Automatic Inspection Planning for Optimizing the Surface Coverage in Industrial Inspection

Mahsa Mohammadikaji

Vision and Fusion Laboratory
Institute for Anthropomatics
Karlsruhe Institute of Technology (KIT), Germany
mahsa.mohammadikaji@kit.edu

Technical Report IES-2017-01

Abstract: Optical product inspection plays an important role in today's industrial manufacturing. Therefore, design of optimized solutions corresponding to the industrial requirements are essential for efficient product quality assurance. To configure an inspection setup one requires to determine the position and orientation of the cameras and illuminations as well as the optical configurations. This problem is commonly known as inspection planning. Today's optical inspection setups are mostly being designed based on trial and error requiring a lot of engineering experience and experimental work. As the design space is high dimensional, the empirical designs typically lead to suboptimal solutions and compromise between contrary requirements. In today's industry we are missing a generic automatic method to translate the inspection requirements into optimized inspection solutions. In this report we propose an optimization framework to automatically propose optimized setup solutions, by minimizing the number of acquisitions which fulfil the inspection requirements. As an example, we consider maximizing the surface coverage for the inspection of a cylinder head in a laser triangulation setup. We characterize the design space and propose different approaches to solve the problem. We finally demonstrate the planning results which successfully cover hard to reach areas on the object.

1 Introduction

A fast, automated, and precise quality inspection process is of high importance in today's industrial production. Automated inspection first made its way into industry with Coordinate Measuring Machines (CMMs) [WPH06], the tactile probes which scan the product in a number of key points. Although CMMs deliver precise measurements, they are being more and more replaced by optical inspection techniques due to two main reasons: the very low inspection rate and their requirement to touch the object [NJ95]. The optical inspection techniques, on the other hand, offer fast touchless scans even with lower hardware costs.

The very benefits that optical methods offer are however not without an extra cost: the inspection planning. If the configuration of CMMs needed a pre-selection of the target key points on the surface, designing an optical inspection setup which would deliver the same measurement quality is not a trivial task [CBL02]. Apart from choosing among the existing inspection techniques such as fringe projection, laser triangulation, deflectometry, interferometry, and etc. (see [BLF15]), one requires to determine the position and orientation of the cameras and illuminations as well as their optical configurations. This problem is commonly referred to as inspection planning in the literature [SRR03, TTA$^+$95]. As the design space is high dimensional, today's optical inspection setups are mostly being designed based on a trial and error process requiring a lot of engineering experience and experimental work. In addition to the high design costs, the empirical designs typically lead to suboptimal solutions and a compromise between contrary requirements. In today's industry we are missing a generic automatic method to translate the inspection requirements for a given industrial product into optimized inspection solutions. Especially for precise inspection of geometrically complex products, such as an engine block in figure 1.1, the need for an automatic inspection planning is more evident.

In this report we try to address this question by proposing an optimization method for planning the inspection of a cylinder head in a laser triangulation setup. The

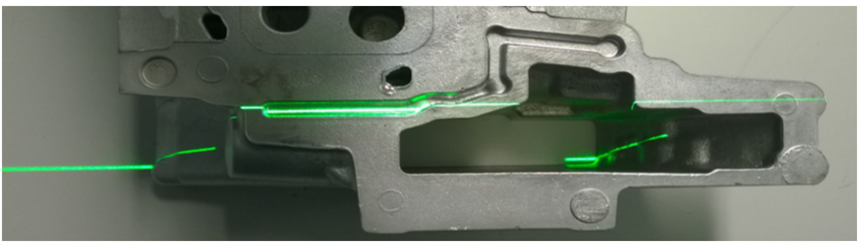

Figure 1.1: Cylinder head illuminated with a laser line

input to this optimization framework is the CAD model of the product with the desired region of interest and the inspection requirements for each region. For the automatic setup planning, we associate each design configuration to a fitness value which enables us to optimize the designed cost function to come up with optimal solutions according to the industry requirements.

The content of this report is organized as follows. In section 2 we formally define the problem of inspection planning as an optimization problem. Section 3 discusses the parameter space of a typical laser triangulation setup. In section 4, we further search for the solution of the optimization problem using two different approaches, the greedy and the combinational approach, with each of them using the Particle Swarm Optimization [BK07] method for the global optimization. In section 5 the achieved planning results will be discussed.

