Normalized Cross-Correlation for Spherical Images

Lorenzo Sorgi[1] and Kostas Daniilidis[2]

[1] Italian Aerospace Research Center
[2] University of Pennsylvania
{sorgi,kostas}@grasp.cis.upenn.edu

Abstract. Recent advances in vision systems have spawned a new generation of image modalities. Most of today's robot vehicles are equipped with omnidirectional sensors which facilitate navigation as well as immersive visualization. When an omnidirectional camera with a single viewpoint is calibrated, the original image can be warped to a spherical image. In this paper, we study the problem of template matching in spherical images. The natural transformation of a pattern on the sphere is a 3D rotation and template matching is the localization of a target in any orientation. Cross-correlation on the sphere is a function of 3D-rotation and it can be computed in a space-invariant way through a 3D inverse DFT of a linear combination of spherical harmonics. However, if we intend to normalize the cross-correlation, the computation of the local image variance is a space variant operation. In this paper, we present a new cross-correlation measure that correlates the image-pattern cross-correlation with the autocorrelation of the template with respect to orientation. Experimental results on artificial as well as real data show accurate localization performance with a variety of targets.

1 Introduction

Omnidirectional cameras [1,2] are now a commodity in the vision and robotics community available in several geometries, sizes, and prices and their images are being used for creating panoramas for visualization, surveillance, and navigation. In this paper we will address the problem of template matching in omnidirectional images. Our methods can be applied to any images which can be exactly or approximately mapped to a sphere. This means that our input images will be spherical images while the template can be given in any form which can be warped so that it has a support on the sphere. The main challenge in template matching is to be able to detect the template under as many geometric or illumination transformations as possible. This can be done by either comparing invariants like moment functions between the template and the image or with statistical techniques [3,4]. It is basic knowledge that we can compute affine invariants and that we can compute an affine transformation from combinations of image moments [5,6] or from Fourier descriptors [7]. Little work has been done on the computation of 3D-rotations from area-based features [8,9,10]. Even if template matching in planar image processing is certainly a problem studied in several different ways, there is hardly any work addressing this problem in omnidirectional imaging.

In this paper we consider 3D rotations, the natural transformations on the sphere, and we try to achieve a normalized cross-correlation to account for linear illumination

T. Pajdla and J. Matas (Eds.): ECCV 2004, LNCS 3022, pp. 542–553, 2004.

changes. A straightforward implementation of normalization in cross-correlation would be space variant, and as such quadratic in the number of image samples and cubic in the tessellation of the rotation space. The contribution of this paper is the development of a matcher that performs in the Fourier domain a sequence of two cross-correlations in order to guarantee the same performance as the normalized cross-correlation and simultaneously to avoid any space-variant operation. By introducing the concept of axial autocorrelation, we have been able to construct a new normalization that preserves the linearity of intensity transformations. The new measure consists of the image-pattern cross-correlation and its subsequent cross-correlation with the axial autocorrelation of the pattern.

The paper is split into 3 parts: Section 2 introduces the preliminary mathematics, Section 3 presents the matcher, and Section 4 analyzes results with real omnidirectional imagery.

2 Mathematical Preliminaries

This section collects the basic mathematical tools required for the formulation and solution of the spherical template matching problem. The proposed method in particular will make use of the analysis of real or complex valued functions defined on S^2, the two-dimensional unit sphere, and their harmonic representation. Throughout the paper we make use of the traditional spherical reference frame, according to which any point on S^2 is uniquely represented by the unit vector $\omega(\theta, \phi) = (\cos\phi \sin\theta, \sin\phi \sin\theta, \cos\theta)$. We denote with $\Lambda(R_{\alpha,\beta,\gamma})$ the linear operator associated with a rotation $R \in SO(3)$ and we use the ZYZ Euler angle parameterization for rotations.

