



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

EndoMetric: Near-Light Monocular Metric Scale Estimation in Endoscopy

Raúl Iranzo^{*[0009-0000-6726-923X]}, Víctor M. Batlle^{*[0000-0002-6837-934X]}, Juan D. Tardós^[0000-0002-4518-5876], and José M. M. Montiel^[0000-0002-3627-7306]

Instituto de Investigación en Ingeniería de Aragón (I3A), Universidad de Zaragoza,
María de Luna 1, 50018 Zaragoza, Spain.

{riranzo,vmbatlle,tardos,josemari}@unizar.es

Abstract. Geometric reconstruction and SLAM with endoscopic images have advanced significantly in recent years. In most medical fields, monocular endoscopes are employed, and the algorithms used are typically adaptations of those designed for external environments, resulting in 3D reconstructions with an unknown scale factor.

For the first time, we propose a method to estimate the real metric scale of a 3D reconstruction from standard monocular endoscopic images, under unknown varying albedo, without relying on application-specific learned priors. Our fully model-based approach leverages the near-light sources embedded in endoscopes, positioned at a small but nonzero baseline from the camera, in combination with the inverse-square law of light attenuation, to accurately recover the metric scale from scratch. This enables the transformation of any endoscope into a metric device, which is crucial for applications such as measuring polyps, stenosis, or assessing the extent of diseased tissue.

Keywords: Endoscopy · Metric scale · Size estimation

1 Introduction

In current endoscopic procedures, endoscope navigation, localization, and tissue measurement are performed manually. Recent advances in Visual Simultaneous Localization and Mapping (VSLAM) for endoscopy [17,18,10] offer the promise of live 3D reconstructions, that will enable autonomous or assisted navigation and robotized interventions. Most specialties use just monocular endoscopes to reduce bulk and cost. However, using a moving monocular camera, the absolute scale of the environment is unobservable, and the 3D reconstructions and trajectories obtained have an unknown scale factor. This also introduces scale drift which significantly reduces map accuracy.

However, endoscopes are equipped with light sources attached to the camera, which introduce significant illumination variations in the scene. Our key insight is to leverage these illumination changes through near-light photometry to accurately recover the true metric scale of monocular reconstructions. Photometry

* These authors contributed equally to this work.

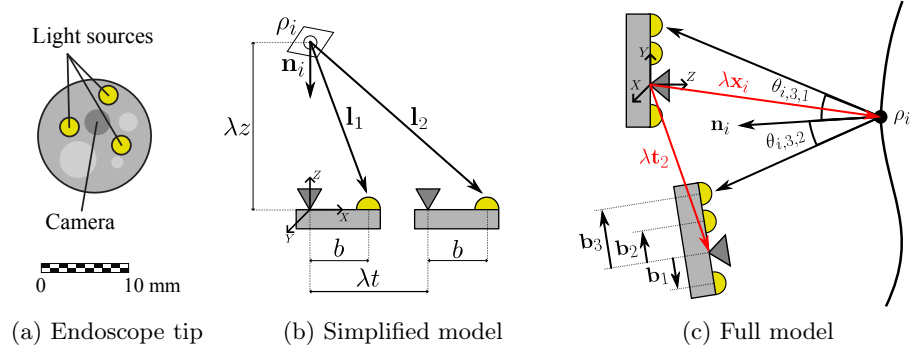


Fig. 1. Our method estimates the metric scale factor λ by leveraging a near-light illumination model applied to multi-view images captured with a monocular endoscope.

is scale-dependent due to two factors: the inverse-square decay of illumination with distance from the light source, and the angle between the incident light and the surface normal when the light source is positioned at a small, but nonzero, baseline from the camera’s optical center (Figure 1).

We demonstrate how true scale can be recovered just from the images captured by a standard monocular endoscope in two steps: first, obtaining an up-to-scale reconstruction with SfM or VSLAM, and then performing photometric optimization to recover scale, gains, and per-point albedo. Our contributions are:

- A scale-dependent near-light photometric model applicable to any monocular, up-to-scale multi-view reconstruction.
- A photometric optimization method to estimate true metric scale, per-point albedo, and camera gain.
- An initialization technique to enhance convergence and avoid local minima.
- Simulations and real experiments demonstrating the scale accuracy achievable in endoscopy.