2 Inspection Planning Problem

Let us assume the parameter vector $\mathbf{c}$ determines all the parameters to specify a particular measurement constellation. This means parameter $\mathbf{c}$ includes all the information regarding the positioning and optical configuration of the setup. Similarly, we can define a sequence of k measurements $\{\mathbf{c}_1, ..., \mathbf{c}_k\}$. Let's further assume there is a cost function $f(\mathbf{c}_1, ..., \mathbf{c}_j) \in \mathbb{R}$ which associates each sequence of measurements to a fitness value corresponding to the quality of the measurement with a maximum value of f_{opt}. In a real measurement setup there are of course always constraints. Therefore, not every sequence of measurement can be realizable. Let's quantify all the constraints as a function $r(\mathbf{c}_1, ..., \mathbf{c}_j)$ which takes positive values if the constraints are met.

With the above definitions, the planning problem can be defined as finding the minimum number of N measurements $\{\mathbf{c}_1^*, ..., \mathbf{c}_N^*\}$ which optimize the cost function and at the same time meet the constraints. Therefore

$$\begin{aligned} & f(\mathbf{c}_1^*, ..., \mathbf{c}_N^*) = f_{opt}, \\ \text{s.t.} \quad & r(\mathbf{c}_1^*, ..., \mathbf{c}_N^*) > 0. \end{aligned}$$

For dimensional inspection, the cost function f usually encompasses either of the surface coverage, the scan resolution, or the measurement uncertainty. In our previous works [MBI$^+$17, MBI$^+$16] we have studied different useful cost functions for this problem.

One can easily verify that the optimization problem defined above is at least as hard as the "Set Cover" problem which is known to be NP-hard [Cor07]. The planning problem would be exactly equivalent to the set cover problem when we only consider a discrete set of constellations C for possible solutions, and the surface coverage as the cost function with no further constraints. In this scenario each constellation can cover a subset of the surface. Let us assume the union of all measurable areas by any of the constellations in C is the total area A. In the planning problem, we look for the minimum number of constellations which cover A. In the general case however, the problem is more complex. Often the number of valid constellations are not finite and the cost function is based on complex quantities such as the measurement uncertainty. The optimum value f_{opt} is also often unknown.

For large and complex object surfaces, one is often not able to search the whole parameter space and therefore, needs to resort to approximations. For the set cover problem there is a simple greedy approximation which is actually shown to be the best possible polynomial approximation to the problem [LY94]. In the context of inspection planning, the idea of the greedy approximation is to choose a constellation which covers the most of the surface (or achieves the best fitness values) and continue choosing next best constellations in the same manner. Later in section 5 we discuss the results of the greedy and non-greedy approach.

2.1 Simulations

Typical optimizations require evaluation of many different constellations which are of course not possible to be evaluated in a real setup. Therefore inspection planning always relies on simulations. The planning cost function is consequently based on evaluating the simulation results which is in general a multi-modal non-derivable function.

Evaluation of different cost functions may require different levels of realism in the simulations. The coverage of the surface can be, for instance, estimated by means of fast rasterization-based [AMHH08] simulations (milliseconds per frame), whereas evaluation of the resulting measurement uncertainty requires physically correct image formation simulations using advanced techniques [DBB16]. In such simulations, the emitted photons from the light sources are traced as they get differently scattered by the objects in the scene, all the way to reaching the camera sensor and being converted to intensity values. In our previous work [IBM+17] we have discussed and compared the results of different simulation techniques for simulation of the images of a laser triangulation setup.

3 Design Space

In this section we discuss the degrees of freedom for the planning of a laser triangulation inspection. In such an inspection setup, we have at least one camera and one laser line projector which illuminates a profile on the object (see figure 1.1). The camera captures images of the illuminated profile which are further processed for obtaining the 3D information [MBI+16]. To scan the whole surface, the laser and camera can follow an arbitrary trajectory around the object and capture image frames all along the trajectories. Figure 3.1 depicts the cylinder head CAD model in an arbitrary constellation. For now we assume that the optical parameters of the laser and the camera (e.g. laser power, shutter time, objective f-number, etc.) are already set and are not supposed to be optimized. We solely consider the determination of the geometrical degrees of freedom of the setup. Positioning the camera and the laser as two rigid bodies has a total of 12 degrees of freedom, which can arbitrarily change along the trajectories for each single image frame. To cover a typical product with an average area comparable with the cylinder head, one needs thousands of frames. Obviously this general parametrization leads to an enormous optimization complexity.

