

Bayesian surface estimation for white light interferometry

Michael Hissmann* Fred A. Hamprecht†

2004/07/05

Abstract

In conventional white light interferometry (WLI) surface estimation, data acquisition is followed by a preprocessing step in which the 3D (space-time) data is reduced to a 2D surface or height map. Finally, a postprocessing step eliminates the height outliers that arise from the preprocessing under ordinary experimental conditions.

We introduce a Bayesian approach that unifies pre- and postprocessing by considering simultaneously both the full 3D data set and knowledge concerning the surface smoothness. This knowledge is coded into the prior probability of local height configurations. The surface is estimated as the mode of the marginal posterior at each pixel. An adept formulation of the prior allows for the exact computation of the estimate, obviating the need to sample from the posterior using Markov Chain Monte Carlo (MCMC) methods. A complete surface can thus be obtained in 3 – 30 s.

A quantitative comparison with (adaptive) median filtering shows that all three approaches decimate outliers, but that the Bayesian estimation leads to smaller average absolute errors. For slow scanning speeds and good raw data, this is due to a reduced tendency to oversmooth; while for poor input data, the enhancement is explained by the more complete exploitation of the observations.

1 Introduction

1.1 Significance of white light interferometry

White light interferometry (WLI) has become an important visual inspection tool in fields which require high precision surface height maps. With its short coherence length, WLI fills the precision gap between mechanical ball testing devices ($> \mu\text{m}$ scale) and laser interferometers ($< \mu\text{m}$ scale). While mechanical devices produce 1D height profiles only, interferometers can acquire a 2D height map in one take, which both opens new analytic possibilities and raises new questions concerning the standardized characterization of surfaces.

*M. Hissmann is with the Robert Bosch GmbH, Stuttgart, Germany (e-mail: michael.hissmann@de.bosch.com).

†F. A. Hamprecht is with the Interdisciplinary Center for Scientific Computing, University of Heidelberg, Germany (e-mail: fred.hamprecht@iwr.uni-heidelberg.de).

The development of white light interferometers for industry was first guided by the requirements of semiconductor wafer inspection. In recent years, especially in the automotive sector, the precision of mechanical component manufacture has exceeded that of mechanical testers. Machined surfaces that are to be inspected typically exhibit a roughness with a standard deviation above $\lambda/2 \approx 0.4 \mu\text{m}$ (λ : center wavelength of the interferometer) and can usually be considered optically rough.

For in-line quality control during the manufacturing process, the requirements often include robust data acquisition and processing at high speed, whereas a height resolution near or beyond the roughness limit is not needed.

In this paper, we propose a new approach for the denoising of WLI data which, although of general applicability, is particularly well suited for the scenario depicted above.

1.2 Data acquisition

For our purpose, we consider a white light interferometer as a conventional two-arm Michelson interferometer with a broad-band light source. One arm is delimited by the surface under investigation, the other by the reference object, usually a plane mirror. The change of the interference pattern caused by a systematic variation of the length of one of the optical paths is recorded in full spatial detail by means of a camera. The recorded dataset consists of a time series for each camera pixel. Each time series contains noise as well as an oscillatory signal centered around the surface height at that pixel, which we aim to detect. WLIs have been proposed as measuring devices in [1], [2] and [3]. Recent technical developments address curved surfaces of high industrial relevance [4], [5].

The ideal WLI signal has the following form for each pixel (cf. [1] for details):

$$I(z) = \bar{I} + A(z - z_0) \cos(\varphi(z)). \quad (1)$$

Here, \bar{I} denotes the incoherent intensity at a pixel and $A(z)$ the envelope of the interference signal. Its width is inversely proportional to the spectral width of the light source. For commonly used LED sources, it is approximately a Gaussian centered around z_0 , the height relative to the reference surface that is to be estimated. The cosine term with phase $\varphi(z)$ comes from the fast interference modulation and is sometimes denoted “inner oscillation” of the signal. Especially for rough surfaces, which the backscattered intensity follows an exponential probability distribution [6], the amplitude is often very low. In addition, the signal is degraded with photonic and electronics noise, see figure 1.

1.3 Conventional height map estimation

The scanning procedure yields a 3D data set, consisting of a time series for each pixel in a 2D plane. The task then is to embed a surface into the data, thus collapsing each pixel’s height dimension into a unique height value. This surface is subject to restrictions imposed by both the data and prior knowledge concerning the smoothness.

In conventional WLI processing, the final height map is obtained in three stages:




Figure 1: Real (above) and ideal (below, cf. eq. (1)) white light interferogram in a single pixel. Scanning velocity is $0.56 \mu\text{m}/\text{frame}$.

