Calibration from statistical properties of the visual world

Etienne Grossmann[1*], José António Gaspar[2], and Francesco Orabona[3]

[1] Tyzx, Inc., Menlo Park, USA
[2] ISR, Instituto Superior Técnico, Lisbon, Portugal
[3] Idiap Research Institute, Martigny, Switzerland

Abstract. What does a blind entity need in order to determine the geometry of the set of photocells that it carries through a changing lightfield? In this paper, we show that very crude knowledge of some statistical properties of the environment is sufficient for this task.
We show that some dissimilarity measures between pairs of signals produced by photocells are strongly related to the angular separation between the photocells. Based on real-world data, we model this relation quantitatively, using dissimilarity measures based on the correlation and conditional entropy. We show that this model allows to estimate the angular separation from the dissimilarity. Although the resulting estimators are not very accurate, they maintain their performance throughout different visual environments, suggesting that the model encodes a very general property of our visual world.
Finally, leveraging this method to estimate angles from signal pairs, we show how distance geometry techniques allow to recover the complete sensor geometry.

1 Introduction

This paper departs from traditional computer vision by not considering images or image features as input. Instead, we take signals generated by photocells with unknown orientation and a common center of projection, and explore the information these signals can shed on the sensor and its surrounding world.

We are particularly interested in determining whether the signals allow to determine the geometry of the sensor, that is, to calibrate a sensor like the one shown in Figure 1. Psychological experiments [1] showed that a person wearing distorting glasses for a few days, after a very confusing and disturbing period, could learn the necessary image correction to restart interacting effectively with the environment. Can a computer do the same when, rather than distorted images, it is given the signals produced by individual photocells? In this situation, it is clear that traditional calibration techniques [2, 3] are out of the question.

Less traditional non-parametric methods that assume a smooth image mapping and smooth motion [4] can obviously not be applied either. Using controlled-light stimuli

* This work was partially supported by TYZX, Inc, by the Portuguese FCT POS_C program that includes FEDER funds, and by the EU-project URUS FP6-EU-IST-045 062

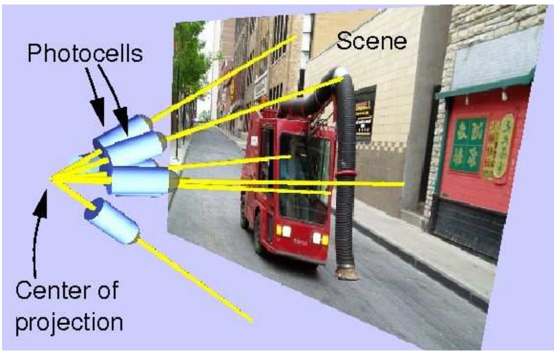

Fig. 1. A discrete camera consists of a number of photocells (pixels) that measure the light traveling along pencil of lines.

or known calibration, matches could be obtained, allowing to use match-based non-parametric techniques [5]. In this study however, we wish to exclude known calibration objects and other controlled stimuli.

Our approach is inspired from the work of Pierce and Kuipers [6], who measure the dissimilarity, or distance, between sensor elements that are not necessarily light sensors. The elements are then embedded in a metric space using metric scaling [7], which also determines the dimension of the space. A relaxation method then improves this embedding, so that the Euclidean distance between sensor elements better matches the dissimilarity between the sensor inputs. Getting close to the problem addressed in the present paper, the authors use this method to reconstitute the geometry of a rectangular array of visual sensors that scans a fronto-parallel image.

Going further, Olsson et al. [8] use the information distance of [9] as a more appropriate method to measure the distance between visual or other sensor elements. They also show how visual sensors -the pixels of the camera of a mobile robot- can be mapped to a plane, either using the method of [6], or their own, that embeds sensor elements specifically in a square grid.

The works of Olsson et al. and of Pierce and Kuipers are very interesting to computer vision researchers, but they cannot calibrate an arbitrary discrete camera, since the embedding space is either abstract or fixed to a grid. In both cases, it lacks an explicit connection to the geometry of the sensor.

Grossmann et al [10] partially fill this gap by showing that the information distance can be used to estimate the angular separation between pairs of photocells, and from there, estimate the geometry of a sensor of limited angular radius.

Because the present work exploits statistical properties of the light-field of the world surrounding a light sensor, it is also related to research on the statistical properties of real-world images. In that area, a model of image formation is used, but images, rather than sequences, are studied. That research has put in evidence fundamental properties, in terms of local, global and spectral statistics, of real-world images, and found ways to exploit these properties for computer vision tasks, such as classification [11], image restoration [12] and 3D inference [13]. Although these results are of great interest, they are not directly applicable in our case, mainly because we lack images.

Moreover, these statistics are about planar images, which is a hindrance in our case: first, we do not want to exclude the case of visual sensor elements that are separated by more than 180 degrees, such as the increasingly popular omnidirectional cameras. Also, the local statistical properties of perspective images depend of the orientation of the image plane with respect to the scene, except in special constrained cases such as the fronto-parallel "leaf world" of Wu et al. [14]. Defining images on the unit sphere thus appears as a natural way to render image statistics independent of the sensor orientation, at least with proper assumptions on the surrounding world and/or the motion of the sensor.

