euspen's 26th International Conference & Exhibition, Krakow, PL, June 2026
www.euspen.eu

Dense Data Fusion and Uncertainty Quantification in Optical Metrology: The Dust3r++ Framework

Ladji Fofana[1, 2], Katarina Josic[1, 2], Louis-Ferdinand Lafon[1], Charyar Mehdi-Souzani[3], Nabil Anwer[2], Hichem Nouira[1]

[1] *Laboratoire national de métrologie et d'essais/Conservatoire national des arts et métiers, 1 Rue Gaston Boissier, 75015 Paris, France*
[2] *Université Paris-Saclay, ENS Paris-Saclay, LURPA, 91190, Gif-sur-Yvette, France*
[3] *Université Paris-Saclay, ENS Paris-Saclay, LURPA, Université Sorbonne Paris Nord, 91190, Gif-sur-Yvette, France*

ladji.fofana@lne.fr

Abstract

Dense data fusion has become a cornerstone of optical dimensional metrology, but the reliability of reconstructed 3D models is still obstructed by measurement noise, changing acquisition conditions, and the absence of systematic uncertainty evaluation. In this work, we present Dust3r++, a dense fusion framework that couples Optimal Transport (OT) with non-rigid warping to merge multi-view data in a stable and consistent manner. A distinctive feature of Dust3r++ is its uncertainty propagation strategy: uncertainties on fused correspondences are inferred from OT dual variables, warp parameter covariances are obtained from the Hessian of the fitting objective, and both are propagated locally to the reconstructed surface. This leads to dense 3D reconstructions enriched with spatially resolved confidence bounds, bringing the results in line with the principles of the GUM. Tests carried out on several industrial reference parts show that Dust3r++ delivers higher robustness and accuracy than state-of-the-art approaches, while also providing, for the first time, a practical framework for dense uncertainty quantification in image-based 3D inspection. The approach lays the groundwork for traceable and real-time optical metrology within Industry 4.0.

Keywords: Dense data fusion, Optimal Transport (OT), Non-rigid warping, 3D reconstruction, Optical dimensional metrology, Uncertainty propagation, GUM, Spatial confidence bounds, Industrial inspection, Industry 4.0

1. Introduction

Multi-view photogrammetry is increasingly used in dimensional metrology as a non-contact measurement technique for automated inspection in manufacturing [1]. By combining measurements acquired from multiple viewpoints, it allows the estimation of the geometry of complex mechanical parts with a high spatial sampling density [2], beyond what can be achieved from isolated point-cloud measurements. Such systems are well suited to in-line inspection scenarios in the context of Industry 4.0. In this work, we consider an experimental setup developed at LNE (Laboratoire National de Metrology et d'Essais) for stereo-based optical dimensional measurement. The acquisition system consists of an industrial robot ensuring repeatable positioning of a calibrated stereo rig for passive image acquisition, complemented by a structured-light sensor used for reference measurements (Figure 1).

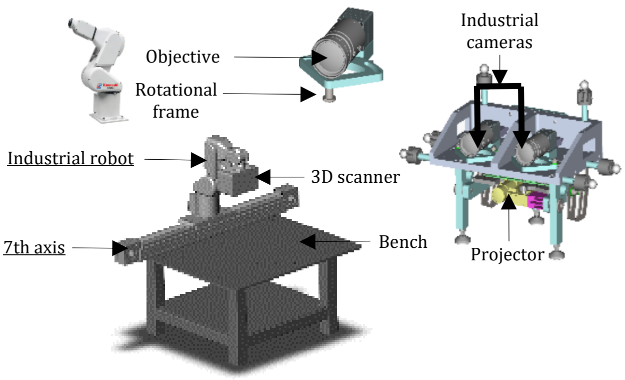

Figure 1. Structured light 3D stereovision scanning system

Despite recent progress, three limitations remain for metrological use [3]: (i) sensitivity to measurement noise and varying acquisition conditions, (ii) susceptibility to outliers and calibration errors in classical correspondence-based fusion [3], and (iii) the lack of uncertainty evaluation and propagation in dense reconstruction pipelines [2][10][11]. To overcome these limitations, this paper introduces Dust3r++, a dense data-fusion framework based on OT and non-rigid warping, providing reconstructed surfaces associated with spatially resolved measurement uncertainties in accordance with the GUM.

