

Radiometric Calibration of a Helmholtz Stereo Rig

(corrected version)

Zsolt Jankó

Computer and Automation Research Institute
Hungarian Academy of Sciences
Budapest, Hungary
janko@sztaki.hu

Ondřej Drbohlav and Radim Šára

Center for Machine Perception
Czech Technical University
Prague, Czech Republic
{drbohlav,sara}@cmp.felk.cvut.cz

Abstract

Helmholtz stereopsis guarantees unbiasedness by BRDF of the search for inter-image correspondences. In a practical setup, calibrated pixel sensitivity and corrected light anisotropy are required for the method to work well. In this paper a simple method for joint light-camera radiometric calibration is proposed. Such calibration is shown to be an ill-posed numerical problem. In the worst case of a single image pair, it can be regularized along epipolar lines only. The calibration problem becomes regularizable everywhere in the case of more pairs and cameras in a general position. A general, simple and fast calibration procedure is proposed that includes a regularizer for the single-pair case. It is shown in a ground-truth experiment that the accuracy of reconstructed surface normals improves by an order of magnitude after radiometric calibration.

1 Introduction

Recently, a number of papers appeared on applications of a symmetry of bi-directional reflectance distribution function (BRDF) called the *Helmholtz reciprocity* [1–3] in stereoscopic matching [4–7], surface registration [8] and shape from reflectance [9].

Helmholtz reciprocity predicts $f(\mathbf{i}, \mathbf{e}) = f(\mathbf{e}, \mathbf{i})$, where $f(\mathbf{i}, \mathbf{e})$ is the surface BRDF (the ratio of the outgoing radiance to the incident irradiance), \mathbf{i} is the direction of the incident light ray and \mathbf{e} the direction of the outgoing ray [1]. A simple stereoscopic setup employing Helmholtz reciprocity consists of a pinhole camera and a point source of isotropic light that can swap their positions in space. *By the swap each individual sensor ray given by an image pixel becomes the same ray emitting light and vice versa*, see Fig. 1. As usual in stereo, two images are taken: one before the swap and one after it. The images look natural, except, as one immediately notices, with all specularities staying fixed relative to surface features [9]. This is because of the reciprocity: The otherwise unknown BRDF does nothing more than just multiplication by an unknown scalar of




Fig. 1: Helmholtz rays $\mathbf{i}_1, \mathbf{i}_2$ at surface point \mathbf{x} in cameras $R_1 = (S_1, C_1)$ and $R_2 = (S_2, C_2)$. The $\rho(\mathbf{i})$ is the radiance along \mathbf{i} when \mathbf{i} acts as a source ray and $\sigma(\mathbf{i})$ is its sensitivity when it acts as a sensor ray. Surprisingly, in Helmholtz stereo it is the ratio $\rho(\mathbf{i})/\sigma(\mathbf{i})$ for each \mathbf{i} that is recoverable from calibration *and* can be used to correct raw image measurements.

every pair of measurements coming from a given surface point. This is in contrast to a standard setup with a fixed light source in which a non-isotropic BRDF causes large non-linear changes in appearance induced by camera motion.

In their paper [9] Magda et al. first formulated the Helmholtz reciprocity condition in a way that can be employed in computer vision. Two pinhole cameras located at $\mathbf{c}_1, \mathbf{c}_2$, co-located with point light sources are assumed. In other words one takes image I_1 by Camera 1 when Camera 2 acts as a light and then takes image I_2 by Camera 2 under Camera 1 illuminating. It is shown that image irradiance values $I_1(\mathbf{i}_1), I_2(\mathbf{i}_2)$ along optical rays $\mathbf{i}_1, \mathbf{i}_2$ (cf. Fig. 1) must satisfy

$$\mathbf{n}^\top \left(\frac{I_1(\mathbf{i}_1)}{\|\mathbf{c}_1 - \mathbf{x}\|^2} \mathbf{i}_1 - \frac{I_2(\mathbf{i}_2)}{\|\mathbf{c}_2 - \mathbf{x}\|^2} \mathbf{i}_2 \right) = 0 \quad (1)$$

at a surface point \mathbf{x} with normal vector \mathbf{n} . *The expression (1) is valid for any BRDF.* It has two roles: it is a constraint on image consistency for correspondence search *and* a constraint on local shape under the hypothesis of correspondence. Helmholtz reciprocity thus has the potential of having great impact on the accuracy of stereopsis requiring correspondence search based on comparing image values.