1. The interferometer is translated perpendicularly to the surface under investigation, while its CCD-sensor records the optical interference pattern in a (time / height) series of frames.
2. In the preprocessing stage, the center of the interference pattern (central fringe) is estimated, separately, for each pixel's time series and stored in a 2D matrix, the height map.
3. In the postprocessing stage, outliers in the height map are detected, and possibly corrected, using prior knowledge. Most tools available for gray-value image processing can be adopted for this task.

In this paper, we propose an approach completely different in nature (section 2), and compare it to alternative methods detailed in section 3. Results are presented in section 4.

2 Bayesian surface estimation

We use a Bayesian approach to amalgamate the last two steps in the scheme above: We aim to reconcile *all* locally available evidence (that is, the entire noisy time series of a pixel, along with those of its neighbors) with our prior knowledge in a single step. This is in contrast to all previous work, which lets

a postprocessing stage make do with what reduced information is passed on by the preprocessing stage.

2.1 Background

Extending the approach by Hartvig and Jensen [7] from spatially resolved binary to spatially resolved time series data, we aim to accommodate information from both observations and prior knowledge in a Bayesian framework as follows:

Consider a set \mathcal{S} of pixel sites and a pixel index $j \in \mathcal{S}$. The possible height values h^j for each pixel in the height map are from a finite, discrete set $\tilde{\mathcal{H}}$, which for convenience we map to $\mathcal{H} = \{1, \dots, h_{\max}\}$. A 2D height map $\mathbf{h} = \{h^j\}$ is a point in the product space $\mathcal{H}^1 \times \dots \times \mathcal{H}^{|\mathcal{S}|}$.

During the scan process, the camera acquires a 3D data set composed of intensities for all pixels and all scan steps / height values. We denote these 3D recording $\mathbf{x} = \{\mathbf{x}^j\}$ with each \mathbf{x}^j a 1D time series of h_{\max} intensities. We assume the scanning process is stable in the sense that the height-to-index mapping $\tilde{\mathcal{H}} \rightarrow \mathcal{H}$ is linear, so that we can calculate height values directly from the framerate and interframe distance of our interferometer. We will not address interpolation between height values of two adjacent camera frames, although especially for higher speeds, interpolation to a finer discretization is a useful and viable extension (cf. [8], [9] for interpolation strategies in preprocessing).

We look for a 2D height map estimate $\hat{\mathbf{h}}$, given as the *marginal posterior mode estimate* (MPME) of the conditional probability density $\mathbb{P}(\mathbf{h}|\mathbf{x})$. According to the Bayes paradigm, for each pixel j ,

$$\mathbb{P}(h^j|\mathbf{x}^j) \sim f(\mathbf{x}^j|h^j) P(h^j) \quad (2)$$

$\mathbb{P}(h^j|\mathbf{x}^j)$ is the *a posteriori* probability of a height value h^j given the data \mathbf{x}^j . The *likelihood* $f(\mathbf{x}^j|h^j)$ describes the probability of the observations given a height value and is determined by the technical characteristics of the interferometer. It plays the role of a data term, linking our estimate to the observations. The *a priori* probability $P(h^j)$ carries our prior knowledge on the surface, e. g. in the form of local smoothness constraints.

Let \mathcal{C} indicate the union of a pixel and its immediate neighborhood. For convenience, we index the elements of \mathcal{C} with $0, \dots, k = |\mathcal{C}| - 1$, and reserve the index 0 for the central pixel. Throughout the later examples, $k = 8$ is used and \mathcal{C} includes the 8 nearest neighbors of the central pixel. We may write:

$$\mathbb{P}(\mathbf{h}^{\mathcal{C}}|\mathbf{x}^{\mathcal{C}}) \sim f(\mathbf{x}^{\mathcal{C}}|\mathbf{h}^{\mathcal{C}}) P(\mathbf{h}^{\mathcal{C}}) \quad (3)$$

with $\mathbf{h}^{\mathcal{C}}$ the configuration of height values and $\mathbf{x}^{\mathcal{C}}$ all observations in \mathcal{C} .