The present article elaborates and improves over our previous work [10]. We innovate by showing that the correlation, like the information distance, can be used to provide geometric information about a sensor. Also, we use a simpler method to model to relation between angles and signal statistics.

More important, we go much beyond [15] in showing that this model generalizes well to diverse visual environments, and can thus be considered to be a reliable characteristic of our visual world. In addition, we show that the presented calibration method performs much better, for example by allowing to calibrate sensors that cover more than one hemisphere.

1.1 Proposed approach

The present work relies on statistical properties of the data streams produced by pairs of sensor elements that depend only on the *angular separation* between the photocells. For example, if the sampled lightfield is a homogeneous random field defined on the sphere [16], then the covariance between observations depends only on the angular separation between the sampled points.

This assumption does not hold in general in our anisotropic world, but it does hold, e.g. if the orientation of the sensor is uniformly distributed amongst all unitary transformations of the sphere, that is, if the sensor is randomly oriented, so that each photocell is just as likely to sample the light-field in any direction.

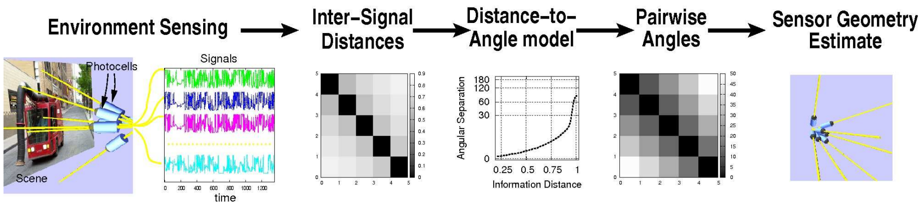

Fig. 2. The process of estimating the geometry of an unknown discrete camera.

This assumption of homogeneity -or isotropy- of the sampled lightfield is of great practical utility, in conjunction with a few other assumptions of good behavior: in this work, we only use statistics that converge properly (e.g. in probability or more strongly) when signal lengths tend to infinity.

Perhaps more importantly we are only interested in statistics that have an expectancy that is a strictly monotonous function of the angular separation of the pair of photocells. That is, if x, y are two signals (random variables) generated by two photocells separated by an angle θ, and $d(x, y)$ is the considered statistic, then the expectancy of $d(x, y)$ is a strictly monotonous function of θ, for $0 \leq \theta \leq \pi$. The importance of this last point is that this function can be inverted, resulting in a functional model that links the value of the statistic to the angle.

The statistic-to-angle graph of such statistics is the a-priori knowledge about the world that we leverage to estimate the geometry of discrete cameras. In the present work, we use discrepancy measures based on the correlation or conditional entropy, defined in Section 3. In Section 4, we show how to build the considered graph.

Having obtained angle estimates, we recover the sensor geometry, in Section 5.1, by embedding the angles in a sphere. This is done using simple techniques from distance geometry [17]. Experimental results are presented in Section 5.2. Finally, Section 6 presents some conclusions and possible directions for future research. The calibration process considered in the present work is outlined in Figure 2. The statistic-to-angle modeling produces the crucial functional relation used in the third-from right element of Figure 2.

2 Discrete camera model and simulation

Before entering into the details of our methodology for estimating the sensor geometry, we define the discrete camera and explain how to simulate it using an omnidirectional image sensor.

We define a discrete camera [10] as a set of N photocells indexed by $i \in \{1, \ldots, N\}$, pointing in directions $X_i \in \mathbb{R}^3$ and having a unique center of projection. These photocells acquire along the time t, brightness measurements $x(i, t)$ in the range $\{0, \ldots, 255\}$. The directions of the light rays, contrarily to conventional cameras, are not necessarily organized in a regular grid. Many examples of cameras can be found under these definitions. One example is the linear camera, where all the X_i are co-planar. Another example is the conventional perspective camera which comprises a rectangular grid of photocells that are enumerated in our model by a single index i,

$$\left\{ X_i \mid X_i \sim K^{-1} \left[\begin{smallmatrix} i\%W \\ \lfloor i/W \rfloor \\ 1 \end{smallmatrix} \right], \; 0 \leq i < HW \right\}$$

where W, H are the image width and height, K is the intrinsic parameters matrix, $\%$ represents the integer modulo operation and $\lfloor . \rfloor$ is the lower-rounding operation. Cameras equipped with fisheye lenses, or having log- polar sensors, can also be modeled again by setting X_i to represent the directions of the light-rays associated to the image pixels. In the same vein, omnidirectional cameras having a single projection center, as the ones represented by the unified projection model [18], also fit in the proposed model. In this paper we use a calibrated omnidirectional camera to simulate various discrete cameras.

2.1 Image sensor

We simulate a discrete camera with known Euclidean geometry by sampling a calibrated panoramic image with unique projection center at fixed locations. Since the camera is calibrated, it is straightforward to locate the position (u, v) in the panoramic image corresponding to the 3D direction X of a photocell that is part of the simulated discrete camera. In the present work, we use bilinear interpolation to measure the graylevel value at non-integer coordinates (u, v).