The set of spherical harmonics $\{Y_m^l(\omega), l \geq 0, \mid m \mid \leq l\}$ (see [11]), forms a complete orthonormal basis over $L^2(S^2)$; thus any square-integrable function on the unit sphere $f(\omega) \in L^2(S^2)$ can be decomposed as series of spherical harmonics. The Spherical Fourier Transform and its inverse are defined [11] as:

$$f(\omega) = \sum_{l \in \mathbb{N}} \sum_{|m| \leq l} \hat{f}_m^l Y_m^l(\omega) \tag{1}$$

$$\hat{f}_m^l = \int_{\omega \in S^2} f(\omega) \overline{Y_m^l}(\omega) d\omega. \tag{2}$$

We assume for any spherical image to be square integrable on S^2 and bandlimited in order to use the sampling theorem for spherical function introduced by Driscoll and Healy [12]:

Theorem 1 (Driscoll and Healy). *Let $f(\omega)$ be a band-limited function, such that $\hat{f}_m^l = 0, \forall l \geq B$,, where B is the function bandwidth. Then for each $\mid m \mid \leq l < B$,*

$$\hat{f}_m^l = \frac{\sqrt{2\pi}}{2B} \sum_{j=0}^{2B-1} \sum_{k=0}^{2B-1} a_j^{(B)} f(\theta_j, \phi_k) \overline{Y_m^l(\theta_j, \phi_k)}$$

where the function samples $f(\theta_j, \phi_k)$ are chosen from the equiangular grid: $\theta_j = j\pi/2B, \phi_k = k\pi/B$, and $a_j^{(B)}$ are suitable weights.

Under a rotation $R(\alpha, \beta, \gamma)$, decomposed using the Euler angles parameterization, each harmonic of degree l is transformed into a linear combination of only those Y_m^l with the same degree $\mid m \mid \leq l$, [12]:

$$\Lambda(R_{\alpha,\beta,\gamma})Y_m^l(\omega) = \sum_{|k|\leq l} Y_k^l(\omega) D_{k,m}^l(R) \tag{3}$$

where $D_{m,k}^l(R)$ is defined as:

$$D_{m,k}^l(R) = e^{-im\alpha} d_{m,k}^l(\cos\beta) e^{-ik\gamma} \tag{4}$$

and for a definition of $d_{m,k}^l$, the irreducible unitary representation of SO(3), we refer to [13].

From (3) we can easily show that the effect of a rotation on a function $f(\omega) \in L^2(S^2)$ is a linear transformation associated with a semi-infinite block diagonal matrix in the Fourier domain:

$$h(\omega) = \Lambda(R_{\alpha,\beta,\gamma}) f(\omega) \Leftrightarrow \hat{h}_m^l = \sum_{|k|\leq l} \hat{f}_k^l D_{m,k}^l(R) \tag{5}$$

We finally add a lemma we first proved in [14] because it will be part of the first stage of our matcher.

Lemma 1. *Given $f(\omega), h(\omega) \in L^2(S^2)$, the correlation between $f(\omega)$ and $h(\omega)$ defined as*

$$C(\alpha, \beta, \gamma) = \int_{S^2} f(\omega) \Lambda(R_{\alpha,\beta,\gamma}) h(\omega) \; d\omega \tag{6}$$

can be obtained from the spherical harmonics $\hat{f}_m^l$ and $\hat{h}_m^l$ via the 3-D Inverse Discrete Fourier Transform as

$$C(\alpha, \beta, \gamma) = IDFT\{\sum_l \hat{f}_m^l \overline{\hat{h}_k^l} d_{m,h}^l(\pi/2) d_{h,k}^l(\pi/2)\}. \tag{7}$$

3 Spherical Pattern Matching

In this section we present normalized correlation and an approach on how to reduce the associated computational cost. We will introduce first the axial autocorrelation function for spherical functions and its invariant properties.

3.1 Autocorrelation Invariance

Let $f(\omega) \in L^2(S^2)$ be a function defined on the unit sphere. We define the axial autocorrelation function, as the autocorrelation of the function $f(\omega)$ computed rotating the function itself only around the Z axis, that points out of the North Pole η:

$$AC_{f,\eta}(\gamma') = \int_{S^2} f(\omega) \cdot \overline{\Lambda(R_{0,0,\gamma'}) f(\omega)} d\omega. \tag{8}$$

Expanding the function $f(\omega)$ as series of spherical harmonics we can rewrite relation (8) as:

$$AC_{f,\eta}(\gamma') = \int_{S^2} \sum_{l,m} \hat{f}_m^l Y_m^l \overline{\sum_{l',m',k} D_{m',k}^{l'}(0,0,\gamma') \hat{f}_k^{l'} Y_{m'}^{l'}} d\omega.$$

where $D_{m,k}^l(0,0,\gamma') = e^{-jk\gamma'}\delta_{m,k}$. Using the spherical harmonics orthogonality, we can finally simplify

$$AC_{f,\eta}(\gamma') = \sum_k \sum_l \left|\hat{f}_k^l\right|^2 e^{ik\gamma'} = IDFT\{\sum_l \left|\hat{f}_k^l\right|^2\}. \tag{9}$$