One can however make meaningful simplifications to reduce the complexity. For instance, we can assume the scan trajectory to be a linear motion along a particular axis but allow the object instead to freely position under the sensor. The effect of laser distance to the surface can be neglected as lasers can be later focused to any particular distance. We may also parametrize the space in a way that the laser and camera both look towards a common point so that the laser always remains in the camera field of view.

Figure 3.2 illustrates the proposed 9-dimensional degrees of freedom for parametrization of one acquisition. Each acquisition is defined as completely scanning the object along the predefined scan direction. Therefore, instead of planning for each single frame we plan for a number of N acquisitions. In the proposed design space, we dedicate four parameters to sensor placement

(ϕ, θ, τ, d) and five to the object $(\alpha, \beta, \gamma, \Delta x, \Delta z)$. Similar to spherical coordinates, camera placement is determined using polar angle θ, azimuthal angle ϕ, and distance d. The laser holds a triangulation angle τ to the camera, with positive values corresponding to bright field illumination and negative values for dark field illumination. The distance of the laser to the origin is set to a predefined value and does not change during the optimization. The rotation angle ϕ rotates both the camera and the laser to always keep the laser line aligned with the image rows. The object can be freely positioned under the sensor by a 3D rotation using rotation angles α, β, γ, as well as a translational vector. The translation has however only two degrees of freedom because during the acquisition the object is completely scanned in one direction and the translational components parallel to the scan direction do not introduce any new constellations. In the proposed parametrization we consider the y-axis as the scan direction and consider it invariable. Due to rotational degrees of freedom of the sensor and the object, variations of the scan direction within the xy plane will be redundant and do not count as an extra degree of freedom.

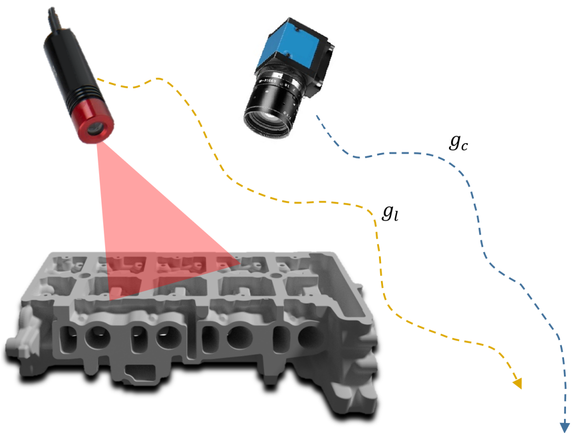

Figure 3.1: Camera and laser can move along arbitrary trajectories g_c and g_l to scan the whole surface.

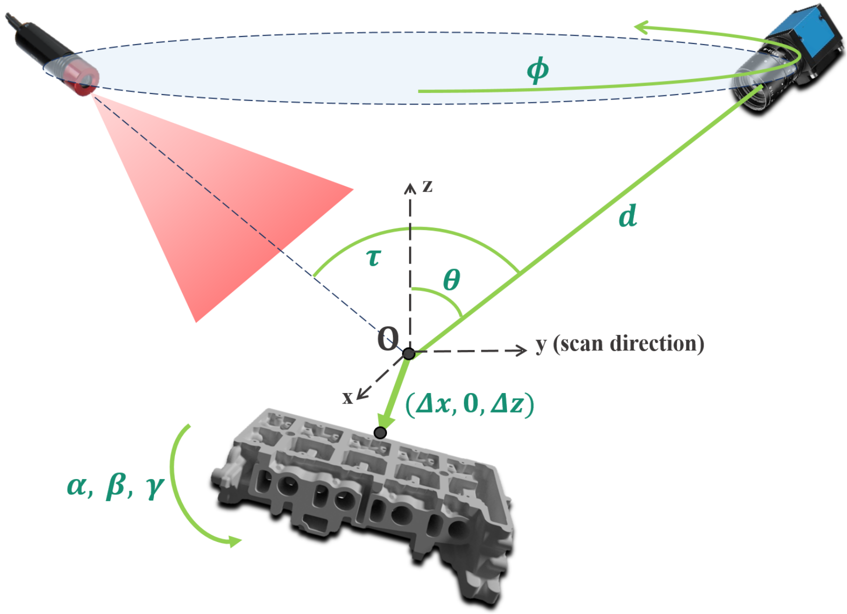

Figure 3.2: Proposed 9D design space for one acquisition

4 Optimization

After definition of the degrees of freedom, in this section we go back to the main optimization problem in equation 2. The main difficulty with the optimization is that we not only want to optimize the constellation parameters but also minimize the number of acquisitions. For a fixed number of M acquisitions, one has an optimization problem with 9M parameters for which we can use standard optimization algorithms. To find the optimal value for M one needs another optimization on possible values of M. The pseudo-code in algorithm 4.1 can be a potential solution to the problem.

