Optimising Light Source Positions to Minimise Illumination Variation for 3D Vision

Tom Botterill
Department of Computer Science
University of Canterbury
Christchurch, NZ
tom.botterill@canterbury.ac.nz

Steven Mills
Department of Computer Science
University of Otago
Dunedin, NZ

Richard Green
Department of Computer Science
University of Canterbury
Christchurch, NZ

Thomas Lotz
Department of Mechanical Engineering
University of Canterbury
Christchurch, NZ

Abstract

Machine vision lighting systems must be designed to evenly illuminate scenes, so that objects' appearances do not change depending on their position. This is challenging when imaging 3D scenes where the camera's field of view or depth of field are large. This paper describes a scheme to design effective lighting systems in such situations: lighting across the scene is modelled, and positions of lights are optimised to minimise variation in illumination levels. The scheme is used to design a lighting system for an agricultural robot, which reduces variation in the appearance of crops due to uneven lighting. The lighting for a compact medical imaging device is redesigned, which will result in a 53% reduction in the range of lighting levels in the scene.

1. Introduction

Machine vision systems are deployed in a wide range of environments for many different applications, and often these systems need additional lighting. For these systems to work effectively, lighting systems must be designed so that variations in the appearance of objects imaged in different parts of the scene are minimised [1, 2, 3].

Many contemporary systems image objects with similar sizes and positions, for example items on a conveyor belt, and measure only a small number of characteristics of these objects (many of these systems are reviewed by [4]). For most of these systems, the depth of field (DOF) and field of view (FOV) are relatively small, and effective illumination can be obtained by placing a light source near to the camera (Figure 1(a)). Modern systems however, often image a variety of 3D objects distributed across a range of depths, for example vision systems on mobile platforms such as road vehicles [5] and robots [6]. Other modern systems operate in confined spaces, for example underneath a tractor [7], or inside a compact medical imaging device [8]. In these situations, where the DOF or FOV are large, lighting becomes challenging, and simple configurations of lights lead to large variations in illumination levels across the scene (Figures 1(b) and 1(c)).

In this paper we describe a method for designing lighting systems to give even illumination across complex 3D scenes with large DOF or FOV, and without requiring exact knowledge of the structure or position of the objects being imaged. The positions of a set of light sources are optimised to minimise the range of lighting levels, hence minimising the variation of the appearance of objects. The system is used to design the lighting system on an agricultural robot (Figure 2), and to redesign the lighting for a medical imaging system. The systems designed substantially reduce the variation in lighting in the scenes which are imaged.

This paper is organised as follows: the following section describes previous work on the automatic design and control of lighting systems, Section 3 describes the scene illumination model used, and Section 4 describes the cost function and optimisation scheme used to position light sources. Section 5 describes and analyses three lighting systems designed using our scheme, and Section 6 concludes with an overview of our findings.

2. Background

A considerable amount of research has been undertaken on the optimal positioning of sensors and lights on robots,

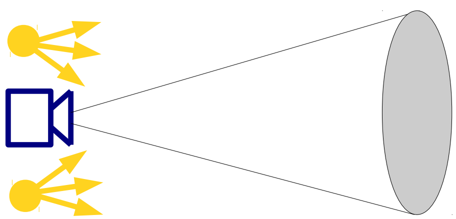

(a) When the FOV and DOF are small, an even lighting field can be obtained easily by positioning light sources close to the camera.

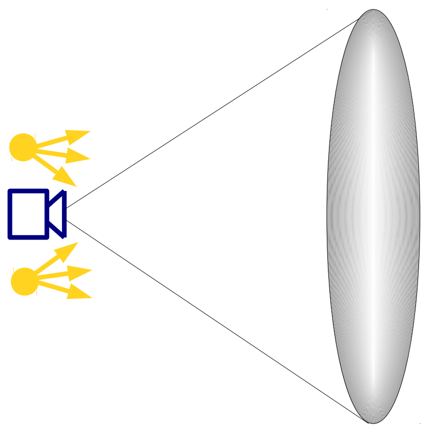

(b) When the FOV is large, typical light sources positioned near to the camera lead to an unevenly-illuminated scene, as points near the centre of the FOV are closer to the light, and are considerably brighter than points at the edge.

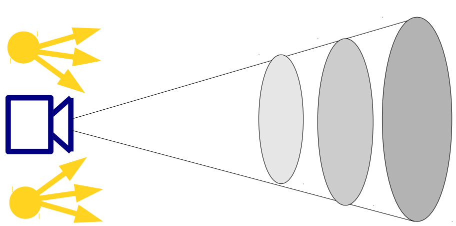

(c) When the DOF is large, light sources positioned near to the camera lead to large variation in light intensity at different depths.

