DOUBLY-ROTATED MATCHED FILTERING

Matthias Mühlich, Thorsten Dahmen, and Til Aach

Institute of Imaging and Computer Vision, RWTH Aachen University, 52056 Aachen, Germany
phone: +(49) 241 80-27860, fax: +(49) 241 80-22200, email: web: www.lfb.rwth-aachen.de,
{muehlich,dahmen,aach}@lfb.rwth-aachen.de

ABSTRACT

Using the concept of steerable filters, it is possible to represent the rotation of a known template – and thus arbitrary orientations of lines or edges – as a linear combination of a fixed number of base filters, thus allowing a computationally efficient detection of 1D structures. In spite of their usefulness in many image processing areas, such structures exhibit one important drawback: Due to the aperture problem, it is impossible to localize such features in images exactly. Junctions or corners, on the contrary, allow to find exact point matches in image sequences, but existing steerable filter approaches are not capable of modeling such templates.

In this paper, we show how to extend rotated matched filtering using steerable filters such that it can represent multi-oriented image structures. As an example, we examine one special type of double-oriented feature which is especially important for applications like registration or camera calibration: the checkerboard pattern.

1. INTRODUCTION: FEATURE DETECTION USING MATCHED FILTERING

In [3], Jacob and Unser presented an approach for detecting edges and lines in images. Their approach is based on *steerable filters* which were introduced by Freeman and Adelson in [1]; these filters allow a computationally effecient solution with a sound theoretical basis in *matched filtering*. We will review both concepts, (rotated) matched filtering and steerable filters, in the first two introductory sections. Then, we will proceed to image structures which can only be described with more than a single adjustable orientation and model such structures as point-by-point product of single-oriented structures (section 3). In section 4, we will extend the theory of steerable filters to *multi-steerable filters*. This allows us to define the counterpart to standard rotated matched filtering, namely the concept of *multiply rotated matched filtering*, in the following section 5. One specific and particularly interesting feature template, e.g. for camera calibration patterns, is the checkerboard template which we will examine in section 6, before we conclude this paper with a summary.

We start the discussion of feature detection using steerable filters with an introduction to matched filtering [5]: The detection of a known template in a signal can be formulated as maximizing the correlation between image patch and template, which can also be expressed as inner product. Assume that we want to detect a known signal $f_0(\mathbf{x})$ in a received signal $f(\mathbf{x})$. The signal model behind the derivation of the optimal filter h_{opt} for this task is that the unknown signal is composed of the known signal and additive white Gaussian noise $\eta(\mathbf{x})$:

$$f(\mathbf{x}) = f_0(\mathbf{x}) + \eta(\mathbf{x}) \ .$$

It can be shown that if we perfom a convolution of the unknown signal $f(\mathbf{x})$ and an arbitrary filter $h(\mathbf{x})$, we obtain the signal-to-noise ratio

$$\frac{S}{N} = \frac{E}{N_0}\rho^2$$

where E is the energy of $f_0(\mathbf{x})$, N_0 is the spectral power density of the noise, and $\rho \in [-1,1]$ is the cross-correlation coefficient of $f_0(\mathbf{x})$ and $h(-\mathbf{x})$. Hence, for the optimal filter $h_{\text{opt}}(\mathbf{x})$, we require $\rho = 1$ and obtain $h_{\text{opt}}(\mathbf{x}) = c \cdot f_0(-\mathbf{x})$ where c is a constant. We observe that the optimal filter is the mirrored version of the known signal. In the following, we will omit the subscript 'opt' and set the constant c to 1:

$$h(\mathbf{x}) = f_0(-\mathbf{x}) \ . \tag{1}$$

For our application of finding certain templates in images, we can define the sought template as a bivariate signal $f_0(\mathbf{x})$ that we want to find in an image $f(\mathbf{x})$. However, as we seek to detect the features even if they are arbitrarily rotated, we have to expand the concept of matched filtering towards *rotated* matched filtering. In order to simplify notation in the following, we introduce the rotation operator $(\cdot)^{\phi}$ that rotates a bivariate function by the angle ϕ:

$$f^{\phi}(r,\theta) = f(r,\theta-\phi) \ . \tag{2}$$

(Note that, in slight abuse of notation, we will always denote the image function as f, regardless whether it is represented in Cartestian or polar coordinates.) This allows the definition of the principle of *rotated matched filtering*: Let a desired feature be represented by a template $f_0(\mathbf{x})$. A measure of how strong this feature is present in an image $f(\mathbf{x})$ at a fixed position $\mathbf{x}_0$ is $A_{\max}(\mathbf{x}_0)$:

$$\begin{aligned} A_{\max}(\mathbf{x}_0) = A\left(\hat{\phi}|\mathbf{x}_0\right) &= \max_{\phi}\left\{\left[f(\mathbf{x}) \circledast f_0^{\phi}(\mathbf{x})\right]_{\mathbf{x}_0}\right\} \\ &= \max_{\phi}\left\{\left[f(\mathbf{x}) * h^{\phi}(\mathbf{x})\right]_{\mathbf{x}_0}\right\} \end{aligned} \tag{3}$$

with '$\circledast$' and '$*$' denoting correlation and convolution, respectively. Each point $\mathbf{x}_0$ in the image is assigned an estimated orientation angle $\hat{\phi}(\mathbf{x}_0)$. At this rotation angle, the cross-correlation of the image patch centered at $\mathbf{x}_0$ and the rotated template is maximized. The corresponding filter kernel is $h(\mathbf{x}) = f_0(-\mathbf{x})$ according to (1). The feature is said to be present at those points $\mathbf{x}_0$ (and aligned with $\hat{\phi}(\mathbf{x}_0)$) where $A_{\max}(\mathbf{x}_0)$ has a (sufficiently prominent) local maximum.

2. STEERABLE FILTERS

Jacob and Unser [3] define a similar criterion and then propose to use templates which can be represented by *steerable filters* [1]. This approach reduces the computational load considerably because such templates allow to represent *arbitrary* rotations by a weighted linear combination of a limited number of base filters. The weights (or interpolation functions) $\beta_d(\phi)$ are determined exclusively by the filter orientation.

However, the class of steerable filters used in [3] is highly restrictive: The authors can only model their templates as a sum of derivatives of the Gaussian function $g(\mathbf{x})$:

$$h(\mathbf{x}) = \sum_{k=1}^{D} \sum_{i=0}^{k} \alpha_{k-i,i} \frac{\partial^{k-i}}{\partial x^{k-i}} \frac{\partial^{i}}{\partial y^{i}} g(\mathbf{x}) \ . \tag{4}$$

The advantage of this class of templates is that all base filters are Cartesian-separable, but unfortunately, the number of base filters is not the minimal number of base filters required to steer a template [1]. Secondly, this class only allows a single steering parameter – the angle ϕ – and therefore cannot model general multi-oriented structures.

Dealing with rotations suggests using polar coordinates. An expansion of $h(r,\theta)$ to a Fourier series is

$$h(r,\theta) = \sum_{p=-P}^{P} a_p(r) \exp(\mathrm{j}p\theta) \ . \tag{5}$$

This derivation directly leads to the formula for the interpolation coefficients in the original work of Freeman and Adelson [1], but it has never been fully exploited for filter design in literature. Some authors focus on Cartesian-separable filters only, while others [4] are interested in *phase-invariant* behaviour [2] which means that the filter response should not depend on the signal orthogonal to some orientation; most importantly, lines and edges should lead to the same energy of the filter response. In contrast to these well-established approaches, we emphasize that the Fourier coefficients $a_p(r)$ in (5) directly represent some *signal model* and therefore, we will directly base our novel steerable filter design approach on the angular function and its Fourier coefficients. With the Fourier coefficients being responsible for the steerability property, it seems logical to optimize these coefficients directly provided that some desired signal model is available. Assuming polar separability, we set

$$h(r,\theta) = a(r)\, h_{\mathrm{ang}}(\theta)$$

with radial function

$$a(r) = \begin{cases} 1 & r \le r_{\max} \\ 0 & \text{else} \end{cases} \tag{6}$$

and idealized angular edge function

$$h_{\mathrm{ang}}^{\mathrm{ideal}}(\theta) = \begin{cases} 1 & 0 \le \theta < \pi \\ -1 & -\pi \le \theta < 0 \end{cases} \ .$$