Algorithm 4.1 Inspection planning algorithm - combinational

Result: Minimum number of acquisition M^*, constellation parameters $\mathbf{p}^*$

$\mathrm{M} = 0$, $\mathbf{p} = [\,]$
$\mathrm{lastFitness} = 0$, $\mathrm{nextFitness} = 0$
repeat
    `/* try the solution to M+1 acquisitions */`
    $\mathrm{M} = \mathrm{M} + 1$
    $\mathrm{lastFitness} = \mathrm{nextFitness}$
    $[\mathrm{nextFitness}, \mathbf{p}] = \mathit{optimizeMultipleAcquisitions}(\mathrm{M})$
until $(\mathrm{nextFitness} - \mathrm{lastFitness} \leq \mathrm{T}$ or $\mathrm{surfaceIsCovered})$;
$\mathrm{M}^* = \mathrm{M}$, $\mathbf{p}^* = \mathbf{p}$

In the above algorithm, one goes through a $9(\mathrm{M}+1)$-dimensional optimization in each iteration to obtain the optimized fitness value and the corresponding parameter vector $\mathbf{p}$, and compares the gain of having $\mathrm{M}+1$ optimized acquisitions compared to M optimized acquisitions. If the gain of allowing one more acquisition is zero or insignificant we stop the optimization and return the result of the previous step. Of course one can modify the linear search in this algorithm into a binary search to significantly increase the performance; however, even for moderate values of M like 20, the high-dimensional optimization ($9 \times 21 = 189$ free parameters) imposes a great overhead. This is not only a problem from the performance point of view (the algorithm is in the complexity class NP), but it makes it significantly less probable for a heuristic optimization algorithm to find the global optimum of a non-convex non-derivable high-dimensional cost function.

As mentioned earlier, the best polynomial approximation to the set cover problem is the greedy approach. In this algorithm, one already uses the previously optimized constellations and adds only a single constellation to the previous ones. This keeps the optimization overhead constant with the increase of the acquisitions; however, this may result in more acquisitions than the theoretical optimal value. For the set cover problem, the approximation factor is shown to be bounded by a logarithmic upper-bound [LY94]. Algorithm 4.2 contains the pseudo-code of the greedy approximation to the inspection planning algorithm.

Algorithm 4.2 Inspection planning algorithm - greedy approximation

Result: Minimum number of scans M^* with parameters $\mathbf{p}^*$
$\mathrm{M} = 0$, $\mathbf{p} = [\,]$
$\mathrm{lastFitness} = 0$, $\mathrm{nextFitness} = 0$
repeat

```
    lastFitness = nextFitness
    /* 9-dimensional optimization                    */
    [nextFitness, c] = findNextBestScan(p)
    /* append to previous scans                      */
    append(p, c)
    M = M + 1
```

until $(\mathrm{nextFitness} - \mathrm{lastFitness} \leq \mathrm{T}$ or surfaceIsCovered$)$;
$\mathrm{M}^* = \mathrm{M}$, $\mathbf{p}^* = \mathbf{p}$

4.1 Particle Swarm Optimization

Using either of the algorithms for solving the inspection planning problem, we need a global optimizer. As mentioned earlier, the cost function is based on evaluating simulations, taking any complex form with multiple local optimums. Therefore, derivative-free randomized heuristic search algorithms seem to be the right choice for the global optimization problem.

Among the many existing global optimizaters in the literature (see [Wei09] for a good review), the Particle Swarm Optimization (PSO) has been widely applied in many fields. PSO is a variant of swarm-based intelligent methods [BDT99] whose main idea is to imitate the behavior of biological species which live in colonies such as birds, ants, and bees. We observe that these species exhibit a high intelligence in their social activities such as searching for food, although they exhibit simple individual acts. The PSO algorithm is especially inspired by the way birds communicate when searching a field for food [Wei09]. The approach of PSO is very intuitive. For optimizing a cost function $f(\mathbf{x})$, one initiates a number of K random particles in the parameter space (the swarm size). The particles are not only individual searchers in the paramter space (similar to many parallel simulated-annealing [Wei09] optimizers), but they also communicate with each other and share the results of their local searches which guides them for the rest of the search.