Figure 1. Simple lighting configurations result in uneven illumination when the FOV or DOF are large.

Figure 2. Lighting system deployed on a vine imaging robot, and a typical vine image.

however typically only one or two lights are moved or directed to better illuminate some known object of interest. These schemes include the system by Cowan and Bergman [9] for computing the regions in which a light and camera may be placed while imaging an object. The set of light and camera positions where the object is fully visible, is close enough to make good use of the camera's resolution, and is illuminated so that the surfaces have brightnesses within the dynamic range of the camera is identified. Similarly, Yi et al. [10] position a camera and light with respect to an object of even colour so that the proportion of edges which have sufficient contrast to be visible to the camera is maximised. The cost function is multimodal, and is optimised by using a global search to alternately optimise the position of the light or camera, while fixing the other. Similarly, for the task of positioning a camera to maximise object visibility in a cluttered environment, Triggs and Laugier [11] also find that the objective function is complex and multimodal, and again use a global search to find a good solution. On mobile platforms, lighting may need to be be controlled dynamically; for example Rife and Rock [12] control the power and direction of an Autonomous Underwater Vehicle's spotlights to illuminate a target (a jellyfish) which is followed, while minimising power usage.

These schemes use models of the objects being imaged to position the lights; alternatively Marchand [13] proposes dynamically moving a light and camera to increase a heuristic measure of the amount of contrast in the image. The system moves the light and camera until the object is fully visible and clearly illuminated.

In this paper we aim to create an even lighting field by positioning multiple light sources; an alternative approach suitable for planar scenes is to design a diffuser or Fresnel lens which creates the same effect. Chen and Uang [14] use a genetic algorithm to design a Fresnel lens for a reading light which gives an even level of illumination over a planar lighting field. The positions of grooves in the lens are optimised to maximise efficiency plus a measure of uniformity.

3. Lighting and camera model

A wide range of light sources are manufactured for use in machine vision, and in addition custom made systems, often LED-based, are frequently used. Each light source has a radiation characteristic function (RCF), which describes how the intensity of the emitted light varies with direction. RCFs are available from manufacturers' datasheets; some examples for "bare" (unlensed) high power LEDs [15, 16], a commercially available "wide angle brick" machine vision light [17], and a wide angle ring light [18], are shown in Figure 3.

One of these light sources, positioned at $\mathbf{L}$, and pointing in direction $\mathbf{D}$, illuminates an object at position $\mathbf{X}$ in proportion to

$$\frac{f(\text{angle}(\mathbf{D}, \mathbf{X} - \mathbf{L}))}{\|\mathbf{X} - \mathbf{L}\|^2} \tag{1}$$

where 'angle' denotes the angle between the two vectors, and f is the RCF.

Models for the appearance of 3D scenes are commonly used in computer graphics [19]. A camera measures the light reflected from an object, which is a function of the object's colour, orientation, and reflectance properties. In general, the light reflected varies depending on the angle between the surface normal and camera, however many objects are approximately Lambertian reflectors, in which case the perceived colour is independent of the viewing angle. If the point $\mathbf{X}$ lies on a Lambertian surface with normal $\mathbf{n}$, then the measured light is proportional to

$$\mathbf{n} \cdot \frac{(\mathbf{X} - \mathbf{L})}{\|\mathbf{X} - \mathbf{L}\|}. \tag{2}$$

In a scene illuminated by N point light sources, positioned at $\mathbb{L} = \{\mathbf{L}_1, \mathbf{L}_2, ..., \mathbf{L}_N\}$, and pointing in the directions $\mathbb{D} = \{\mathbf{D}_1, \mathbf{D}_2, ..., \mathbf{D}_N\}$, the radiance of $\mathbf{X}$ is proportional to the sum of the contributions from each light source:

$$R(\mathbf{X}; \mathbb{L}, \mathbb{D}) \propto \sum_{i=1}^{N} \frac{f(\text{angle}(\mathbf{D}_i, \mathbf{X} - \mathbf{L}_i))}{\|\mathbf{X} - \mathbf{L}_i\|^2} (\mathbf{n} \cdot \frac{(\mathbf{X} - \mathbf{L})}{\|\mathbf{X} - \mathbf{L}\|}). \tag{3}$$

For each application, objects of interest fall within some region of space Ω. To minimise the variation of an object's appearance with position, we should minimise the range of values for R for each point $\mathbf{X} \in \Omega$. For some applications where scene objects vary in both position and orientation, the surface normal $\mathbf{n}$ is not known. In this situation we assume that an object at $\mathbf{X}$ faces approximately towards the light sources which provide the majority of its illumination, so that $\mathbf{n}.(\mathbf{X} - \mathbf{L})/\|(\mathbf{X} - \mathbf{L})\| \approx 1$.

