

Reconstruction of refractive index maps using photogrammetry

A. Miller ^a, A.J. Mulholland^b, K.M.M. Tant^a, S.G. Pierce ^c, B. Hughes^d and A.B. Forbes^d

^a Department of Mathematics and Statistics, University of Strathclyde, Glasgow, UK; ^b Department of Engineering Mathematics, University of Bristol, Bristol, UK; ^c Department of Electronic and Electrical Engineering, University of Strathclyde, Glasgow, UK; ^d National Physical Laboratory, Teddington, London, UK

ABSTRACT

Large volume metrology is a key enabler of autonomous precision manufacturing. For component positioning, the optical-based metrology technique of photogrammetry could be used more widely if its accuracy was improved. These positional measurements are subject to uncertainties which can be greater than manufacturing tolerances. One source of uncertainty is due to thermal gradients, which cause the refraction of the light rays in large-scale industrial environments. This paper uses light-based sensor data to reconstruct a heterogeneous spatial map of the refractive index in air. We use this reconstructed refractive index map to discount the refractive effects and thereby reduce the uncertainty of this positioning problem. This new inverse problem employs Voronoi tessellations to spatially parameterize the refractive index map, the Fast Marching Method to solve the forward problem of calculating the light rays through this medium, and a Bayesian approach in the inversion. Using simulated data, this methodology leads to positioning improvements of up to 37%.

ARTICLE HISTORY

Received 6 November 2019
Accepted 6 June 2021

KEYWORDS

Photogrammetry; metrology; refractive index; reversible jump Markov Chain Monte Carlo; Fast Marching Method; Voronoi tessellations

2010 MATHEMATICS SUBJECT CLASSIFICATIONS

74J25; 65C05; 68T40

1. Introduction

Almost half of UK exports are from manufacturing [1] and so it is crucial that the UK continues to manufacture high-value goods quickly, to a high standard and at a low cost to sustain this income. Manufacturing is increasingly using robots for the tasks of positioning and assembling of parts [2,3]. Industries are interested in autonomous manufacturing as it can potentially reduce costs and increase productivity [4]. A key enabler of autonomous manufacturing is large volume metrology. Robots require sensors to allow them to accurately position themselves and the component they are working on, in a similar manner to a local GPS system. Large volume or large-scale metrology typically involves using light rays that gather data on the distance and/or the direction between measurement apparatus and an object being measured [5].

CONTACT A. Miller  ainsley.miller@strath.ac.uk

 Supplemental data for this article can be accessed here. <https://doi.org/10.1080/17415977.2021.1946533>

© 2021 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group
This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way.

There have been many advancements in indoor positioning systems (IPS) and some of these systems use light rays; these systems are called optical-based metrology systems [6]. Optical-based metrology systems might be distance based, angle based, or focused on surface form measurement [7]. This paper will consider the optical-based metrology system of photogrammetry. Photogrammetry systems generate a set of photographs taken from different spatial positions to produce positional measurements of objects of interest [8]. It is possible to obtain robust and dependable data about the surface, shape, size and position of an object using photogrammetry [7]. By using two or more images of the object of interest from different angles, it is possible to find the three-dimensional coordinates using optical triangulation [5,8]. However, due to thermal fluctuations in the air, the light may refract and this can lead to significant positioning errors in the triangulation calculation [9].

A photogrammetry system's ability to measure accurately depends on many factors including the ability to take into account atmospheric effects such as temperature, pressure, refractive index and humidity [10,11]. If we extend the measured volume to that of an industrial setting, we must address light refraction due to temperature variations within the volume; a factory environment does not have a constant temperature as machines and operators emit heat. This therefore means that the first-order approximation which says that rays of light move in straight lines are not applicable [5]. This bending of light rays introduces errors in the received measurements (that is, the angle measurements of the photogrammetry system) and consequently, as the volume measured increases, these errors grow and can severely affect the accuracy of the co-ordinate positioning [10]. Accounting for the refractive index effects is a key step to improving the accuracy and quantifying uncertainties.

The literature demonstrates that sub-millimetre errors in the position of a target cause significant problems [12–15]. The 'Large volume unified metrology for industry and novel applications and research (LUMINAR)' project is a recently completed international collaboration that focused on tackling refraction effects in industrial settings [13]. In aircraft manufacturing, the errors due to positional measurements can be vital in ensuring aircraft safety and quality and indeed for full aircraft measurements an error of 0.26 mm can be significant [14,16]. Under typical working conditions, the VICON T160 photogrammetry system can have millimetre scale average positional errors [15]. Simulated light rays propagating through a laboratory-sized spatial domain (3 m × 4 m) with a Gaussian temperature profile with a maximum temperature of 370 K leads to a positional error of approximately 0.25 μm [17].

In this paper, we propose that this long-standing problem in large volume metrology is tackled via the reconstruction of the refractive index map of a 2D plane from simulated photogrammetry sensor data. There is much work in the literature already where statistical methods have been used to create a spatial map of some property of interest. In disciplines such as environmental science and neuroscience, a nonparametric regression method, namely multivariate adaptive regression splines (MARS), has allowed researchers to create spatial maps of, for example, forest fires and neurological events [18–21] and it has been shown that MARS is able to cope with large and complex data sets [19]. Furthermore, this technique is flexible and adaptive and as a result, it can be applied to linear and non-linear problems [19]. There has been some success with this technique in reducing the positional uncertainties of the data points by reducing the discretisation of the geometry

[18]. However, due to the differences in the application, the discretization in [18] is still much larger than in the metrological setting central to this paper.

As far as the authors are aware, this paper presents for the first time, a statistical inverse problem approach (so that is one that involves the use of a forward mathematical model of the underlying physics) to reduce the uncertainty and error associated with this use of photogrammetry. To tackle this problem empirically (without a forward mathematical model) is very difficult due to the complex interaction between the temperature (refractive index) maps and the light rays. The aforementioned LUMINAR project was a large-scale collaborative project focusing on eliminating the refractive effects in photogrammetry metrology for high-precision manufacturing using empirical means but the progress made was limited. Photogrammetry is an angular metrology method and the positional errors that arise are extremely sensitive to spatial variations in the refractive index. Given the many degrees of freedom in this problem, it is not surprising that previous empirical approaches have struggled to make meaningful inroads. Using a mathematical model within an iterative framework does of course require computational resource and so there is a timeliness to this methodology that takes advantage of the current availability of affordable and fast computational resource.

This paper presents a Bayesian inverse problem approach, namely the reversible-jump Markov chain Monte Carlo (rj-MCMC) method, as a way of inverting photogrammetry sensor data to obtain information on the refractive index map the light rays are propagating through. No a priori assumptions are made on the initial refractive index map in this work and the initial model randomly assigns refractive indices drawn from an uninformative uniform distribution to the irregularly partitioned domain (a Voronoi tessellation) [22]. The only prior information exploited is the size of the spatial domain (in two dimensions) and the location of the cameras and reflectors. The inversion method employed in this paper has proved useful in applications of other sensing modalities in other application areas. For example, it has been used in ultrasonic tomography for improved flaw detection in non-destructive testing [22,23]. The data used in that setting is time of flight from the ultrasonic sensor (as opposed to the angular direction to a target in photogrammetry) and the material map is locally anisotropic (whereas the refractive index is locally isotropic in the photogrammetry setting). The approach has also been used by geoscientists to obtain a spatial map of the wave speed within the Earth's crust [24–27].

2. Experimental positioning errors due to temperature variations in the volume

An experiment was carried out to determine the magnitude of the positioning errors of the Vicon T160 photogrammetry system when positioning a cluster of retroreflectors in a spatially heterogeneous refractive index volume created by applying a localized heat source to the volume. In order to do so it is also necessary to measure the temperature profile in the working geometry.

Two Vicon T160 cameras (16 MP with standard Vicon 18 mm focal length lens) (Vicon Motion Systems, LA, USA) were mounted on tripods and these tripods were positioned approximately 3.5 m away from a group of reflectors as shown in Figure 1. There were five retroreflectors all of which were steel and spherical with a 3.81 cm diameter. We mounted the reflectors on to a plate attached to the top of a tripod that was approximately 1 m high.

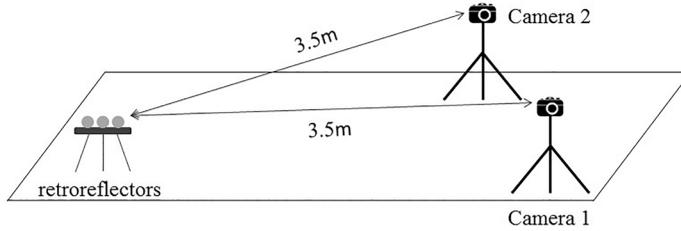


Figure 1. Baseline set up of the experiment which has two Vicon T160 cameras (16 MP with standard Vicon 18 mm focal length lens) tracking a cluster of stationary reflectors over time. Each camera is positioned approximately 3.5 m from the reflectors.

The cameras will be denoted as c_1 and c_2 and the photogrammetry system records the estimated position (\hat{s}) of the cluster of the reflectors as the centroid of them. Photogrammetry systems use the method of triangulation to estimate the position of objects of interest [9,28] and so in a two camera setting the intersection point of the light ray paths emitted from the cameras yields the estimated position \hat{s} . For the experiments below the two cameras and the cluster of retroreflectors remained in a fixed spatial position.