Under the action of a rotation $R_{\alpha,\beta,\gamma}$ the North Pole η is transformed into the point $\omega' = (\beta,\alpha)$. Following the same procedure as above, we can show that the axial autocorrelation of the rotated function $h(\omega) = \Lambda(R_{\alpha,\beta,\gamma})f(\omega)$, computed around the axis ω', is related to the function $AC_{f,\eta}$ by a simple translation:

$$AC_{h,\omega'}(\gamma') = AC_{f,\eta}(\gamma' - \gamma). \tag{10}$$

The proof of the previous identity is again straightforward if we rewrite the function $AC_{h,\omega'}(\gamma')$ using the expansion (1):

$$\begin{aligned} AC_{h,\omega'}(\gamma') &= \int_{S^2} \Lambda(R_{\alpha,\beta,\gamma})f(\omega) \cdot \overline{\Lambda(R_{\alpha,\beta,\gamma'})f(\omega)} d\omega = \\ &= \int_{S^2} \sum_{l,m,k} D_{m,k}^l(\alpha,\beta,\gamma)\hat{f}_k^l Y_m^l \cdot \overline{\sum_{l',m',k'} D_{m',k'}^{l'}(\alpha,\beta,\gamma')\hat{f}_{k'}^{l'} Y_{m'}^{l'}} d\omega = \\ &= \sum_{l,m,k,k'} d_{m,k}^l(\beta) d_{m,k'}^l(\beta) \hat{f}_k^l \overline{\hat{f}_{k'}^l} e^{-ik\gamma} e^{ik'\gamma'}. \end{aligned}$$

Using the orthogonality of the matrices $d^l(\beta)$:

$$\sum_m d_{m,k}^l(\beta) d_{m,k'}^l(\beta) = \delta_{k,k'}$$

we can finally simplify the $AC_{h,\omega'}(\gamma')$ expression:

$$AC_{h,\omega'}(\gamma') = \sum_k \sum_l \left|\hat{f}_k^l\right|^2 e^{ik(\gamma'-\gamma)} = AC_{f,\eta}(\gamma' - \gamma).$$

3.2 Normalized Cross-Correlation

Let $I(\omega)$ be a spherical image and $P(\omega)$ the template that is to be localized. If a pattern that matches exactly with the template is present in the image $I(\omega)$ at the position ω_0, the linear filtering that maximizes the signal to noise ratio in $\omega = \omega_0$, can be expressed as a cross-correlation. Such a cross-correlation function $C : SO(3) \to \mathbb{R}$ for functions defined on the unit sphere (6) may be computed fast as a 3D Inverse Fourier transform of

a linear combination of the template and the image spherical harmonics (lemma 1). The difficulty is in normalizing this expression to account for variations of the local signal energy. A normalized cross-correlation coefficient would be written as

$$NC(R) = \frac{\int_{S^2}(I(\omega) - \bar{I}_w)(\Lambda(R_{\alpha,\beta,\gamma})P(\omega) - \bar{P})d\omega}{\sqrt{\int_{PW}\left|I(\omega) - \bar{I}_w\right|^2 d\omega \int_{PW}\left|P(\omega) - \bar{P}\right|^2 d\omega}}, \tag{11}$$

where PW is the image window defined by the support of the pattern, and the over bar means the mean value of the signals in the PW region. Hereafter, we use the overbar to represent the local mean of the image instead of the complex conjugate as we did in subsection 3.1.

Digital spherical images are processed after a projection to the uniformly sampled $\{\theta, \phi\}$ plane, where the distribution of samples is related to the chosen equiangular grid. With this representation the number of samples inside the overlapping window PW, and the shape of the window itself are space-variant. This means that to perform the normalization process shown in (11), we must refresh the shape of the window PW for every possible rotation which yields to an extremely high computational cost. In particular if N_i^2 is the number of image samples and N_r^3 is the dimension of the discretized rotational space $\{\alpha, \beta, \gamma\}$, then the computational cost of (11) is $O(N_i^2\, N_r^3)$. Observe that the cost is proportional to the number of sphere samples as well as the number of samples of α, β, γ because the entire sphere has to be revisited for every rotation due to the space variant pattern window.