To find the marginal posterior probability of a height value for the central pixel given all observations in the surroundings, we sum over all possible height configurations of the neighbors:

$$\mathbb{P}(h^0|\mathbf{x}^{\mathcal{C}}) = \sum_{h^1=1}^{h_{\max}} \dots \sum_{h^k=1}^{h_{\max}} \mathbb{P}(\mathbf{h}^{\mathcal{C}}|\mathbf{x}^{\mathcal{C}}) \quad (4)$$

If the aperture is chosen such that the subjective speckle size matches the camera's pixel dimensions, the likelihoods at neighboring pixels are approximately,

but not entirely, independent. Assuming independence, we can factorize the likelihood and obtain

$$\mathbb{P}(h^0|\mathbf{x}^C) \sim \sum_{h^1=1}^{h_{\max}} \cdots \sum_{h^k=1}^{h_{\max}} f(\mathbf{x}^C|\mathbf{h}^C)P(\mathbf{h}^C) \quad (5)$$

$$\sim f(x_0|h_0) \sum_{h^1=1}^{h_{\max}} \cdots \sum_{h^k=1}^{h_{\max}} \prod_{j=1}^k f(x^j|h^j)P(\mathbf{h}^C) \quad (6)$$

The vast number of possible height configurations in the neighborhood is reflected in the nested summations above. It is obvious that a direct evaluation is computationally prohibitive, even for small neighborhoods: for $h_{\max} = 1000$ height steps in an 8-pixel neighborhood ($k = 8$), we face of the order of $|\mathcal{H}|^{k+1} = 10^{27}$ operations to find a height estimate for each pixel, provided the likelihoods have been precalculated.

However, the estimate can be transferred to the realm of the feasible for a limited class of prior probabilities, by using a manipulation proposed by Hartvig and Jensen [7] and detailed below.

In our application, we wish to cast our prior knowledge, that the surfaces under investigation exhibit smoothness on a macroscopic scale, into a prior probability. A simple way of expressing this knowledge is to assume that those configurations, in which no neighbor has a vertical distance larger than δ from the central height value, have a larger probability than the others:

$$P(\mathbf{h}^C) = \begin{cases} q_1 & \text{if } h^1, \dots, h^k \in [-\delta + h^0, \delta + h^0] \\ q_0 & \text{otherwise} \end{cases} \quad (7)$$

q_0 and q_1 obey the following normalization condition (which neglects minor boundary effects):

$$1 = q_1(2\delta + 1)^k \cdot h_{\max} + q_0(h_{\max}^{k+1} - (2\delta + 1)^k \cdot h_{\max}) \quad (8)$$

The enormity of the configuration space leads to small and very small numerical values for q_1 and q_0 . Since we are only interested in the mode and to avoid numerical instabilities, both can be rescaled.

The algorithm proposed thus far depends on two parameters, δ and the ratio q_0/q_1 . We suggest choosing δ to be of the same order as the surface roughness. The best ratio q_0/q_1 can be found through screening a small range of values typical for the specific surface. See section 4 and table 1 for the appropriate values for the sample piece depicted in figure 2.

Using the prior in eq. (7), we can write

$$\begin{aligned} \mathbb{P}(h^0|\mathbf{x}^C) \sim & f(\mathbf{x}^0|h^0) \left(\sum_{h^1=1}^{h_{\max}} \cdots \sum_{h^k=1}^{h_{\max}} \prod_{j=1}^k f(\mathbf{x}^j|h^j) q_0 \right. \\ & \left. + \sum_{h^1=h^0-\delta}^{h^0+\delta} \cdots \sum_{h^k=h^0-\delta}^{h^0+\delta} \prod_{j=1}^k f(\mathbf{x}^j|h^j) (q_1 - q_0) \right) \end{aligned} \quad (9)$$

It is now possible to make use of the following equality:

$$\sum_{h^1} \cdots \sum_{h^k} \prod_{j=1}^k f(\mathbf{x}^j|h^j) = \prod_{j=1}^k \sum_h f(\mathbf{x}^j|h), \quad (10)$$

which can be proven easily via induction over k .

We can now rearrange $\mathbb{P}(h^0|\mathbf{x}^c)$ to obtain a computable formula:

$$\mathbb{P}(h^0|\mathbf{x}^c) \sim f(\mathbf{x}^0|h^0) \left(q_0 \prod_{j=1}^k \sum_{h=1}^{h_{\max}} f(\mathbf{x}^j|h) + (q_1 - q_0) \prod_{j=1}^k \sum_{h=h^0-\delta}^{h^0+\delta} f(\mathbf{x}^j|h) \right) \quad (11)$$

The number of operations required for each pixel is significantly reduced to the order of $k \cdot h_{\max} = 8000$ in our example. The simple a priori density of eq. (7) allows further speed-up by implementing a sliding summation and other efficient routines. Our final implementation reaches practicable execution times (see section 4).