Note that the radial function $a(r)$ can be designed separately since it does not influence steerability; the disc chosen in (6) is mainly used for its simplicity. Even a dependency on the Fourier coefficient, i.e. $a_p(r)$ instead of a common $a(r)$ would be possible. Polar separability of the whole template simplifies the upcoming derivations, but is not required mathematically. The Fourier approximation of (odd) order P to the angular function $h_{\mathrm{ang}}^{\mathrm{ideal}}(\theta)$ rotated by ϕ is then

$$h_{\mathrm{ang}}^{\phi}(\theta) = \frac{4}{\pi} \sum_{p=1,3,\ldots}^{P} \frac{1}{p} \sin(p(\theta-\phi)) = \sum_{k=1}^{D} \beta_k(\phi) h_{\mathrm{ang}}^{\phi_k}(\theta) \ . \tag{7}$$

Here, the right hand side is the steerable representation where $D = P+1$ is the number of required base filters, $h_{\mathrm{ang}}^{\phi_k}$ are the base filters chosen as rotated versions of the template function and β_k are the interpolation coefficients; for equidistant sampling angles ϕ_k, the formula for β_k can be found in [1]. We now illustrate how such edge representations can be used to model complex image templates.

(a) (b)

Figure 1: Double-oriented features. (a): a checkerboard pattern with orientation angles aligned with the coordinate axes, i.e. $\phi_1 = 0°$ and $\phi_2 = 90°$. (b): checkerboard with $\phi_1 = 20°$ and $\phi_2 = 130°$. For all templates, we define the scaling to go from -1 (black) to 1 (white).

3. IDEALIZED MULTIPLE-ORIENTED FEATURES

In contrast to single-oriented structures – like the edges and lines discussed in [3] –, we are now interested in features that contain more than just one orientation. Such templates can be used to obtain point matches in image pairs or image sequences. One especially interesting template is the checkerboard pattern; for instance, images of checkerboards for camera calibration with the well-known Matlab Camera Calibration Toolbox from Bouguet (`http://www.vision.caltech.edu/bouguetj/calib_doc`, [6]) lead to patterns like the ones shown in fig. 1. We stress that perspective distortions lead to *two* independent orientations – and not just rotations of the pattern in fig. 1a as a whole – if a checkerboard is imaged under arbitrary observation angles.

When dealing with two orientations, the conventional steerable filter approach is insufficient for these features since it is confined to steering only *one* orientation angle. Our novel idea is now to generate multi-oriented features as point-by-point products of single-oriented features. To simplify notation in the following, we introduce the *multiple rotation operator*:

$$f_{\mathrm{multi}}^{\phi_1,\ldots\phi_M} = \prod_{m=1}^{M} f_m^{\phi_m}$$

which means that the function $f_{\mathrm{multi}}^{\phi_1,\ldots,\phi_M}$ consists of the product of M individual functions and each of these functions is rotated by the angle ϕ_m according to our rotation operator $(\cdot)^{\phi}$ defined in (2). We illustrate this concept for the creation of a checkerboard pattern: The pattern $f_{\mathrm{check}}(\mathbf{x})$ in fig. 1a is black (-1) in the first and third quadrant and white (1) in

Figure 2: Creation of a checkerboard pattern as product of two individually rotated idealized edges. If black corresponds to -1 and white to 1, this construction principle does not only hold for $\phi = 0^\circ$ and $\phi_2 = 90^\circ$ (left), but also for *arbitrary* angles like $\phi_1 = 20^\circ$ and $\phi_2 = 130^\circ$ (right).

the second and fourth quadrant. Evidently, $f_{\text{check}}(\mathbf{x})$ can easily be composed of the point-by-point product of two edges where the second edge is rotated by $\frac{\pi}{2}$:

$$f_{\text{check}}^{0,\frac{\pi}{2}}(\mathbf{x}) = f_{\text{edge}}^{0}(\mathbf{x})\, f_{\text{edge}}^{\frac{\pi}{2}}(\mathbf{x}) = \text{sign}(y)\,\text{sign}(-x) = \begin{cases} -1 & \text{in } 1^{\text{st}} \text{ and } 3^{\text{rd}} \text{ quadrant} \\ 1 & \text{in } 2^{\text{nd}} \text{ and } 4^{\text{th}} \text{ quadrant} \end{cases}.$$

The key element of our approach is now that an appropriate scaling of our idealized single-oriented templates (here: to the range $[-1,1]$) means that exactly the same construction principle also holds if we rotate both edges by arbitrary angles, i.e.

$$f_{\text{check}}^{\phi_1,\phi_2} = f_{\text{edge}}^{\phi_1} f_{\text{edge}}^{\phi_2} \tag{8}$$

is a *general* formula for generating checkerboard patterns with two *arbitrary* orientations.