2. Dust3r++ dense fusion method

We propose a traceable dense data fusion and 3D reconstruction framework called Dust3r++, which has been developed by the authors and is currently being used by LNE for passive optical imaging. This research framework is not a commercial software product. It is designed to ensure reproducibility, reliability and metrological traceability of the multi-view reconstruction process.

2.1. Measurand definition

The measurand, in the sense of the Vocabulary of Metrology (VIM), is defined as the true surface geometry of the observed object [4], expressed in a common three-dimensional coordinate system. In this work, the geometry of the surface is represented by a dense collection of three-dimensional points, reconstructed from several acquisitions. Therefore, the measurement process is to estimate this surface geometry from the observations derived from the image. In common photogrammetric frameworks, surface geometry is parametrically defined implicitly by dense point or depth estimates obtained from the corresponding images [5]. These quantities are related to the physical geometry by means of calibrated imaging models and are used for the reconstruction of a three-dimensional surface [6].

2.2. Optical Measurement model and input quantities

An optical measurement model implements a probabilistic estimation process by treating observed geometric values X and corresponding ideal geometric values Y as two random variables that are distributed according to an unknown data distribution [7]. Common state-of-the-art approaches use a pre-trained predictor like a vision transformer (ViT) as measurement model [8]. By considering a measurement model g that is obtained during the learning phase by minimising the empirical risk for a large image pairs dataset, g^* minimises the expected risk for the underlying data distribution. In a regression setting, uncertainty is represented by a Bayesian risk formulation [9]. The total uncertainty associated with the estimate at a given observation x is defined as the conditional expected squared error:

$$U_t(x) = E[(\hat{g}(X) - Y)^2 \mid X = x] \text{ (1) [10]}$$

The data (average) uncertainty is related to the internal variability of the measurement process and given by $U_d(x) = E[(g^*(X) - Y)^2 \mid X = x]$. Since the optimal estimator satisfies $g^*(X) = E[Y \mid X]$, the data uncertainty is reduced to the conditional variance $Var[Y \mid X = x]$, which depends only on the measurement conditions. The model uncertainty is then defined as the difference $U_m(x) = U_t(x) - U_d(x)$ and reflects the imperfection of the pretrained measurement model with respect to the ideal estimator [13].

2.3. Estimation of the surface geometry

The estimation process is purely probabilistic [7] and is formulated by modelling the visual observations X and the corresponding geometric ideal Y as two randomly-distributed random variables. The pre-trained estimator $\hat{g}$ is obtained by minimising the empirical risk for the training set, while the ideal g^* minimises the Bayesian risk for the underlying dataset [9]. For dense image fusion, the overall uncertainty evaluation process from image acquisition to surface geometry definition is summarised in the optical measurement chain illustrated in (Figure 2).

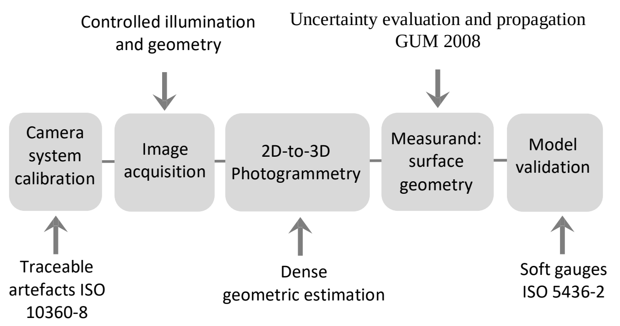

Figure 2. Optical Measurement Chain for Dense 3D Reconstruction with GUM-Based Uncertainty Propagation [11]

The reconstructed surface geometry constitutes the measurement result. In accordance with GUM, this result must be associated with an evaluated measurement uncertainty characterizing the dispersion of values that could reasonably be attributed to the measurand. As illustrated in (Figure 2), uncertainty evaluation and propagation are therefore an integral part of the measurement chain, enabling uncertainty-aware quantification of 3D dense reconstructions for dimensional inspection [10][11].