Even more: As a free bonus, one obtains a constraint on local shape. All this makes the method very appealing.

In [4, 10] Zickler et al. implemented a working Helmholtz stereopsis. Winner-take-all (WTA) correspondence search based on rank deficiency constraint derived from (1) was used. Normals and depths were estimated in a calibrated setup. The joint constraints on the solution surface and its normal provided by (1) were not yet used. Light anisotropy and radiometric camera calibration was hinted as nontrivial but left undescribed.

In [5] Zickler et al. considered a geometrically uncalibrated multi-view scheme and proposed a stratified 3D surface reconstruction method. A new matching constraint is derived that requires neither geometric nor radiometric calibration on the assumption that camera photoresponse is linear and spatially uniform and the lights are isotropic. WTA correspondence search based on (1) is used. Under distant lights one has a bilinear calibration-estimation problem solvable up to the $GL(3)$ group. Assuming affine cameras, the problem is disambiguated up to a free affine coordinate frame. Metric reconstruction is possible using a standard method. At least four observed points and four Helmholtz cameras are needed. Surface is integrated from the recovered normal field after enforcing integrability. Good results were obtained on a plastic mannequin head.

Later, Zickler et al. observed that (1) is a first-order nonlinear PDE [6]. The solution to the surface reconstruction problem is found by integration along epipolar lines starting from all possible boundary conditions and then imposing a coherence constraint across epipolar lines using dynamic programming. Infinitely-distant lights are assumed. Good results are obtained on a weakly textured plastic mannequin head up to small scanline artifacts showing up as wrinkles on the surface.

In [7] Tu and Mendonça attempted to employ the whole potential of the constraint (1): depth and orientation are constrained *simultaneously, during the search*. This is possible by observing that the ratio $\mathbf{n}^\top \mathbf{i}_1 / \mathbf{n}^\top \mathbf{i}_2$ is invariant to the projection of the surface normal \mathbf{n} onto the epipolar plane (spanned by $\mathbf{i}_1, \mathbf{i}_2$). These projections are easy to estimate in disparity space during the correspondence search. Dynamic programming is used to solve the matching problem. Small scanline artifacts are visible, too. A better formulation of the optimization problem would probably result in less numerical problems and recovering more detail.

Only qualitative results are reported in all previous work. Upon closer look, it becomes clear that—unless one maintains uniform pixel sensitivity and perfectly isotropic illumination—the recovered normals must be rather imprecise. As we will show the measurement vector which is supposed to be perpendicular to the normal vector \mathbf{n} in (1) can differ by as much as 15° from perpendicularity if radiometric calibration is not done (Fig. 5).

Non-uniform camera photoresponse is relatively easy to calibrate [11] but implementing an isotropic point illuminant is very difficult, if not impossible. Calibrating light anisotropy is possible but that would be insufficient in a Helmholtz setup: the relative orientation of the anisotropy with respect to the camera must be calibrated as well, which, again, is rather difficult. In this paper we propose calibrating cameras and illuminants of the Helmholtz setup jointly. This improves the accuracy over individual light and camera calibrations, since less degrees of freedom are in fact necessary to be calibrated.

2 Radiometric Calibration

For the sake of the subsequent analysis we need to be precise in formulating the problem. Consider a set of lights S_i and perspective cameras C_i with projection centers \mathbf{c}_i , both types of devices represented as ray bundles that span the 3-space (Fig. 1). We assume the cameras are geometrically calibrated (their 3×4 projection matrices \mathbf{P}_i are known). A quadruple (S_1, S_2, C_1, C_2) is called a *Helmholtz (stereo)-pair* if $S_1 = C_1$ and $S_2 = C_2$, i.e. if cameras and lights are co-located. The pair $R_i = (S_i, C_i)$, $i = 1, 2$ will be called a *Helmholtz camera* and the pair (R_1, R_2) the *Helmholtz pair*. The elements of R_i are *rays* which we can identify with Helmholtz camera pixels. We will express the rays as unit-length vectors oriented towards the camera. The relation between pixels and rays is given from geometric camera calibration. The vectors may therefore be thought of as image points in homogeneous coordinates.