Note that any prior that reduces to a fixed number of constants independent of the likelihood, can be used to transform the nested a posteriori expression. While less restrictive priors from this class, investigated in [7], may better reflect our prior knowledge on the surface, they also reduce the achievable computational simplification. Also note that we can compute our estimate exactly, avoiding the sampling of the posterior by means of Gibbs or Metropolis sampling that is customary in spatial Bayesian inference [10], [11].

Likelihood functions There are two major paths to a likelihood function $f(\mathbf{x}^j|h^j)$:

On one hand, the likelihood can be obtained from a physical modeling of the WLI signal formation process. This strategy is the most accurate, but is hampered by our restricted knowledge on the optical scattering processes and the technical inaccuracies of the interferometer's components. In addition, a proper modeling will take into account the unknown phase of the inner oscillation, which leads to a large computational burden.

On the other hand, there is the empirical route to a likelihood function. This phenomenological approach mimics the characteristics of a properly derived likelihood function. An ideal function would feature a strong response in the vicinity of the true height value and a weak response to the noise elsewhere in the time series. The methods proposed and established by various authors (an overview is given by [12], [13]) for interferogram preprocessing are a good approximation.

For the results presented in the next sections, we chose the second path and used a quasi-likelihood obtained from low-pass filtered finite differences of the time series in a pixel.

3 Methods

3.1 Data acquisition

We recorded interference patterns of a part of a turned steel piece, 19.5 mm in diameter, that features three circular steps of $20 \mu\text{m}$ height difference each, produced using a new tool, and additionally smaller tracks, obtained using a worn




Figure 2: Photograph (full view) of the piece (diameter: 19.5 mm) used in the subsequent analyses.

turning tool (see picture, figure 2). This surface exhibits various characteristics that are typically found in the industrial application of WLI.

To challenge the algorithms, we recorded interferograms at scan speeds of $14\ \mu\text{m/s}$ (which, at a framerate of 50 Hz, corresponds to the Nyquist frequency of the inner oscillation of the interferogram), $28\ \mu\text{m/s}$, $56\ \mu\text{m/s}$, $84\ \mu\text{m/s}$ and $112\ \mu\text{m/s}$ (which corresponds to an 8-fold subsampling). Two exemplary results are shown in figures 3 and 4.

These results were obtained with a C++ implementation of the algorithm. On a 1.2 GHz P IIIm machine, the processing time varied between 31s for $14\ \mu\text{m/s}$ and 2.8 s for $112\ \mu\text{m/s}$, including the raw data input from hard disk.

3.2 Height map estimation

We compare the Bayesian method introduced here with other established approaches detailed in the following. For all evaluations, a 3×3 pixel neighborhood or filter mask was chosen.

Preprocessing only A pixelwise detection of interference patterns with no postprocessing yields, in the case of rough surfaces, a height map which is heavily contaminated by outliers that render further analysis difficult or impossible. These results are included only to allow for an unbiased comparison of the remaining methods.

Median filter mask Filtering the noisy height map obtained from preprocessing with a median filter mask is simple and a common approach. As we are dealing mainly with high amplitude outliers, the median filter seems appropriate. The rank order filter provides optimal robustness against outliers (see e. g. [14]) and preserves edges well. On the other hand, this filter does not adapt

