
SUPERPOINT FEATURES IN ENDOSCOPY *

O. L. Barbed, J. Morlana, J.M. Martínez Montiel, A. C. Murillo
University of Zaragoza
Zaragoza, Spain
leon@unizar.es

F. Chadebecq
University College London
London, UK

ABSTRACT

There is often a significant gap between research results and applicability in routine medical practice. This work studies the performance of well-known local features on a medical dataset captured during routine colonoscopy procedures. Local feature extraction and matching is a key step for many computer vision applications, specially regarding 3D modelling. In the medical domain, handcrafted local features such as SIFT, with public pipelines such as COLMAP, are still a predominant tool for this kind of tasks. We explore the potential of the well known self-supervised approach SuperPoint [4], present an adapted variation for the endoscopic domain and propose a challenging evaluation framework. SuperPoint based models achieve significantly higher matching quality than commonly used local features in this domain. Our adapted model avoids features within specular regions, a frequent and problematic artifact in endoscopic images, with consequent benefits for matching and reconstruction results. Training code and models available <https://github.com/LeonBP/SuperPointEndoscopy>

Keywords deep learning · self-supervision · local features · endoscopy.

1 Introduction

Endoscopic procedures are a frequent medical practice. The endoscope guided by the physician traverses hollow organs or body cavities, such as the colon. Improvements in quality and efficiency of this kind of procedures can benefit numerous patients and broaden screening campaigns reach. In endoscopy, as in plenty other medical imaging tasks, computer vision has potential to help in numerous aspects, such as assistance for diagnosis [33] or 3D modelling [13]. Unfortunately, there is still a significant gap between research results and applicability into the clinic, as discussed for example in [3]. This study emphasizes the need for unsupervised methods that can fully exploit *in the wild* medical data, which is in itself an already scarce resource. To move forward, it is often key to consider challenging and realistic evaluations of current techniques, to determine where specific adaptations are needed.

Our work is motivated by the automated acquisition of 3D models of the endoluminal scene, that can facilitate augmented reality applications or assistance for navigation or patient monitoring. A core step in 3D reconstruction techniques, such as structure from motion (SfM) or Simultaneous Localization and Mapping (SLAM), is local feature detection and matching. Many broadly used SfM or SLAM frameworks still rely on hand-crafted local feature computation [17], although deep learning based techniques are boosting the state of the art. SuperPoint [4] is one of the seminal works in this topic and has inspired many follow up works discussed next. This promising research stream of learning based local features is recently being exploited in the endoscopic image domain [12], since evaluations and benchmarks on local feature detection and matching are typically focused on conventional images and mostly rigid scenes [10].

Endoscopic images captured during routine procedures present many challenges (such as challenging textures, frequent artifacts and scene deformation) that hinder local feature extraction. Fig. 1 shows matches on two pairs of 1 second apart frames from a real colonoscopy where general purpose hand-crafted features (SIFT) can not tackle scenarios that our learned model features do. SIFT concentrates a lot on specular artifacts, while our adapted SuperPoint model

* *Citation:* O.L. Barbed, F. Chadebecq, J. Morlana, J.M.M. Montiel, A.C. Murillo. SuperPoint Features in Endoscopy. MICCAI Workshop on Imaging Systems for GI Endoscopy, International Workshop on Graphs in Biomedical Image Analysis (2022). Pages 45-55. DOI:10.1007/978-3-031-21083-9_5.




Figure 1: Feature extraction (red circle) and matching (green line) on endoscopy samples.

achieves more and better distributed matches, key for good 3D reconstructions. The main contributions of this work are: 1) A thorough study of SuperPoint effectiveness on *in the wild* endoscopic images, compared to typically used hand-crafted local features, including the proposed framework to evaluate these aspects in endoscopic data captured during daily medical practice; 2) our Superpoint adaptation to the endoscopic domain that improves its performance.

2 Related Work

Endoscopic image registration for 3D reconstruction and mapping in minimally-invasive surgery. Endoscopic image registration is an open problem essential to image-guided intervention. Current efforts are directed at developing benchmarks and techniques able to tackle this challenging domain [2]. Learning-based approaches have shown their efficiency for general image registration, but they remain difficult to adapt to minimally-invasive imaging constraints, largely due to a lack of robust feature detection and matching in these scenarios. Most vision-based approaches for 3D reconstruction in medical domains still rely on hand-crafted features [7, 6]. Some works avoid the need for image registration by directly estimating and fusing key frames depth map [21] or combining them with camera pose estimates [16]. These pipelines get the input frames in a temporally consistent way. Our work is focused on a more general problem of feature extraction without any temporal information given to the model.