The cameras were recalibrated each time the experiments were carried out and, during each calibration of the cameras, 2000 frames were captured using the active wand; the active wand is a T-shaped calibration artefact that has 5 LED sensors. The active artefact minimizes the number of data outliers captured as the number of data points that have an obstructed view is reduced as the LEDs are actively lit up [15]. In fact, using an active calibration leads to an improvement in the accuracy of 13% compared with using a passive calibration artefact [15]. We used the active wand to set the geometry's origin, and to do so we placed the active wand on the ground where both cameras could see it. The origin remained in a fixed position throughout all experimentation.

The first experiment was carried out at room temperature and the estimated position (\hat{s}) of the retroreflector centroid was monitored over approximately 8 min and 20 s (time step: 0.02 seconds, number of time steps: 25,000). The estimated position of the reflectors is output as a series of (x, y, z) co-ordinates. In addition to recording the co-ordinates, the temperature of the volume was also recorded throughout using a PCB with eight temperature sensors. The temperature sensors were MCP9808 Precision I2C, which have a maximum accuracy of $\pm 0.5^\circ \text{C}$ and a typical accuracy of $\pm 0.25^\circ \text{C}$ [29]. On average, the ambient temperature of the laboratory was between 21°C and 22°C and it took approximately 2 s for the temperature sensor to collect one set of readings. From the experimental measurements, the minimum temperature (T_{\min}) was determined to be 21°C as this was the modal temperature of the laboratory with no heat source applied.

After collecting these baseline readings, we introduced a localized heat source to the volume. The heat source used in this experiment was a 1.6 kW halogen heater. We switched the heater on until the laboratory reached a steady state. We placed the heat source on the ground between camera c_1 and the cluster of retroreflectors. The heat source was elevated approximately 70 cm as shown in Figure 2. The mean temperature of the volume when there was a localized heat source applied to the volume was 51°C and the maximum temperature was 55°C . The temperature occurring from adding the heat source is consistent with the normal temperature fluctuation range in industrial production. The authors of

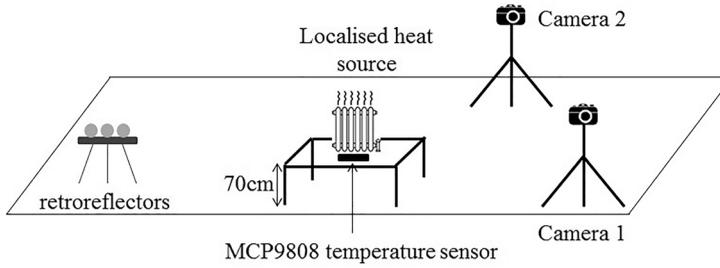


Figure 2. Experimental Set up: The 2 Vicon cameras and the retroreflectors remain in the same position throughout the experiment. A localized heat source is placed 70 cm above the ground and is directly between camera c_1 and the retroreflectors. There is also an MCP9808 temperature sensor placed in front of the heat source measuring the temperature of the volume.

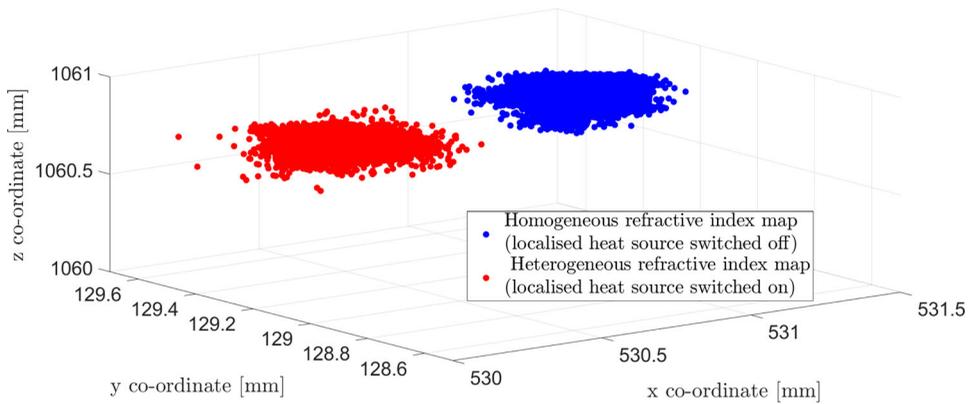


Figure 3. Point cloud showing the (x, y, z) co-ordinates of the reflectors before the heat was applied to the volume (right cluster) and after (left cluster).

[30] reported that in a steel plant the mean radiant temperature is 58.2°C with a standard deviation of 1.24°C .

In order to analyse the impact that the heat source has on the estimated positions of the cluster of reflectors the data was taken into MATLAB for analysis. We repeated the experiment five times and calculated the mean point at each of the 25,000 iterations. From this, a point cloud of the estimated position of the cluster of reflectors was plotted for the control case (localised heat source switched off) and the heterogeneous (localised heat source switched on) in Figure 3.

The mean estimated position of the cluster of reflectors (denoted by s^*) for the homogeneous case is taken to be the true position of the reflectors. The mean estimated position of the cluster of reflectors for the heterogeneous case (denoted by \hat{s}) is then also calculated. Then the error in the estimated position is the Euclidean distance between these two points, $\varepsilon_c = \|s^* - \hat{s}\|_2$ which was found to be $\varepsilon_c = 679\ \mu\text{m}$. This set of experiments show that the Vicon T160 photogrammetry system is sensitive enough to detect significant positional errors in the retroreflectors caused by heat fluctuations in the working volume. Importantly, the size of these errors is sub-millimetre over a relatively small distance (3.5 m) and such positional errors could be significant in high-precision manufacturing

scenarios. The next section will introduce a methodology to recover the spatial refractive index map and this map will be used to reduce this error.

3. Generation of simulated data using COMSOL

The effect of the positional errors obtained by simply using Snell's Law is far smaller than those observed in practice and to better account for these measurement errors, one must include transverse gradients [31]. COMSOL includes the effects of temperature variations that are transverse to the direction of the ray path. To test its accuracy in this regard, we replicated the example presented in Appendix E of the American Society of Mechanical Engineers (ASME) Standard [31] and it was found that COMSOL obtained good agreement.

Using simulated photogrammetry data has the advantage that complete ground-truth is known about the camera and reflector positions and the temperature variations in the volume. In addition, it allows us to isolate the error due to these temperature fluctuations alone. However, it is prudent to use a different model to generate this simulated data than that used in the forward solver in the inverse implementation. To that end, the ray tracing package in COMSOL Multiphysics was used [32].

One limitation of COMSOL is that it is not possible to define the start and the end point of a ray. Since COMSOL solves an initial value problem for an Eikonal equation [33, Equation (2)], we prescribe the start point, the initial directional vector and the length of the ray. Therefore, we devised the following method for generating camera and reflector positions. This is unusual as in an industrial setting the camera and reflector positions would be known a priori. Starting from a known refractive index map m^* and a spatial position s^* , C rays were released each with a distinct direction but the same length. We used the end point of each ray to define the camera positions c_k . We then recorded the angles $\theta_{k,s}^*$ that each ray from s^* enters the camera c_k . One could use a shooting-matching method to iterate to get each ray to join a prescribed camera to a prescribed reflector but that is a major task for a large number of camera-reflector pairings and adds very little to the framework being presented here and to the error analysis upon which the method is being judged. This type of calculation could of course be performed in practice but is beyond the scope of this paper as the purpose of this paper is to show the potential of this approach. We used COMSOL to create a series of test refractive index maps. The typical domain size of each numerical experiment was the industrial scale and of the order of 40 m in the lateral direction and 60 m in the axial direction and attention was restricted to the 2D horizontal plane case.

4. Methodology for recovering the refractive index map

4.1. Parameterisation of the refractive index map

In this work, we use Voronoi tessellations to parameterize the refractive index map and to produce adaptive and irregular partitionings which facilitate low dimensional (low degrees of freedom) representations of the spatial domain whilst affecting large regions with a single parameter perturbation. The Voronoi tessellation begins with a set of $P^{(j)}$ randomly chosen seed locations denoted by $X^{(j)} = (x_1^{(j)}, \dots, x_p^{(j)})$. Note that we use a transdimensional

approach and so $P^{(j)}$ is not fixed and is thus a function of iteration number j . We introduce the notion of a refractive Voronoi tessellation, where each Voronoi cell is assigned a refractive index $\eta_p^{(j)}$; we denote the set of refractive indices by $H^{(j)} = (\eta_1^{(j)}, \dots, \eta_p^{(j)})$. In the remainder of the paper, we will denote the model $m^{(j)} = (X^{(j)}, H^{(j)}; P^{(j)})$. In two dimensions, the map $m^{(j)}$ has $3P^{(j)} + 1$ degrees of freedom, which arise from the two-dimensional seed coordinates and the refractive index of the $P^{(j)}$ cells, and the number of cells itself.

4.2. Modelling the light ray propagation in a heterogeneous refractive index map

For each camera (c_k , where $k = 1, \dots, C$) and retroreflector (r_i , where $i = 1, \dots, R$) the ray tracing algorithm (forward model) will generate an estimated angle denoted by $\theta_{k,i}$ with the measured (simulated) angle denoted by $\theta_{k,i}^*$; all angles are measured anticlockwise relative to a fixed set of axes.