We assume that for each ray bundle R_i there is a non-uniform radiometric calibration function, which is the radiance $\rho: R_i \mapsto \mathbb{R}^+$ when R_i acts as a source and the *sensitivity* $\sigma: R_i \mapsto \mathbb{R}^+$ when it acts as a sensor. We show that for correcting raw image measurements in Helmholtz stereo it suffices to know *the ratio* $\mu(\mathbf{i}) = \rho(\mathbf{i})/\sigma(\mathbf{i})$ for every image pixel \mathbf{i} in each image.

Let \mathbf{x} be a surface point and \mathbf{n} be the surface normal at \mathbf{x} , cf. Fig. 1. Let \mathbf{i}_1 and \mathbf{i}_2 be the respective optical rays in camera R_1 and R_2 incident with \mathbf{x} . If I_i , $i = 1, 2$ is the Helmholtz image pair, the image value read at pixel $\mathbf{i}_2 \in R_2$ will be

$$I_2(\mathbf{i}_2) = \rho_1(\mathbf{i}_1) \sigma_2(\mathbf{i}_2) f(\mathbf{i}_1, \mathbf{i}_2) \frac{\mathbf{i}_1^\top \mathbf{n}}{\|\mathbf{c}_1 - \mathbf{x}\|^2}, \quad (2)$$

where $\rho_1(\mathbf{i}_1)$ is the radiance along the ray $\mathbf{i}_1 \in R_1$, the $\sigma_2(\mathbf{i}_2)$ is the sensitivity of pixel $\mathbf{i}_2 \in R_2$ (it actually lumps nonuniform pixel sensitivity and lens vignetting), f is the BRDF at \mathbf{x} (including albedo), and \mathbf{c}_1 is the position of light/camera R_1 . By swapping subscripts 1 and 2 we obtain the expression for the other image in the Helmholtz pair.

Similarly as in [9] the pair of relations (2) can be used to establish a stereoscopic matching constraint by exploiting

the reciprocity $f(\mathbf{i}_1, \mathbf{i}_2) = f(\mathbf{i}_2, \mathbf{i}_1)$:

$$\mathbf{n}^\top \left(\frac{\mu_1(\mathbf{i}_1) I_1(\mathbf{i}_1)}{\|\mathbf{c}_1 - \mathbf{x}\|^2} \mathbf{i}_1 - \frac{\mu_2(\mathbf{i}_2) I_2(\mathbf{i}_2)}{\|\mathbf{c}_2 - \mathbf{x}\|^2} \mathbf{i}_2 \right) = 0, \quad (3)$$

where we introduced two ratios $\mu_i: R_i \mapsto \mathbb{R}^+$

$$\mu_i(\mathbf{i}_i) = \frac{\rho_i(\mathbf{i}_i)}{\sigma_i(\mathbf{i}_i)}, \quad i = 1, 2, \quad (4)$$

which may be thought of as the *effective sensitivity* of the ray \mathbf{i}_i in camera $i = 1, 2$ of the Helmholtz pair. It is an easy exercise to see that (3) indeed holds and that every μ_i lives in its camera R_i , even if the camera participates in multiple Helmholtz pairs. This means our subsequent analysis applies to any multi-camera Helmholtz setup.

The expression (3) can be directly used for radiometric calibration using planes of known pose as will be described in Sec. 4. In the following section we analyze a sufficient setup for such calibration.

3 Sufficient Setup

This section analyzes the sufficient conditions under which radiometric calibration is possible, describes a simple calibration setup and analyzes its efficiency. It first deals with the case of a single Helmholtz pair and later generalizes it to more pairs. As it turns out the calibration problem becomes numerically well-posed only if more than a single Helmholtz pair are calibrated. The goal is to calibrate effective sensitivities μ_i in both images of the pair by a procedure that is as simple as possible. The idea is to insert a plane performing a calibrated motion in the object space.

Let there be planes Π_j , $j = a, b, \dots, p$ represented by known homogeneous 4-vectors $[\mathbf{n}_j, d_j]^\top$, in which \mathbf{n}_j is the unit normal vector of plane Π_j . Plane Π_j induces bijection (homography) \mathbf{H}_j transferring rays (pixels) from R_1 (image 1) to R_2 (image 2) of the Helmholtz pair. Given a ray $\mathbf{i}_1 \in R_1$ and the induced Helmholtz ray $\mathbf{i}_2 = \mathbf{H}_j \mathbf{i}_1$ in R_2 , we get the following ratio from (3):