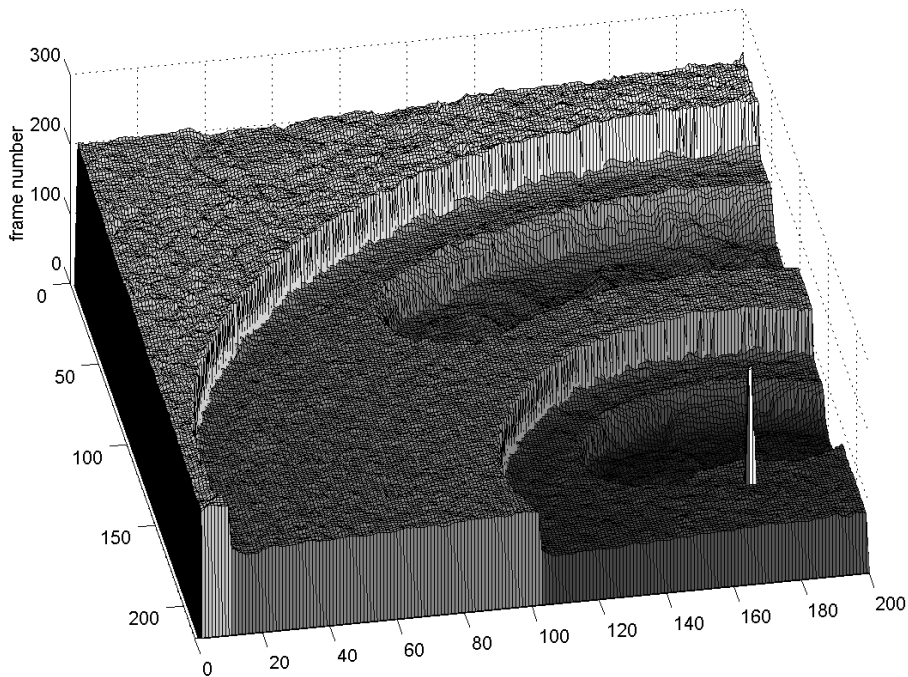


Figure 3: A height map of the sample piece obtained with the Bayesian algorithm at $14 \mu\text{m/s}$. Scale: $0.28 \mu\text{m}/\text{frame}$.

to the quality of the original data and so tends to oversmoothing in areas with little or no outliers.

Outlier-sensitive median filter The basic idea is to replace only suspected outliers with the median of the neighborhood, and to leave all other height values unscathed.

The median filter itself provides a way of detecting outliers: if applied to a central pixel of height h^0 , it yields $\tilde{h}^0 = \text{med}\{h^0, \dots, h^k\}$. For an outlying central height h^0 , the difference $|h^0 - \tilde{h}^0|$ is outside the usual range which reflects the height variations of a rough surface. An optimum distinction requires a threshold that adapts locally to the fraction of outliers.

In a study on breakdown points, Hampel [15] proposed to identify outliers by their larger-than-normal variability, based on

$$|h^0 - \tilde{h}^0| \geq c \text{mad}\{h^0, \dots, h^k\}. \quad (12)$$

Here, $\text{mad}\{\}$ denotes the *median of absolute deviations* (MAD) of the sample. The parameter c can be optimized by simulation [16]. We chose c to minimize the average absolute error of the estimate (see section 4).




Figure 4: A height map of the sample piece obtained with the Bayesian algorithm at $112 \mu\text{m/s}$. Scale: $2.24 \mu\text{m}/\text{frame}$. Note the coarse height steps here.

3.3 Reference height map

No other measurement method delivers height maps that are directly comparable to those from WLI. In spite of first efforts to establish a correspondence between results from WLI and tactile devices [17], their relation is as yet unresolved for real-world surfaces. Instead, we recurred to a partially recursive argument by constructing a reference height map from repeated high-quality WLI measurements. To this end, we performed a series of 25 measurements under optimum conditions (balanced illumination, minimal vibrations, constant temperature and a slower scan speed of $2 \mu\text{m/s}$). The data of one recording was spoiled and had to be removed. The remaining time series were free of deterioration, according to visual inspection of the raw data.

We applied the sliding average of absolute finite differences as preprocessing algorithm, which is a recommended linear filtering algorithm [13], to produce height maps \hat{h}_l . Overall shifts of the maps were detected with the image-wise median value and then corrected by moving the height maps to a common reference height. Tilts and potential higher order deformations were not accounted for.

The pixelwise median height map h_{med}

$$h_{\text{med}} = \text{med}\{\hat{h}_1, \dots, \hat{h}_{24}\} \quad (13)$$

exhibits a small number of invalid pixels, as expected for rough surfaces. These




Figure 5: The reference height map.

sites could easily be detected by the large variation of height values reported from preprocessing. The median absolute deviation over 24 recordings gives a robust estimate of this variation:

$$\text{mad}\{\mathbf{h}_{\text{med}}\} = \text{med}\{|\hat{\mathbf{h}}_1 - \mathbf{h}_{\text{med}}|, \dots, |\hat{\mathbf{h}}_{24} - \mathbf{h}_{\text{med}}|\} \quad (14)$$

By heuristics, we marked those sites having a $\text{mad} \geq 30$ as defective and replaced their values by the median of the immediate neighborhood:

$$\mathbf{h}_{\text{ref}} = \begin{cases} \text{med}_{3 \times 3} \mathbf{h}_{\text{med}} & \text{for } \text{mad}\{\mathbf{h}_{\text{med}}\} \geq 30 \text{ fr.} \\ \mathbf{h}_{\text{med}} & \text{otherwise} \end{cases} \quad (15)$$

The reference map \mathbf{h}_{ref} is depicted in figure 5.