2 Related work

Feature-based monocular SLAM [4,7] and SfM [22] recover up-to-scale geometry via bundle adjustment, ignoring photometry. Photometric monocular SLAM [8,27] estimates geometry, albedo, and camera gains but assumes constant illumination and cannot recover scale. In contrast, we first estimate up-to-scale geometry via bundle adjustment and then determine metric scale, albedo, and camera gain using non-linear optimization of near-light photometric errors.

Our method relates to photometric stereo, originating from [24], which used orthographic images from a fixed camera with distant, switchable light sources to recover surface normals, but not scale. The first method to recover metric scale was [15], leveraging inverse-square illumination decay from multiple point lights at known positions. Near-field photometric stereo, where the camera and

light sources are close to the object, was formalized in [19], with semi-calibrated approaches introduced by [20], requiring known light positions but unknown intensities. Using an endoscope for calibrated near-light photometric stereo was first explored in [5], combining three colored light sources into a single shot for true-scale reconstruction. In summary, near-light photometric stereo can generate real-scale reconstructions from three or more images with different point light sources. However, for endoscopy, the device must be modified to control the lights or use colored lights, which is impractical in clinical settings.

Other hardware modifications use structured light to overlay a metric scale on the image for polyp measurement [26,21]. Some works train a deep network to classify polyps into two size classes (smaller or larger than 10 mm) [14,13] or to predict dense depth with true scale for polyp measurement [6,23]. However, these methods require task-specific learning and do not generalize to other medical applications. In contrast, we achieve metric-scale reconstruction and estimate the camera’s metric trajectory using a standard endoscope without hardware modifications or task-specific learning, relying solely on near-light photometry.

Near-light shape-from-shading was used in [25] to recover a metric-scale reconstruction from a single perspective image, but relying on strong assumptions of known light intensity and a constant, known albedo. Similar methods [11,3] assumed that the light source is located at the optical center, i.e., the baseline is zero. As a result, it becomes impossible to disambiguate scale from albedo. This yields only up-to-scale reconstructions due to three unknown factors: illumination power, camera gain, and surface albedo. Multi-view near-light photometry to jointly recover scale, gain, and per-point albedo was first proposed in [9], but only validated through simplistic simulations with up-to-scale ground truth geometry. We take a step further by demonstrating that the method can work from the near-light images alone in realistic simulations and in real colonoscopies.

3 Fundamentals

To analyze some properties of near-light photometry we first consider a simplified two-view problem [9] (Figure 1b). It assumes a moving monocular camera with a single point light source at a distance b from the optical center, observing a Lambertian point with albedo ρ , which lies along the camera’s optical axis at a depth λz , with its normal pointing towards the camera. The second camera is translated λt to the right. We aim to compute the unknown scale λ .

Assuming the light intensity L_0 and camera gain g are constant, and no gamma compression, the intensity of the point i in each image k will be [15]:

$$I_{i,k}(\lambda) = \frac{\rho_i g L_0}{\pi} \frac{\mathbf{n}_i \cdot \mathbf{l}_k}{\|\mathbf{l}_k\|^3} = \frac{\rho'_i}{\pi} \frac{\mathbf{n}_i \cdot \mathbf{l}_k}{\|\mathbf{l}_k\|^3} \quad (1)$$

where $\rho'_i = \rho_i g L_0$ is a scaled albedo. The intensities in both images are:

$$I_{i,1} = \frac{\rho'_i}{\pi} \frac{\lambda z}{(b^2 + \lambda^2 z^2)^{3/2}} \quad , \quad I_{i,2} = \frac{\rho'_i}{\pi} \frac{\lambda z}{((\lambda t + b)^2 + \lambda^2 z^2)^{3/2}} \quad (2)$$

We can eliminate ρ'_i to get a second-order equation on λ :

$$b^2 + \lambda^2 z^2 = c \left((\lambda t + b)^2 + \lambda^2 z^2 \right) \quad (3)$$

where $c = (I_{i,2}/I_{i,1})^{2/3}$ is a known constant determined from the intensities measured in the images. This equation allows to obtain λ , except when $b = 0$, in which case λ simplifies away, and cannot be solved.