A possible remedy to this problem has been introduced in [14] where the matching is performed between the image and template gradients:

$$C_G(R) = \int_{S^2} \langle \nabla_T[I(\omega)], \nabla_T[\Lambda(R_{\alpha,\beta,\gamma})P(\omega)]\rangle d\omega. \tag{12}$$

This approach, also based on space invariant operations, even if much faster than the normalized cross-correlation still does not guarantee robustness against any possible intensity linear transformation of the template intensity the same way a normalized cross-correlation would do. Taking advantage of the properties of the axial autocorrelation, we can instead modify the standard normalized cross-correlation matching and decrease the processing time without losing the normalization properties. Matching will be accomplished in two steps, the first of which is the image-template cross-correlation (6) followed by a 1D normalized cross-correlation performed using the axial autocorrelation as kernel. Let us suppose that a rotated version $P_r(\omega) = \Lambda(R_{\tilde{\alpha},\tilde{\beta},\tilde{\gamma}})P(\omega)$ of the template $P(\omega)$ is present in our image $I(\omega)$ (as shown Fig. 1.a), then for any value of $(\tilde{\alpha}, \tilde{\beta}, \tilde{\gamma})$ we can state that the following relation holds:

$$C(R)|_{\alpha=\tilde{\alpha},\beta=\tilde{\beta}} = C(\gamma) = AC_{P,\eta}(\gamma - \tilde{\gamma})$$

where $C(R)$ is the cross-correlation, computed with $\alpha = \tilde{\alpha}$ and $\beta = \tilde{\beta}$ and $AC_{P,\eta}(\gamma)$ is the template axial autocorrelation (Fig. 1.b).

The basic idea is then to perform the localization task using the normalized cross-correlation between the axial autocorrelation of the pattern and the image-pattern cross-correlation, instead of using the normalized correlation (11) between the pattern and the

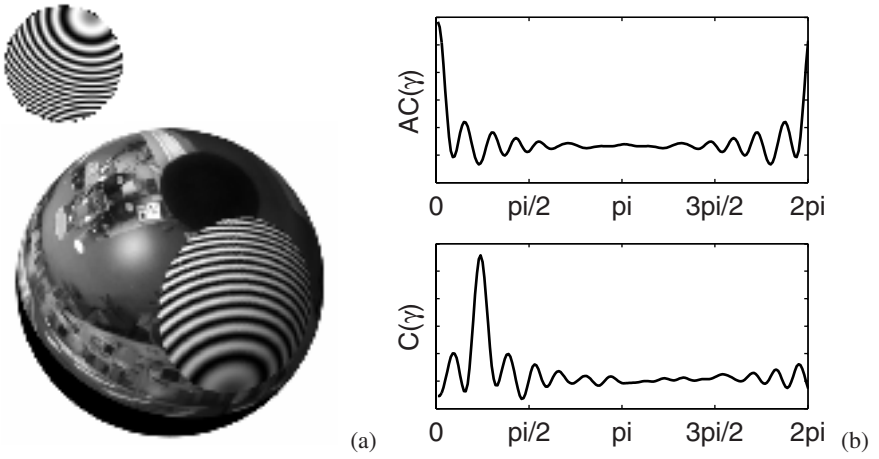

(a) (b)

Fig. 1. On the left (a) a catadioptric image mapped on the sphere and an artificial pattern to be matched. (b) On the right the axial autocorrelation of the pattern $AC_{P,\eta}$ and the pattern-image cross-correlation $C(R)|_{\alpha=\tilde{\alpha},\beta=\tilde{\beta}}$, computed in the point $(\tilde{\alpha},\tilde{\beta})$ where the pattern is located.

image; this means that the matching is not performed on the pattern function but on its autocorrelation, (fig. 13):

$$M(\alpha,\beta) = \max_{\gamma' \in (0,2\pi]} \int_{\gamma} C(R_{\alpha,\beta,\gamma}) \cdot AC_{P,\eta}(\gamma - \gamma') d\gamma \tag{13}$$

We can show that if the pattern has undergone a linear intensity transformation then a linear transformation relates also their axial autocorrelations:

$$P'(\omega) = a \cdot P(\omega) + b \qquad \leftrightarrow \qquad AC_{P',\eta}(\gamma) = A \cdot AC_{P,\eta}(\gamma) + B$$

where $A = a^2$ and $B = 2ab \int_{S^2} P(\omega) d\omega + 4\pi b^2$. Besides since $AC_{P,\eta}$ and $C(R)|_{\alpha,\beta}$ are one-variable (γ) discrete functions and the interval $\gamma \in [0, 2\pi)$ is uniformly sampled, the normalization process is thus reduced to a one-dimensional, space invariant operation.