3. Dust3r++ dense fusion model

Dust3r++ is an extension of the Dust3r (Dense Unconstrained Stereo Reconstruction) [12] measurement model and estimation principles described in (section 2). The goal is to estimate a consistent representation of the geometry measurand from dense, multiple-view geometric observations, explicitly taking into account sources of uncertainty that arise during the fusion process. As shown in (Figure 3), Dust3r++ optimizes the pre-trained state-of-the-art dense image fusion model (Dust3r) with an OT based correspondence formulation that is also refined through non-rigid warping [13].

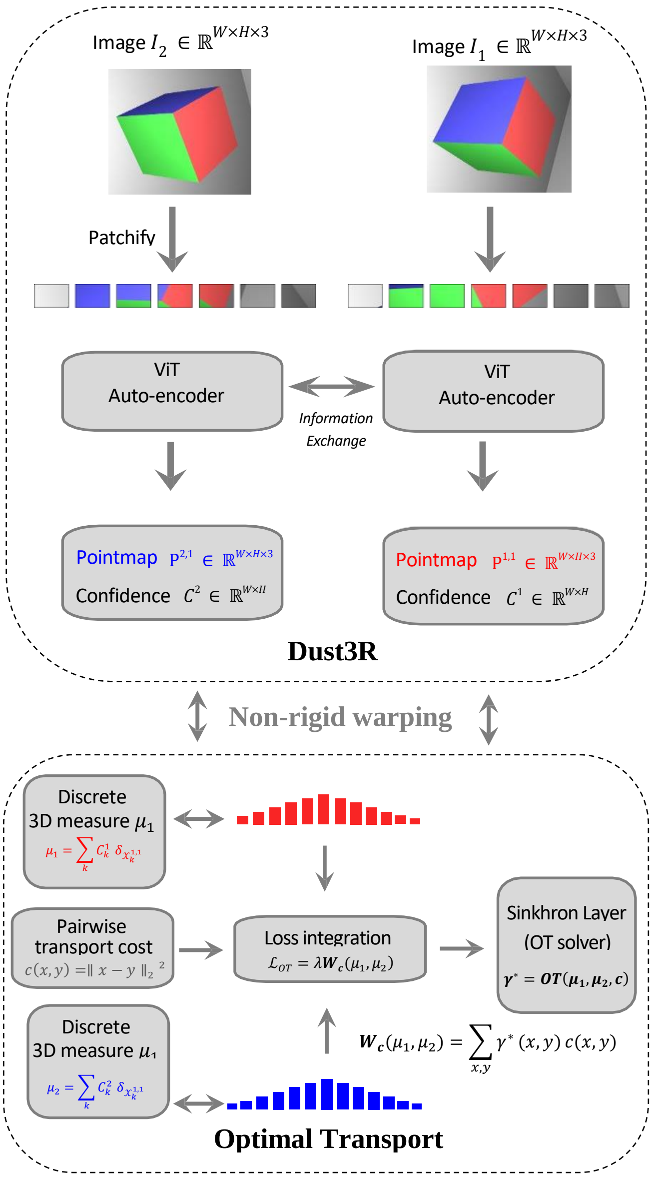

Figure 3. Dust3r++ Overview: Dust3r Features Coupled with Optimal Transport (Sinkhorn) Fusion

3.1. Optimal transport-based correspondences

For each view I, the pretrained dense image fusion model produces a dense point map P^i, and an associated confidence map $C^{i,k}$. Each view is represented by a discrete geometric measure defined by:

$$\mu_{\mathrm{i}} = \sum_k C_k^{\mathrm{i}} \ \delta_{\chi_k^{\mathrm{i},1}} \quad (2)$$

where δ denotes the Dirac measure located at the estimated point positions. Dense multi-view fusion is formulated as the OT problem between these discrete measures. Given two views i and j, the transport cost between points $P^{i,k}$and $P^{j,k}$ is defined as the squared Euclidean distance $\left\|P^{i,k} - P^{j,k}\right\|^2$. The optimal transport plan γ^* is obtained by minimising the total transport costs under the marginal constraints imposed by μ_i and μ_j. This global formulation provides reliability for outliers and missing correspondences, as geometrically inconsistent points naturally have a low transport cost, without requiring an explicit exclusion of outliers. The OT formulation also gives rise to two variables related to the optimum solution. In Dust3r++, these two variables are interpreted as indicators of correspondence, which reflects the sensitivity of transport costs to changes in input measures. They are therefore the first source of uncertainty information associated with dense correspondences.