4. INTRODUCING MULTI-STEERABLE FILTERS

The concept of generating multi-oriented features using point-by-point products of single-oriented features directly gives us the idea of approximating the single-oriented features by single-steerable filters. Steering each of these single-steerable filters individually should result in an approximation of the multi-oriented feature with corresponding steering angles. We define:

A *multi-steerable filter* $h_{\text{multi}}^{\phi_1,\ldots,\phi_M}(\mathbf{x})$ is the point-by-point product of single-steerable filters $h_m^{\phi_m}(\mathbf{x})$, each characterized by a set of D_m rotated versions of itself as base templates and corresponding interpolation coefficients $\beta_k(\phi_m)$ which only depend on the steering angle ϕ_m:

$$h_{\text{multi}}^{\phi_1,\ldots,\phi_M}(\mathbf{x}) = \prod_{m=1}^{M} h_m^{\phi_m}(\mathbf{x}) = \prod_{m=1}^{M} \sum_{k=1}^{D_m} \beta_k(\phi_m)\, h_m^{\phi_k}(\mathbf{x}) \,. \tag{9}$$

Exchanging product and summation shows that:

- these multi-steerable filters can again be represented as linear combination of weighted base functions and
- the base functions and interpolation coefficients can be computed as products of M single-steerable base functions and coefficients, respectively.

As an example, we show this for the double-steerable case:

$$\begin{aligned} h_{\text{multi}}^{\phi_1,\phi_2}(\mathbf{x}) &= \sum_{k=1}^{D_1} \beta_k(\phi_1)\, h_1^{\phi_k}(\mathbf{x}) \cdot \sum_{\ell=1}^{D_2} \beta_\ell(\phi_2)\, h_2^{\phi_\ell}(\mathbf{x}) \\ &= \sum_{k=1}^{D_1} \sum_{\ell=1}^{D_2} \underbrace{\beta_k(\phi_1)\,\beta_\ell(\phi_2)}_{\beta_{k\ell}(\phi_1,\phi_2)} \cdot \underbrace{h_1^{\phi_k}(\mathbf{x})\, h_2^{\phi_\ell}(\mathbf{x})}_{h_{\text{multi}}^{\phi_k,\phi_\ell}(\mathbf{x})} \,. \end{aligned} \tag{10}$$

The base functions $h_{\text{multi}}^{\phi_k,\phi_\ell}(\mathbf{x})$ are point-by-point products of the (single-)steerable filter base functions. Analogously, the interpolation functions $\beta_{k\ell}(\phi_1,\phi_2)$ are products of the individual interpolation functions for both factors.

In general, the number of multiple base filters – and hence the number of necessary convolutions – is equal to the number of combinations of the single base filters. If two single-steerable filters require D_1 and D_2 rotated versions for single-steering, we obviously obtain $D_1 \cdot D_2$ combinations. Fortunately, double-oriented features like the checkerboard pattern can be composed of two *equal* sets of individual filters. This reduces the number of distinct combinations of base filters: Double-steerable filters constructed from two identical sets of D single-steerable base filters require only $\frac{D(D+1)}{2}$ multiple base filters.

5. THE PRINCIPLE OF MULTIPLY ROTATED MATCHED FILTERING

In the same way as (single-)steerable filters provided an elegant and computationally efficient implementation for the detection of single-oriented features like edges or lines in [3], we will now show that the much larger class of multi-oriented features allow a similar treatment with multi-steerable filters. As a first step, we consequently need an extension of the principle of rotated matched filtering towards *multiple* rotated matched filtering and define: A multi-oriented feature has M independent orientation angles $\phi_1,\ldots,\phi_M$ and is represented by a template $f_{\text{multi}}^{\phi_1,\ldots,\phi_M}(\mathbf{x})$. A measure of how strong this feature is present in an image $f(\mathbf{x})$ at a fixed position $\mathbf{x}_0$ is $B_{\max}(\mathbf{x}_0)$:

$$\begin{aligned} B_{\max}(\mathbf{x}_0) &= \max_{\phi_1,\ldots,\phi_M} \left\{ \left[f(\mathbf{x}) \circledast f_{\text{multi}}^{\phi_1,\ldots,\phi_M}(\mathbf{x}) \right]_{\mathbf{x}_0} \right\} \\ &= \max_{\phi_1,\ldots,\phi_M} \left\{ \left[f(\mathbf{x}) * h_{\text{multi}}^{\phi_1,\ldots,\phi_M}(\mathbf{x}) \right]_{\mathbf{x}_0} \right\} . \end{aligned}$$