In this paper a lighting system is designed using these approximation, and is shown to perform well in practice. If the approximations were not appropriate for the objects being imaged however, a model-based approach could be used instead. For scenes with unknown 3D structure, lighting configurations could be designed for and evaluated against large numbers of simulated 3D scenes, with the best-performing configurations being selected.

4. Minimising lighting variation

The aim of this work is to design lighting systems which minimise the relative range of illumination levels across the scene, so that the best use of the camera's sensitivity is made, and to minimise the variation in appearance of objects. The relative range of illumination levels is given by:

$$C_{\text{range}}(\mathbb{L}, \mathbb{D}) = \frac{R_{\text{brightest}}(\mathbb{L}, \mathbb{D}) - R_{\text{darkest}}(\mathbb{L}, \mathbb{D})}{R_{\text{brightest}}(\mathbb{L}, \mathbb{D})} \tag{4}$$

$$\text{where } R_{\text{brightest}}(\mathbb{L}, \mathbb{D}) = \max_{\mathbf{X} \in \Omega} R(\mathbf{X}; \mathbb{L}, \mathbb{D}), \tag{5}$$

$$R_{\text{darkest}}(\mathbb{L}, \mathbb{D}) = \min_{\mathbf{X} \in \Omega} R(\mathbf{X}; \mathbb{L}, \mathbb{D}) \tag{6}$$

so to minimise lighting variation we optimise $\mathbb{L}$ and $\mathbb{D}$ to minimise C_{range}. Light positions $\mathbb{L}$ are generally constrained to some surface or region (i.e. attached to the bottom of a tractor, or inside a medical imaging device).

One practical and versatile implementation of this cost function is to use a discretised Ω, Ω_d, consisting of G points distributed evenly throughout Ω. e.g. in a grid. The lighting level R is then evaluated for each $\mathbf{X} \in \Omega_d$, and the relative range is then computed from the minimum and maximum lighting level.

The problem with evaluating C_{range} on a grid is that valleys form where points of minimum or maximum brightness change from one grid point to another, and the gradient is discontinuous at the valley bottoms (where gradient descent optimisation algorithms simply get stuck). To smooth C_{range}, hence making it more suitable for nonlinear optimisation, we observe that:

$$C_{\text{range}}(\mathbb{L}, \mathbb{D}) = \|\mathbf{R}\|_\infty \tag{7}$$

where $\mathbf{R}$ is a vector of relative lighting intensities at each of the points in Ω_d with coefficients:

$$R_i = \frac{R_{\text{brightest}} - R(\mathbf{X}_i; \mathbb{L}, \mathbb{D})}{R_{\text{brightest}}} \text{ for } \mathbf{X}_1, \mathbf{X}_2, ..., \mathbf{X}_G \in \Omega_d, \tag{8}$$

and $\|\cdot\|_\infty$ is the L_∞ norm, i.e. the maximum coefficient of the vector. The L_∞ norm has the properties that

$$\|\mathbf{R}\|_\infty = \lim_{p \to \infty} \|\mathbf{R}\|_p \text{ and } \|\mathbf{R}\|_\infty \leq \|\mathbf{R}\|_p \leq G^{\frac{1}{p}} \|\mathbf{R}\|_\infty, \tag{9}$$

where the L_p norm is given by $\|\mathbf{R}\|_p = (\sum_{i=1}^{G} |R_i|^p)^{\frac{1}{p}}$. Instead of minimising C_{range}, we minimise

$$C_{L_p}(\mathbb{L}, \mathbb{D}) = \|\mathbf{R}\|_p \tag{10}$$

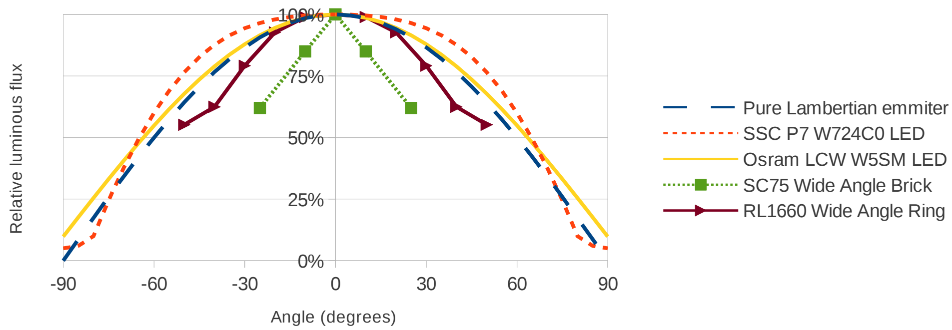

Figure 3. RCFs for some wide angle machine vision lights and LEDs.