The Fast Marching Method (FMM) [34,35] is used to calculate the ray paths through the Voronoi tessellation. Using COMSOL for the forward model would be computationally prohibitive, would be difficult to automate and embed within the inversion algorithm, and would lead to the ‘inverse crime’ of using the same forward model as that used to generate the synthetic data. The FMM is a fast means of calculating time fields and this enables the inversion algorithm to make many forward model calculations and have many model realisations.

Using Fermat’s principle [36], the time taken for a ray of light following path $r(Q)$, between camera and reflector, where Q is the length of the ray measured from a fixed point, is defined by

$$t = \frac{1}{c_\infty} \int_{c_k}^{r_i} \eta(r(Q)) \, dQ \quad (1)$$

where c_∞ is the speed of light in a vacuum, $\eta(r(Q))$ is the refractive index along the path which takes the shortest time between c_k and r_i , and dQ is a small displacement along the ray. Following this, the optical path length S between the camera and the retroreflector is

$$S = \int_{c_k}^{r_i} \eta(r(Q)) \, dQ \quad (2)$$

where $S = c_\infty t$. Using the calculus of variations

$$\Delta S = \Delta \int_{c_k}^{r_i} \eta(r(Q)) \, dQ = 0 \quad (3)$$

where Δ is the change in optical path length and this means that the path $r(Q)$ satisfies Fermat’s principle. The Fast Marching Method (FMM) solves a boundary value problem of an associated Eikonal equation to calculate the travel time field [33]. Once the travel time field is calculated, the ray that takes the shortest time to travel between the camera c_k and the reflector r_i can be found using Equation (3).

4.3. Misfit function between the modelled and measured data

The misfit function at iteration j is given by

$$\gamma^{(j)}(\eta_p^{(j)}, \underline{x}_p^{(j)}) = \frac{\sum_{(k,i) \in \Phi} \left(\left| \theta_{k,i}^* - \theta_{k,i}^{(j)}(\eta_p^{(j)}, \underline{x}_p^{(j)}) \right| \right)^2}{N_{\text{rays}}}, \quad (4)$$

where $\gamma^{(j)}$ is measured in degrees, $\theta_{k,i}^*$ is the angle that the light ray enters the camera c_k from the retroreflector r_i and the single vertical lines denote the absolute value of the quantity therein. The set of camera-reflector pairs is denoted by Φ , where $\Phi : \mathcal{C} \rightarrow \mathcal{R}$ maps the set of cameras $\mathcal{C} = \{c_1, c_2, \dots, c_C\}$ to the set of reflectors $\mathcal{R} = \{r_1, r_2, \dots, r_R\}$ and $N_{\text{rays}} = |\Phi|$. The inclusion of the number of rays in the denominator in Equation (4) was to allow for any variations in the number of camera-reflector pairs as this could alter as an object moves through a volume and lines of sight alter. However, in this paper, this was held constant on grounds of simplicity and therefore this parameter has not affected the method and the results that follow.

Given the discretized ray path the angle $\theta_{k,i}^{(j)}$ is calculated using, what we will refer to as, the δ method. A schematic of the discretized ray path is found in Figure 4. In this schematic, the small squares show the Voronoi cell seeds \underline{x}_p . The camera c_1 and reflector r_1 are both shown by a larger square. The dashed line shows the straight ray path between the camera c_1 and the reflector r_1 . The thick solid line is the discretized ray path of the light ray through this Voronoi tessellation and the $b_j, j = 1, \dots, B$ are the interface points at adjacent Voronoi cells. We then find the angle $m_j, j = 1, \dots, B$ between the camera c_1 and the points $b_j, j = 1, \dots, B$. We define $B_p = [pB]$ as the first p percent of points in the ray where $[x]$ is the integer closest to x . The mean of these angles in the first p percent of the interface points is

$$\delta_p = \frac{(\sum_{j=1}^{B_p} m_j)}{B_p}. \quad (5)$$

To determine how many percentage point we should use a series of refractive index maps based on Voronoi tessellations were imported into COMSOL and synthetic data $\theta_{k,i}^*$ was generated. The agreement between the known angle $\theta_{k,i}^*$ and the estimated angle at model iteration zero, $\theta_{k,i}^{(0)}$, from the FMM applied to these known refractive index material maps was examined. The angle $\theta_{k,i}^{(0)}$ was calculated from the gradient between the starting point of the ray (at the camera c_k) and each subsequent point of the discretized ray in the first δ % of the ray points. The mean of these gradients was found for each ray and this was converted into an angle (measured in degrees), where $-180 \leq \theta_{k,i}^{(0)} \leq 180$. The results of this are shown in Figure 5 for the cases where $\delta = 10, 30, 50, 70, 90$ and 100 % have been plotted versus the known angle $\theta_{k,i}^*$. The mean absolute error was also found for the cases where $\delta = 10, 30, 50, 70, 90$ and 100 % and the resulting values were $29.58^\circ, 25.09^\circ, 19.86^\circ, 17.22^\circ, 5.49^\circ$ and 3.08° , respectively. Based on these investigations the most accurate algorithm was to use all the points of the ray to calculate $\theta_{k,i}^{(0)}$.

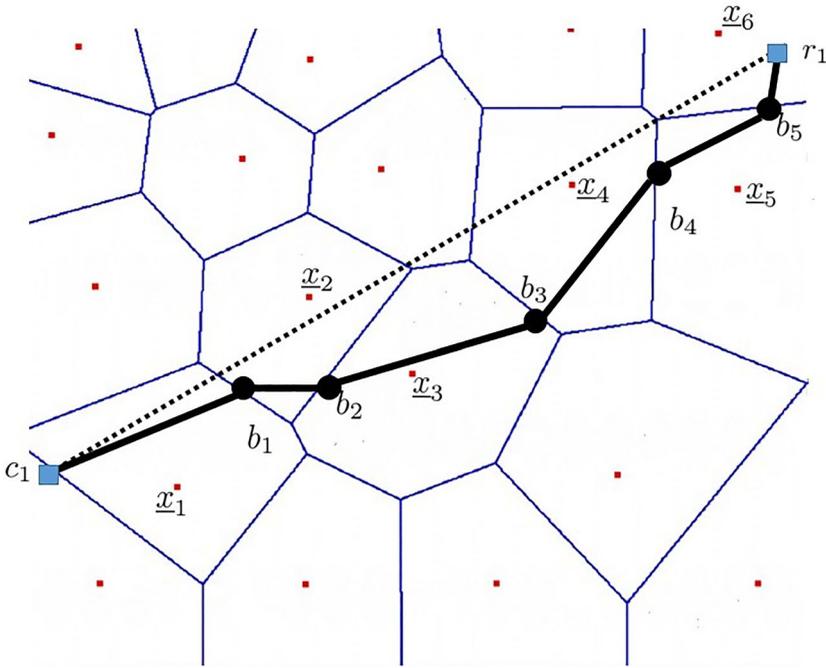


Figure 4. A schematic showing the discretized ray path of a light ray travelling through a Voronoi tessellation. The small squares show the Voronoi cell seeds, x_p , and the camera c_1 and reflector r_1 are both depicted by larger squares. The dashed line shows the straight ray path between the camera c_1 and the reflector r_1 . The thick solid line is the discretized ray path of the light ray through this Voronoi tessellation, and the $b_j, j = 1, \dots, B$ are where the ray meets the interface points between adjacent Voronoi cells.

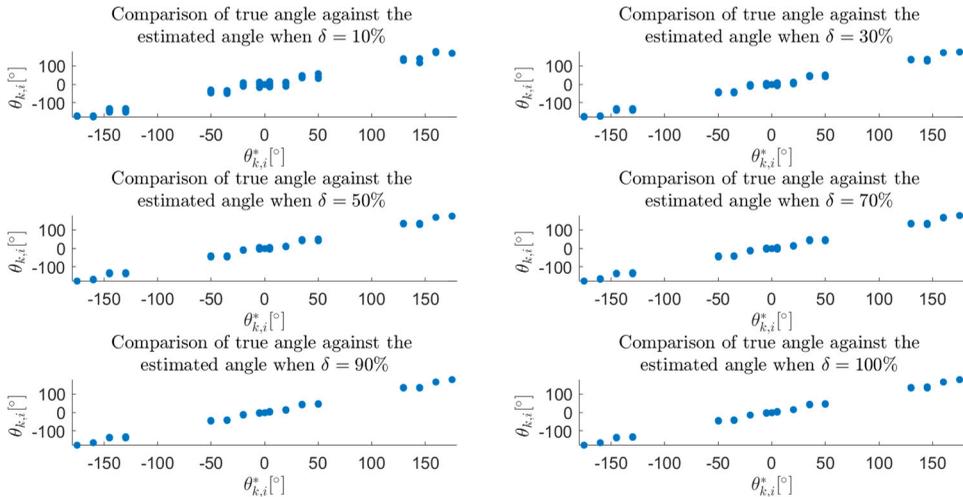


Figure 5. Scatter-plots of $\theta_{k,i}^{(0)}$ [°] against $\theta_{k,i}^*$ [°] when $\delta = 10\%, 30\%, 50\%, 70\%, 90\%$ and 100% , where $\theta_{k,i}^*$ is the true angle from the camera (from synthetic data) to the receiver and $\theta_{k,i}^{(0)}$ is the estimated angle from the FMM. The correlation coefficient is given for each percentage of the FMM ray used, δ , and the red line denotes $\theta_{k,i}^{(0)} = \theta_{k,i}^*$.