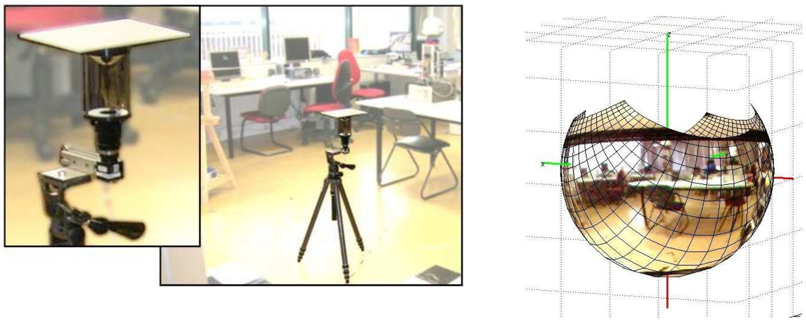

Fig. 3. Left: The camera used to sample omnidirectional images (image mirrored). **Right:** A calibrated omnidirectional image mapped to a sphere.

Images are acquired by a VStone catadiopric camera consisting of a perspective camera fitted to a hyperbolic mirror, shown in Figure 3, left. This system is modeled as single projection center camera [18] with a $360° \times 210°$ field of view and a $\sim 45°$ blind spot at the south pole (Fig. 3, right). The mirror occupies a 453×453 pixel region of the image. The angular separation between neighboring pixels in the panoramic image is usually slightly smaller than 0.5°. Also, some mild vignetting occurs, that could be corrected. Apart for these minor inconveniences, simulating a discrete camera by an omnidirectional camera presents many advantages: no other specialized hardware is needed and each omnidirectional image can be used to simulate many discrete camera "images", as in Fig. 4, right. With respect to perspective cameras, the available field of view allows to study very-wide-angle discrete cameras.

3 Distances between pairs of signals

In this section, we define the measures of distance between signals, correlation and information distance, that will later be used to estimate angles.

3.1 Correlation distance

We call correlation distance between signals $x(t)$ and $y(t)$, $1 \leq t \leq T$, the quantity

$$d_c(x, y) = \frac{1}{2}\left(1 - C(x, y)\right),$$

where $C(x, y)$ is the correlation between the signals. It is easy to verify that $d_c(.,.)$ is a distance.

For the task considered in this paper, it is natural to prefer the correlation distance over the variance or the (squared) Euclidean distance $\|x - y\|^2$, because both vary with signal amplitude (and offset, for the latter), whereas $d_c(.,.)$ is offset- and scale-invariant.

3.2 Information distance

Given two random variables x and y (in our case, the values produced by individual pixels of a discrete camera) taking values in a discrete set $\{1, \ldots, Q\}$, the *information distance* between x and y is [9]:

$$d(x, y) = H(x|y) + H(y|x) \; = \; 2H(x, y) - H(y) - H(x), \tag{1}$$

where $H(x, y)$ is the Shannon entropy of the paired random variable (x, y), and $H(x)$ and $H(y)$ are the entropies of x and y, respectively. It is easy to show that Eq. (1) defines a distance over random variables. This distance is bounded by $H(x, y) \leq \log_2 Q$, and is conveniently replaced thereafter by the *normalized information distance* :

$$d_I(x, y) = d(x, y) / H(x, y), \tag{2}$$

which is bounded by 1, independently of Q [9].

It should be noted that estimating the information distance is non-trivial: naively replacing unknown probabilities $p_x(q)$ by sample frequencies $\hat{p_x}(q) = |\{t | x(t) = q\}| / T$, where T is the signal length and $|.|$ denotes the set cardinal, yields a biased estimator $\hat{H}(x)$. This estimator has expectancy

$$E\left\{\hat{H}\right\} = H - \frac{Q-1}{2T} + \frac{1 - \sum_q \frac{1}{p_x(q)}}{12T^2} + O\left(\frac{1}{T^3}\right). \tag{3}$$

This expression shows the slow convergence rate and strong bias of $\hat{H}(x)$. We somewhat alleviate these problems by first, correcting for the first bias term $(Q-1)/2T$, i.e. applying the Miller-Madow correction; and by re-quantizing the signal to a much smaller number of bins, $Q = 4$. Extensive benchmarking in [15] has shown these choices to be beneficial.

4 Estimating angular separation from inter-signal distance

As explained earlier, our a-priori knowledge of the world will be encoded in a graph mapping a measure of discrepancy between two signals, to the angular separation between the photocells that generated the signals. We now show how to build this graph, and assess its effectiveness at estimating angles.