which is smooth and differentiable, so can be minimised using any gradient descent-based optimisation, for example Conjugate Gradient Descent (CGD) [20]. Choosing $p = 32$ provides a suitable balance between L_p being a good approximation to L_∞ (typically 2% greater), while still being smooth and numerically stable to compute. Choosing $G \approx 1000$ avoids significant artefacts or valleys in the cost function.

Although C_{L_p} is smooth, it is also multimodal, with many false minima caused by the shape of Ω, and by stable configurations of less than N lights. With typically $4N$ parameters, the parameter space is generally too large for a full global search, so to find good solutions, multiple runs of CGD optimisation are used, with starting points chosen randomly from throughout the parameter space. The best parameter set found usually has lights in a regular configuration, for example in a ring. Intermediate solutions often have regular structure in some regions but not others, so to quickly find solutions where all lights are arranged regularly, a genetic algorithm is used to choose some starting points from a population of the best parameter sets found so far.

5. Experimental results and case studies

This section describes three situations in which the proposed scheme has been used to design lighting layouts. Firstly, the system is used to design the lighting system for a tractor-mounted robot; secondly, the lighting system in a medical imaging system is redesigned; and thirdly, the use of the system to design lighting layouts to illuminate planar scenes is analysed.

5.1. Lighting vines viewed in the confined space under a tractor

The original motivation for this system was to design an illumination system for a tractor-mounted robot which images vines [7]. The robot has three cameras in a trinocular stereo rig, each fitted with 110° FOV lenses. Dormant woody vines (which are approximately Lambertian) are imaged at a depth of between 35cm and 65cm from the cameras. Even illumination is needed to identify dead, diseased or damaged branches, and to image features such as buds and crossing points.

Lighting is provided by N sets of four SSC P7 white wide-angle LEDs [15], with each set mounted on an aluminium heat sink. Lights are constrained to the same planar surface as the cameras. Cameras are set up so that points at depths between 35cm and 65cm from the central forward-pointing camera are visible to all three cameras; this set of points forms the set Ω. Either the lights' positions alone are optimised, while keeping lights forward-pointing, or both the lights' positions and their directions are optimised. When optimising light directions, a barrier term is used to constrain the optimisation so that the direction of lights is within 45° of the camera direction. The optimisation is run for one hundred thousand CGD runs for the hardest cases (N=100), although C_{range} decreases by less than 1% in total after ten thousand runs, indicating good convergence. For small values of N, good configurations are found in a few tens of runs.

Lighting configurations found are shown in Figure 4. Lights with fixed directions are positioned in a ring around the central camera. When both light directions and positions are optimised, the ring is larger, with inward pointing lights. With large numbers of lights, many solutions with very similar cost are found, these all involve having the lights scattered around an even larger ring.

Figure 5 shows the range of illumination obtained for each N. Each lighting configuration is compared to a circle of lights with a diameter matching the outer width of the FOV, and a simple grid of lights. Grids of lights, or a single light source, perform very poorly in comparison with optimised layouts or circular layouts. For small N, a circular layout provides reasonable performance. When light directions are optimised, a considerably more even field can be obtained, with a range of illumination levels of just 20%

