey. IEEE Transactions on Industrial Electronics 55(1), 348–363 (2008)
8. Kumar, N., Zhang, L., Nayar, S.: What is a good nearest neighbors algorithm for finding similar patches in images? In: Forsyth, D., Torr, P., Zisserman, A. (eds.) ECCV 2008, Part II. LNCS, vol. 5303, pp. 364–378. Springer, Heidelberg (2008)
9. Moser, B.: Similarity measure for image and volumetric data based on Hermann Weyl's discrepancy measure. IEEE Transactions on Pattern Analysis and Machine Intelligence 33(11), 2321–2329 (2011)
10. Ngan, H.Y., Pang, G.K., Yung, N.H.: Automated fabric defect detection-a review. Image and Vision Computing 29(7), 442–458 (2011)
11. Olonetsky, I., Avidan, S.: Treecann - KD tree coherence approximate nearest neighbor algorithm. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Sato, Y., Schmid, C. (eds.) ECCV 2012, Part IV. LNCS, vol. 7575, pp. 602–615. Springer, Heidelberg (2012)
12. Sanby, C., Norton-Wayne, L., Harwood, R.: The automated inspection of lace using machine vision. Mechatronics 5(23), 215–231 (1995)
13. Stübl, G., Bouchot, J.-L., Haslinger, P., Moser, B.: Discrepancy norm as fitness function for defect detection on regularly textured surfaces. In: Pinz, A., Pock, T., Bischof, H., Leberl, F. (eds.) DAGM/OAGM 2012. LNCS, vol. 7476, pp. 428–437. Springer, Heidelberg (2012)
14. Sureka, H., Narayanan, P.J.: Mixed-resolution patch-matching. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Sato, Y., Schmid, C. (eds.) ECCV 2012, Part VI. LNCS, vol. 7577, pp. 187–198. Springer, Heidelberg (2012)
15. Yazdi, H.R., King, T.: Automatic machine vision for lace inspection. In: Proceedings of Machine Vision Applications, Architectures, and Systems Integration V, vol. 2908, pp. 109–117 (1996)