For this purpose, we use the 31-pixel planar discrete camera (or "probe") shown in Fig. 4, left. This probe design allows to study the effect of angular separations ranging from 0.5 to 180 degrees and each sample provides 465=31(31-1)/2 pixel pairs. In

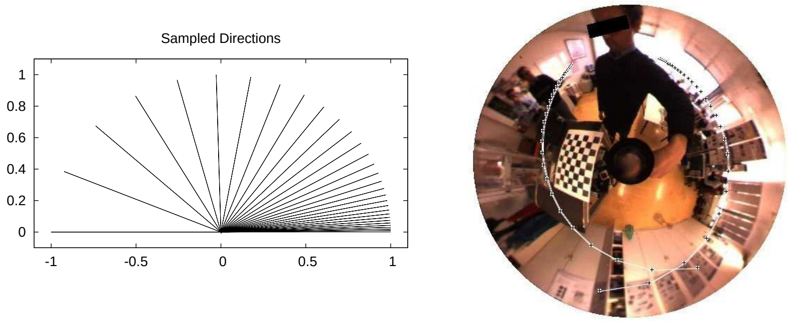

Fig. 4. Left: Geometry of a discrete camera consisting of a planar array of thirty one (31) pixels, spanning 180° in the plane. The first two pixels are separated by 0.5°, the separation between consecutive photocells increases geometrically (ratio $\simeq 1.14$), so that the 31^{st} photocell is antipodal with respect to the first. **Right:** Two instances of the linear discrete camera, inserted in an omnidirectional image. Pixels locations are indicated by small crosses connected by white lines.

the "tighter" part of the discrete camera layout, there exists a slight linear dependence between the values of consecutive pixels due to aliasing.

The camera is hand-held and undergoes "random" general rotation and translation, according to the author´s whim, while remaining near the middle of the room, at 1.0 to 1.8 meters from the ground. We acquired three sequences consecutively, in very similar conditions and joined them in a single sequence totaling 1359 images, i.e. approximately 5 minutes of video at ˜4.5 frames per second.

To simulate the discrete camera, we randomly choose an orientation (i.e. half a great circle) such that all pixels of the discrete camera fall in the field of view of the panoramic camera. Figure 4 shows two such choices of orientations. For each choice of orientation, we produce a sequence of 31 samples $x(i,t)$, $1 \leq i \leq 31$, $1 \leq t \leq 1359$, where each $x(i,t) \in \{0,\ldots,255\}$. Choosing 100 different orientations, we obtain 100 discrete sensors and 100 arrays of data $x_n(i,t)$, $1 \leq n \leq 100$. Appending these arrays we obtain 31 signals $x(i,t)$ of length to 135900.

We then compute, for each pair of pixels (indices) $1 \leq i,j \leq 31$, the correlation and information distances, $d_c(i,j)$ and $d_I(i,j)$. Joining to these the known angular separations $\theta_{i,j}$, we obtain a set of pairs $(\theta_{i,j}, d(i,j))$, $1 \leq i,j \leq 31$.

From this dataset, we build a constant by parts model of the expectancy of the distance, knowing the angle. For the correlation distance, we limit the abscissa to values in $[0, 1/2]$. After verifying and, if needed enforcing, the monotonicity of this model, we invert it, obtaining a graph of angles as a function of (correlation or information) distances. Strict monotonicity has to be enforced for the correlation-based data, owing to the relatively small number of data points used for each quantized angle.

Figure 5 shows the resulting graphs. This figure shows one of the major issues that appear when estimating the angular separation between pixels from the correlation or information distance: the graphs become very steep for large values of the distance, indicating that small changes of the distance result in large changes in the estimated angle. On the other hand, for small distance values, the curves are much flatter, suggesting

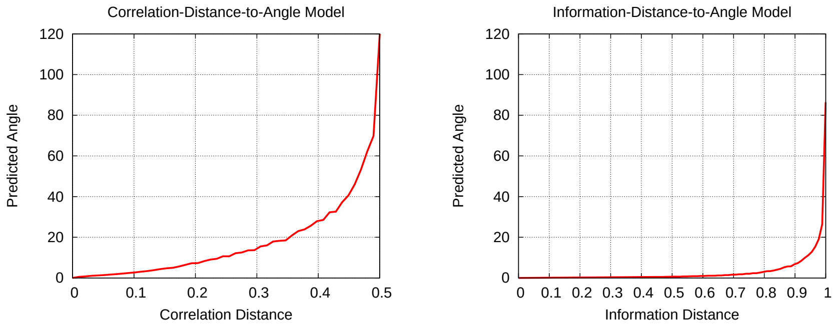

Fig. 5. Models relating correlation (left) or information distance (right) to angular separation between photocells. These models were build from simulated signals produced by the linear probe of Fig. 4, left. Signals of length $T = 135900$, acquired indoors were used.

that small angles can be determined with greater accuracy. Both trends are particularly true for the information distance.

4.1 Experimental validation

We now assess how well angles can be estimated from the graphs obtained in the previous section. For this purpose, we use 100 sets of 31 signals $x_n(i,t)$, $1 \leq n \leq 100$, $1 \leq i \leq 31$, $1 \leq t \leq 1359$ acquired in the same conditions as above. We compute the correlation and information distances of pairs of signals $d_c(n,i,j)$ and $d_I(n,i,j)$ and, using the models in Fig. 5, angular estimates $\hat{\theta}_c(n,i,j)$ and $\hat{\theta}_I(n,i,j)$.

Figure 6 shows the precision and accuracy of the estimated angles. This figure shows that the estimated angles are fairly accurate for angular separations smaller than $5°$, but degrades sharply for greater values. As could be expected from our comments at the beginning of the section, the curves confirm that the information distance yields better estimates of small angles, while correlation distance does best (but still not very well) for larger angles.