In analogy to (single) rotated matched filtering, we define $h_{\text{multi}}^{\phi_1,\ldots,\phi_M}(\mathbf{x}) = f_{\text{multi}}^{\phi_1,\ldots,\phi_M}(-\mathbf{x})$ to be the corresponding filter kernel and maximize the cross-correlation of the image patch centered at $\mathbf{x}_0$ and the multiply rotated template, thus assigning estimated orientation angles $\hat{\phi}_1(\mathbf{x}_0),\ldots,\hat{\phi}_M(\mathbf{x}_0)$ to each image point $\mathbf{x}_0$. But in contrast to the detection of linear structures in [3], we are not limited by the aperture problem anymore; this also allows *exact* localization of the sought features next to the estimation of orientation angles. Additionally, it is possible to achieve sub-pixel accuracy with refinement steps like fitting a paraboloid to the maximum. However, the individual rotation of more than one orientation angle has introduced two new issues which are not known from standard matched filtering:

- Some multi-oriented features, for instance the checkerboard pattern discussed here, change their mean value if both angles are rotated independently. Hence, the mean value of the image patch influences the optimal orientation angles and the optimal correlation value $B_{\max}$.
- In general, varying the angles also changes the filter energy which can potentially be a problem for the optimization of $B_{\max}$ with repect to the angles: The highest correlation is then biased towards angles for which the filter energy is high.

(a) (b)

Figure 3: As double steerable filters are not necessarily zero-mean filters, it is often advisable to subtract the mean from the image patch. The best fitting checkerboard model for (a) would otherwise be all white, whereas subtracting the mean in (b) allows to find the correct checkerboard pattern.

Fortunately, the second issue is no problem for the idealized checkerboard template because its absolute value (and also its energy) is 1 everywhere. Later on, we will see that checkerboard templates approximated with real multi-steerable filters still do not require energy normalization because they only show slight energy variation unless the Fourier approximation order P is extremely low.

The first issue, on the other hand, requires some careful considerations for checkerboard pattern detection. Let us assume that the "light" area of a modified checkerboard pattern in an image patch is on level 1 and the "dark" area is on level 0.5 as illustrated in fig. 3a. The double-rotated checkerboard filter would then give maximum output for *equal* orientation angles, i.e. $\hat{\phi}_1 = \hat{\phi}_2$, which corresponds to an averaging filter: maximum correlation for "all white" – mathematically correct, but a result which we most probably do not want to obtain. Subtracting the mean value from the image patch in a preprocessing step (see fig. 3b) solves this problem and the checkerboard detector finds the correct angles $\phi_1 = 20°$ and $\phi_2 = 80°$. Additionally, the correlation $B_{\max}(\mathbf{x}_0)$ will be much smaller than for black-white patterns due to the lower contrast.

6. THE CHECKERBOARD PATTERN DESIGNED WITH MULTI-STEERABLE FILTERS

As discussed in section 3, the ideal checkerboard pattern can be composed as a point-by-point product of two rotated ideal edges and in section 2, we created a steerable edge filter by expanding the ideal edge $h_{\text{edge}}^{\text{ideal}}(\mathbf{x}) = \text{sign}(y)$ into an truncated angular Fourier series, thus achieving steerability at the expense of approximation. Combining these two results means that we can design a double-steerable filter respresenting checkerboard patterns if we plug the weights β_k derived for the (single-)steerable approximated edge functions (7) into the double-steerable filter equation (10).

A considerable advantage of our design approach for single-steerable filters using truncated Fourier series (in contrast to other steerable filter designs) is that we have a very flexible control over the ratio between approximation quality versus computation time by increasing or lowering the number of Fourier coefficients, see fig. 4. We observe that increasing the order P, i.e. spending more computation time, allows to adjust the sharpness of the edges to a desired quality level, which can be important for a good feature localization. For comparison, we also illustrate checkerboard templates generated by using a different single-steerable filter design. With edges designed with the approach of [3] in place of the ones designed by Fourier series approximations, we obtain fig. 5: In comparison to our design approach, the edges look blurred and increasing the order D in (4) does not help much, especially when we take into account that the number of base filters for single steerable filters based on derivatives of Gaussians rises quadratically with the number of Fourier coefficients, in contrast to a linear increase for our steerable filter design. Again, we emphasize that the radial function (6) for our filter can be chosen arbitrarily. For instance, $a(r) = c\,e^{-r^2}$ would lead to filters that look roughly similar to the derivative-based approach – but with sharp edges instead of blurred ones.