We can conclude that in a multi-view near-light scenario, metric scale is observable only when the baseline between camera and light source is non-zero, even if light intensity, gain and albedo are unknown. Also, we can expect scale accuracy to degrade when the distance to the surface is too large compared with the camera-light baseline. The remainder of the paper proposes a practical method for obtaining the scale in real endoscopic settings and studies its accuracy.

4 EndoMetric

Our proposed method, EndoMetric, works in two steps: first obtain an *up-to-scale multi-view reconstruction* of the scene using any SLAM or SfM method, and then solve an optimization problem for *metric scale estimation* that impose albedo consistency. The key for scale estimation is leveraging a *near-light photometric model* that takes into account the baseline of the light sources with respect to the camera and the inverse-square law of illumination decline with distance.

4.1 Up-to-scale Multi-view Reconstruction

Classical multiview geometry can produce, from a sequence of calibrated monocular images (at least two), the geometry of n scene points $\{\mathbf{x}_i\}_{i=1}^n$ and m camera poses $\{\mathbf{T}_k = (\mathbf{R}_k, \mathbf{t}_k)\}_{k=1}^m$ up to an unknown scale factor λ , by solving bundle adjustment, i.e. the non-linear optimization of the re-projection errors of the points matched along the sequence. We use the well known COLMAP [22] to compute multiview geometry from endoscopic sequences. Photometric models need not only the sparse scene geometry, but also the surface normals at each scene point. We propose to compute the normals \mathbf{n}_i by fitting a plane to the p neighbors of each scene point.

4.2 Near-light Photometric Model

We assume a calibrated mobile camera with r fixed point light sources at known positions relative to the optical center $\{\mathbf{b}_j\}_{j=1}^r$. All lights share the same intensity L_0 , uniformly distributed in all directions, with a calibrated lens vignetting. Points near specular reflections are discarded, and the surface is assumed to be Lambertian with an unknown, varying albedo.

We have m grayscale images $\{I_k\}_{k=1}^m$ taken from m poses while observing n scene points. For a given point i in an image k the image intensity depends on

the point albedo ρ_i , the camera gain g_k and the light intensity L_0 . The observed albedos are always multiplied by the camera gain, g_k , and L_0 . As these values are not provided by currently available endoscopes, we define two new variables that are observable from the images:

$$\rho'_i = \rho_i g_1 L_0 \quad , \quad g'_k = \frac{g_k}{g_1} \quad (4)$$

where ρ'_i is a scaled albedo (that may be greater than 1) and g'_k represents the gain change with respect to the first image.

Illumination depends on the incidence angle and the inverse-square of the distance between light and point. Then, the perceived radiance depends on the unknown scale factor λ of the multi-view reconstruction as (Figure 1c):

$$\mathcal{I}_{i,k}(\rho'_i, g'_k, \lambda) = \left(\frac{\rho'_i g'_k}{\pi} \sum_{j=1}^r \frac{\cos \theta_{i,j,k}(\lambda)}{\|\lambda \mathbf{x}_i - (\mathbf{R}_k \mathbf{b}_j + \lambda \mathbf{t}_k)\|^2} V(\mathbf{x}_i) \right)^{1/\gamma} \quad (5)$$

where $\cos \theta_{i,j,k}(\lambda) = \mathbf{n}_i \frac{\lambda \mathbf{x}_i - (\mathbf{R}_k \mathbf{b}_j + \lambda \mathbf{t}_k)}{\|\lambda \mathbf{x}_i - (\mathbf{R}_k \mathbf{b}_j + \lambda \mathbf{t}_k)\|}$, $V(\mathbf{x}_i)$ is the calibrated lens vignetting and γ is gamma compression.

Note that, if all light-camera baselines \mathbf{b}_j were zero, λ^2 could be extracted from the denominator and all intensities would be proportional to ρ'_i / λ^2 , producing a fundamental ambiguity: you can multiply the scale by any constant c just multiplying all albedos by c^2 . So, also in the general case, a non-zero baseline is needed to break the ambiguity and estimate true scale and per-point albedo.