We present now in detail a new algorithm for template matching on the sphere and the associated computational cost of each step. We assume that the rotational space can be sampled with N_r^3 samples and that the spherical image sampling N_i^2 is lower than N_r^2.

1. Initialize a $N_r \times N_r$ square localization map $M(\alpha,\beta)$. The dimension N_r determines the Euler angle estimation precision, the rotational space will be in fact discretized in N_r^3 points.
2. Compute the spherical harmonics of both the image and the pattern. Using the algorithm described in [12] this operation require a computational cost $O(N_i \log^2 N_i)$, where N_i^2 is the number of samples of the uniformly sampled (θ,ϕ)-plane.

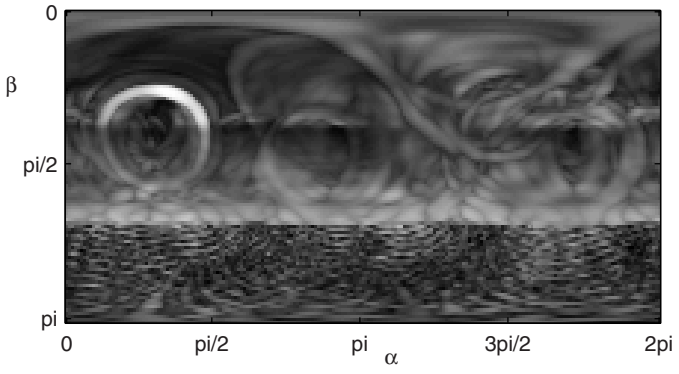

Fig. 2. The localization map $M(\alpha, \beta)$), relative to fig. 1.a.

3. Compute the axial autocorrelation of the template $AC_{P,\eta}$. This function can be obtained as 1D IDFT, so using the standard FFT requires only $O(N_r \log(N_r))$
4. Compute the cross-correlation $C(R)$ between the template and the image. Using Lemma 1 the computational cost of this step is $O(N_r^3 \log(N_r))$.
5. Build the localization map M: for every point (α, β) the normalized 1D circular cross-correlation between $AC_{P,\eta}$ and $C|_{\alpha,\beta}$ is computed, and the maximum value is assigned to the map. $M(\alpha, \beta) = max\{AC \otimes C|_{\alpha,\beta}\}$, (equation 13). Computing every normalized cross-correlation via 1D FFT the global computational cost of this step is again $O(N_r^3 log(N_r))$.
6. Localization of the pattern: exhaustive search of the max of the map $M(\alpha, \beta)$. Computational cost $O(N_r^2)$.
7. The estimation of the pattern orientation γ is performed computing the shift factor between $AC_{P,\eta}$ and $C|_{\alpha,\beta}$. Computational cost $O(N_r)$.

The final computational cost is $O(N_r^3 \log(N_r))$ where N_r^3 is the sampling of the rotation space.

4 Experiments

In this section we will present some results of our algorithm applied to artificial as well as real images. Before pattern matching, an omnidirectional image is projected into the uniformly sampled (θ, ϕ) space, required to compute the SFT using the sampling theorem (in any test has been chosen a square space of 700×700 pixels). The spherical Fourier transform has been performed assuming for the images to be band-limited, and we will use a subset of their spectrum in the third and fourth steps of the algorithm, considering a number of spherical harmonic up to the degree $l = 40$.