3.2. Non-rigid warping

Despite the global enforcement of correspondence by OT, residual geometrical differences remain due to measurement noise, calibration residuals and imperfections in the pre-trained fusion model [3]. We therefore use a non-dynamic warp model as a corrective mechanism in the estimation process. Let p_k denote a fused 3D point obtained after the OT step, and let $W(p_k;\ \theta)$ be a parametric warping function with parameters θ. The parameters θ are estimated by minimizing a least-squares objective of the form $\Sigma_{\mathrm{k}}\|W(p_k;\ \theta) - \bar{p}_{\mathrm{k}}\|^2$, enforcing geometric consistency in a common reference frame. The uncertainty associated with the correction model is quantified by the covariance of the estimated parameters, given by $Cov(\theta) = H^{-1}(\theta^*)$, where H is the Hessian of the objective at the optimum value. This uncertainty is combined with correspondence-related uncertainty inferred from the OT formulation and propagated to the reconstructed surface geometry.

4. Uncertainty propagation to the reconstruction surface

The reconstructed surface geometry is the measurement result and associated with an evaluated measurement uncertainty in accordance with the GUM. We denote by $\hat{y}$ a given reconstructed surface point obtained from the fused correspondences and the non-rigid correction model. The uncertainty associated with $\hat{y}$ results from two contributions: the uncertainty on the dense correspondences obtained through the OT formulation, and the uncertainty on the estimated warp parameters θ. Uncertainty propagation is performed locally by linearizing the surface reconstruction model with respect to these quantities. Therefore, the covariance of the reconstructed surface point is expressed as:

$$Cov(\hat{y}) = J_p Cov(p) J_p^T + J\theta\ Cov(\theta) J\theta^T \quad (3)$$

where p represent the correspondence-related quantities, $Cov(p)$ represents their covariance inferred from the OT formulation, $Cov(\theta)$ is the covariance of the warp parameters obtained from the inverse Hessian of the fitting objective, and J_p and $J\theta$ are the corresponding Jacobians. This formulation yields a spatially resolved uncertainty associated with the surface geometry, corresponding to a model-based (Type A) uncertainty evaluation consistent with the GUM.

5. Experiment and validation

In this work, the experimental validation is conducted on 16 real-world automotive data measured on the LNE optical metrology test bench. These mechanical parts exhibit complex geometries, varying surface orientations, and limited texture. For clarity, a brake caliper is shown here as an example (Figure 4) illustrating the way we represent image pairs for multi-view data fusion.

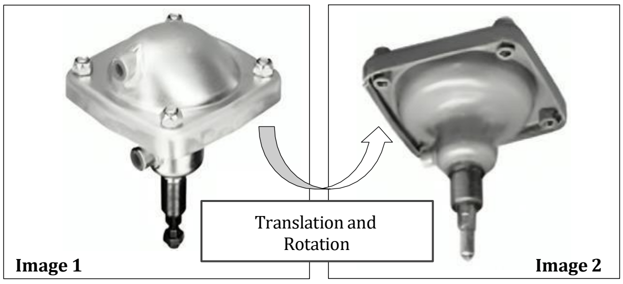

Figure 4. Image Pair Configuration for Dense Fusion of a Brake Caliper

5.1. Dense fusion performances

The robustness of dense fusion models is evaluated by using standard image fusion metrics (reprojection error RE, inlier ratio IR and F1-score). The RE quantifies the geometric consistency of the estimated correspondences, while the IR measures their reliability with respect to outliers and missing data. The two quantities are combined by their harmonic average to give the F1 score, which is a compact indicator of reconstruction accuracy.