4.3 Metric Scale Estimation

Given the up-to-scale reconstruction, our near-light photometric model, and the original images $I_1 \dots I_m$, our goal is to recover the scale factor λ and, as a side product, the point albedos ρ'_i and camera gain changes g'_k . This can be achieved by minimizing the photometric error with respect to the model:

$$\arg \min_{\{\lambda, \rho'_1 \dots \rho'_m, g'_2 \dots g'_m\}} \sum_{i,k} \|\mathcal{I}_{i,k}(\rho'_i, g'_k, \lambda) - I_{i,k}\|_e^2 \quad (6)$$

where $I_{i,k}$ is the intensity of point i observed in image I_k . A robust cost function is used to reduce the influence of spurious data. This nonlinear optimization is solved using the Levenberg-Marquardt method implemented in Ceres [1].

4.4 Initial Guess for the Scale

To avoid local minima and achieve faster convergence it is crucial to find good initial values for the optimization variables (λ , ρ'_i , and g'_k). Indeed, these variables are closely related, therefore, instead of estimating their initial values separately as in [9], we propose to estimate them jointly.




Fig. 2. Method accuracy depends on surface distance, performing best when it matches the light-camera baseline. *Left:* Scale error increases from 1% to 5% with distance. *Right:* Greater distances weaken the photometric cost function’s minimum.

We perform an exhaustive search for the scale parameter λ over a logarithmic space Λ . For each trial value $\hat{\lambda}$, we estimate the albedo values $\hat{\rho}_i'$ solving (5) for each point in the first image. Then, to estimate the relative gains of the rest of images, we perform a robust regression. First, we undo the gamma compression I^γ to work in linear space. Next, we find the gain value \hat{g}_k that minimizes the difference between the real value $I_{i,k}^\gamma$ and the estimated value $\mathcal{I}_{i,k}^\gamma$ of the points using the robust cost function. Finally, we select the value of $\hat{\lambda}$ with the lowest residual according to (6).

5 Experiments

5.1 Datasets

Simulation dataset. Real endoscopic images with a ground-truth metric scale are difficult to obtain. Therefore, we assessed the accuracy of our method through simulations. Our endoscope has a monocular fisheye camera and three surrounding light sources with a ~ 3 mm baseline (Figure 1a). We used a real 3D mesh from [12], a Lambertian illumination model, and Gaussian pixel noise of 4 gray levels. Examples are available in the supplementary video.

EndoMapper dataset [2]. We validate our method using real endoscopy videos to assess its performance in real-world conditions. To our knowledge, EndoMapper is the only dataset that provides both photometric calibration of the endoscope and metric-scale annotations of polyp sizes estimated by endoscopists. It uses a calibrated Olympus endoscope with three light sources. We interpret the provided light spread as vignetting, since both were jointly estimated, and we use isotropic light sources positioned according to the manufacturer’s datasheet.

5.2 Impact of Distance to the Surface

In our *simulation dataset*, the endoscope is positioned to face a polyp at varying distances from the surface (3 to 20 mm), capturing images from slightly different

Table 1. Accuracy at near-field ranges (5mm). The study highlights the importance of precise multi-view geometry \mathbf{T}_k and the benefits of our initial guess.

| | | Optimized | | | Initial guess | \mathbf{T}_k | \mathbf{n}_i | % error | | |
|-------------------|---|-----------|---------|-----------|------------------|----------------|----------------|-----------|---------|------|
| | | λ | ρ' | g' | | | | λ | ρ' | g' |
| EndoMetric | | ✓ | ✓ | ✓ | ✓ | SfM | SfM | 0.95 | 5.48 | 3.37 |
| Ablation study | A | ✓ | ✓ | ✓ | ✓ | SfM | GT | 0.81 | 4.52 | 4.04 |
| | B | ✓ | ✓ | GT | ✓ | SfM | GT | 0.80 | 3.99 | – |
| | C | ✓ | ✓ | GT | ✓ | GT | GT | 0.17 | 3.44 | – |
| | D | ✓ | ✓ | ✓ | X | SfM | SfM | 38.21 | 18.43 | 3.88 |
| | E | ✓ | ✓ | ✓ | [9] | GT | GT | 1.23 | 4.25 | 4.57 |

viewpoints, which could be easily achieved in practice by bending the endoscope tip. Four images are used at a time to reconstruct the 3D shape with COLMAP and estimate its scale. Figure 2L presents the average scale error as a function of distance to the surface. Since different image sets may yield varying results, we report the mean and standard deviation over five experiments.