We now turn to the generalization ability of the models in Fig. 5. For this purpose, we use 100 31-uplets of signals of length 2349, taken from an out- and indoor sequence, four images of which are shown in Fig. 7. In this sequence, and contrarily to the previous sequence, the camera remains mostly horizontal. Also, the scene is usually farther away and more textured. A lot of saturation is also apparent.

Following the previous procedure, we estimate angles from these new signals and show the precision and accuracy statistics in Figure 8.

The striking resemblance between Figures 8 and 6 indicates that the models in Fig. 5 generalize pretty well to outdoors scenes. We surmise that the fact that the correlation distance yields more accurate estimates outdoors than indoors is due to the extra texture, which increases the correlation distance for small angles, and corrects the bias in angular estimates observed near the origin of the top left curve of Fig. 6.

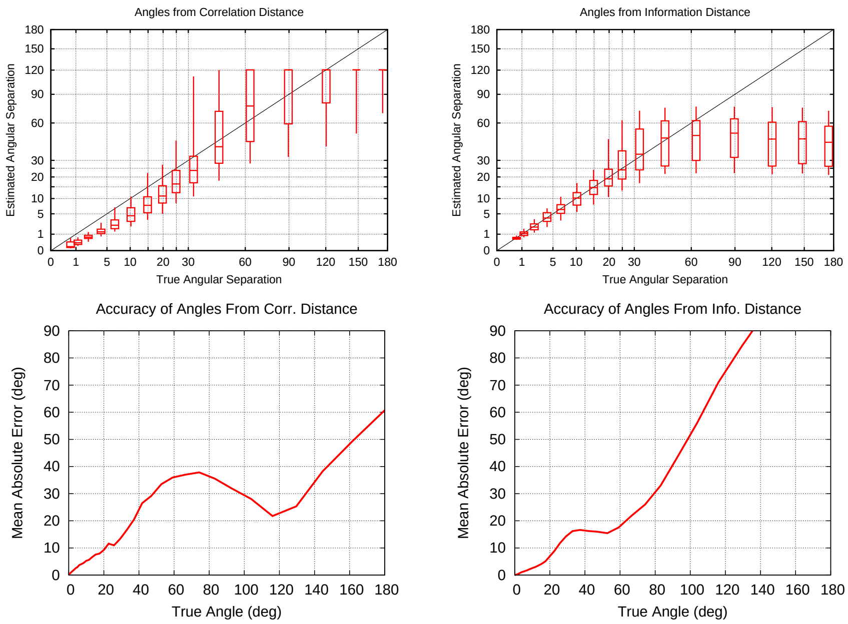

Fig. 6. Precision and accuracy of angles estimated from correlation (left) or information distance (right). The boxplots at the top show the 5$^{\text{th}}$ percentile, first quartile, median, third quartile and 95$^{\text{th}}$ percentile of the estimated angles, plotted against the true angles. The bottom curves show the mean absolute error in the estimated angles. These statistics were generated from 100 planar probes (Fig. 4, left) and signals of length $T = 1359$. The angles were estimated using the models of Fig. 5. The signals were acquired in the same conditions as those used to build the models.

5 Calibrating a discrete camera

Having seen the qualities and shortcomings of the proposed angle estimators, we now show how to use them to calibrate a discrete camera.

To stress the generalization ability of the angle estimators, all the reconstructions produced by the above method are obtained from the in- and outdoors sequence of Fig. 7, rather than from the indoors sequence used to build the distance-to-angle models.

5.1 Embedding points in the sphere

The last step we take to calibrate a discrete camera requires solving the problem:

Problem 1) Spherical embedding problem: Given angle estimates θ_{ij}, $1 \leq i, j \leq N$, find points X_i on the unit sphere, separated by angles approximately equal to θ_{ij}, i.e. $X_i^\top X_j \simeq \cos\theta_{ij}$, for all i, j.

This problem can be reduced to the classical problem of distance geometry [17]:

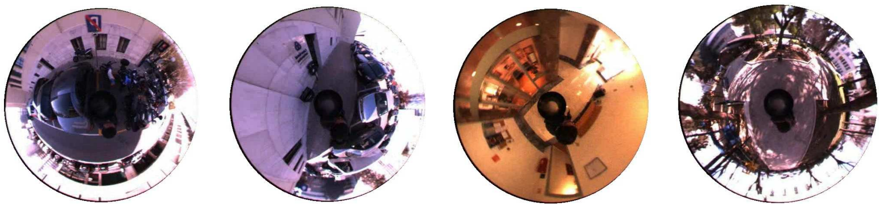

Fig. 7. Four images from a sequence of 2349 images acquired indoors and outdoors at approximately 4.5FPS.