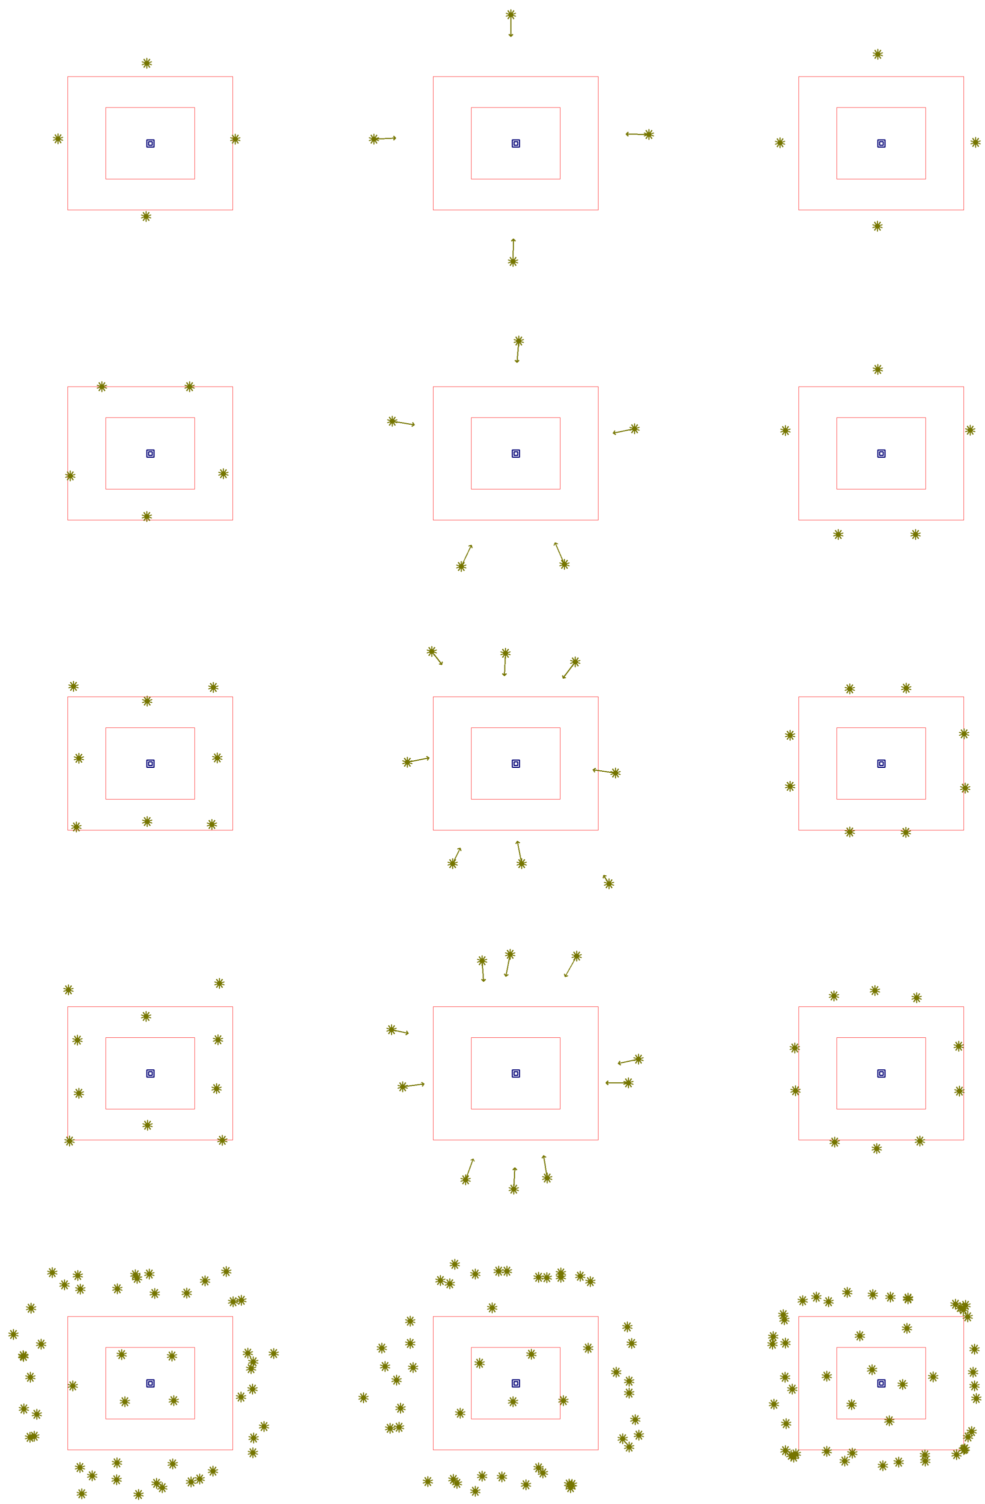

Figure 4. Lighting layouts for the vine imaging robot for $N = 4, 5, 8, 10, 50$. View from behind the front-pointing camera (small double-square), with the FOV at 35cm and 65cm marked with rectangles. Left: position front-pointing lights to minimise relative range; Middle: optimise positions and directions; Right: minimise range with occlusion.

for five lights, although this comes at the expense of efficiency, as the lights are positioned further from the region of interest (the total illumination is 39% lower than for the fixed-direction case). A measure of efficiency could be incorporated into the cost function if required.