Every particle in this algorithm is composed of three vectors: its current position in the search space $\mathbf{x}_i$, its best individually found position so far $\mathbf{b}_i$, and its velocity $\mathbf{v}_i$. The particles are also aware of the best position $\mathbf{g}_i$ found by the rest of the particles so far. This might be actually the globally best found position, or the best position that the particle has so far heard from those who have communicated with it. The algorithm updates the particles at each iteration by updating their velocities and positions according to

$$\begin{aligned} \mathbf{v_i}^{t+1} &= \mathbf{v_i}^{t} + \epsilon_1(\mathbf{b}_i^t - \mathbf{x}_i^t) + \epsilon_2(\mathbf{g}_i^t - \mathbf{x}_i^t), \\ \mathbf{x}_i^{t+1} &= \mathbf{x}_i^t + \mathbf{v_i}^t. \end{aligned}$$

The random factors ϵ_1 and ϵ_2 weight tendencies of the particle to search further towards the local and global optimum. A particle also has a tendency to keep its previous direction, therefore we also add the $\mathbf{v_i}^t$ term to the velocity update rule. The position is simply computed as moving from the previous position along the updated velocity.

Since the introduction of PSO, there have been a few standardizations proposed which give recommendations on choosing the parameters of the algorithm like the swarm size, the communication structure, and weightings. In this work, we have orientated the PSO implementation based on the standardization given in 2007 [BK07].

5 Coverage Planning Results

The cylinder head object contains hard to reach areas such as deep intake and exhaust manifolds. Therefore, maximizing the surface coverage with minimized number of scans is a non-trivial problem which is of high interest for the industry. In this section, we present the inspection planning results with the goal of maximizing the surface coverage. For evaluating the surface coverage, we use a fine mesh model of the object as shown in figure 5.1. The granularity of the mesh elements must correspond to the required inspection resolution. The surface coverage can be then evaluated as the area of all the patches which have been measured with at least one point.

For the optimization, we have used the proposed parameter space in figure 3.2 with a simplification of setting the camera distance d to a constant value of 0.5 meter. This choice can be justified by the fact that the object translation component along the z-axis introduces very similar effects to changing the camera distance. However, in the future we intend to get the results with the full degrees of freedom. As optimization constraints, we assume valid ranges for each of the parameters, based on the degrees of freedom of the real physical setup. For the current results we have bounded the parameters to

$$\begin{array}{lll} \phi \in [0^\circ, 360^\circ), & \theta \in [0^\circ, 80^\circ], & \tau \in [-80^\circ, -10] \cup [10^\circ, 80], \\ \alpha, \beta, \gamma \in [0^\circ, 360^\circ), & \Delta x \in [-0.4, 0.4]\, m, & \Delta z \in [-0.5, 0.5]\, m. \end{array}$$

Many parameters can already take their full range, such as ϕ. Others such as the triangulation angle τ must be constrained to deliver meaningful measurements.

Figure 5.2 compares the results of the greedy vs. the combinational approach. This chart displays the optimized measurable area for each number of acquisitions. The orange line determines the full area of the object, which is however not fully measurable because some areas are either completely unreachable or require constellations which violate the optimization constraints. The blue dotted graph displays the improvements of the greedy surface coverage planning vs. the number of acquisitions. One can see that the algorithm makes rather big improvements at the beginning; however, the contribution of the next acquisitions gradually reduces until it falls below the threshold for the 30$^{\text{th}}$ acquisition. For comparison, we have also applied the combinational planning algorithm for