Table 1 Dense Fusion Performance Metrics (RE, IR, Chamfer + deviations)

Algorithms	Mean inlier ratio $\bar{\tau}$ (%)	Mean reprojection error $\bar{e}$ (pixel)	Mean Chamfer distance $\bar{d} \pm \Delta\bar{d}$ (µm)
Dust3r++	**91**	$\mathbf{1.45}$	$\mathbf{245.63 \pm 1.42}$
Dust3r [12]	88	1.50	287.57 ± 5.05
Mast3r [14]	86	1.50	344.89 ± 6.16
RoMa [15]	74	1.85	352.08 ± 10.68

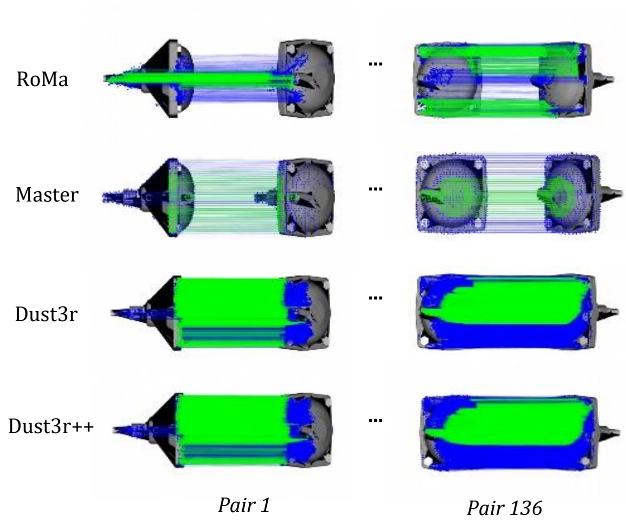

Figure 5. Dust3r++ vs State-of-the-Art Dense Matching on 136 pairs

Compared to state-of-the-art methods, our Dust3r++ method maintains stable reprojection errors and significantly higher inlier ratios across distant viewpoints, demonstrating improved robustness of the dense corfrespondence estimation. This robustness is consistent with the global OT formulation, which naturally downweights inconsistent correspondences instead of relying on hard rejection rules.

5.2. Dust3r++ vs 3D reconstruction benchmarks

Surface reconstruction errors on the brake caliper are evaluated against a traceable CMM reference (ZEISS CONTURA G2 7.10.6; 700 × 1000 × 600 mm³; MPE_E = 1.8 + L/300 µm, MPE_P = 1.8 µm, ISO 10360) [16]. Dust3r++ is compared with COLMAP, AliceVision and OpenMVG/MVS [17] under identical evaluation conditions. Results are reported in Table 2.

Table 2 Comparative reconstruction errors with respect to a CMM

Method	Mean Signed Dist $\bar{d}$(mm)	Mean Std Dev $\bar{\sigma}$(mm)
COLMAP	4,35	0,77
AliceVision	3,30	0,27
OpenMVG/MVS	3,29	0,80
Dust3r++	0,25	0,18
CMM (reference)	0	-

Dust3r++ achieves lower errors and reduced variability, and the reconstructed geometry is closer to the CMM reference, which indicates that it is suitable for dimensional metrology applications below 300 µm using optical systems with a dense image matching algorithm. Figure 6 shows the spatial distribution of signed point-to-surface deviations obtained with Dust3r++.

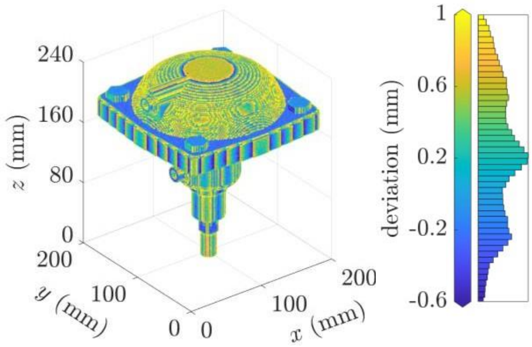

Figure 6. Dust3r++ Reconstruction Error Maps

6. Conclusion and perspectives

In this work, we introduced Dust3r++, a dense multi-view fusion framework with uncertainty-aware surface reconstruction. The proposed approach combines Optimal Transport-based correspondence estimation with a non-rigid warping model to allow a reliable reconstruction of complex industrial geometry, while explicitly propagating uncertainty to the geometry measurand in accordance with the GUM. Experimental validation on a real-world industrial part demonstrates high reconstruction accuracy and lower uncertainty values compared to state-of-the-art matching approaches and 3D reconstruction benchmarks, with a spatial distribution consistent with the observed reconstruction errors. The results contribute to bridging the gap between dense 3D reconstruction and traceable dimensional inspection. Future work will focus on addressing in-line implementation, real-time constraints and extending the dissemination of uncertainty to dimensional properties.