The results of this simulation were used to position five forward-pointing light sources around the robot's central camera, as shown in Figure 2. A vine was imaged using an earlier four light configuration, the new five light configuration, and using the ambient light in a well lit room. The background was removed (by filtering on hue), and pixel intensities were examined. Figure 6 shows that the distribution of intensities on the vine is tighter with the new lighting configuration than with either of the other configurations—the standard deviation of intensities is 38% lower than for the old lighting configuration for this example. With the new optimised lighting configuration, more of the intensity variation seen in the image is due to vine colour, rather than variation in lighting.

5.1.1 Minimising lighting variation due to occlusion

When imaging 3D scenes, some lights may be occluded, increasing the range of illumination levels. While all lights could in theory be occluded, this is unlikely on the vine imaging robot due to the sparse structure of the vines. We have investigated the effects of occlusion on lighting design by considering the range of illumination levels when the brightest light at any point may be occluded. This is implemented by modifying $\mathbf{R}$ in Equation 10 to subtract the contributions of the brightest light at each point. The optimisation is run for various values of N again: similar ring-like configurations are found (Figure 4).

5.2. Illumination in a breast cancer screening system

We also used our system to design a new lighting configuration for the Digital Image Elasto-Tomography (DIET) breast cancer screening system [8]. The DIET system uses a ring of five inward-pointing cameras to image a breast as it is mechanically actuated. A 3D model is constructed, then observed surface motion is used to reconstruct the breast's internal stiffness, hence identifying tumours. Currently, lighting is provided by a ring light on each camera, however the light levels on the breast vary significantly. For accurate tracking of skin surface motion, even lighting is desirable. Our optimisation scheme is used to design new lighting configurations, which are compared with the existing configuration (Figure 7). Light positions are constrained to the hemisphere containing the cameras. The approximate position of the breast surface is known, and is incorporated into the cost function as per Equation 3. Lambertian reflectance is assumed again, which is a good approximation for skin when incidence angles are low citeIgarashi-etal-2007. With multiple light sources, illumination is provided predominately at low incidence angles (by the closest lights), and specular effects which are present for some skin are ignored.

Positioning four lights at the outer edges of the unit, and one overhead, provides a more even lighting field than with the original configuration, with a 22% range of lighting intensities on the breast, compared with a 47% range for the original. Increasing the number of lights to 10 reduces the range to 11%. The new design will be incorporated into a prototype which is currently under construction.

5.3. Illuminating a planar scene

Another common situation is where a camera images objects on a plane, for example on the ground under a vehicle, or on a conveyor belt. If the FOV is large, then illuminating the scene with a single light source will not generally give an even light field. Figure 8 shows lighting configurations found for a camera with a 90° square FOV, imaging objects on a plane. Again, ring-like lighting configurations are found for small N, but for larger N, some lights are also positioned near to the camera. A single wide angle LED light source leads to 89% variation in illumination across the scene, however optimising the positions of four similar lights results in a range of 31%, ten result in a range of 12% and 25 light sources reduces that to just 2%.

6. Conclusions

Robots and machine vision systems are increasingly deployed in environments where illumination is challenging, including in confined spaces, and where target objects move in three dimensions; in these situations careful design of lighting systems is important to ensure even illumination. This paper has described a method for designing such lighting systems, by constructing a simplified model of the illumination of objects, then optimising positions and directions of multiple point light sources to minimise the range of lighting intensities across the scene. A lighting system designed for a tractor-mounted robot is shown to reduces the variation in the appearance of objects due to uneven lighting. When illuminating a planar scene imaged with a wide FOV lens, an optimised lighting configuration of 25 light sources provides a 98% lower range of illumination levels, compared with a single light source. A lighting system for a medical imaging device is designed, which will reduce the variation in lighting by approximately 53%.