$$\kappa(\mathbf{i}_1, \mathbf{i}_2 | \Pi_j) = \frac{\mu_1(\mathbf{i}_1)}{\mu_2(\mathbf{i}_2)} = \frac{\mathbf{n}_j^\top \mathbf{i}_2 \|\mathbf{c}_1 - \mathbf{x}\|^2 I_2(\mathbf{i}_2 | \Pi_j)}{\mathbf{n}_j^\top \mathbf{i}_1 \|\mathbf{c}_2 - \mathbf{x}\|^2 I_1(\mathbf{i}_1 | \Pi_j)}, \quad (5)$$

where $\mathbf{x} \in \Pi_j$ and the Π_j in $I(\cdot | \Pi_j)$ is used to stress the fact we are working with the images of plane Π_j . Given cameras R_i , $i = 1, 2$, calibration plane Π_j , and the measurements $I_1(\cdot | \Pi_j)$, $I_2(\cdot | \Pi_j)$ the ratio $\kappa(\mathbf{i}_1, \mathbf{i}_2 | \Pi_j)$ is known. Note that the radiometric calibration plane may have arbitrary BRDF, including albedo, which means we can re-use the images of our geometric camera calibration plane.¹

¹As long as the geometric calibration pattern is not too dark causing loss of radiometric resolution due to image quantization.




Fig. 2: The pair $\mathbf{i}_1^1, \mathbf{i}_1^2$ of the propagation chain $\chi(\mathbf{i}_1^1)$ on which the effective sensitivity μ_1 of Camera 1 is defined up to scale (left). Propagation chains almost never meet (right).

The formula (5) gives just the ratio of effective sensitivities *across the cameras*. We shall see now one more calibration plane suffices to factor out the ratio to its components μ_1 and μ_2 . Suppose we have two calibration planes Π_j , $j = a, b$, each inducing an inter-image homography \mathbf{H}_j . Given an initial ray \mathbf{i}_1^1 in Camera 1, its intersection with plane Π_a induces the ray $\mathbf{i}_2 = \mathbf{H}_a \mathbf{i}_1^1$ in Camera 2 for which the ratio $\kappa(\mathbf{i}_1^1, \mathbf{i}_2 | \Pi_a)$ is known, follow Fig. 2 (left). We use the ray \mathbf{i}_2 to transfer \mathbf{i}_1^1 to \mathbf{i}_1^2 in Camera 1 via the intersection of \mathbf{i}_2 with Π_b : $\mathbf{i}_1^2 = \mathbf{H}_b^{-1} \mathbf{H}_a \mathbf{i}_1^1$. This and (5) give us the ratio $\kappa(\mathbf{i}_1^2, \mathbf{i}_2 | \Pi_b)$, which, in turn, allows us to get another ratio

$$r_1(\mathbf{i}_1^1, \mathbf{i}_1^2) = \frac{\kappa(\mathbf{i}_1^1, \mathbf{i}_2 | \Pi_a)}{\kappa(\mathbf{i}_1^2, \mathbf{i}_2 | \Pi_b)} = \frac{\mu_1(\mathbf{i}_1^1)}{\mu_1(\mathbf{i}_1^2)}, \quad (6)$$

which lives in Camera 1. Given $\mathbf{i}_1^k, \mathbf{H}_a, \mathbf{H}_b, I(\cdot | \Pi_a), I(\cdot | \Pi_b)$ the ratio $r_1(\mathbf{i}_1^k, \mathbf{i}_1^{k+1})$ is known at every ray (pixel) $\mathbf{i}_1^k \in R_1$ (Camera 1). The homography $\mathbf{H}_1 = \mathbf{H}_b^{-1} \mathbf{H}_a$ mapping \mathbf{i}_1^k to \mathbf{i}_1^{k+1} is a bijection $R_1 \mapsto R_1$ induced by calibration plane motion $\Pi_a \mapsto \Pi_b$ as observed in Camera 1. The ratio r_2 involving μ_2 is based on $\mathbf{H}_2 = \mathbf{H}_b \mathbf{H}_a^{-1}$ in Camera 2.