The method achieves an error of approximately 1% at distances up to 8 mm, which are typical in endoscopic procedures. This confirms its effectiveness in near-field conditions, where the distances between the camera, light sources, and surface are of the same order of magnitude. The scale error rises to 5% as we move away from the surface, since the photometric cost function exhibits a less marked minimum (Figure 2R).

5.3 Impact of Multi-View Reconstruction Accuracy

Focusing on near-field conditions, Table 1 presents an ablation study using only images from our *simulation dataset* captured at approximately 5 mm from the surface. Results show that scale accuracy remains similar whether ground-truth surface normals (row A) or camera gain values (row B) are provided. However, the error decreases significantly in the scenario with the most ground-truth information (row C), suggesting that improving the underlying multi-view reconstruction could enhance the accuracy of our scale estimation method.

5.4 Impact of Initial Guess

Our photometric cost function is non-linear and can exhibit multiple local minima, as illustrated in Figure 2R. To ensure robustness and generality, we propose a method for computing the initial seed for non-linear optimization. Without this step, our method may get trapped in a sub-optimal solution, resulting in large errors, as shown in Table 1 (row D). Previous work [9] initialized albedo and gains to constant values and assumed known camera poses and scene geometry. We tested this configuration with our software and achieved an error of 1.23% in estimating the real scale under these ideal conditions (row E). In contrast, despite estimating camera poses and scene geometry automatically, our method obtains a smaller error of 0.95%, thanks to a better initial guess.


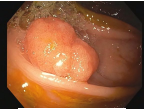



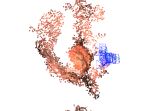
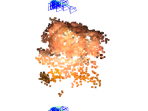


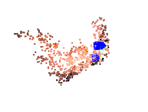
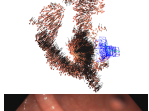

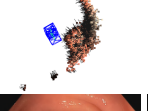
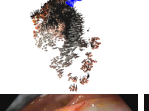

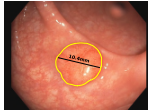
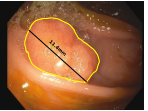
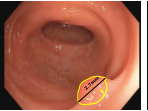
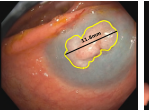

| | Seq_041 Polyp A | Seq_34 Polyp A | Seq_058 Polyp A | Seq_041 Polyp D | Seq_022 Polyp A |
|-------------|---|---|---|--|---|
| Image |  |  |  |  |  |
| 3D |  |  |  |  |  |
| Normals |  |  |  |  |  |
| Result |  |  |  |  |  |
| Our method | 10.4mm | 11.4mm | 2.7mm | 11.8mm | 4.2mm |
| Endoscopist | 10mm | 10mm | 3mm | 8–10mm | 4–5mm |
| Discrepancy | 0.4mm (4%) | 1.4mm (14%) | 0.3mm (10%) | 2.8mm (31%) | 0.3mm (7%) |

Fig. 3. Results in EndoMapper dataset [2].

5.5 Real Polyps Measurement

In the *EndoMapper dataset*, some sequences include metadata with annotations derived from the endoscopist’s speech recorded during the procedure. We selected five polyps where the practitioner estimated the lesion size. The selection criteria included a clear view of the polyp, a distance to the surface of 5–15mm, and smooth endoscope tip motion. We use four seconds of video at 5 FPS—yielding a total of 20 non-contiguous frames—and process them with COLMAP to reconstruct the 3D shape, whose scale is computed by our method.¹ We apply Segment Anything [16] to identify the polyp in a single frame. The size of the polyp is determined by measuring the longest diameter of the 3D points within the segmented region. Figure 3 presents an input image, the reconstructed 3D shape, the estimated surface normals, and our diameter measurement within the polyp boundaries. On average, our measurements deviate from the endoscopist’s estimation by 1.0 mm (13%), demonstrating the potential of the method for standardizing polyp size assessment.