Problem 2) Euclidean embedding problem: Given distance estimates $D_{ij}, 1 \leq i, j \leq N$, find points Y_i in a metric vector space, such that, for all i, j, $\|Y_i - Y_j\| \simeq D_{ij}$

Indeed, by defining an extra point $Y_0 = (0, 0, 0)$, and distances $D_{ij} = \sqrt{2 - 2\cos\theta_{ij}}$ for $i, j \neq 0$ and $D_{oi} = 1$, the mapping of the first problem to the second is immediate. Solutions to both problems (with exact equality, rather than approximate) were published in 1935 [19][4]. Schoenberg´s Theorem 2 [19] states that if the matrix C with terms $C_{ij} = \cos\theta_{ij}$ is positive semidefinite with rank $r \geq 1$, then there exist points on the unit $(r-1)$ −dimensional sphere that verify $X_i^\top X_j = C_{ij}$ for all i, j. This result directly suggests the following method for embedding points in the 2-sphere:

1. Build the matrix C with terms $C_{ij} = \cos\theta_{ij}, 1 \leq i, j \leq N$.
2. Compute, using the SVD decomposition, the rank-3 approximation $\tilde{C} = UU^\top$ of C, where U is $N \times 3$.
3. Define $X_i = (U_{i1}, U_{i2}, U_{i3}) \,/\, \|(U_{i1}, U_{i2}, U_{i3})\|$.

One should note that this very simple algorithm is not optimal in many ways. In particular, it does not take into account that the error in the angles θ_{ij} is greater in some cases than in others. It is easy to verify that the the problem is not directly tractable by variable-error factorization methods used in computer vision.

Noting that the error in the estimated angles is approximately proportional to the actual angle suggests an embedding method that weighs less heavily large angular estimates. One such method is Sammon´s algorithm [20], which we adapt and modify for the purpose of spherical embedding from our noisy data. In this paper, we minimize the sum

$$\sum_{i,j} w_{i,j} \left(X_i^\top X_j - C_{ij} \right)^2, \text{ where } w_{ij} = \begin{cases} \max\left\{0, \frac{1}{1-C_{ij}} - \frac{1}{1-C_o}\right\} & \text{if } C_{ij} \neq 1 \\ \frac{1}{\eta} & \text{otherwise.} \end{cases}$$

To reflect the fact that big angles are less well estimated, we set $C_0 = 0.9$, so that estimates greater than $\operatorname{acos}(0.9) \simeq 25°$ be ignored. The other parameter, η is set to 1, allowing the points X_i to stray a little bit away from the unit sphere. Our implementation is inspired by the second-order iterative method of Cawley and Talbot (`http://theoval.sys.uea.ac.uk/~gcc/matlab/default.html`). For initialization, we use an adaptation of [21] to the spherical metric embedding problem, which will be described in detail elsewhere.

[4] Schoenberg cites previous work by Klanfer and by Menger, to which we did have access.

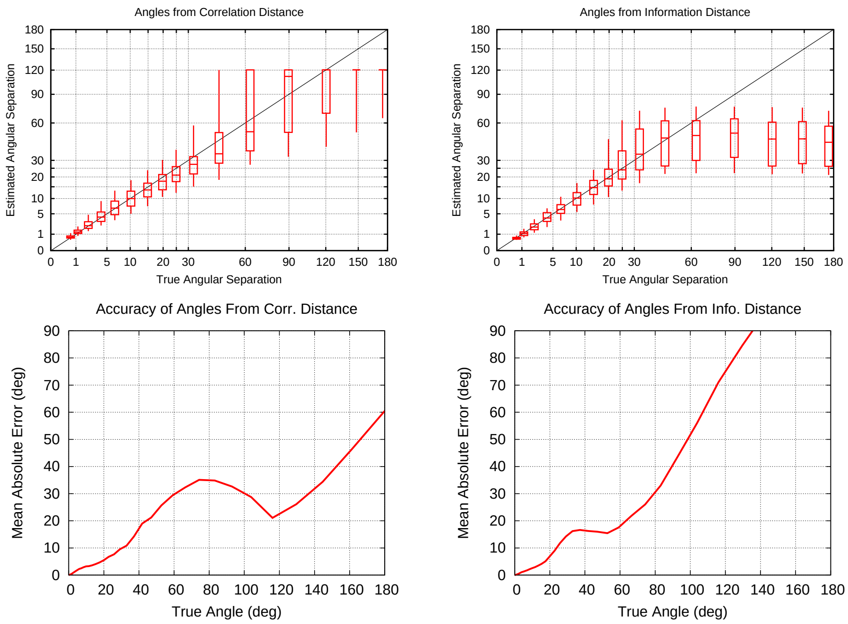

Fig. 8. Precision and accuracy of angles estimated in the same conditions as in Fig. 6, except that signals extracted from an indoors-and-outdoors sequence (Fig. 7) were used. These figures show that the models in Fig. 5 generalize fairly well to signals produced in conditions different from that in which the models were produced. In particular, the angles estimated from the correlation distance are improved w.r.t. those of Fig. 6 (see text).

5.2 Sensor calibration

We now evaluate the results of this embedding algorithm on data produced by the angle-estimating method of Sec. 4. For this purpose, we produce sequences of pixel signals in the same conditions as previously, using the outdoors and indoors sequence shown in Figure 7, except that the sensor shape is different. The information and correlation distances between pixels is then estimated from these signals, the angular separation between the pixels is estimated using Sec. 4, and the embedding method of Sec. 5.1 is applied to these angle estimates.

Figure 10 shows the results of our calibration method on sensors covering more than a hemisphere, which thus cannot be embedded in a plane without significant distortion. It should be noted that, although the true sensor is each time more than hemispheric, the estimated calibration is in both cases smaller. This shrinkage is a known effect of some embedding algorithms, which we could attempt to correct.