Given an initial ray $\mathbf{i}_1^1 \in R_1$ the homography \mathbf{H}_1 maps it to the successor \mathbf{i}_1^2 (via \mathbf{i}_2^1 given by the intersection of \mathbf{i}_1^1 with Π_a), which, in turn, gets mapped to \mathbf{i}_1^3 (via \mathbf{i}_2^2 given by the intersection of \mathbf{i}_1^2 with Π_a), etc. We obtain discrete *oriented propagation chains* $\chi(\mathbf{i}_i^1), i = 1, 2$ induced by \mathbf{H}_i starting at $\mathbf{i}_i^1 \in R_i$.² It is clear from Fig. 2 the rays of the discrete chain all lie in the same epipolar plane given by the starting ray and camera locations $\mathbf{c}_1, \mathbf{c}_2$. Every ray (pixel) in R_i gives rise to one such chain. Clearly, one obtains the function μ_i up to a scalar constant along every chain $\chi(\mathbf{i}_i^k)$ by propagating an unknown value of $\mu_i(\mathbf{i}_i^k)$ from the starting point $\mathbf{i}_i^k \in R_i$. Alternatively, for every pair of rays $\mathbf{i}_i^k, \mathbf{i}_i^{k+1}$ in R_i related by $\mathbf{i}_i^{k+1} = \mathbf{H}_i \mathbf{i}_i^k$ one has the following constraint on

²Note the chains of *pixels* in both Helmholtz images are induced by the same set of spatial points. They are therefore *not* independent.

$$\mu_i(\mathbf{i}_i^{k+1}) = \mu_i(\mathbf{i}_i^k) \cdot r_i(\mathbf{i}_i^k, \mathbf{i}_i^{k+1}), \quad (7)$$

in which $r_i(\mathbf{i}_i^k, \mathbf{i}_i^{k+1})$ is known from geometric calibration and radiometric calibration measurement according to (6). From (7) it is clear the solution μ_i is recoverable up to a set of positive scalar multipliers, *one per chain* χ . We shall see that in a single Helmholtz pair setup, the chains do not generally communicate, hence there is no constraint binding these multipliers together.

If we express the system (7) for $i = 1, 2$ in terms of all the raw measurements I_1, I_2 and the geometric terms, we can see the solutions for μ_1 and μ_2 are not independent, since exactly the same set of measurements participates in recovering each of them. It therefore suffices to study the solvability for one fixed i , say $i = 1$.

Let the calibration plane undergo motion such that the chain converges towards the baseline vector $\mathbf{b} = \mathbf{c}_1 - \mathbf{c}_2$ in R_i . The chain never reaches \mathbf{b} as the successive steps get ever smaller, see Fig. 2 (right). The image epipole is therefore a point of convergence (a fixed point of \mathbf{H}_i), which means it is not a ‘point of exchange’ for the image chains and therefore there is no new constraint on μ_i that would relate the individual chains approaching the epipole. This situation does not improve with more calibration planes. Note even the chains along a single epipolar line are generally disjoint, hence it is impossible to recover the function μ_i without additional constraints. The solution can be recovered by regularization assuming the function μ_i is smooth.

If $p > 2$ radiometric calibration planes are used, the situation stays essentially the same except the number of constraints of the type (7) increases dramatically, since every ray \mathbf{i} has $\binom{p}{2}$ successors (one per pair of planes) instead of just one, and the oriented propagation chains become *oriented propagation graphs* along every epipolar line. The graphs remain essentially independent as in the case of chains, however.

All entities in (7) are positive. By taking logarithms,

$$\lambda_i(\mathbf{i}_i^{k+1}) = \lambda_i(\mathbf{i}_i^k) + \delta_i(\mathbf{i}_i^k, \mathbf{i}_i^{k+1}), \quad (8)$$

where $\lambda(\cdot) = \log \mu(\cdot)$ and $\delta(\cdot, \cdot) = \log r(\cdot, \cdot)$. Similarly as in (7) the solution λ is recoverable up to an additive scalar per propagation chain χ . Due to the discrete nature of the image one has to interpolate the function λ_i at $\mathbf{i}_i^{k+1} = \mathbf{H}_i \mathbf{i}_i^k$. The interpolation coefficients are known and can be aggregated in a sparse matrix \mathbf{A}_i of size $m \times n$, where n is the number of rays in R_i having a successor in R_i (i.e. up to image boundary) and m is the number of constraints (8). There are $m = n \binom{p}{2} - n_0$ constraints if p calibration planes are used. The $n_0 \sim \sqrt{n}$ is the number of constraints unavailable at the image boundary. The collection of all m constraints is then expressed as