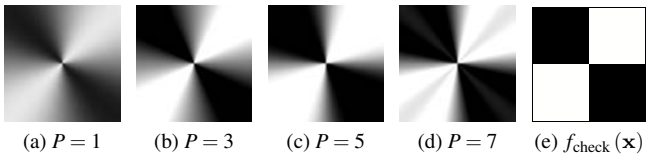

(a) $P=1$ (b) $P=3$ (c) $P=5$ (d) $P=7$ (e) $f_{\text{check}}(\mathbf{x})$

Figure 4: Checkerboard patterns created by two steerable Fourier expansions of edge-functions for Fourier coefficients $p = 1,3,\ldots,P$. The higher P is chosen, the better the steerable filter approximates the idealized template shown in (e). A radial weighting can be added if desired.

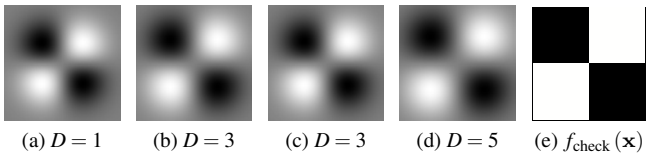

(a) $D=1$ (b) $D=3$ (c) $D=3$ (d) $D=5$ (e) $f_{\text{check}}(\mathbf{x})$

Figure 5: Double-steerable checkerboard templates created by multiplying two perpendicularly aligned single-steerable edges which were design according to the criteria found in [3]. These templates are less suited for feature localization than those shown in fig. 4.

Finally, we regard some checkerboard patterns that are steered to $\phi_1 = 0$ and different angles ϕ_2 and examine the energy variations. In fig. 6, four different realizations for ϕ_2 and the energy variations as a function of the angle ϕ_2 are shown. We observe that if we choose the order P to be high enough, we can neglect the energy variations.

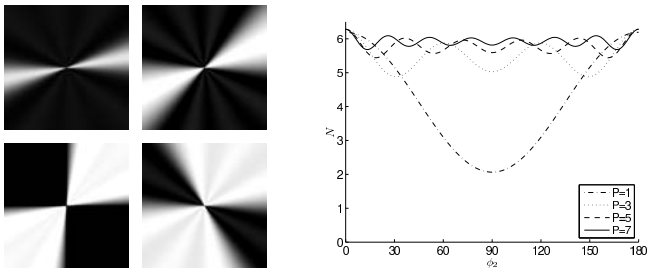

Figure 6: Differently steered checkerboard patterns created by two steerable Fourier expansions of angular sign-functions up to the order $P = 5$ (as in fig. 4c). The angle ϕ_1 is fixed at 0° whereas ϕ_2 varies. Left: $\phi_2 = 20°$, $\phi_2 = 45°$, $\phi_2 = 90°$, and $\phi_2 = 120°$. Right: Filter energy N as a function of ϕ_2 for all filters from fig. 4.

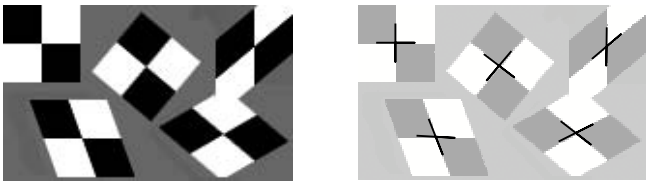

(a) (b)

Figure 7: Detection of several multi-oriented checkerboard patterns: (a) Test image (90×150 pixels) with five differently distorted checkerboards. (b) Pixel-based estimation result. The original image in the background is brightened.

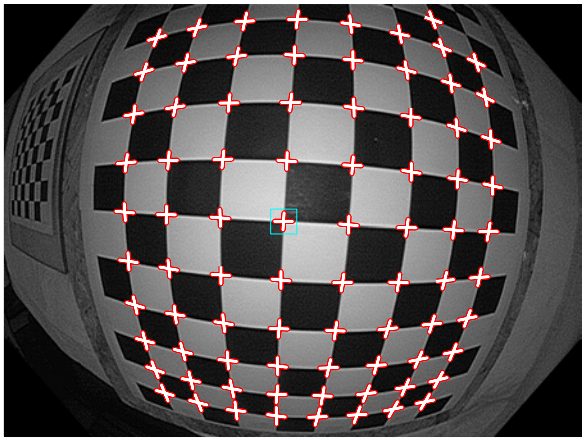

Figure 8: Checkerboard calibration image taken with an endoscope camera and detected checkerboard crossings.