Local feature detection and description for image registration. Image registration in general settings is a long studied problem [15]. A key aspect in this work is learning-based methods for image registration in endoscopy. Early learning-based approaches solely focused on feature description. Advanced training loss and strategies significantly improved feature descriptor performances, e.g., by relying on triplet loss which aims at maximizing descriptor discrepancy between close but negative pairs of matches [18]. Similar results have been achieved by sampling more negative pairs as proposed in [29]. The learning-based feature detection problem has been less investigated. Former approaches learn to detect co-variant features and aim at reproducing and eventually improving hand-crafted feature detectors [5, 11]. Unlike these approaches, [24] learns in an unsupervised way to rank keypoints according to their repeatability. The repeatability constraint is now generally combined with peakiness constraints for improving the robustness of the detector [19, 32].

State-of-the-art registration approaches directly integrate feature extraction and description in a single framework. It has been shown that such approaches significantly improve matching results over classical hand-crafted feature-based registration methods [20]. Preliminary approaches such as LIFT [31] aim at reproducing the different stages of classical image registration pipelines. The need for Structure-from-Motion labels to train supervised methods makes these approaches impractical for applications such as endoscopy. More recent unsupervised approaches such as SuperPoint aim at jointly detecting and describing image landmarks [4]. The training is an iterative process that starts by learning from a synthetic dataset of random 2D shapes. The next iterations learn from the problem-specific dataset generated by applying random homographies to source images and using self-supervision from the previous iteration. It remains among the most efficient feature detection and description methods, and is still being considered in recent comparatives [10]. The R2D2 network [22], based on the L2-Quad architecture, jointly estimates a reliability and repeatability map together with a dense descriptor map. Despite their efficiency, methods jointly detecting and describing features are difficult to train and do not generalize well to different application domains [10]. To overcome these limitations, [30] propose to rely on a describe-to-detect strategy which takes advantage of the efficiency and performances of learning-based descriptor models. Recently, [14] propose a descriptor training strategy based on the formulation of a novel landmark tracking loss. While results demonstrate the efficacy of the proposed method, its computational cost remains generally high. Recent image matching trends propose dense matching as an intermediate step to local matching [34] and incorporating attention for the matching stages [23, 9, 28]. However, these approaches rely on 3D reconstruction ground truth for training, which is often not available for recordings acquired during routine medical practice.

3 SuperPoint in endoscopy

Local feature matching is typically divided in four steps: feature detection, descriptor computation, matching and, often, outlier filtering. Our goal is to evaluate and improve existing methods on the first two steps for *in the wild* endoscopy imagery. The well-known SuperPoint, a seminal work regarding self-trained deep learning solutions for feature detection and description, is the base for our study. We next describe the Superpoint model variations used and the matching strategy applied. More implementation details in the supplementary material.

3.1 SuperPoint models considered

SuperPoint Base. Original SuperPoint model [4]. For this and the following model, we use the implementation by [8], which allows us to use the original model weights as well as training new models. SuperPoint follows the known encoder-decoder architecture, but with two parallel decoders (detection and description heads). SuperPoint processes a single image ($I \in \mathbb{R}^{H \times W}$. H and W are the height and width, respectively) as input and produces two outputs: *detection*, image location of each keypoint extracted, and *description*, one descriptor for each keypoint. The *detection* head maps I into a tensor $\mathcal{X} \in \mathbb{R}^{H/8 \times W/8 \times 65}$. The depth of 65 corresponds to a cell of 8×8 pixels in I plus an additional channel called dustbin or “no interest point”. After performing a softmax over the third dimension (we refer to it as `softmd()`), the dustbin is removed and the rest is reshaped to recover I ’s dimensions ($\text{d2s}(X) : \mathbb{R}^{H/8 \times W/8 \times 64} \rightarrow \mathbb{R}^{H \times W}$). The result is interpreted as a probability heatmap of the keypoints in the image. The *description* head maps I into a tensor $\mathcal{D} \in \mathbb{R}^{H/8 \times W/8 \times 256}$. The depth of 256 is the descriptor size, associated with a whole cell of 8×8 pixels in I . Bi-cubic interpolation is used to upsample \mathcal{D} into having H and W as the first two dimensions. The descriptors are L2-normalized. SuperPoint is trained by contrasting the outputs of an image and a warped version of itself via a known homography and pre-computed pseudo-labels of image keypoints. The **loss function** is