References

[1] Yamada, K., Nakano, T., Yamamoto, S.: Robot vision robust to changes of lighting conditions using a wide-

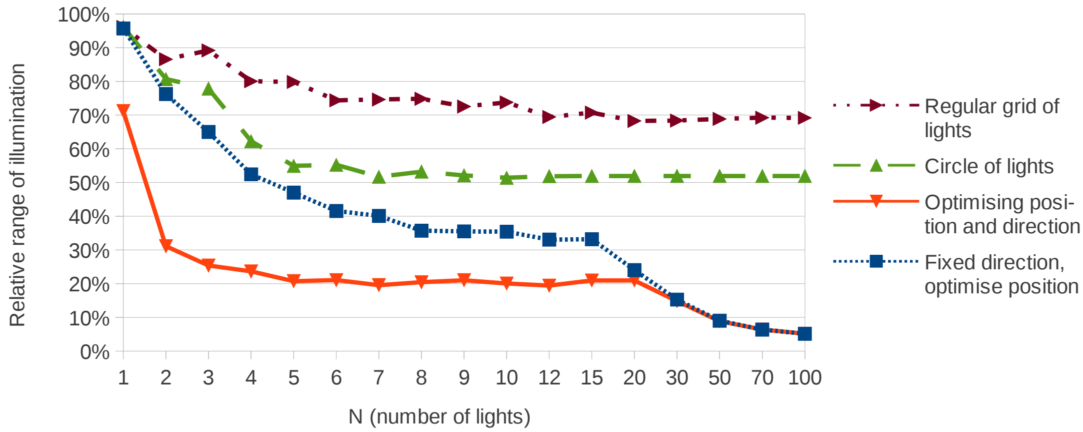

Figure 5. Relative range of illumination obtained for the vine imaging robot, with varying N. Lower values are best.

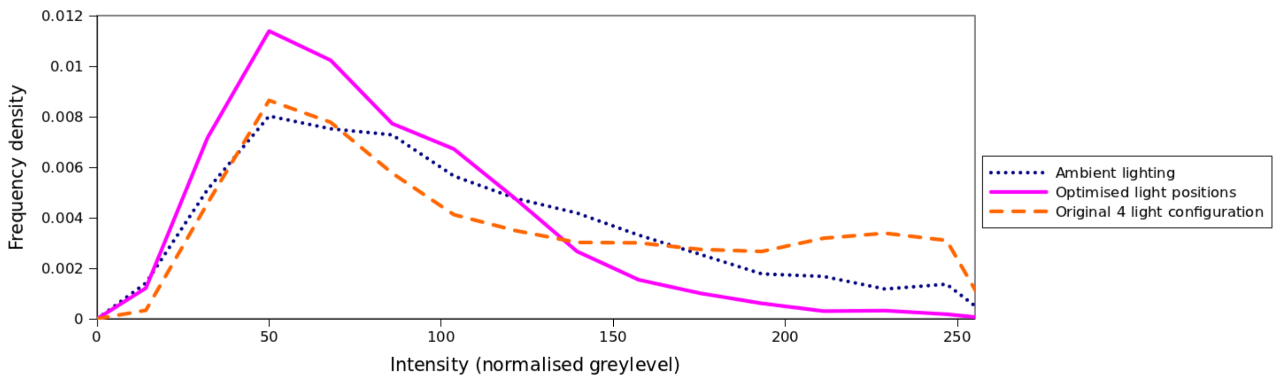

Figure 6. Variation in image intensity on vine pixels, with old and new lighting configurations, and under ambient light.

dynamic-range visual sensor. Electrical engineering in Japan **120** (1997) 34–40 1

[2] Chen, Y., Chao, K., Kim, M.: Machine vision technology for agricultural applications. Computers and Electronics in Agriculture **36** (2002) 173–191 1

[3] Klančar, G., Kristan, M., Karba, R.: Wide-angle camera distortions and non-uniform illumination in mobile robot tracking. Robotics and Autonomous Systems **46** (2004) 125–133 1

[4] Malamas, E., Petrakis, E., Zervakis, M., Petit, L., Legat, J.: A survey on industrial vision systems, applications and tools. Image and vision computing **21** (2003) 1

[5] Li, Z., Li, L., Zhang, Y.: Future research in vehicle vision systems. Intelligent Systems **24** (2009) 62-65 1

[6] Dudek, G., Jenkin, M.: Computational principles of mobile robotics. Cambridge University Press (2010) 1

[7] Botterill, T., Mills, S., Green, R.: Design and calibration of a hybrid computer vision and structured light 3D imaging system. In: Proc. Int. Conf. Automation, Robotics, and Applications (ICARA). (2011) 1, 4

[8] Peters, A., Milsant, A., Rouzé, J., Ray, L., Chase, J., Houten, E.: Digital image-based elasto-tomography: proof of concept studies for surface based mechanical property reconstruction. JSME International Journal Series C. (2004) 1, 6