¹ On an i7 10700k 3.8 GHz CPU, the running time is 45 s for COLMAP reconstruction, 15 s for the initial scale estimation (Python prototype) and 0.4 s for minimizing eq. (6) with Ceres. In future work, we will replace COLMAP with a real-time SLAM method, and optimize the initial scale estimation in C++.

6 Conclusions

We have presented, for the first time, a method to obtain 3D reconstructions with real metric scale from a conventional monocular endoscope, under unknown varying albedo, solely based on physical principles. The method does not require any application-specific learning, prior knowledge, or hardware modifications — only a calibrated light-camera setup. Our simulations demonstrate that accurate metric scale can be recovered under practical conditions. Our experiments on the EndoMapper dataset show that the method produces polyp measurements closely matching those estimated visually by an endoscopist. This provides a preliminary, yet solid, proof that the method bridges the simulation-to-real gap and effectively estimates metric scale in real images. A quantitative evaluation of its accuracy will require a dataset with ground-truth polyp size annotations.

Near-field photometry paves the way for real-scale visual SLAM with monocular endoscopes. This will be critical in the short term for accurate measurements, and in the long term, for autonomous robotic exploration and surgery.

Acknowledgments. This work was supported by the EU-H2020 grant 863146: ENDOMAPPER, Next Generation EU *Programa Investigo -116-69*, the Spanish grants PID2021-127685NB-I00 and FPU20/06782, and Aragón grant DGA_T45-17R.

Disclosure of Interests. The authors have no competing interests to declare.

References

1. Agarwal, S., Mierle, K., The Ceres Solver Team: Ceres Solver (10 2023), <https://github.com/ceres-solver/ceres-solver>
2. Azagra, P., Sostres, C., Ferrandez, Á., Riazuelo, L., Tomasini, C., Barbed, O.L., Morlana, J., Recasens, D., Batlle, V.M., Gómez-Rodríguez, J.J., Elvira, R., López, J., Oriol, C., Civera, J., Tardós, J.D., Murillo, A.C., Lanas, A., Montiel, J.M.M.: Endomapper dataset of complete calibrated endoscopy procedures. *Scientific Data* **10**(1), 671 (2023)
3. Batlle, V.M., Montiel, J.M.M., Tardós, J.D.: Photometric single-view dense 3D reconstruction in endoscopy. In: *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*. pp. 4904–4910 (2022)
4. Campos, C., Elvira, R., Gómez-Rodríguez, J.J., Montiel, J.M.M., Tardós, J.D.: ORB-SLAM3: An accurate open-source library for visual, visual-inertial, and multi-map SLAM. *IEEE Transactions on Robotics* **37**(6), 1874–1890 (2021)
5. Collins, T., Bartoli, A.: 3D reconstruction in laparoscopy with close-range photometric stereo. In: *Int. conf. on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. pp. 634–642. Springer (2012)
6. Du, S., Zhang, Q., Zhang, Z., Cai, C., Li, X., Qian, D.: Polyp size estimation by generalizing metric depth estimation and monocular 3D reconstruction. In: *IEEE International Symposium on Biomedical Imaging (ISBI)*. pp. 1–5 (2024)
7. Elvira, R., Tardós, J.D., Montiel, J.M.M.: CudaSIFT-SLAM: multiple-map visual SLAM for full procedure mapping in real human endoscopy. *arXiv preprint arXiv:2405.16932* (2024)