$$(\mathbf{A}_i - \mathbf{E}) \mathbf{l}_i = \mathbf{d}_i, \quad (9)$$

where \mathbf{l}_i is the stacked vector of all values $\lambda_i(\cdot)$ at all rays in R_i , the vector \mathbf{d}_i stacks the elements $\delta_i(\cdot, \cdot)$, and \mathbf{E} is an identity matrix. The interpolation kernel acts as a weak regularizer that binds the solutions λ along close individual chains together. In the worst case of parallel epipolar lines such constraint need not suffice in the direction *orthogonal* to epipolar lines: The system (9) breaks down to independent under-constrained subproblems, one per epipolar line. Radiometric calibration that is independent for each epipolar line need not be a problem as long as the stereo correspondence search is also going to be done on epipolar lines independently.

Consider now the case of more than a single Helmholtz pair. If one camera R_i participates in several such pairs, there is a set of constraints on μ_i per pair.³ The chains become interconnected, forming a dense graph whose nodes cover the entire image domain. Therefore, the problem of isolated chains along epipolar lines disappears, as long as the systems of epipolar lines are independent. Two calibration planes then suffice to calibrate by simply solving the system of equations (9) which is now over-constrained. Note that the solution μ_i will be automatically positive definite by exponentiation of λ_i . It will be obtained up to an overall positive multiplicative scalar.

4 Implementation and Experiments

In this section, the actual calibration procedure is first described, including a ground-truth experimental setup. It is then shown how radiometric calibration improves the accuracy of recovered surface normals.

So far we used \mathbf{i} to denote a spatial direction vector associated with a given pixel in (one of) the Helmholtz cameras. From now on the \mathbf{i} will stand for an image pixel. We hope this will not lead to confusion.

Instead of solving the large system (9) directly, we approximate the solution by a mixture of a small number of given positive semi-definite kernels $\ker(\mathbf{i})$

$$\mu_i(\mathbf{i}) = \sum_{l=1}^L \alpha_i^l \ker(\mathbf{i} - \mathbf{i}_0^l), \quad (10)$$

whose *control points* \mathbf{i}_0^l are fixed locations in the image plane. This choice has three reasons: (1) it significantly improves the computational complexity of the radiometric calibration, (2) it regularizes the problem in case just two calibration planes are used with a single Helmholtz pair and (3) it filters the solution giving smooth μ .

Let camera projection matrices \mathbf{P}_i be given. A radiometric calibration object of known geometry is inserted in the scene. Helmholtz images of the object are acquired. The

³Recall the radiometric calibration function $\mu_i: R_i \mapsto \mathbb{R}^+$ for every given i stays the same in all pairs.




Fig. 3: The measurement setup with light S and camera C (left). The motorized stage has two degrees of freedom: rotation ϕ and vertical translation t . Raw calibration image from one camera (right). The bar at right shows the range of image values.

surface points \mathbf{x}^s , $s = 1, 2, \dots, S$ are sampled from the known parametric model of the calibration object. From this data the coefficients α_i^l are found as follows. The 3D points \mathbf{x}^s with normals \mathbf{n}^s are projected to images via camera projection matrices \mathbf{P}_i . The values of μ_i at the projected points \mathbf{i}_i^s are expressed in the form (10) in each of the images, with the coefficients α_i^l remaining the only unknowns. The interpolated values of μ_i are plugged in equation (3). One 3D point \mathbf{x}^s thus gives one equation linear in α_i^l . In total, S equations are obtained for $n \times L$ unknowns (L unknowns for each of the n cameras). The system is then solved under the additional constraints $\alpha_l \geq 0$ and $\sum_{l^*} \alpha_{i^*}^{l^*} = 1$, where l^* are indices of control points $\mathbf{i}_0^{l^*}$ that lie within the respective image boundary. This is a constrained linear least-squares fitting problem which can be solved very fast.

The stereo rig used in the experiment, shown in Fig. 3, consisted of a digital radiometric camera⁴ C and a custom-built light source⁵ S , both fixed in space. The camera-light distance was 60cm. The object was placed on a precise motorized rotation stage 1m away from the camera. The rotation stage was carefully manually centered to achieve reciprocity.⁶ The Helmholtz images were taken by rotating the stage around its axis by $\phi = 180^\circ$. Fig. 3 (right) shows one raw image acquired from the rig. Camera dark field was subtracted from all acquired images in the experiment.

Geometric calibration of the cameras was done using a planar glass calibration panel with a high-accuracy calibration pattern that was moved by a known distance of $t = 50\text{mm}$. This gave us camera projection matrices \mathbf{P}_i .