4 Results

Four algorithms were compared in the benchmarking:

1. Sliding average of absolute finite differences, the standard preprocessing for rough surface WLI,
2. the same sliding average, followed by a 3×3 median filter mask post-processing,
3. the sliding average followed by the adaptive 3×3 median filter mask (section 3.2), and

Table 1: Optimum parameters for the algorithms under comparison by the measure of eq. (16). The median filter has no parameter other than the mask size.

scanning speed	h_{\max}	pre-processing	adaptive median	Bayesian estimation	
		window size	threshold c	δ	q_0/q_1
14 $\mu\text{m/s}$	289 fr	9 fr	15.4	6 fr	10^{-2}
28 $\mu\text{m/s}$	144 fr	5 fr	10.7	4 fr	10^{-2}
56 $\mu\text{m/s}$	72 fr	3 fr	0.5	4 fr	10^{-4}
84 $\mu\text{m/s}$	49 fr	2 fr	0	5 fr	10^{-4}
112 $\mu\text{m/s}$	37 fr	2 fr	0	3 fr	10^{-4}

Table 2: Absolute error per pixel (\mathcal{E}_{pp}) for the algorithms under comparison. For the reasons given in section 4, the signal to noise ratio of the raw data is lowest for 84 $\mu\text{m/s}$.

scanning speed	preprocessing only	median filter	adaptive median	Bayesian estimation
14 $\mu\text{m/s}$	0.55 μm	0.69 μm	0.56 μm	0.45 μm
28 $\mu\text{m/s}$	0.71 μm	0.78 μm	0.72 μm	0.68 μm
56 $\mu\text{m/s}$	1.27 μm	1.02 μm	1.02 μm	0.89 μm
84 $\mu\text{m/s}$	4.84 μm	2.43 μm	2.43 μm	1.59 μm
112 $\mu\text{m/s}$	2.69 μm	1.64 μm	1.64 μm	1.17 μm

4. the Bayesian algorithm with a 3×3 neighborhood, as described in section 2.

The implementation of the sliding average algorithm and the setting for its parameter (size of sliding mean window) follows ref. [13].

The benchmarking data (see section 3.1) were obtained under real-world conditions with a suboptimal illumination, and sequences of 25 recordings were acquired during normal business hours in our laboratory.

We computed a scalar error estimate to optimize the respective parameters of all algorithms, which can be found in table 1, and to provide a coarse quality measure. The 25 height maps calculated with each algorithm were rescaled to the scale of the reference map and any overall height shift was removed by comparison of the image-wise median value with that of the reference map. The mean absolute difference per pixel \mathcal{E}_{pp} between the 25 estimated height maps and the reference map is not a robust estimator and thus penalizes left-over outliers:

$$\mathcal{E}_{\text{pp}} = \frac{1}{25} \cdot \frac{1}{|\mathcal{S}|} \cdot \sum_{l=1}^{25} \sum_{j=1}^{|\mathcal{S}|} |h_{\text{ref}}^j - \hat{h}_l^j| \quad (16)$$

Table 2 gives the results for optimized parameters. For 84 and 112 $\mu\text{m/s}$, the optimal adaptive median filter is *not* adaptive, the reasons being that over-smoothing does not occur with large interframe distances and that outliers be-




Figure 6: Distribution of errors in a height map obtained from preprocessing (see section 3) of interferograms acquired at $0.56 \mu\text{m}/\text{frame}$. Logarithmized occupancies of the histogram are plotted as a function of estimation error and variability of a height estimate in repeated trials, see text for details.

come frequent. Not surprisingly, the error generally increases with higher scan speeds, both due to the more difficult localization of the interference pattern in the subsampled signals, and the coarser discretization of permissible height values. Also, the higher density of erroneous pixels leads to more and larger clusters that cannot be corrected with the 3×3 filter masks used here.

The results for $84 \mu\text{m}/\text{s}$ are, however, particularly inaccurate. The reason is that this sampling rate corresponds almost exactly to four entire periods of the interferogram, thus leading to a weak response in a preprocessing based on finite differences. This is not the case with the data obtained at $112 \mu\text{m}/\text{s}$, which therefore leads to better height estimates.

For a more detailed analysis of the estimation error, we present 2D histograms in which the distribution of estimation errors is given as a function of error magnitude and the difficulty of obtaining a correct estimate for a pixel. The latter was measured by the pixelwise spread of the estimates \hat{h}_i obtained by preprocessing only. This spread was, in turn, quantified by the median absolute deviation (MAD, see section 3.2) of the estimates.

These histograms can be computed for each estimator, and we will refer to those as `histogram(preprocessing only)`, `histogram(median filter)`, etc.