Acknowledgment

This work is part of the project 23IND08 DI-Vision, which has received funding from the European Partnership on Metrology, co-financed from the European Union's Horizon Europe Research and Innovation Programme and by the Participating States.

References

[1] Alberto Mendikute, José Yagüe-Fabra, Mikel Zatarain, Álvaro Bertelsen, and Ibai Leizea. Self-calibrated in-process photogrammetry for large raw part measurement and alignment before machining. Sensors, 17:2066, 09 2017. doi: 10.3390/s17092066.

[2] Shuai Liu, Mengmeng Yang, Tingyan Xing, and Ran Yang. A survey of 3d reconstruction: The evolution from multi-view geometry to nerf and 3dgs. Sensors, 25:5748, 09 2025. doi: 10.3390/s25185748.

[3] Young Kim, Christian Theobalt, James Diebel, J. Kosecka, Branislav Micusik, and Sebastian Thrun. Multi-view image and tof sensor fusion for dense 3d reconstruction. pages 1542 - 1549, 11 2009. doi: 10.1109/ICCVW.2009.5457430.

[4] Philip Potts. Glossary of analytical and metrological terms from the international vocabulary of metrology (2008). Geostandards and Geoanalytical Research, 36, 09 2012. doi: 10.1111/j.1751-908X.2011.00122.x.

[5] Yangming Lee. Three-dimensional dense reconstruction: A review of algorithms and datasets. Sensors, 24:5861, 09 2024. doi: 10.3390/s24185861.

[6] Michael A., Jean-José Orteu, and Hubert Schreier. Volumetric Digital Image Correlation (VDIC), pages 1-16. 02 2009. ISBN 978-0-387-78746-6. doi: 10.1007/978-0-387-78747-38.

[7] Wojciech Toczek and Janusz Smulko. Risk analysis by a probabilistic model of the measurement process. Sensors, 21:2053, 03 2021. doi: 10.3390/s21062053.

[8] Ruiming Jia, Xin Chen, Jiali Cui, and Zhenghui Hu. Mvs-t: A coarse-to-fine multi-view stereo network with transformer for low-resolution images 3d reconstruction. Sensors, 22:7659, 10 2022. doi: 10.3390/s22197659.

[9] Harvey Mitchell. Bayesian Decision Theory, pages 311-331. 11 2025. ISBN 978-3-662-71022-7. doi: 10.1007/978-3-662-71023-414.

[10] Eyke Hüllermeier and Willem Waegeman. Aleatoric and epistemic uncertainty in machine learning: an introduction to concepts and methods. Machine Learning, 110, 03 2021. doi: 10.1007/s10994-021-05946-3

[11] Miroslaw Wojtyla and Wojciech Plowucha. Measurement models - theory and practice of uncertainty evaluation of geometrical deviation measurements. Metrology and Measurement Systems, pages 1-21, 11 2025. doi: 10.24425/mms.2025.154671

[12] Wei Zhang, Yihang Wu, Songhua Li, Wenjie Ma, Xin Ma, Qiang Li, and Qi Wang. Review of feed-forward 3d reconstruction: From dust3r to vggt, 07 2025.

[13] Zhengyu Su, yalin wang, Rui Shi, Wei Zeng, Jian Sun, Feng Luo, and Xianfeng Gu. Optimal mass transport for shape matching and comparison. IEEE Transactions on Pattern Analysis and Machine Intelligence, 37, 11 2015. doi: 10.1109/TPAMI.2015.2408346.

[14] Vincent Leroy, Yohann Cabon, and Jerome Revaud. Grounding image matching in 3d with mast3r, 06 2024.

[15] Johan Edstedt, Qiyu Sun, Georg Bökman, Mårten Wadenbäck, and Michael Felsberg. Roma: Revisiting robust losses for dense feature matching, 05 2023.

[16] R. Koteras, M. Wieczorowski, P. Znaniecki, Acceptance and reverification of cmm in industrial conditions, Advances in Science and Technology Research Journal 12 (2018) 80–88. doi:10.12913/22998624/80987.

[17] E. Stathopoulou, M. Welponer, F. Remondino, Open-source image-based 3d reconstruction pipelines: Review, comparison and evaluation, The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences XLII-2/W17 (2019) 331–338. doi:10.5194/isprs-archives-XLII-2-W17-331-2019.