⁴Vosskuehler COOL-1300, 12 bit monochromatic CCD, 2/3", Peltier cooled, 1280 × 1024 square pixels, with a quality 1"-chip manual lens LINOS MeVis-C 50mm 1:1.8, negligible vignetting.

⁵A clear 150W incandescent bulb powered by a stabilized DC current, enclosed in a black housing with a round hole of 6cm in diameter, 14cm distant from the filament.

⁶The rotation axis bisects the camera-light line at 90°.




Fig. 4: The solution for the effective sensitivity μ_1 with high-lighted isocontours at values $[0.85 : 0.05 : 1.5]$. The bar at right shows the range of the values. The other Helmholtz camera was positioned to the right of this one.

The panel was then covered with a sheet of standard white paper and re-used for radiometric calibration.

Each of the two radiometric calibration planes is sampled in a uniform rectangular grid of $S = 130$ points.⁷ Bi-cubic kernel $\ker(x, y) = \sum_{u=0}^2 a_u x^u \cdot \sum_{v=0}^2 b_v y^v$ was used and Matlab[®] function `lsqlin` for solving constrained linear least squares was used to recover the interpolation parameters of each μ_i , $i = 1, 2$ as described in the previous section. Fig. 4 shows the reconstructed effective sensitivity μ_1 for Camera 1.

The next goal was to assess the influence of radiometric calibration on the accuracy of recovered normals. Additional known plane Π_v at $t = 24\text{mm}$ was inserted into the calibrated space. The plane was sampled and the points \mathbf{x}^s projected to images, as above. Each pair of projected points gave us one vector \mathbf{m} according to (3):

$$\mathbf{m} = \frac{\mu_1(\mathbf{i}_1) I_1(\mathbf{i}_1)}{\|\mathbf{c}_1 - \mathbf{x}^s\|^2} \mathbf{i}_1 - \frac{\mu_2(\mathbf{i}_2) I_2(\mathbf{i}_2)}{\|\mathbf{c}_2 - \mathbf{x}^s\|^2} \mathbf{i}_2. \quad (11)$$

This vector is supposed to be orthogonal to the known plane normal \mathbf{n} . The minimal angular difference from orthogonality was measured as the error function. The left column of Fig. 5 shows the error without considering radiometric calibration (with $\mu_i \equiv 1$) and the second column with the calibrated μ_i . Note the different ranges on the horizontal axes. Both results are shown on three planes: the two radiometric calibration planes (second and fourth row) and the independent cross-validation plane (third row). The following table shows the improvement of RMS error in angular degrees due to radiometric calibration in the form mean-value \pm RMS. Important are the RMS values. Results on validation plane Π_v are shown in bold. The error remaining after calibration is at the precision limit of the calibration setup (calibration panel non-flatness).

⁷The number varied for the different planes according to joint visibility in the images.