4.4. Recovery of the refractive index map via a Bayesian approach

To sample the model space and construct an approximation of the posterior distribution, we employ the reversible jump Markov Chain Monte Carlo (rj-MCMC) [37] method for the optimization step. The rj-MCMC is a stochastic iterative process used to create samples from the posterior distribution which is the unknown probability distribution describing the likelihood of each Voronoi tessellation being the reconstructed refractive index map. The approach relies on Bayes' Theorem [38] which leads to

$$p\left(m^{(j)} \mid \theta_{k,i}^*\right) \propto p\left(\theta_{k,i}^* \mid m^{(j)}\right) p\left(m^{(j)}\right), \quad (6)$$

where $p(m^{(j)})$ is a probability density function representing the prior knowledge of the model $m^{(j)}$. The likelihood of observing the data $\theta_{k,i}^*$ given a particular model $m^{(j)}$ is $p(\theta_{k,i}^* \mid m^{(j)})$ and this is proportional to

$$\exp\left(\frac{-\gamma^{(j)}\left(\eta_p^{(j)}, \mathbf{x}_p^{(j)}\right)}{\left(\zeta^{(j)}\right)^2}\right), \quad (7)$$

where $\zeta^{(j)}$ is the noise parameter and $\gamma^{(j)}$ (measured in degrees) is the misfit function defined in Equation (4). The posterior distribution which describes the probability of $m^{(j)}$ being the refractive index map m^* given the data $\theta_{k,i}^*$ is denoted by $p(m^{(j)} \mid \theta_{k,i}^*)$. The likelihood function is used to calculate $p(\theta_{k,i}^* \mid m^{(j)})$.

The initial prior probability density function $p(m^{(0)})$ is derived from any information which is known a priori such as the locations of the cameras and the reflectors, the dimensions of the space and the likely range of refractive indices. The algorithm consists of the steps outlined in Figure 6. We start by constructing an initial refractive index map model $m^{(0)}$ by randomly selecting the parameters for the Voronoi tessellation and the noise parameter from a uniform distribution [39]. Following this, at each iteration, the Voronoi tessellation model of the refractive index map $m^{(j)}$ is perturbed to make a new model $m^{(j+1)}$. There are five ways to perturb the model: cell birth, cell death, seed move, cell refractive index change or system noise change [24]. The cell birth step is when an additional Voronoi cell seed is added to the model. The cell death step is the opposite of the birth step and this is where a Voronoi cell seed is removed from the current model. The cell move step is where a randomly selected Voronoi seed $\mathbf{x}_p^{(j)}$ is moved. The refractive index change step is where a refractive index, $\eta_p^{(j)}$, is changed and lastly the random noise change is when the value of $\zeta^{(j)}$ is changed in the likelihood function [Equation (7)]. A change in $\zeta^{(j)}$ is the only perturbation which does not require the rays to be re-traced through the new model $m^{(j)}$. To determine which type of perturbation of the map will be conducted at each iteration the algorithm randomly generates, with equal probability either a zero or a one (Step 1). If the algorithm generates a zero then the perturbation will alter the map geometry. That is, with equal probability, either a perturbation involving a cell birth, a cell death or a cell move. If the algorithm generates a one then, with equal probability, the perturbation will be a refractive index change or a system noise change. In the case of a cell birth, the position of the new cell seed is randomly selected (from a uniform distribution across the spatial domain). Then the refractive index for the new cell is drawn from

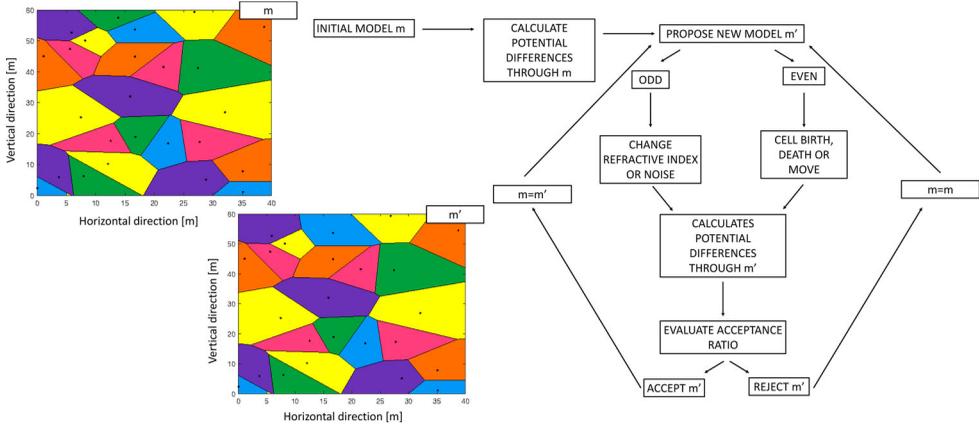


Figure 6. Workflow of the inversion method. In this example, the geometry of model m is perturbed in a death step by removing a Voronoi cell seed at (14,57) m.

a Gaussian centred on the refractive index that was assigned to that point in the domain (that is, it is determined by the cell in the previous Voronoi cell that the new seed lies in). There is a specific standard deviation prescribed on the refractive index for birth cases, which is 5×10^{-5} . For a cell death model perturbation a cell seed is randomly chosen with equal probability from within the set of current seeds and then removed from that set. For perturbations involving refractive index changes, movement of Voronoi cells and noise, the proposal distributions are Gaussian centred on the existing value. For a perturbation involving the movement of a Voronoi cell seed, a cell is chosen (with equal probability) from amongst the set of current seeds and its new location is chosen arbitrarily (so each point in the domain is as likely to be chosen). For a perturbation involving a change in refractive index, the proposal distributions are Gaussian and centred on the existing value with a standard deviation of 9×10^{-4} . The standard deviations of these Gaussian distributions are very important. If the standard deviation is too big then large jumps in parameter values can occur and the algorithm can take a long time to converge. Conversely, if the standard deviation is too small then the algorithm samples the parameter space slowly which can slow down convergence. Therefore, the standard deviation can be used to tune the algorithm so that it explores the parameter space efficiently.

The rj-MCMC is an iterative process and at each iteration the algorithm decides whether to replace the existing model $m^{(j)}$ with the new perturbed one, $m^{(j+1)}$, using the Metropolis-Hastings criterion [40,41]. The probability of acceptance depends on the ratio of the probabilities arising from the posterior distributions of $m^{(j)}$ and $m^{(j+1)}$ via

$$p(\text{accept}) = \min \left(1, \frac{p(m^{(j+1)} | \theta_{k,i}^*)}{p(m^{(j)} | \theta_{k,i}^*)} \times \frac{q(m^{(j)} | m^{(j+1)})}{q(m^{(j+1)} | m^{(j)})} \right), \quad (8)$$

where $q(m^{(j)} | m^{(j+1)})/q(m^{(j+1)} | m^{(j)})$ is the ratio of the proposal distributions. Here $q(m^{(j+1)} | m^{(j)})$ is the probability of moving to model $m^{(j+1)}$ from $m^{(j)}$ and this is called the forward step, whilst $q(m^{(j)} | m^{(j+1)})$ is the reverse step. For a non-transdimensional move (no cell birth or death), the ratio of the proposal distributions is equal to 1 [23,24]. The

Metropolis Hastings algorithm proceeds by randomly sampling the model space, sometimes accepting the perturbations and sometimes remaining in place. If a perturbation results in a model with a higher likelihood than the existing point, it is always accepted. If a perturbation results in a lower likelihood model, the probability of rejection is dependent on the relative drop in probability. Thus, the algorithm will more densely sample high likelihood regions and less frequently sample low likelihood regions, and so intuitively returns an ensemble of samples which well represents the underlying posterior distribution.

The steps of the inversion process are summarized in the workflow shown in Figure 6. We refer to the first set of iterations of the model (of the Markov chain) as the burn-in period. The burn-in period is required to allow the Markov chain to converge to the estimated posterior distribution from which samples are taken. From the literature, we desire an acceptance rate after the burn-in period of between 25% and 50% [25,42]. From Equations (7) and (8), it can be seen that $\zeta^{(j)}$ can be changed to alter the probability of accepting an inferior model. This parameter automatically adjusts as the inversion algorithm iterates and facilitates efficient exploration of the parameter space.

The main output of the inversion process is an ensemble of maps which approximate the posterior distribution on refractive indices throughout the domain. By considering moments of these distributions, we are effectively integrating over the dynamic parameterisations, and so the method is inherently self-smoothing and provides self regularized solutions driven by the data [24,25]. Another of the advantages of using a probabilistic method is that we can automatically access higher moments of the posterior distribution of refractive index maps and so one can observe where the method is struggling to find an optimal solution (high variance). In addition, we can make statements about the level of uncertainty associated with any metrology related statements.