$$\mathcal{L}_{SP}(\mathcal{X}, \mathcal{X}', \mathcal{D}, \mathcal{D}'; Y, Y', S) = \mathcal{L}_p(\mathcal{X}, Y) + \mathcal{L}_p(\mathcal{X}', Y') + \lambda \mathcal{L}_d(\mathcal{D}, \mathcal{D}', S), \quad (1)$$

where \mathcal{X} and \mathcal{X}' are the raw detection head outputs for image I and warped image I' , respectively. Their associated detection pseudo-labels are Y and Y' . \mathcal{D} and \mathcal{D}' are the raw description head outputs. $S \in \mathbb{R}^{H/8 \times W/8 \times H/8 \times H/8}$ is the homography-induced correspondence matrix. \mathcal{L}_p is the detection loss, which measures the discrepancies between the detection outputs and the pseudo-labels. \mathcal{L}_d is the description loss, that forces descriptors that correspond to the same region in the original image to be similar, and different to the rest. λ is a weighting parameter.

E-SuperPoint. Specularities are very frequent artifacts in endoscopic images [27], and feature extractors often tend to detect features in the contour or within these image specularities. Features on specularities are not well suited for rigid model estimation, suffer from bad localization, and they turn out to be unreliable in downstream tasks such as tracking and 3D reconstruction. Although they can be masked out later, as we see in our experiments, they account for a too large portion of the features and matches. Thus, we aim to prevent them from happening in the first place, to encourage the detectors to focus on other regions.

We fine-tune the original model using endoscopic images (resized to 256×256) from routine medical practice recordings (dataset detailed in Sec. 4). Pseudo-labels are obtained with the original SuperPoint model. The pseudo-label is set to zero where the confidence value is lower than a threshold of 0.015. Non-maximum suppression is applied over windows of 9×9 pixels, and only the top 600 points are finally saved. We fine-tune the model for 200000 iterations with learning rate of $1e-5$ and batch size of 2. We use sparse loss for more efficient convergence [8], and the rest of parameters are the same as they describe. For testing we set the detection threshold to 0.015 and non-maximum suppression over 3×3 windows.

Our modification of the SuperPoint model adds a new term to the training loss, our specularity loss \mathcal{L}_s . The purpose of \mathcal{L}_s is to account for all the keypoints that are extracted on top of specularities, and is close to zero when there are no keypoints on those locations. **The final loss** is:

$$\mathcal{L}_{ESP}(I, I', \mathcal{X}, \mathcal{X}', \mathcal{D}, \mathcal{D}'; Y, Y', S) = \mathcal{L}_{SP}(\dots) + \lambda_s \mathcal{L}_s(\mathcal{X}, I) + \lambda_s \mathcal{L}_s(\mathcal{X}', I'), \quad (2)$$

where we add to the original \mathcal{L}_{SP} the value of our specularity loss \mathcal{L}_s , once per image, weighted by the scale factor λ_s . \mathcal{L}_s is defined as

$$\mathcal{L}_s(\mathcal{X}, I) = \frac{\sum_{h,w=1}^{H,W} [\text{m}(I)_{hw} \cdot \text{d2s}(\text{softmd}(\mathcal{X}))_{hw}]}{\epsilon + \sum_{h,w=1}^{H,W} \text{m}(I)_{hw}}, \quad (3)$$

where `softmd()` and `d2s()` are softmax and reshape functions from the original SuperPoint, and $\epsilon = 10^{-10}$. The subscript X_{hw} refers to the value of X at row h and column w . $\text{m}(I)$ is a weighting mask: it is > 0 for pixels near a specularity and 0 otherwise. The mask comes from post-processing I with three operations: a binary threshold of

$I_{hw} > 0.7$, a dilation of this binary output with a 3×3 kernel size, and a Gaussian blur of the mask with 9×9 kernel size and $\sigma = 4$. The threshold of 0.7 was chosen empirically, after observing that higher values missed too many specularities and lower values were discarding too many valid regions. To balance the new loss component \mathcal{L}_s , we set the weighting parameter $\lambda_s = 100$ so the losses have similar magnitudes for better optimization. Testing parameters remain the same.