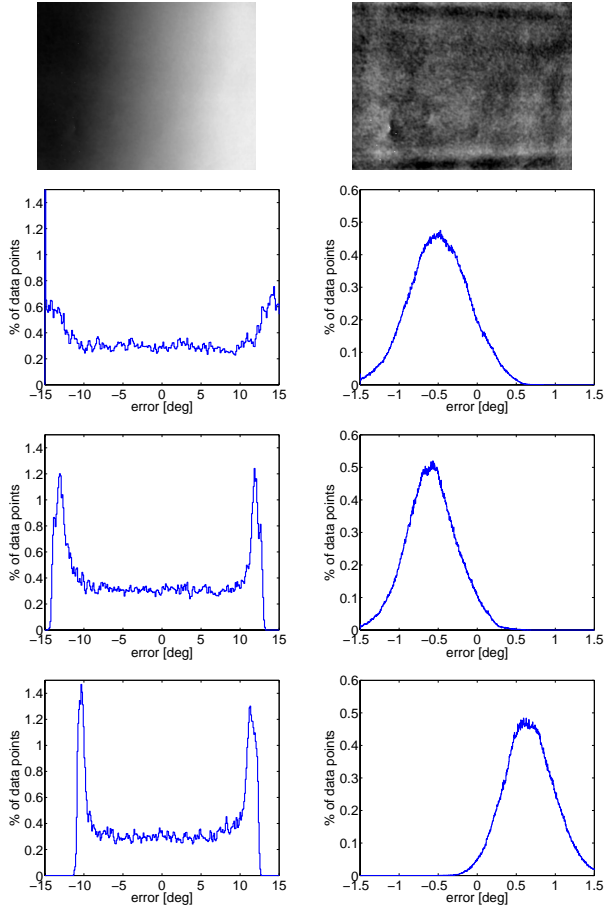


Fig. 5: Error in degrees before (left column) and after (right column) radiometric correction, error ranges are $[-15, 15]$ and $[-1.5, 1.5]$, respectively. The top row shows spatial error distribution in the jointly visible area. The 3rd row shows results on an independent cross-validation plane Π_v , the 2nd and 4th rows show results on calibration planes Π_a , Π_b .

plane	before [deg]	after [deg]
Π_a	-0.96 ± 10.4	-0.49 ± 0.39
Π_v	-1.1 ± 9.1	-0.58 ± 0.33
Π_b	1.1 ± 8.1	0.67 ± 0.32

5 Conclusions and Future Work

In this paper we have shown that for the success of accurate Helmholtz stereopsis radiometric calibration is as essential as geometric camera calibration. We have observed the radiometric calibration problem is numerically ill-posed. A simple regularized radiometric calibration method was proposed which is fast and accurate without requiring a specially equipped photometric lab. An experiment testing the necessary condition of Helmholtz stereopsis showed that the accuracy of the constraint improves dramatically after calibration. RMS error of 0.33° was achieved using a pla-

nar calibration object undergoing a single translation.

All previous work including ours assumed (but not verified) the setup is truly Helmholtzian. An experimental procedure that would let us tune the setup to be as close to Helmholtzian as possible is left for future work. So is the evaluation of accuracy on more complex, curved objects. The improved accuracy will undoubtedly improve the discriminatory power of the Helmholtz condition (1) when used in an image consistency statistic for stereoscopic matching.

Acknowledgments

This work was supported by the following grants: T038355 and M28078 of Hungarian Scientific Research Fund (OTKA), LN00B096 of the Czech Ministry of Education, and CONEX GZ 45.535 of the Austrian Ministry of Education. The authors wish to thank Vladimír Smutný from CMP for his help with the experimental setup.

References

- [1] H. von Helmholtz, *Optique physiologique*. Sceaux, France: Éditions Jaques Gabay, I: 231–2, 1989.
- [2] F. J. J. Clarke and D. J. Parry, Helmholtz reciprocity: Its validity and application to reflectometry, *Lighting Research & Technology*, 17(1): 1–11, 1985.
- [3] W. Snyder, Z. Wan, and X. Li, Thermodynamic constraints on reflectance reciprocity and Kirchoff’s law, *Applied Optics*, 37(16): 3464–3470, 1998.
- [4] T. Zickler, P. N. Belhumeur, and D. J. Kriegman, Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction, in *Proc ECCV*, 3: 869–884, 2002.
- [5] T. E. Zickler, P. N. Belhumeur, and D. J. Kriegman, Toward a stratification of Helmholtz stereopsis, in *Proc CVPR*, 1: 548–555, 2003.
- [6] T. E. Zickler, J. Ho, D. J. Kriegman, J. Ponce, and P. N. Belhumeur, Binocular Helmholtz stereopsis, in *Proc ICCV*, 2: 1411–7, 2003.
- [7] P. Tu and P. R. S. Mendonça, Surface reconstruction via Helmholtz reciprocity with a single image pair, in *Proc CVPR*, 1: 541–7, 2003.
- [8] P. Tu, P. R. S. Mendonça, J. Ross, and J. Miller, Surface registration with a Helmholtz reciprocity image pair, in *Proc IEEE Wkshp on Color and Photometric Methods in Computer Vision*, 2003.
- [9] S. Magda, D. J. Kriegman, T. Zickler, and P. N. Belhumeur, Beyond Lambert: reconstructing surfaces with arbitrary BRDFs, in *Proc ICCV*, 2: 391–8, 2001.
- [10] T. Zickler, P. N. Belhumeur, and D. J. Kriegman, Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction. *Int J Comp Vision*, 49(2): 215–27, 2002.
- [11] B. M. Ratliff, M. M. Hayat, and T. J. Scott, Radiometrically accurate scene-based nonuniformity correction for array sensors, *J Opt Soc Am Ser A*, 20(10): 1890–9, 2003.