4.5. Quantifying the measurement uncertainty

To quantify whether the recovered refractive index map can be used to obtain a better estimate of objects of interest than the homogeneous refractive index map a reflector was placed in the domain in a fixed position s^* . Let $\hat{s}_h(\theta_{k,s}^*)$ be the estimated position of this reflector when the domain is assumed to be homogeneous; that is, the rays are assumed to be straight. Let the resulting error in the position of the reflector be denoted by $\varepsilon_h = \left\| s^* - \hat{s}_h(\theta_{k,s}^*) \right\|_2$. We can estimate the position of the reflector using the recovered refractive index map (either the mean, median or maximum a-posteriori of the posterior distribution). Let $\hat{s}_\alpha(\theta_{k,s}^*, \hat{m})$ be the estimated position of the reflector. The resulting error in position using the recovered refractive index map is denoted by $\varepsilon_\alpha = \left\| s^* - \hat{s}_\alpha(\theta_{k,s}^*, \hat{m}) \right\|$. In addition, the relative improvement of the error in the position of the object of interest when compared to the homogeneous case is defined by $\varepsilon_{\text{imp}} = (\varepsilon_h - \varepsilon_\alpha) / \varepsilon_h$. Current positioning techniques used by other researchers assume that the volume has a constant refractive index and that there is no ray bending. The estimated position of the target using these techniques is therefore \hat{s}_h . In this paper, the improvement in the position compared to these other methods is therefore ε_{imp} . The goal of this paper is to carry out numerical experiments and observe whether the error is reduced using the inversion algorithm. We aim to reconstruct a refractive index map such that $\varepsilon_\alpha \ll \varepsilon_h$ and to reduce the uncertainty.

In order to calculate \hat{s}_h , a homogeneous spatial domain was created in COMSOL (the mean refractive index of the above map was chosen but any value will suffice of course). Then from each camera c_k a single straight ray with angle $\theta_{k,s}^*$ was released; the rays propagate until they reach the edge of the domain. We then input the set of discrete points of each ray into MATLAB (version 2017b) [43]. All the intersection points of these rays were obtained and then the centre of mass of the resulting point cloud was found and this produced an estimate of \hat{s}_h . We repeated this process for the recovered refractive index map, \hat{m} (typically the mean, or median, or mode of the posterior distribution) and found an estimate for the reflector position \hat{s}_α .

5. Results from simulated data

5.1. Numerical experiment A

Numerical experiment A had a refractive index map consisting of a Voronoi tessellation with one hundred Voronoi cells. We assign the refractive index in each seed via

$$T(x, y) = (T_{\max} - T_{\min}) \exp\left(-\frac{(x^2 + y^2)}{\sigma^2}\right) + T_{\min}, \quad (9)$$

where x and y are the horizontal and vertical coordinates of the Voronoi cell seed, $T(x, y)$ is the temperature at point (x, y) in the domain, σ is the standard deviation of the temperature distribution measured in metres, T_{\max} is the maximum temperature in the domain and similarly, T_{\min} is the minimum temperature. In this case, T_{\max} was 363.15 K (90 °C), T_{\min} was 295.15 K (22 °C), and σ was 20 m to create a plausible industrial setting. COMSOL requires the refractive index values in order to define the material properties. These temperatures were converted to refractive indices using the Ciddor equation [44, Equation (5)]; all other parameters in the equation were set at the National Institute of Standards and Technology (NIST) average values [45]. The Ciddor equation (which models the refractive index in air) is a popular method of calculating the refractive index of air, other methods include the Edlen equation [44,46]. To add this equation to the COMSOL simulation, the *component* option in the *Model Builder* window was accessed and a new *variable* was added in the *definitions* section. Once in this menu, we input the right-hand side of Equation (9) into the *Expression* box and named this variable T_domain . Following this, we created and defined the rays propagating within this geometry COMSOL allows the user to prescribe the start point of a ray, its initial directional vector and its length. See Section 2 of the supplementary file for full details.

The position of s^* was (23, 37) m and $C = 20$ rays were released from s^* . As a validation step we used the known refractive index map m^* in the ray tracing algorithm in COMSOL to estimate the position of the target using the aforementioned triangulation procedure and this was denoted by \hat{s}^* . The error in \hat{s}^* is negligible with $\varepsilon^* = 1.34 \times 10^{-6}$ mm. This result gives confidence in our implementation of the model in COMSOL as it shows that if we precisely recovered the known map m^* then the error would be negligible.

As described in subsection 4.5 the light rays were then traced from c_k with initial angle $\theta_{k,s}^*$ but this time the refractive index map was assumed to be homogeneous. Figure 7 shows these ray paths. The ray co-ordinates were output and taken into MATLAB and using the method of triangulation an estimate for the position of the reflector (\hat{s}_h) was found. The

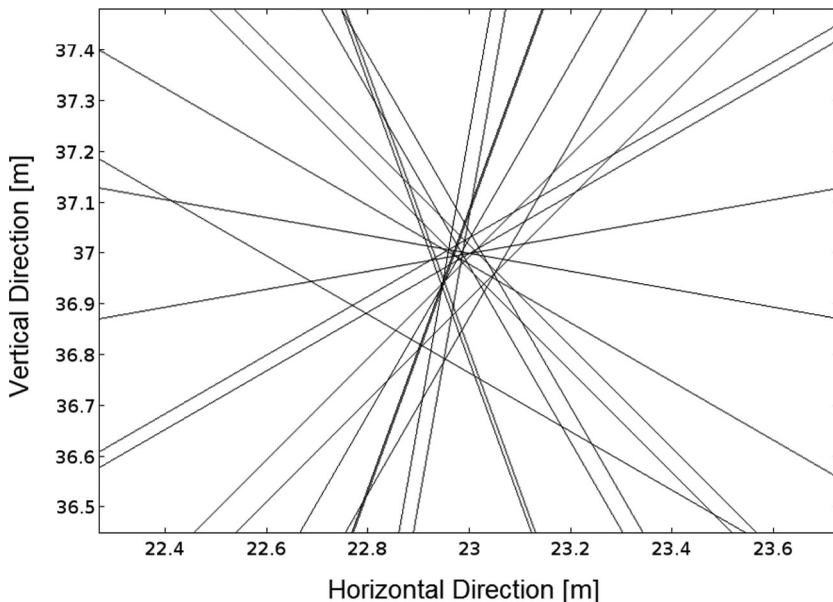


Figure 7. Numerical experiment A: Image showing the straight rays propagating through a homogeneous refractive index map from cameras c_k with initial angle $\theta_{k,s}^*$ towards the retroreflector at known position s^* ; the rays do not all intersect at a single point with this homogeneous refractive index map assumption.

cluster of intersection points is shown in Figure 8 along with the true position of the reflector s^* (shown by the square) and the estimated position of the reflector \hat{s}_h (illustrated by the filled circle). For illustration purposes, s^* has been translated such that it is at the origin. This translation has also been applied to \hat{s}_h . We found that the positional error was $\varepsilon_h = 279.3 \mu\text{m}$.

5.2. Numerical experiment A: the inversion process

For each numerical experiment, the inversion process ran for 100,000 iterations (realizations) with a burn-in period of 20,000. Due to the significant number of iterations in the rj-MCMC optimization, the results are not produced in real time and can take several hours to complete. The computational time is of course subject to the particular hardware system it is run on and the code is not currently optimized for speed. The scenarios considered here started from a uniform initial model (so a spatially constant refractive index map). In a commercial setting, once this system was in a steady state, the initial model would be the previously acquired refractive index map so it would be starting in the neighbourhood of the desired solution. This would significantly reduce the number of iterations and hence the computational time. The major computational overhead is associated with the choice of optimization method employed and it may be, in the case where sufficient prior information is available, that a suitable deterministic, gradient based method could be found and this overhead significantly reduced. Throughout the running of the inversion algorithm, we checked the misfit function and system noise for convergence by using Geweke's Convergence Diagnostic (GCD) [47]. GCD assumes that if the Markov Chain has

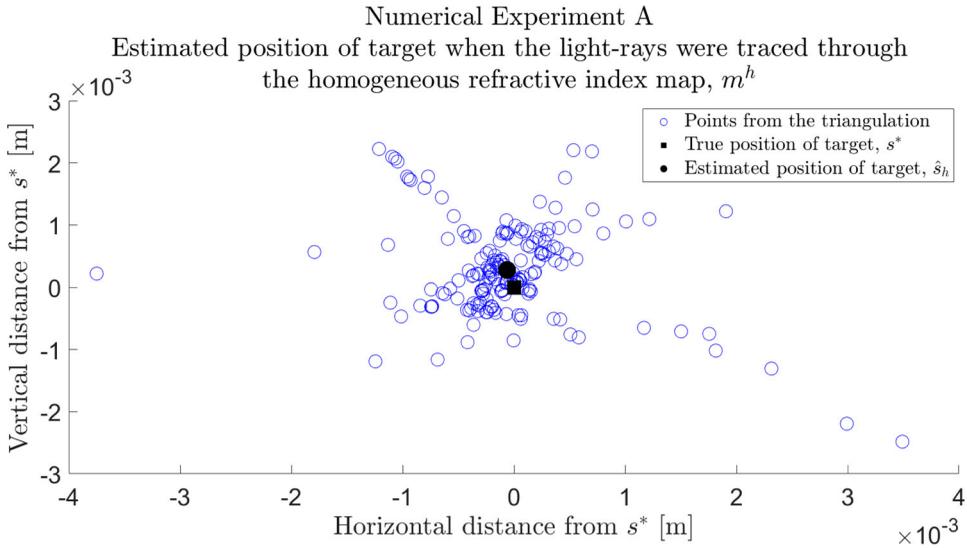


Figure 8. Numerical experiment A: This scatter-plot illustrates the cluster of light ray intersection points when the light rays are traced from cameras c_k with initial angle $\theta_{s,k}^*$ through a homogeneous refractive index map. The mean of the light ray intersection points yields \hat{s}_h (illustrated by the filled circle). The square shows the true position of the object of interest (s^*). The true position of the reflector, s^* has been translated such that it is at the origin. This translation has also been applied to \hat{s}_h .

converged then the mean and variance of the posterior samples should no longer change. The GCD confirmed that the misfit function had converged in all the numerical experiments shown and the GCD Z-score values obtained for all numerical experiments were approximately 0.7. Since this value is between -2 and 2 then we can assume that the data has converged [47]. While not used in this paper, it is possible to have additional convergence criteria using the Geweke diagnostic formulation by examining the stationarity of refractive index values at specific points in the spatial domain and the noise parameter. From intensive testing of the algorithm, it was determined that 100,000 iterations provided the best trade-off between convergence, accuracy and computational time.