Figure 6 shows `histogram(preprocessing only)`. The strong peak in the lower left corner indicates that the vast majority of pixels has a variability of about

5-10 frames and is estimated with an absolute error of 5-10 frames. The horizontal ridge contains all pixels of high MAD, which are nonetheless estimated accurately sometimes. The vertical ridge stems from pixels which give outlying height estimates in a number of measurements smaller than the break-point of the MAD-estimator.

For a direct comparison of the different methods we propose to inspect the difference of two histograms.

Figure 7 shows the small-error range of histogram(median filter) – histogram(preprocessing only). The median filter eliminates outliers; however, these are outnumbered by far by the remaining pixels and do not contribute significantly to the distribution shown. Instead, the figure shows the impact of the median filter on “ordinary” pixels, most of which are affected by the filtering operation, and mostly adversely: excessive smoothing leads to a deterioration of height estimates from the preprocessing, as manifested by the large depletion and accumulation, respectively, in the absolute error ranges around 4 and 9 frames. In summary, the median filter smoothes in regions in which it is not necessary and thus “destroys” the microstructure of the surface.

Inspection of histogram(median) (not shown) reveals that the bulk of the distribution resides in an error range of 3 – 5 frames, which hints at a surface roughness of the magnitude of $1.5 - 3 \mu\text{m}$ on the scale of 3×3 pixels, that is, on a horizontal scale of about $100 \mu\text{m}$.

The adaptive median filter based on Hampel’s outlier detector reduces over-smoothing; the effect is already visible in table 2 for slow scanning speeds, and is also pronounced in histogram(adaptive median filter) – histogram(preprocessing), see figure 8.

Note that the color scales in figures 7 through 9 are different: the adaptive median filter changes less height values than the median filter.

Figure 9 shows that Bayesian surface estimation introduces the least bias. Since the histograms are sparsely populated for large MAD-values, only a summary analysis is reasonable:

Figure 10 shows the (cumulative) “marginal distribution” for $\text{MAD} \geq 20$ frames (it is not a proper distribution since the histogram difference images attain values below zero and integrate to zero). The profiles show that Bayesian denoising reduces errors also in the high-variability regime.

5 Conclusion

Bayesian inference can accommodate both experimental observations and smoothness assumptions in the estimation process, without imposing a pipeline structure in which the postprocessing loses sight of the original data. If the smoothness assumptions can be expressed in a simple prior on the local height configurations, full Bayesian inference without recourse to stochastic sampling is possible.

The procedure introduced in section 2 gives estimates that are more accurate than the (adaptive) median filtering we compare to, in the sense of minimizing the average absolute error. The difference in accuracy is more pronounced in the case of weak signal to noise ratios and fast scanning speeds, cf. the last two lines in table 2. The computational cost of the method presented here is of the order of seconds and increases with the scan length h_{max} .

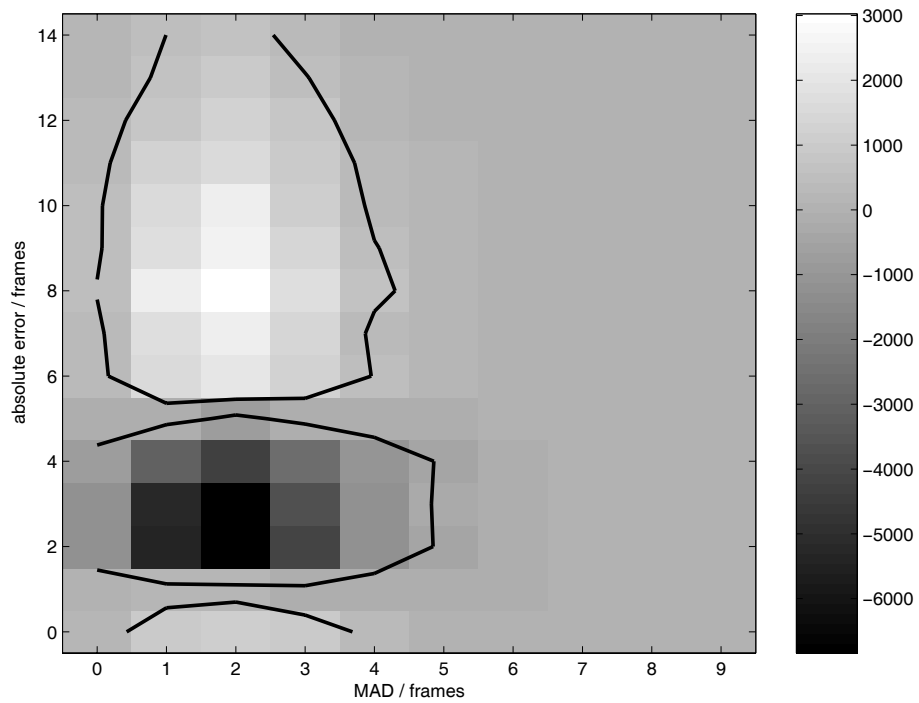


Figure 7: Error histogram of height maps obtained from standard median filtering minus error histogram of height maps obtained from preprocessing. In short: $\text{histogram}(\text{median filter}) - \text{histogram}(\text{preprocessing})$. Scale: $0.56 \mu\text{m}/\text{frame}$.