All the real pictures used to test the algorithm performance have been captured by a Nikon Coolpix 995 digital camera equipped with unique effective viewpoint catadioptric system (parabolic convex mirror and orthographic lens [1]). In [15], it is shown that

such a projection is equivalent to the projection on the sphere followed by a stereographic projection from the North Pole of the sphere to the equator plane. The coordinate transformation

$$u = \cot(\theta/2)\cos(\phi) \tag{14}$$

$$v = \cot(\theta/2)\sin(\phi) \tag{15}$$

allows us to map the catadioptric images directly on the unit sphere

$$I(\theta, \phi) = I(\cot(\theta/2)\cos(\phi), \cot(\theta/2)\sin(\phi)). \tag{16}$$

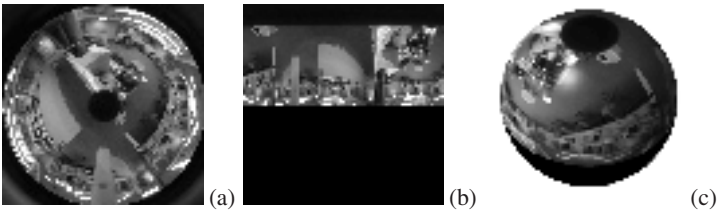

(a) (b) (c)

Fig. 3. On the left a 1400×1400 pixels catadioptric test-image. In the middle the same image presented in the equiangular (θ, ϕ) plane and mapped on the sphere (on the right). The southern emisphere has been padded with zero-value samples.

Table 1. Euler angle estimation for the letters (Fig. 4). For every letter, it is indicated the real position (RP), the estimated position (EP1) using the original templates and the estimated position relative to the scaling-offset transformed version (Fig. 4.b), (EP2). In the last case we provide the indication of the different offset and scaling factors applied to each letter.

		A sc.=0.6 offs.=25	M sc.=0.8 offs.=10	P sc.=0.9 offs.=-10	U sc.=0.5 offs.=-15	V sc.=0.7 offs.=20
α	RP	0°	45°	90°	135°	180°
	EP1	0.89°	45.47°	91.82°	136.4°	180.98°
	EP2	0.89°	45.75°	91.82°	136.4°	180.98°
β	RP	60°	60°	60°	60°	60°
	EP1	58.85°	58.85°	58.85°	58.85°	58.85°
	EP2	58.85°	58.85°	58.85°	58.85°	58.85°
γ	RP	0°	0°	270°	26°	120°
	EP1	0.89°	0.89°	271.29°	25.98°	120.77°
	EP2	0.89°	0.89°	271.29°	27.77°	120.77°

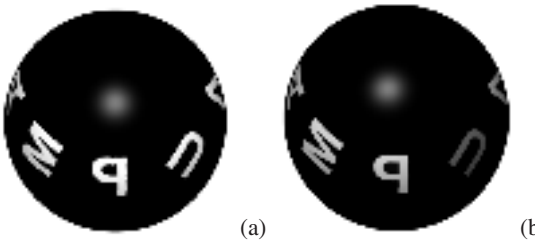

(a) (b)

Fig. 4. Two artificial images with five different letters. In (b) different scaling factors and offsets have been applied to the letters' intensity to test the matcher invariance to this class of transformation.

Table 2. Euler angle estimation for the real pattern (Fig.5). For every pattern is indicated the real position (RP) and the estimated position obtained with the proposed cascading normalized correlation (NC) compared with the analogous result obtained applying the gradient based correlation (GB). NC_s and GB_s refer to the matching performed using noisy images.

		Logo	Field	Paint	Park	Car	Door
α	RP	92.57°	236.69°	77.18°	321.59°	257.27°	272.71°
	NC	93.61°	236.25°	77.56°	321.85°	257.66°	273.70°
	GB	95.39°	233.23°	81.13°	340.60°	257.66°	229.12°
	NC_n	95.07°	237.04°	77.32°	320.7°	257.32°	270.0°
	GB_n	315.63°	237.04°	143.24°	249.71°	152.66°	215.12°
β	RP	64.28°	51.45°	51.45°	64.32°	66.89°	64.32°
	NC	62.40°	51.70°	51.70°	62.41°	65.97°	62.41°
	GB	68.30°	52.35°	71.32°	62.41°	67.88°	65.97°
	NC_n	63.38°	48.17°	50.70°	63.38°	65.91°	63.38°
	GB_n	78.59°	53.24°	70.98°	40.56°	42.16°	118.42°
γ	RP	0°	26.37°	131.85°	80.25°	120.38°	28.66°
	NC	-0.89°	27.63°	133.50°	81.53°	120.77°	29.57°
	GB	-0.71°	28.45°	125.70°	95.60°	120.36°	38.54°
	NC_n	1.28°	29.37°	134.04°	85.53°	121.28°	29.37°
	GB_n	247.18°	206.62°	209.15°	97.60°	95.28°	246.12°

Naturally, the range of this mapping is limited by the field of view of the original catadioptric system: $\theta \in [0, 102^\circ), \phi \in [0, 360^\circ)$. The portion of the sphere that is not covered by the mapping has been padded with zero-value samples, (fig 3). We remark that the mapping of the catadioptric images onto the sphere is necessary to compute the spherical harmonics and perform the matching in the transformed space. Working directly with the original images in fact, would oblige us to deal with the negative effects of the distortion due to the template rotations. Such distortions of objects' shape in catadioptric images are extremely stressed and would certainly lead us to results of low quality and as shown above high processing time.