The inputs into the algorithm were the position of the cameras c_k , the position of the reflectors r_i , the angle that the light enters the camera $\theta_{k,i}^*$, a mapping, Φ , detailing which camera sees which reflector, the range of refractive indices the model can choose for each Voronoi cell at each perturbation and the range of the number of Voronoi cells for the model (with the lowest value always set to 5 Voronoi cells). The algorithm naturally looks for the smallest number of Voronoi cells to represent the map. It is also possible to prescribe an initial refractive index map (initial model) for each inversion, which takes the form of a Voronoi tessellation. In this paper, three choices for this initial map were trialed: the known refractive index map m^* , a uniform refractive index map and a random refractive index map.

In the case of the known and uniform initial models, their corresponding input files contained the seed positions of the Voronoi tessellation and the velocity assigned to each cell. For each numerical experiment, we assumed that the uniform initial model had the same number of Voronoi cells as the known refractive index map m^* . When a random

refractive Voronoi tessellation was chosen as the initial model, its seed co-ordinates and the refractive indices assigned to each cell were drawn from uniform priors (with limits determined by the spatial domain of interest and the underlying physics). The number of Voronoi seeds P in this initial model was drawn from the uniform distribution over the integers in a specified range. Each numerical experiment was repeated three times to allow the algorithm to start from each of the initial models (known refractive index map, uniform refractive index map and random refractive index map).

At each iteration j of the inversion, the corresponding Voronoi tessellation is stored within the algorithm and once the inversion is finished the mean, median and MAP of the posterior value of each pixel in the refractive index map is calculated over j realisations and these values produce the recovered refractive index maps. In addition to these recovered maps, the inversion process also outputs the values of the posterior distribution of the noise parameter, $\zeta^{(j)}$, the number of Voronoi cells at each iteration $P^{(j)}$ and the values of the misfit function $\gamma^{(j)}$.

5.3. Numerical experiment A: results

We will firstly examine the case where the known refractive index map m^* is used as the initial Voronoi tessellation $m^{(0)}$. The inversion used 100,000 iterations with a burn-in period of 20,000 iterations. The misfit function, $\gamma^{(j)}$, fluctuated around an initial value of $O(10^{-3})$, the associated Voronoi tessellations oscillated around 95 Voronoi cells and ranged from 81 to 118, and the acceptance rate was 40%.

Three recovered refractive index maps were output from the inversion algorithm: the mean, median and MAP of the posterior distribution. We found that the mean of the posterior distribution was most similar to m^* so chose this as the map to calculate ε_α . We chose the mean as it resulted in a smoother refractive index landscape and better results in terms of reducing the error in the position of the object of interest, s^* , and the associated uncertainty in this measurement. The results of the triangulation method are shown in Figure 9. The unfilled circles show the intersection points of the light rays, the square shows the true position of the reflector s^* and the filled circle shows the estimated position of the reflector when the mean recovered refractive index map is used. The true position of the reflector has been translated to the origin. We found that the positional error ε_α was $149.0 \mu\text{m}$ which yields a relative improvement of ε_{imp} of 47%.

The uniform refractive index map was then used as the initial model. The uniform refractive index was set to be a 100 cell Voronoi tessellation with each cell having a single refractive index which was the mean of the refractive indices used. The acceptance rate for this inversion was again 40% and the number of Voronoi cells ranged from 83 to 123. We again found that the mean of the posterior distribution was most similar to the known refractive index map m^* and therefore chose this as the refractive index map to calculate ε_α .

We traced the light rays through the recovered refractive index map in order to find the error in the positioning of the reflector s^* . Using the triangulation method an estimate for \hat{s}_α was found and the results of the triangulation indicated that the error in the estimated position ε_α was $205.8 \mu\text{m}$, which leads to a ε_{imp} value of 26%.

A final inversion was carried out for numerical experiment A wherein a random refractive index map was chosen as the initial model. This refractive index map was a Voronoi

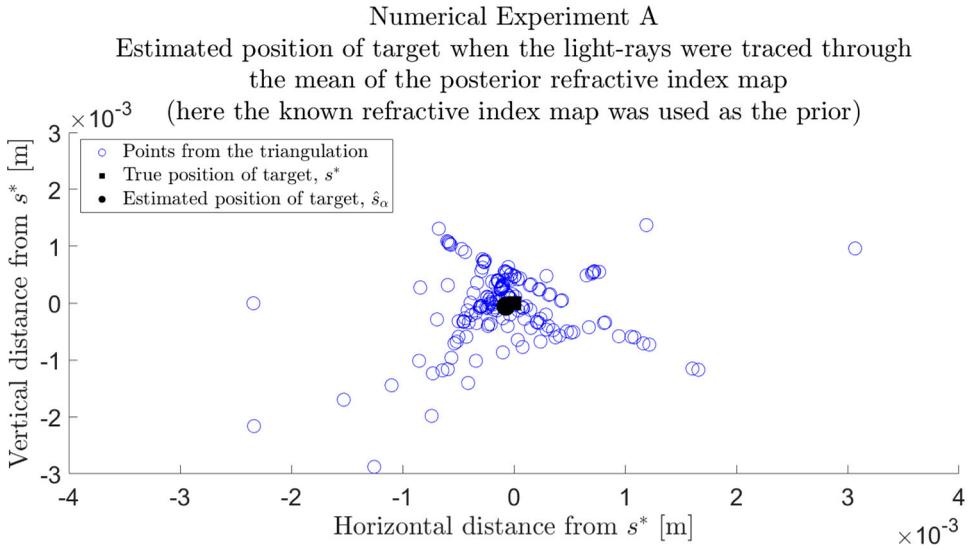


Figure 9. Numerical experiment A: This plot shows the results of the triangulation method when the light rays were traced through the mean recovered refractive index map. The unfilled circles show the intersection points of the light rays, the filled circle shows the estimated position of the reflector (\hat{s}_α) and the square shows the true position (s^*). The true position of the reflector, s^* has been translated such that it is at the origin.

Table 1. Improvement in uncertainty and position of the object of interest s^* when the light rays are traced through various refractive index maps.

Spatial domain	ε (μm)	ε_{imp} (%)
Homogeneous m^h	279.3	NA
Known map as initial model	149.0	47
Uniform map as initial model	205.8	26
Random map as initial model	226.2	19

Note: These results are for numerical experiment A.

tessellation with 111 cells and the number of Voronoi cells ranged between 90 and 124 in the inversion process. The acceptance rate was again 40%. We again found that the mean of the posterior distribution was most similar to m^* so chose this as the refractive index map to calculate ε_α . We traced the light rays through the mean of the posterior recovered refractive index map in order to find the error in the estimated position of the reflector s^* . Using the triangulation method an estimate for \hat{s}_α was found and the results of the triangulation indicated that the error in the estimated position ε_α was 226.2 μm which leads to an ε_{imp} value of 19%. A full summary of the results for numerical experiment A can be found in Table 1.

The only thing that has changed in these implementation of the method was the initial model of the refractive index map for the inversion. The posterior distributions converge to the same similar maps. It is the sensitivity of the triangulation method which results in the changes in ε_{imp} for the known, uniform and random initial models. In summary, in

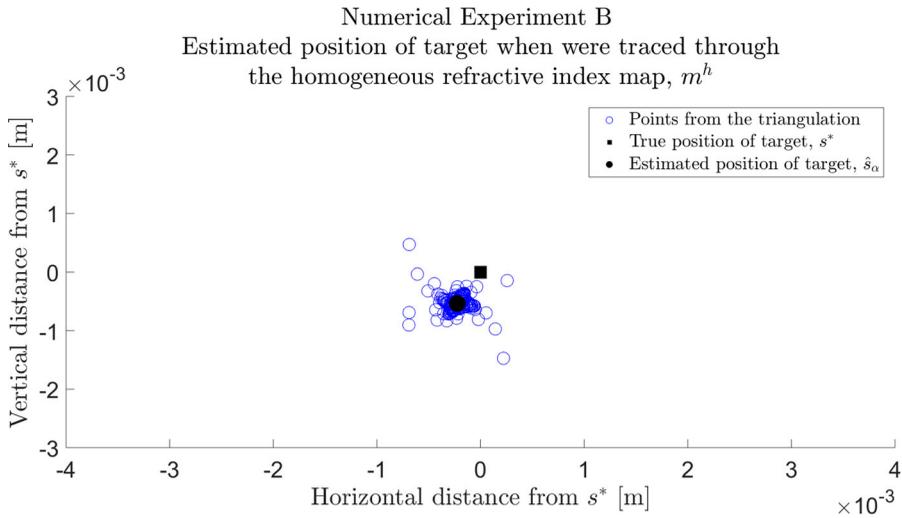


Figure 10. Numerical experiment B: This scatter-plot illustrates the cluster of light ray intersection points when the light rays are traced from cameras c_k with initial angle $\theta_{s,k}^*$ through a homogeneous refractive index map. The mean of the light ray intersection points yields \hat{s}_h (illustrated by the filled circle). The position of the object of interest (s^*) is shown by the square. The true position of the reflector, s^* has been translated to the origin.

numerical experiment A, every recovered refractive index map yielded an improvement in the positioning of s^* .