Figure 8: Histogram(adaptive median filter) – histogram(preprocessing).
Scale: $0.56 \mu\text{m}/\text{frame}$.

In summary, if the measurements are of high quality and a simple post-processing is desired, we recommend the adaptive median filter. For a more accurate description of the microtopology, or if the input data is of low quality, Bayesian surface estimation appears a better choice.

Acknowledgment

The authors would like to thank Chr. Hennig (Seminar for Statistics, ETH Zürich) for pointing out the possibility of using the Hampel detector in an adaptive filtering scheme. Support from Robert Bosch GmbH is gratefully acknowledged.

References

- [1] Thomas Dresel, Gerd Häusler, and Holger Venzke. Three-dimensional sensing of rough surfaces by coherence radar. *Appl. Opt.*, 31(7):919–925, 1992.
- [2] James C. Wyant. How to extend interferometry for rough-surface tests. *Laser Focus World*, 29(9):131–135, 1993.
- [3] Peter de Groot and Leslie Deck. Three-dimensional imaging by sub-Nyquist sampling of white-light interferograms. *Opt. Lett.*, 18(17):1462–1464, 1993.




Figure 9: Histogram(Bayesian estimation) – histogram(preprocessing).
Scale: $0.56 \mu\text{m}/\text{frame}$.

- [4] David J. Aziz. Interferometric measurement of surface roughness in engine cylinder walls. *Opt. Eng.*, 37(5):1429–1434, 1998.
- [5] Peter de Groot and Xavier Colonna de Lega. Valve cone measurement using white light interference microscopy in a spherical measurement geometry. *Opt. Eng.*, 42(5):1232–1237, 2003.
- [6] D. W. Goodman. *Statistical Properties of Laser Speckle Patterns*. 2nd edition, 1984.
- [7] Niels Væver Hartvig and Jens Ledet Jensen. Spatial mixture modeling of fMRI data. *Hum. Brain Mapping*, 11:233–248, 2000.
- [8] Akiko Harasaki, Joanna Schmit, and James C. Wyant. Improved vertical-scanning interferometry. *Appl. Opt.*, 39(13):2107–2115, 2000.
- [9] Min-Cheol Park and Seung-Woo Kim. Direct quadratic polynomial fitting for fringe peak detection of white light scanning interferograms. *Opt. Eng.*, 39(4):952–959, 2000.
- [10] Stuart Geman and Donald Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. 6(6):721–741, 1984.
- [11] Gerhard Winkler. *Image Analysis, Random Fields and Markov Chain Monte Carlo Methods*. Springer, Berlin, Heidelberg, New York, 2nd edition, 2003.




Figure 10: Horizontal integration (sum and cumulative sum) of histogram figure 9 for MAD > 20 frames. Scale: $0.56 \mu\text{m}/\text{frame}$.

- [12] Gunther Bohn. *Hardware-implementierte Algorithmen zur Optimierung des Meßprinzips Kohärenzradar*. PhD thesis, Univ. Erlangen-Nürnberg, 2000.
- [13] Joachim Schraud. *Optimierung der Signalaufnahme und Signalverarbeitung am optischen 3D-Sensor Kohärenzradar*. Diploma thesis, Univ. Erlangen-Nürnberg, 2000.
- [14] David L. Donoho and Peter J. Huber. *The Notion of Breakdown Point*, pages 157–184. Wadsworth Intl., Belmont, Cal., 1983.
- [15] Frank R. Hampel. The breakdown points of the mean combined with some rejection rules. *Technometrics*, 27(2):95–107, 1985.
- [16] Laurie Davies and Ursula Gather. The identification of multiple outliers. *J. Amer. Statist. Assoc.*, 88(423):782–801, 1993.
- [17] Robert Windecker and Hans J. Tiziani. Optical roughness measurements using extended white-light interferometry. *Opt. Eng.*, 38(6):1081–1087, 1999.