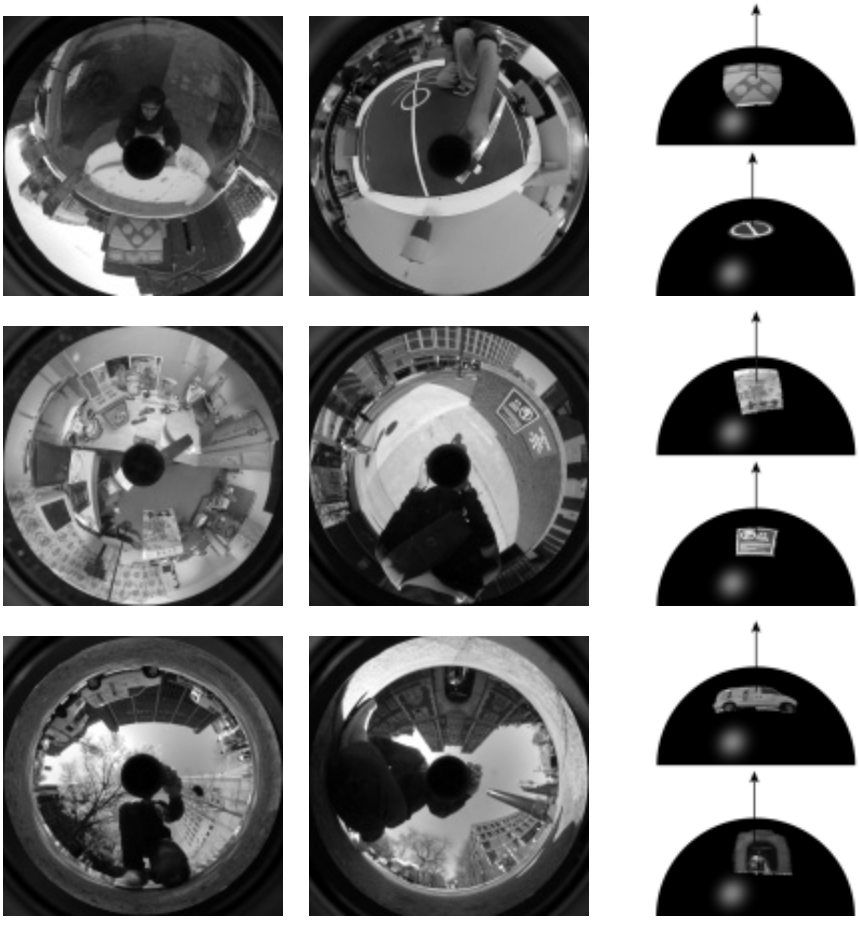

Fig. 5. Six catadioptric images. The patterns are presented mapped onto the sphere and slightly rotated forward to enhance their visibility.

We present first the localization performance of our normalized cross-correlation algorithm applied to an artificial image with five different letters A, M, P, U and V overlapped on a uniform background (Fig. 4.a). For every letter, in the second test, the grayscale value has been modified applying different scaling factors and offsets (Fig. 4.b). Table (1) presents real and estimated position of every letter. In this first test we do not present any comparison with other matching approaches because our main purpose is to show the effective scaling-offset invariance of the proposed method.