5.4. Numerical experiment B

To create a refractive index map that is closer to that observed in an industrial setting a smooth Gaussian distribution was chosen for the known refractive index map m^* using Equation (9). This refractive index map was input into COMSOL following the steps outlined previously. In fact, this geometry was simpler to create in COMSOL as the geometry is not parameterized using a Voronoi tessellation, but discretized using a regular square mesh with a 0.05 m grid spacing. The advantage of numerical experiment A over B is that the known refractive index map is a Voronoi tessellation and so can be input into the inversion process as the initial model. This allows a sensitivity analysis of the misfit function to be conducted as its value is approximately zero. The error is due to the inversion algorithm using a different forward model (a fast marching method and ray tracing solver) to that in COMSOL (numerical solution of the Eikonal differential equation posed as an initial value problem). The parameters used in Equation (9) and in the Ciddor equation were the same as those used in numerical experiment A.

As in the previous numerical experiment the error when the refractive index map is assumed to be homogeneous was calculated first and ε_h was 174.2 μ m. In Figure 10, the unfilled circles show the intersection points from the triangulation, the true position of the reflector s^* is shown by the square, and the estimated position of the reflector \hat{s}_h is shown by the filled circle.

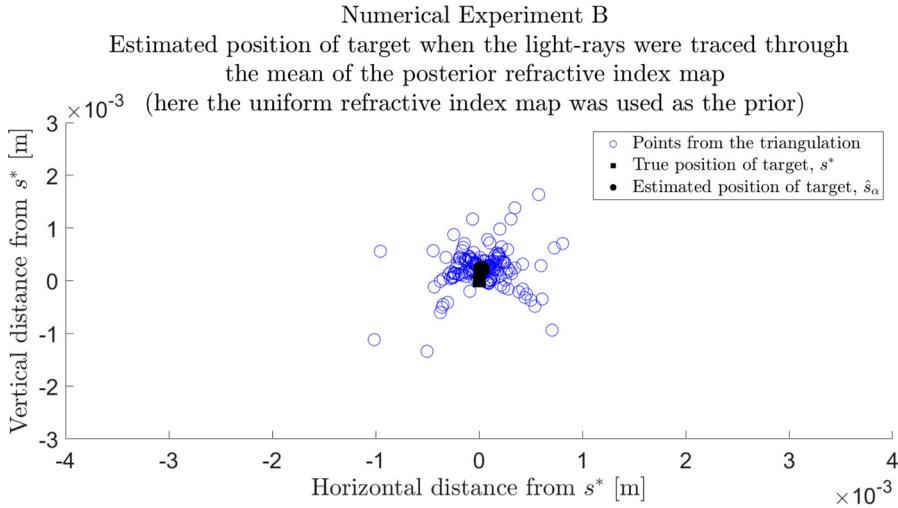


Figure 11. Numerical experiment B: This scatter-plot illustrates the cluster of light ray intersection points when the light rays are traced from cameras c_k with initial angle $\theta_{s,k}^*$ through a homogeneous refractive index map. The mean of the light ray intersection points yields \hat{s}_h (illustrated by the filled circle). The true position of the object of interest (s^*) is shown by the square. The position of the reflector, s^* has been translated such that it is at the origin.

Unlike numerical experiment A it is not possible to express the known refractive index map exactly in terms of a Voronoi tessellation. This means that for numerical experiment B the inversion process only used a uniform and a random initial model.

The uniform initial model used in numerical experiment B is the same as the one used in numerical experiment A. As in all previous numerical experiments, the inversion ran for 100,000 iterations and had a burn-in of 20,000. There is no significant difference in computational time of the inversion algorithm between numerical experiment A and numerical experiment B. The range of Voronoi cells throughout the inversion ranged from 74 to 114 and the acceptance rate was 39.6%.

The mean of the posterior refractive index map was once again used as the spatial map for which the light rays were traced. The triangulation method resulted in a positional error of $\varepsilon_\alpha = 110.2 \mu\text{ m}$ and this is shown in Figure 11. The mean of the light ray intersection points yields \hat{s}_h (illustrated by the filled circle). The square shows the true position of the object of interest (s^*). This leads to an improved positioning of ε_{imp} of 37%.

We then used a random Voronoi tessellation as the initial model and this Voronoi tessellation had 96 cells. The range of Voronoi cells throughout the inversion varied from 80 to 110, which is a smaller range than when the uniform initial model was used for numerical experiment B, and the acceptance rate was 40%. The misfit function starts off with a large $\gamma^{(j)}$ value and then converges as j increases and finally begins to oscillate around 2.5×10^{-3} ; roughly the converged value in numerical experiment A.

The mean of the posterior refractive index map was used as the spatial map for which the light rays were traced. The triangulation method obtained the following results: $\varepsilon_\alpha = 147.2 \mu\text{ m}$. This leads to an ε_{imp} of 15%. The results of numerical experiment B are shown in Table 2.

Table 2. Improvement in uncertainty and position of the object of interest s^* when the light rays are traced through various refractive index maps.

Spatial domain	ε (μm)	ε_{imp} (%)
m^h	174.2	NA
Uniform map as initial model	110.2	37
Random map as initial model	147.2	15

These results are for numerical experiment B.

6. Conclusions

This paper has introduced a new inverse problem that concerns the reconstruction of a refractive index map from photogrammetry data. It is clear from the literature that there is currently no light-based sensor system capable of reconstructing the refractive index map of air. This is a major barrier to the successful deployment of this technology in a manufacturing setting and so this work has taken a first step in rectifying this.

To show the scale of the problem, we presented a physical experiment, which used two Vicon T160 cameras to track a cluster of retroreflectors. We used a 1.6 kW heater as the localized heat source and when this was switched on it generated a heterogeneous refractive index map. We found that the positional error was $679 \mu\text{m}$. Although the photogrammetry equipment is sensitive enough to detect the submillimetre errors in position caused by thermal fluctuations, they are deficient in their ability to correct for these errors, highlighting the importance of this work. To that end, a stochastic, transdimensional approach for refractive index map recovery has been proposed in this paper.

We recovered the refractive index map by solving an inverse problem via a Bayesian approach; namely the reversible jump Markov Chain Monte Carlo method (rj-MCMC). This recovered refractive index map was then used to discount the effect of variations in this refractive index field in the positioning of reflector targets.

Numerical experiment A introduced a Voronoi tessellation with 100 cells. For all three initial models (known, uniform and random) there was a reduction in uncertainty in the position of s^* by 47%, 26% and 19%, respectively. When the inversion algorithm started with the known refractive index map, and a randomly generated refractive index map. Numerical experiment, B, was a more realistic scenario with a single Gaussian distribution representing the heterogeneous temperature distribution. Using a uniform initial model led to an ε_{imp} of 37%, whereas a random initial model resulted in an ε_{imp} of 15%.

Despite the success of this work in a simulated environment, there are of course questions remaining as to whether or not this method can be deployed in an industrial setting. Due to the iterative approach of the rj-MCMC method, the results are not generated in real time and can take several hours to obtain. That being said, the method currently uses uninformative uniform priors, bounded by reasonable estimates on the limits on refractive indices that we would expect to observe in practice. In a practical situation, the previous estimate of the refractive index map could be used as the initial map for the next deployment of the algorithm. Therefore, once the transient phase was overcome, this steady state deployment would increase the feasibility of the approach. The camera-reflector pairs required for this photogrammetry approach are far less expensive than laser tracking systems which are based on interferometry. Many more camera-reflector pairs can therefore

be deployed to increase the aperture and the volume of data. To fully assess the system in practice, the refractive index map in a 3D volume would have to be established at quite fine resolution. At present there is not a technique to do this. Indeed the technology created in this paper could have a second, and perhaps even more important, application as a means of producing temperature maps in a 3D volume.

As this is a new inverse problem then there are questions pertaining to the best choice of spatial parameterization, forward model, misfit function, and optimization routine. To this end, future work will aim to introduce a new method for calculating $\theta_{k,i}^{(j)}$ which avoids using the ray tracing algorithm. The ray tracing algorithm at present is a computational bottleneck and is not robust enough to deal with the small changes in $\theta_{k,i}^{(j)}$ at each model iteration j . We hope that this new method will produce a more sensitive misfit $\gamma^{(j)}$ and reduce computational time, which will allow more model realisations to be achieved.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This work was funded by a studentship with the University of Strathclyde in collaboration with the National Physical Laboratory (NPL), London, UK.