[9] Cowan, C., Bergman, A.: Determining the camera and light source location for a visual task. In: Proc. ICARA (1989) 509–514 2

[10] Yi, S., Haralick, R., Shapiro, L.: Automatic sensor and light source positioning for machine vision. In: Proc. CVPR (1990) 55–59 2

[11] Triggs, B., Laugier, C.: Automatic task planning for robot vision. In: International Symposium on Robotics Research. (1995) 2

[12] Rife, J., Rock, S.: A low energy sensor for AUV-based jellyfish tracking. In: Proc. Int. Symposium Unmanned Untethered Submersible Technology (2001) 2

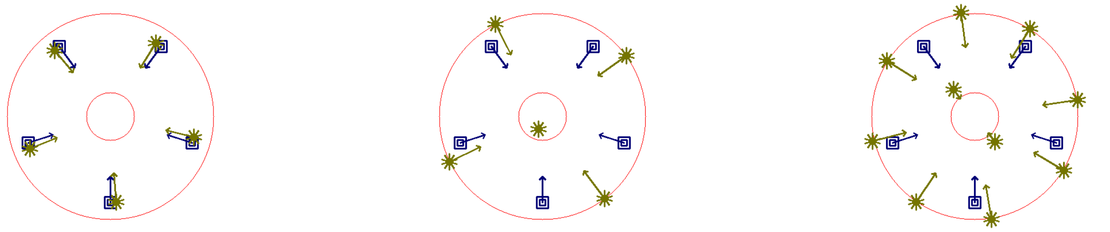

Figure 7. Lighting layouts for the DIET breast cancer screening system; orthographic view from above. The lights and cameras lie on a hemisphere, which is outlined with the large circle (30cm radius). The breast position is marked by the smaller circle. The original lighting configuration is shown on the left; the optimised configuration of five lights is shown in the middle, and an optimised configuration of 10 lights is shown on the right.

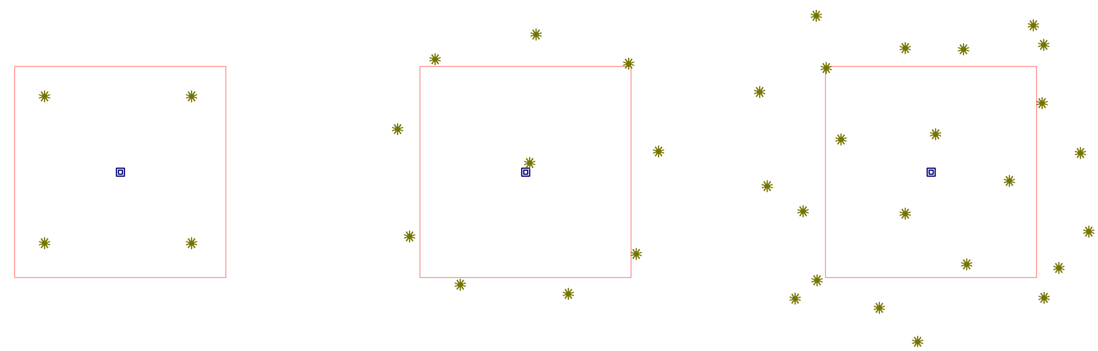

Figure 8. Lighting layouts for a square planar scene, 90° FOV. $N = 4, 10, 25$.

[13] Marchand, E.: Control camera and light source positions using image gradient information. In: Proc. ICARA (2007) 417–422 2

[14] Chen, W.G., Uang, C.M.: Better reading light system with light-emitting diodes using optimized fresnel lens. Optical Engineering **45** (2006) 1–7 2

[15] Seoul Semiconductor: Technical Data Sheet W724C0 (P7 LED). (Online) http://www.seoulsemicon.com/en/product/prd/discontinued.asp, retrieved Feb 2012. 3, 4

[16] Osram: LCW W5SM Series Datasheet. (Online) retrieved Feb 2012. 3

[17] Smart Vision Lights: Distribution datasheet for SC75 Brick light with Wide Lenses. (Online) http://www.smartvisionlights.com/wp-content/uploads/pdf/SC75_Distribution_Wide_Data6.pdf, retrieved Feb 2012. 3

[18] Advanced Illumination: (Rl1660 apposelite LED ring light data sheet) http://www.advill.com/uploads/products/RL1660+.pdf, retrieved Feb 2012. 3

[19] Foley, J.: Computer graphics: principles and practice. Addison-Wesley (1996) 3

[20] Press, W., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P.: Numerical recipes in C. Cambridge University Press (2001) 4