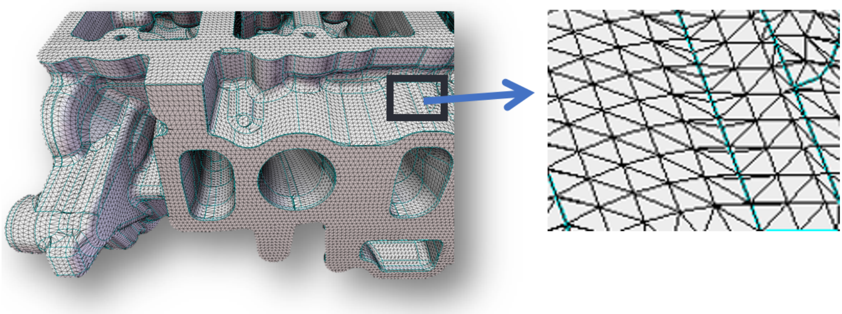

Figure 5.1: Cylinder head fine mesh model for evaluating surface coverage

5, 10, 15, and 20 acquisitions, where we allow the optimizer to combinationally optimize all the acquisitions together. In theory the combinational approach is able to find the global optimum as it looks for every possible combination of constellations. The greedy approach however, might be trapped in local optimum solutions as the constellations are optimized in a non-reversible approach. In practice, however, we see that the results of the combinational approach can even underperform the greedy method. This is due to the fact that it is less probable for a heuristic random optimizer to find the optimum of a high dimensional problem, compared to a problem with significantly less degrees of freedom.

Figure 5.3 illustrates the resulting point cloud of the object after applying the 30 optimized acquisitions obtained by the greedy planning. As it can be seen, the deep cavities corresponding to the intake and exhaust manifolds have been covered.

6 Summary and Future work

In this report we proposed and implemented an optimization framework for automatic optimization of industrial inspection setups. We went through the

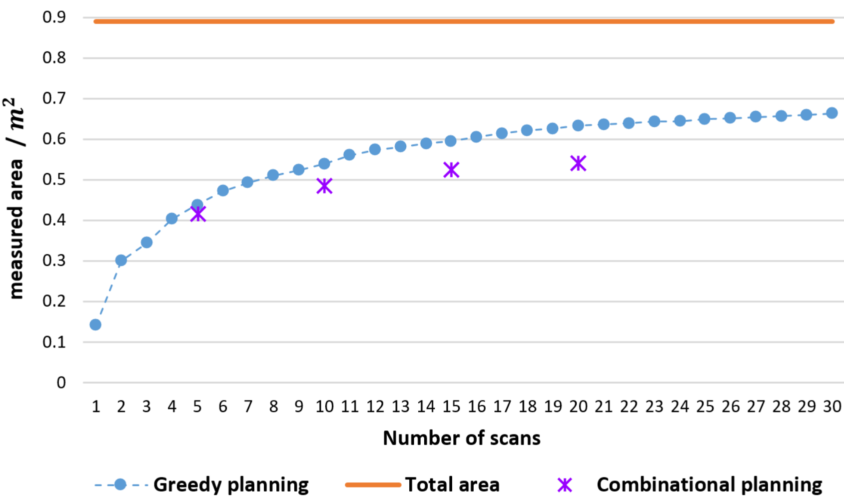

Figure 5.2: Optimized surface coverage vs. number of acquisitions

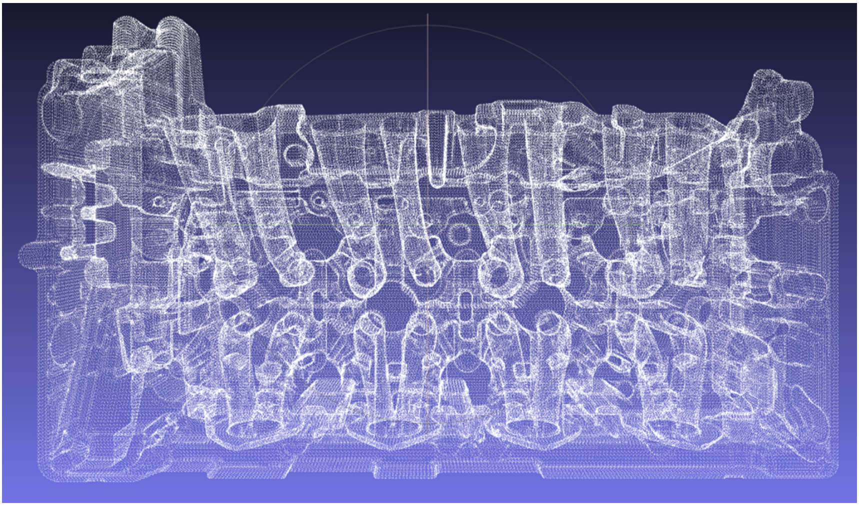

Figure 5.3: Point cloud of the resulting measurement after applying 30 optimized acquisitions obtained by the greedy planning.

parametrization of the design space for inspecting a cylinder head in a laser triangulation setup and compared different approaches for finding the minimized number of acquisitions which maximize the surface coverage.