Figure 11 shows how our method applies to signals produced by a different sensor from the one used to build the distance-to-angle models, namely an Olympus Stylus 300 camera. An 8-by-8 square grid pixels spanning 34 degrees was sampled along a 22822 image sequence taken indoors and outdoors. From this sequence, the estimated angles were generally greater than the true angles, which explains the absence of shrinkage.

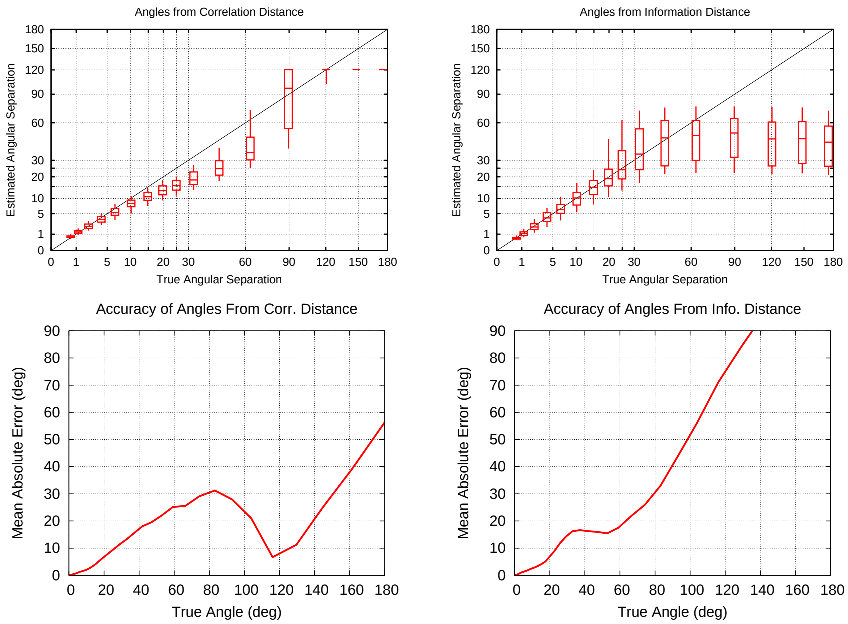

Fig. 9. Precision and accuracy of angles estimated in the same conditions as in Fig. 8, except that the planar probes are constrained to remain approximately horizontal. These figures show that the models in Fig. 5 are usable even if the isotropy assumption of the moving entity is not valid.

The higher angle estimates were possibly due to higher texture contents of the sequence. The estimated angles were also fairly noisy, possibly due to the sequence length, and we surmise that longer sequences would yield better results.

These results represent typical results that researchers reproducing our method may encounter. Results from other experiments will be presented elsewhere.

6 Discussion

In this paper, we have shown that simple models exist that relate signal discrepancy to angular separation, and are valid in indoors and outdoors scenes. This suggests the existence of near-universal properties of our visual world, in line with other work showing statistical properties of natural images. Contrarily to previous works, we consider statistics of the lightfield taken as a function defined on the sphere, rather than the plane, a choice that allows us to consider fields of view greater than 180 degrees.

We addressed the problem of determining the geometry of a set of photocells in a very general setting. We have confirmed that a discrete camera can be calibrated to a large extent, using just two pieces of data: a table relating signal distances to angles; and a long enough signal produced by the camera.

The presented results are both superior and of a much wider scope than that of [15]: we have shown that it is necessary neither to strictly enforce the assumptions that the

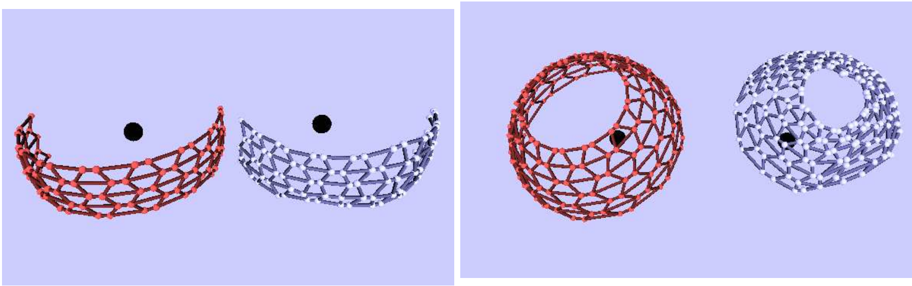

Fig. 10. Calibrations of two different sensors covering more than one hemisphere. On the left, a band-like sensor consisting of 85 photocells, calibrated from correlations (estimated: smaller, true: bigger). On the right, a discrete camera covering more than $180 \times 360^{\circ}$, of 168 photocells, calibrated from the information distance (estimated: smaller, true: bigger). Each ball represents a photocell except the big black balls, representing the optical center.