8. Engel, J., Koltun, V., Cremers, D.: Direct sparse odometry. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **40**(3), 611–625 (2017)
9. Fernandes-Araujo, A., Montiel, J.M.M.: Estimación de escala absoluta para Structure from Motion en endoscopio con fuente de luz cercana. B.s. thesis, EINA, Universidad de Zaragoza (2022), <https://zaguan.unizar.es/record/120688#>
10. Gómez-Rodríguez, J.J., Montiel, J.M.M., Tardós, J.D.: NR-SLAM: Non-rigid monocular SLAM. *IEEE Transactions on Robotics* **40**, 4252–4264 (2024)
11. Gonçalves, N., Roxo, D., Barreto, J., Rodrigues, P.: Perspective shape from shading for wide-FOV near-lighting endoscopes. *Neurocomputing* **150**, 136–146 (2015)
12. İncetan, K., Celik, I.O., Obeid, A., Gokceler, G.I., Ozyoruk, K.B., Almalioglu, Y., Chen, R.J., Mahmood, F., Gilbert, H., Durr, N.J., Turan, M.: VR-Caps: a virtual environment for capsule endoscopy. *Medical image analysis* **70**, 101990 (2021)
13. Itoh, H., Oda, M., Jiang, K., Mori, Y., Misawa, M., Kudo, S.E., Imai, K., Ito, S., Hotta, K., Mori, K.: Binary polyp-size classification based on deep-learned spatial information. *International Journal of Computer Assisted Radiology and Surgery* **16**(10), 1817–1828 (2021)
14. Itoh, H., Roth, H.R., Lu, L., Oda, M., Misawa, M., Mori, Y., Kudo, S.e., Mori, K.: Towards automated colonoscopy diagnosis: binary polyp size estimation via unsupervised depth learning. In: *Int. conf. on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. pp. 611–619. Springer (2018)
15. Iwahori, Y., Sugie, H., Ishii, N.: Reconstructing shape from shading images under point light source illumination. In: *IEEE International Conference on Pattern Recognition*. vol. 1, pp. 83–87 (1990)
16. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollar, P., Girshick, R.: Segment anything. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 4015–4026 (October 2023)
17. Lamarca, J., Parashar, S., Bartoli, A., Montiel, J.M.M.: DefSLAM: Tracking and mapping of deforming scenes from monocular sequences. *IEEE Transactions on Robotics* **37**(1), 291–303 (2020)
18. Ma, R., Wang, R., Zhang, Y., Pizer, S., McGill, S.K., Rosenman, J., Frahm, J.M.: RNNSLAM: Reconstructing the 3D colon to visualize missing regions during a colonoscopy. *Medical image analysis* **72**, 102100 (2021)
19. Mecca, R., Wetzler, A., Bruckstein, A.M., Kimmel, R.: Near field photometric stereo with point light sources. *SIAM Journal on Imaging Sciences* **7**(4), 2732–2770 (2014)
20. Quéau, Y., Wu, T., Cremers, D.: Semi-calibrated near-light photometric stereo. In: *Int. Conf. Scale Space and Variational Methods in Computer Vision (SSVM)*. vol. 6, pp. 656–668. Springer, Kolding, Denmark (June 2017)
21. von Renteln, D., Djinbachian, R., Zarandi-Nowroozi, M., Taghiakbari, M.: Measuring size of smaller colorectal polyps using a virtual scale function during endoscopies. *Gut* **72**(3), 417–420 (2023)
22. Schonberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 4104–4113 (2016)
23. Wang, J., Li, Y., Chen, B., Cheng, D., Liao, F., Tan, T., Xu, Q., Liu, Z., Huang, Y., Zhu, C., et al.: A real-time deep learning-based system for colorectal polyp size estimation by white-light endoscopy: development and multicenter prospective validation. *Endoscopy* **56**(04), 260–270 (2024)
24. Woodham, R.J.: Photometric method for determining surface orientation from multiple images. *Optical engineering* **19**(1), 139–144 (1980)

25. Wu, C., Narasimhan, S.G., Jaramaz, B.: A multi-image shape-from-shading framework for near-lighting perspective endoscopes. *International Journal of Computer Vision* **86**(2), 211–228 (2010)
26. Yoshioka, M., Sakaguchi, Y., Utsunomiya, D., Sonoda, S., Tatsuta, T., Ozawa, S., Teramura, Y., Harada, K., Kinugasa, H., Okada, H.: Virtual scale function of gastrointestinal endoscopy for accurate polyp size estimation in real-time: a preliminary study. *Journal of biomedical optics* **26**(9), 096002–096002 (2021)
27. Zubizarreta, J., Aguinaga, I., Montiel, J.M.M.: Direct sparse mapping. *IEEE Transactions on Robotics* **36**(4), 1363–1370 (2020)