It is very insightful to calculate the actual measurable area on the cylinder head to better evaluate the achieved coverage results. As a future work, we plan to apply methods for calculating ambient occlusion [Mil94] to calculate the area measurable for a given set of inspection constraints. In addition, there is also potentials to combine the greedy and the combinational optimization approaches to benefit from both. In a hybrid approach, one can use the greedy approach to come up with good starting points with less optimization overhead and further apply the combinational approach on the achieved suboptimal results to globally improve the results.

Bibliography

[AMHH08] Tomas Akenine-Möller, Eric Haines, and Naty Hoffman. *Real-Time Rendering*. A. K. Peters, Ltd, Natick, MA, USA, 3 edition, 2008.

[BDT99] Eric Bonabeau, Marco Dorigo, and Guy Theraulaz. *Swarm intelligence: from natural to artificial systems*. Oxford University Press, 1999.

[BK07] Daniel Bratton and James Kennedy. Defining a standard for particle swarm optimization. In *Proc. 2007 IEEE Swarm Intelligence Symposium*, pages 120–127, 2007.

[BLF15] Jürgen Beyerer, Fernando Puente León, and Christian Frese. *Machine Vision: Automated Visual Inspection: Theory, Practice and Applications*. Springer, 2015.

[CBL02] A. Contri, P. Bourdet, and C. Lartigue. Quality of 3D digitised points obtained with non-contact optical sensors. *CIRP Annals - Manufacturing Technology*, 51(1):443–446, 2002.

[Cor07] Thomas H. Cormen. *Introduction to algorithms*. MIT Press, Cambridge, Mass., 2 edition, 2007.

[DBB16] Philip Dutre, Philippe Bekaert, and Kavita Bala. *Advanced global illumination*. CRC Press, 2016.

[IBM+17] S. Irgenfried, S. Bergmann, M. Mohammadikaji, J. Beyerer, C. Dachsbacher, and H. Wörn. Image formation simulation for computer-aided inspection planning of machine vision systems. In Jürgen Beyerer and Fernando Puente León, editors, *Proc. SPIE Optical Metrology*, SPIE Proceedings, page 1033406. SPIE, 2017.

[LY94] Carsten Lund and Mihalis Yannakakis. On the hardness of approximating minimization problems. *Journal of the ACM*, 41(5):960–981, 1994.

[MBI+16] M. Mohammadikaji, S. Bergmann, S. Irgenfried, J. Beyerer, C. Dachsbacher, and H. Wörn. A framework for uncertainty propagation in 3D shape measurement using laser triangulation. In *IEEE International Instrumentation and Measurement Technology Conference (I2MTC)*, pages 6–11, 2016.

[MBI+17] M. Mohammadikaji, S. Bergmann, S. Irgenfried, J. Beyerer, C. Dachsbacher, and H. Wörn. Probabilistic surface inference for industrial inspection planning. In *Proc. 2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1008–1016, 2017.

[Mil94] Gavin Miller. Efficient algorithms for local and global accessibility shading. In Dino Schweitzer, Andrew Glassner, and Mike Keeler, editors, *Proc. SIGGRAPH 1994*, pages 319–326, 1994.

[NJ95] Timothy S. Newman and Anil K. Jain. A survey of automated visual inspection. *Computer Vision and Image Understanding*, 61(2):231–262, 1995.

[SRR03] William R. Scott, Gerhard Roth, and Jean-François Rivest. View planning for automated three-dimensional object reconstruction and inspection. *ACM Computing Surveys*, 35(1):64–96, 2003.

[TTA+95] Konstantinos Tarabanis, Roger Y Tsai, Peter K Allen, et al. The MVP sensor planning system for robotic vision tasks. *IEEE Transactions on Robotics and Automation*, 11(1):72–85, 1995.

[Wei09] Thomas Weise. Global optimization algorithms – theory and application. *Self-Published*, 2009.

[WPH06] A Weckenmann, G Peggs, and J Hoffmann. Probing systems for dimensional micro-and nano-metrology. *Measurement Science and Technology*, 17(3):1–504, 2006.