To further demonstrate the viability of our theory, we generated a test image (90×150 pixels) which contains five differently distorted checkerboard patterns (fig. 7a). Our aim is to find the locations of all center points and the orientation angles at those point using the principle of multiply rotated matched filtering. Here, the order of the edge-templates is only $P = 3$ and their size is 20×20. As the steerable edges only have odd angular Fourier coefficients, we conclude that we need $P+1 = 4$ single-base filters to steer each edge. The number of distinct multiple base filters is $\frac{1}{2}(P+1)(P+2) = 10$. Hence, we need to perform 10 convolutions of these base filters and the image. We used the Levenberg-Marquardt algorithm to optimize both angles for each pixel; the initial angles were chosen as $\phi_1 = 0^\circ$ and $\phi_2 = 90^\circ$. All center points were detected correctly, see fig. 7b.

In addition to this straightforward implementation, we realized a more sophisticated approach: specifically for checkerboard patterns, we implemented an approach which first detects some initial crossing in the center of the image and then follows the edges to find neighboring crossings. Using this concept, we are also able to adapt to varying lighting conditions. Fig. 8 shows a real world image where such extensions become necessary: it is a checkerboard calibration pattern taken with an endoscope camera. With our refined approach based on double-steerable filters, we needed 20.4 seconds on a 3 GHz dual Pentium computer to find 81 crossings in this 1280×1008 pixel image (pure Matlab code, no C parts).

7. SUMMARY AND CONCLUSIONS

We presented a general approach to model multi-oriented image features (like junctions, corners or checkerboard patterns) as a point-by-point product of single-oriented features (like lines and edges). Combined with a new single-steerable filter design approach, we were able to transfer the benefits of steerable filtering – namely high estimation quality with low computational complexity – to such multi-oriented templates. This allowed us to introduce the novel concept of *multiply rotated matched filtering* using *multi-steerable filters*. In comparison to feature detection with (standard) steerable filters, this considerably increased the class of image features which our theory can handle properly.

Additionally, we discussed two important properties which are not known from single-steerable filtering: Both mean value and energy of the filters can now depend on the rotation angles, which can introduce some subtle problems for the optimization of the steering angles. Subtracting the DC component in the image patch and normalizing the template energy can solve the problem, but fortunately, for the checkerboard pattern, we could show that a normalization is not necessary if the approximation order for the edge templates is not too low.

Concluding this paper, we state that we have also successfully implemented other multi-steerable templates, including corners and different types of line junctions (L-, T-, X- and Y-junctions), all characterized by two (or in case of the Y-junction: three) *independently* steerable orientation angles [7]. The checkerboard pattern presented here is just one example for a very general theory which, in our opinion, has a wide applicability for feature detection and beyond.

Matlab demonstration code for double-steerable filters can be downloaded at our homepage: `www.lfb.rwth-aachen.de/en/highlights/multi_steerable_filters.html`.

REFERENCES

[1] W.T. Freeman and E.H. Adelson. The design and use of steerable filters. *IEEE Trans. PAMI*, 13(9):891–906, 1991.

[2] G. Granlund and H. Knutsson. *Signal Processing for Computer Vision*. Kluwer, 1995.

[3] M. Jacob and M. Unser. Design of steerable filters for feature detection using Canny-like criteria. *IEEE Trans. PAMI*, 26(8):1007–1019, 2004.

[4] E.P. Simoncelli and H. Farid. Steerable wedge filters. In *Proc. Int. Conf. Computer Vision*, June 1995.

[5] C.W. Therrien. *Decision, Estimation and Classification*. John Wiley and Sons, 1989.

[6] Z. Zhang. A flexible new technique for camera calibration. In *IEEE Trans. PAMI*, pages 1330–1334, November 2000.

[7] M. Mühlich, T. Dahmen and T. Aach. Design of Multi-Steerable Filters and their Application for the Detection of Corners and Junctions. In *Proc. ICIP 2007*, San Antonio, USA, September 16–19, 2007, *to appear*