ORCID

A. Miller  <http://orcid.org/0000-0002-5295-4848>

S.G. Pierce  <http://orcid.org/0000-0003-0312-8766>

References

- [1] UK Robotics and Autonomous Systems Network. Manufacturing robotics: the next robotic industrial revolution. 2016. [Online]. [cited 2019 Aug 20]. Available from: https://www.ukras.org/wp-content/uploads/2018/10/UK_RAS_wp_manufacturing_web.pdf.
- [2] Pedersen MR, Nalpantidis L, Andersen RS, et al. Robot skills for manufacturing: from concept to industrial deployment. *Robot Comput Integr Manuf.* 2016;37:282–291.
- [3] Tobe F. IFR predicts 15 percent industrial robot growth throughout 2018; 2015. The Robot Report. [Online]. [cited 2019 Aug 20]. Available from: <https://www.therobotreport.com/news/ifr-predicts-15-industrial-robot-growth-through-2018>.
- [4] Park HS, Tran NH. An autonomous manufacturing system for adapting to disturbances. *Int J Adv Manuf Technol.* 2011;56(9-12):1159–1165.
- [5] Estler WT, Edmundson KL, Peggs GN, et al. Large-scale metrology—an update. *CIRP Annals Manuf Technol.* 2002;51(2):587–609.
- [6] Williams DC. *Optical methods in engineering metrology.* London:Chapman and Hall; 1993.
- [7] Peggs GN, Maropoulos PG, Hughes B, et al. Recent developments in large-scale dimensional metrology. *Proc Inst Mech Eng Part B: J Eng Manufacture.* 2009;223(6):571–595.
- [8] Schenk T. *Introduction to photogrammetry.* Department of Civil and Environmental Engineering and Geodetic Science, The Ohio State University, Columbus (OH); 2005.
- [9] Hartley RI, Sturm P. Triangulation. *Comput Vis Image Underst.* 1997;68(2):146–157.
- [10] Holá M, Hrabina J, Sarbort M, et al. Contribution of the refractive index fluctuations to the length noise in displacement interferometry. *Meas Sci Rev.* 2015;15(5):263–267.
- [11] Nasr KM, Hughes B, Forbes AB, et al. Determination of laser tracker angle encoder errors. In *EPJ Web of Conferences*, vol. 77. EDP Sciences; 2014.

- [12] Kyle S, Robson S, Macdonald L, et al. Compensating for the effects of refraction in photogrammetric metrology. In 14th International Workshop on Accelerator Alignment, Grenoble, France; Oct 2016. p. 3–7.
- [13] Lewis A. Large volume unified metrology for industry, novel applications & research (LUMINAR): large volume in industry. London: National Physical Laboratory. [Online]. 2016[cited 2019 Aug 20]. Available from: <http://projects.npl.co.uk/luminar/>.
- [14] Robson S, MacDonald L, Kyle S, et al. Close range calibration of long focal length lenses in a changing environment. *Int Arch Photogramm Remote Sens Spat Inf Sci*. 2016;41:115–122.
- [15] Summan R, Pierce SG, Macleod CN, et al. Spatial calibration of large volume photogrammetry based metrology systems. *Measurement*. 2015;68:189–200.
- [16] Palmateer J. Effect of stratified thermal gradients on measurement accuracy with application to tracking interferometer and theodolite measurement. In 7th International Metrology Conference. *Metrologie*, 1995.
- [17] Miller A, Mulholland AJ, Pierce SG, et al. Positional uncertainty in optical-based metrology systems in large volume manufacturing. *ACTA IMEKO*. 2018;7(2):84–89.
- [18] Kuter N, Kuter S, et al. Investigation of wildfires at forested landscapes: a novel contribution to nonparametric density mapping at regional scale. *Appl Ecol Environ Res*. 2018;16(4):4701–4716.
- [19] Kuter S, Akyurek Z, Weber GW. Retrieval of fractional snow covered area from modis data by multivariate adaptive regression splines. *Remote Sens Environ*. 2018;205:236–252.
- [20] Weber GW, Batmaz İ, Köksal G, et al. Cmars: a new contribution to nonparametric regression with multivariate adaptive regression splines supported by continuous optimization. *Inverse Probl Sci Eng*. 2012;20(3):371–400.
- [21] Çevik A, Weber GW, Murat Eyüboğlu B, et al. Voxel-mars: a method for early detection of Alzheimer's disease by classification of structural brain mri. *Ann Oper Res*. 2017;258(1):31–57.
- [22] Tant KMM, Galetti E, Curtis A, et al. Effective grain orientation mapping of complex and locally anisotropic media for improved imaging in ultrasonic non-destructive testing. *Inverse Probl Sci Eng*. 2020;28(12):1–25.
- [23] Tant KMM, Galetti E, Curtis A, et al. A transdimensional Bayesian approach to ultrasonic travel-time tomography for non-destructive testing. *Inverse Probl*. 2018;34(9):095002.
- [24] Bodin T, Sambridge M. Seismic tomography with the reversible jump algorithm. *Geophys J Int*. 2009;178(3):1411–1436.
- [25] Bodin T, Sambridge M, Gallagher K. A self-parametrizing partition model approach to tomographic inverse problems. *Inverse Probl*. 2009;25(5):055009.
- [26] Galetti E, Curtis A, Baptie B, et al. Transdimensional love-wave tomography of the British isles and shear-velocity structure of the east Irish sea basin from ambient-noise interferometry. *Geophys J Int*. 2016;208(1):36–58.
- [27] Rawlinson N, Sambridge M. The fast marching method: an effective tool for tomographic imaging and tracking multiple phases in complex layered media. *Explor Geophys*. 2005;36(4):341–350.
- [28] Luhmann T, Robson S, Kyle S, et al. Close-range photogrammetry and 3D imaging. 2nd ed. Berlin, Germany: Walter de Gruyter; 2014.
- [29] MCP9808: $\pm 0.5^\circ\text{C}$ Maximum accuracy digital temperature sensor data sheet. [Online]. 2011 [cited 2016 Nov 11]. Available from: <http://ww1.microchip.com/downloads/en/DeviceDoc/25095A.pdf>.
- [30] Karim Fahed A, Ozkaymak M, Ahmed S. Impacts of heat exposure on workers' health and performance at steel plant in Turkey. *Eng Sci Technol Int J*. 2018;21(4):745–752.
- [31] American National Standards Institute B89.4.19-2006. Performance evaluation of laser-based spherical coordinate measurement systems. New York (NY): American Society of Mechanical Engineers; 2006.
- [32] COMSOL Inc. *COMSOL Multiphysics: Version 5.3a*. Burlington (MA), 2018.
- [33] Siddiqi K, Tannenbaum A, Zucker SW. A Hamiltonian approach to the Eikonal equation. In: *Energy Minimization Methods in Computer Vision and Pattern Recognition*, vol. 1654. Berlin/Heidelberg/York: Springer; 1999. p. 1–13.

- [34] Sethian JA. A fast marching level set method for monotonically advancing fronts. *Proc Nat Acad Sci*. 1996;93(4):1591–1595.
- [35] Sethian JA. Theory, algorithms, and applications of level set methods for propagating interfaces. *ACTA Numerica*. 1996;5:309–395.
- [36] Born M, Wolf E. Principles of optics: electromagnetic theory of propagation, interference and diffraction of light. 7th ed. Cambridge, UK: Cambridge University Press; 1999.
- [37] Green PJ. Reversible jump Markov chain monte carlo computation and Bayesian model determination. *Biometrika*. 1995;82:711–732.
- [38] Rice JA. Mathematical statistics and data analysis. 3rd ed. Thomson Brooks/Cole: Belmont, CA; 2007.
- [39] Galetti E, Curtis A. Transdimensional electrical resistivity tomography. *J Geophys Res: Solid Earth*. 2018;123(8):6347–6377.
- [40] Hastings WK. Monte carlo sampling methods using Markov chains and their applications. *Biometrika*. 1970;57(1):97–109.
- [41] Metropolis N, Rosenbluth AW, Rosenbluth MN, et al. Equation of state calculations by fast computing machines. *J Chem Phys*. 1953;21(6):1087–1092.
- [42] Gelman A, Roberts GO, Gilks WR. Efficient metropolis jumping rules. *Bayesian Stat*. 1996;5:599–607.
- [43] The MathWorks, Inc. MATLAB and Statistics Toolbox Release 2017b, 2017 (MA).
- [44] Ciddor PE. Refractive index of air: new equations for the visible and near infrared. *Appl Opt*. 1996;35(9):1566–1573.
- [45] Stone JA, Zimmerman JH. Index of refraction of air: vacuum wavelength and ambient conditions based on the Ciddor equation. National Institute of Standards and Technology [Online]. 2017 [cited 2019 Aug 20]. Available from: <http://emtoolbox.nist.gov/Wavelength/Ciddor.asp>.
- [46] Edlén B. The refractive index of air. *Metrologia*. 1966;2(2):71–80.
- [47] Geweke J. Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments. 196, (MN): Federal Reserve Bank of Minneapolis, Research Department Minneapolis; 1991.