Then our algorithm has been tested using real images and real templates. As described in section 3.2, the matcher uses the pattern placed on the North Pole of the spherical reference frame as kernel to compute the cross-correlation function (6); however due to the structure of the catadioptric system used to take the pictures, it is impossible to

have a real image of the pattern placed on the North Pole. This problem has been overcome with an artificial rotation of the target pattern, selected from an omnidirectional image of the same environment captured with a different tripod orientation (rigid translational+rotational motion). Figure 5 shows six catadioptric images with the relatives selected objects and in the Table 2 we presents the localization performances. For comparison purposes, we performed the same test using the gradient matching introduced in [14], because it is the only matching system that requires the same computational cost of the algorithm presented in this paper and guarantees a good robustness to scaling-offset template transformation. The last test has been performed using a noisy version of the same real images of Fig. 5, obtained with an additive gaussian white noise (zero mean with 0.05 variance). Results are presented in the same Table 2. It is interesting to emphasize that using the axial autocorrelation kernel the template pose estimation is very robust to noisy images. This is due to the fact that the matching is performed by a harmonic analysis that involves a reduced number of transformed coefficients, which are less sensitive to noise.

5 Conclusion

In this paper we presented a pattern matching algorithm for spherical images based on the computation of a new cross-correlation measure: the cross-correlation of the image-pattern cross-correlation with the axial autocorrelation of the pattern. The new cross-correlation measure is invariant to linear intensity transformations and can be computed in $O(N_r^3 log(N_r))$ where N_r^3 is the sampling of the SO(3) Euler angle parameterization. This reduces the sampling cost significantly comparing to the cost associated with an explicit traversal of the image and the rotation space $O(N_i^2\, N_r^3)$. To validate the proposed algorithm we presented results obtained using artificial as well as real images and patterns, with the main goal to check the real maintenance of the matcher's invariance properties.

Just before the submission of the camera ready version of the paper we discovered a way to compute the local signal energy $\sqrt{\int_{PW} \left|I(\omega) - \bar{I}_w\right|^2 d\omega}$ required in the classical normalized cros-correlation coefficient (11) by computing the spherical harmonics of the support mask as well as the spherical harmonics of the square of the image and simultaneously keeping the same computational cost at $O(N_r^3 log(N_r))$. Time limits did not allow obtaining experimental results for this new procedure which will be thoroughly presented in a future report.

Acknowledgements. The authors are grateful for support through the following grants: NSF-IIS-0083209, NSF-IIS-0121293, NSF-EIA-0324977 and ARO/MURI DAAD19-02-1-0383.

References

1. Nayar, S.: Catadioptric omnidirectional camera. In: IEEE Conf. Computer Vision and Pattern Recognition, Puerto Rico, June 17-19 (1997) 482–488

2. Baker, S., Nayar, S.: A theory of catadioptric image formation. In: Proc. Int. Conf. on Computer Vision, Bombay, India, Jan. 3-5 (1998) 35–42
3. Brunelli, R., Poggio, T.: Template matching: Matched spatial filters and beyond. Pattern Recognition **30** (1997) 751–768
4. Turk, M., Pentland, A.: Eigenfaces for recognition. Journal of Cognitive Neuroscience **3** (1991) 72–86
5. Hall, E.: Computer Image Processing and Recognition. Academic Press (1979)
6. Sato, J., Cipolla, R.: Extracting group transformations from image moments. Computer Vision and Image Understanding **73** (1999) 29–42
7. Zahn, C., Roskies, R.: Fourier descriptors for plane close curves. IEEE Trans. Computers **21** (1972) 269–281
8. Lenz, R.: Rotation-invariant operators. In: Proc. Int. Conf. on Pattern Recognition, Paris, France, Sept. 28-3 (1986) 1130–1132
9. Tanaka, M.: On the representation of the projected motion group in 2+1d. Pattern Recognition Letters **14** (1993) 871–678
10. Segman, J., Rubinstein, J., Zeevi, Y.: The canonical coordinates method for pattern deformation: Theoretical and computational considerations. IEEE Trans. Pattern Analysis and Machine Intelligence **14** (1992) 1171–1183
11. Arfken, G., Weber, H.: Mathematical Methods for Physicists. Academic Press (1966)
12. Driscoll, J., Healy, D.: Computing fourier transforms and convolutions on the 2-sphere. Advances in Applied Mathematics **15** (1994) 202–250
13. Talman, J.D.: Special Funcitons. W.A.Benjamin Inc., Amsterdam (1968)
14. Sorgi, L., Daniilidis, K.: Template matching for spherical images. In: SPIE 16° Annual Symposium Electronic Imaging, San Jose, CA, January 18-22 (2004)
15. Geyer, C., Daniilidis, K.: Catadioptric projective geometry. International Journal of Computer Vision **43** (2001) 223–243