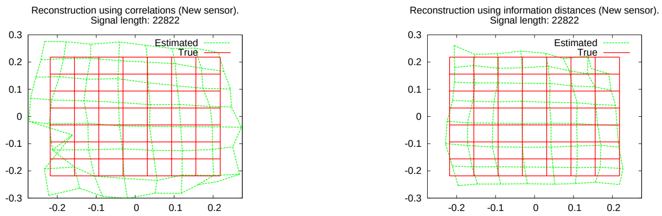

Correlation-Based Information Distance-Based

Fig. 11. Reconstructed and true pixel layouts of a discrete camera consisting of photocells lying on a rectangular grid. The sensor used differs from that with which the models of Fig 5 were built. The reconstructions are obtained by first estimating the pairwise angular distances, then embedding the angles in the sphere (see text). For visualization, the reconstructions are aligned by the usual procrustes method, mapped to the plane by projective mapping with unit focal length. Added line segments show the true pixel neighborhood relations. The left plot is obtained from the correlation distance, and the right from the information distance.

camera directs each pixel uniformly in all directions, nor that statistically similar environments be used to build the statistic-to-angle table and to calibrate the discrete camera. This flexibility reinforces the impression that models such as those shown in Figure 5 have a more general validity than the context of calibration.

We showed also that angle estimators based on correlation and information distance (entropy) have different performance characteristics. It would be very interesting to apply machine learning techniques to leverage the power of many such weak estimators.

Finally a more curious question is worth asking in the future: can the problem of angle estimation be altogether bypassed in a geometrically meaningful calibration procedure? Embedding methods based on rank or connectivity [17, 22], e.g. correlation or information distance, suggest that this is possible.

References

1. Kohler, I.: Experiments with goggles. Scientific American **206** (1962) 62–72
2. Tsai, R.: An efficient and accurate camera calibration technique for 3D machine vision. In: IEEE Conf. on Computer Vision and Pattern Recognition. (1986)
3. Hartley, R., Zisserman, A.: Multiple View Geometry in Computer Vision. Cambridge University Press (2000)
4. Nistér, D., Stewenius, H., Grossmann, E.: Non-parametric self-calibration. In: proc. ICCV. (2005)
5. Ramalingam, S., Sturm, P., Lodha, S.: Towards complete generic camera calibration. In: Proc. CVPR. Volume 1. (2005) 1093–1098
6. Pierce, D., Kuipers, B.: Map learning with uninterpreted sensors and effectors. Artificial Intelligence Journal **92**(169–229) (1997)
7. Krzanowski, W.J.: Principles of Multivariate Analysis: A User's Perspective. Clarendon Press (Statistical Science Series, 1988)
8. Olsson, L., Nehaniv, C.L., Polani, D.: Sensory channel grouping and structure from uninterpreted sensor data. In: NASA/NoD Conference on Evolvable Hardware. (2004)
9. Crutchfield, J.P.: Information and its metric. In Lam, L., Morris, H.C., eds.: Nonlinear Structures in Physical Systems–Pattern Formation, Chaos and Waves. Springer-Verlag (1990) 119–130
10. Grossmann, E., Orabona, F., Gaspar, J.A.: Discrete camera calibration from the information distance between pixel streams. In: Proc. Workshop on Omnidirectional Vision, Camera Networks and Non-classical Cameras, OMNIVIS. (2007)
11. Torralba, A., Oliva, A.: Statistics of natural image categories. Network: Computation in Neural Systems **14** (2003) 391–412
12. Freeman, W.T., Pasztor, E.C., Carmichael, O.T.: Learning low-level vision. International Journal of Computer Vision **40**(1) (2000) 25–47
13. Potetz, B., Lee, T.S.: Scaling laws in natural scenes and the inference of 3d shape. In: NIPS – Advances in Neural Information Processing Systems, MIT Press (2006) 1089–1096
14. Wu, Y.N., Zhu, S.C., Guo, C.E.: From information scaling of natural images to regimes of statistical models. Technical Report 2004010111, Department of Statistics, UCLA (2004)
15. Grossmann, E., Gaspar, J.A., Orabona, F.: Discrete camera calibration from pixel streams. In: Submitted to Computer Vision and Image Understanding. (2008)
16. Roy, R.: Spectral analysis for a random process on the sphere. Annals of the institute of statistical mathematics **28**(1) (1976)
17. Dattorro, J.: Convex Optimization & Euclidean Distance Geometry. Meboo Publishing (2005)
18. Geyer, C., Daniilidis, K.: A unifying theory for central panoramic systems and practical applications. In: proc. ECCV. Volume II., Springer-Verlag (2000) 445–461
19. Schoenberg, I.J.: Remarks to Maurice Fréchet's article "Sur la définition axiomatique d'une classe d'espaces distanciés vectoriellement applicable sur l'espace de Hilbert". Annals of Mathematics **36**(3) (1935) 724–732
20. Sammon, J.W.J.: A nonlinear mapping for data structure analysis. IEEE Transactions on Computers **C-18** (1969) 401–409
21. Lee, R.C.T., Slagle, J.R., Blum, H.: A triangulation method for the sequential mapping of points from n-space to two-space. IEEE Trans. Computers **26**(3) (1977) 288–292
22. Shang, Y., Ruml, W., Zhang, Y., Fromherz, M.P.J.: Localization from mere connectivity. In: MobiHoc '03: Proc. ACM Intl. Symp. on Mobile Ad Hoc Networking & Computing, ACM Press (2003) 201–212