3.2 SuperPoint Matching

SuperGlue [23] is a well-known matching strategy proposed for SuperPoint. However, it requires correspondence ground-truth for training so we can not easily adapt it to endoscopy imagery. We opt to use bi-directional brute force matching, the originally recommended matching for SuperPoint. We also perform a robust geometry estimation with RANSAC to remove outliers, assuming local rigidity for short periods of time. Matching of frames too far apart along the video would need to account for significant deformations, which is out of the scope for this work.

4 Experiments

This section summarizes the main results and insights from our comparison of different SuperPoint models and well-known local features applied in endoscopic data. Implementation details are in the supplementary material, Table 1.

Datasets. A key aspect in this research is to evaluate local feature performance on *in the wild* endoscopic recordings. The model is trained on a set of private videos, and evaluated on two public benchmarks: EndoMapper [1] and Hyper-Kvasir [2]. The supplementary material includes sample images from all sets.

- **Train set.** Endoscopy videos captured across several days of regular medical practice, each video corresponding to a routine procedure on a different patient. We use 11 videos and extract 125000 training frames and another 7179 for validation.
- **EndoMapper test set.** 6 full endoscopies for testing (14191 frames). Sequences 1, 2, 14, 16, 17 and 95. This dataset is similar to the videos used for training.
- **Hyper-Kvasir test set.** 31 short test videos (total of 51925 frames). The labeled videos in “lower-gi-tract/quality-of-mucosal-view/BBPS-2-3”.

Evaluation framework proposed. As often discussed in recent literature, common matching quality metrics, such as repeatability or homography estimation, are not fully representative of local features behaviour in real world settings [10]. This work also shows that hand-crafted features, particularly SIFT, can still surpass more recent deep learning based features regarding accuracy in 3D vision tasks such as image registration. For features and matches to be useful in posterior 3D reconstruction tasks, known desired properties include: good amount of quality matches (reliable and accurate) and matches covering all the scene to better capture the 3D scene information. We propose the following for the evaluation:

- To use an **existing SfM approach**, COLMAP [26, 25], to **pre-compute a pseudo-ground truth** for the relative pose between each pair of frames. COLMAP runs a final global bundle adjustment optimization to recover all relative camera poses. This pseudo-ground truth is used to compute rotation estimation errors and matching quality metrics detailed next.
- A set of **matching quality metrics** to account for: 1) matching quantity and quality, with inliers obtained from Essential (when camera calibration is available, E Inl.) or Fundamental (F Inl.) matrix RANSAC-based estimation, and inliers according to the relative pose provided as pseudo-ground truth (pGT Inl.); 2) scene coverage, with image cell % (out of a 16×16 grid) with at least one inlier (%Gr).

Matching quality evaluation. The following experiments analyze how well each feature can be matched along challenging endoscopic sequences. We extract and match features across pairs of frames 1 second apart from each other from sequences in EndoMapper (1 second = 40 frames for three videos, 1 s = 50 frames for the other three) and Hyper-Kvasir (1 second = 25 frames).

Table 1 shows the performance of different baselines and our adapted model. The changes proposed have a noticeable effect, obtaining improvements in amount of features extracted and inlier matches (both with RANSAC and with the pseudo-GT) and spreading of these matches over the image in both scenarios. This is remarkable because E-SP was not fine-tuned in (b) Hyper-Kvasir data but it still mostly outperforms the rest.

	Feat/Img	E Inl.	%Gr	pGT Inl.	%Gr	Feat/Img	F Inl.	%Gr
SIFT	2350.2	148.2	11.9	71.6	9.3	825.7	151.3	18.6
ORB	2163.0	153.0	8.5	80.9	6.0	361.3	137.2	6.3
SP Base	1333.7	96.4	11.4	57.1	8.2	211.8	51.3	11.1
E-SP	4500.9	278.9	13.2	172.0	9.8	591.3	200.4	11.3

(a) **EndoMapper** *test set* (1080x1080 resolution) (b) **Hyper-Kvasir** (512x512)

Table 1: Matching quality metrics for the two different test sets. E-SP trained on *Train set*. pGT only available for (a) EndoMapper data.

	Feat/Img		E Inliers	
	all feat.	w/o S	all feat.	w/o S
SIFT	2350.2	2006.2 (85.4%)	148.2	129.1 (87.1%)
ORB	2163.0	867.7 (40.1%)	153.0	61.8 (40.4%)
SP Base	1333.7	1035.2 (77.6%)	96.4	73.5 (76.2%)
E-SP	4500.9	4431.3 (98.5%)	278.9	274.7 (98.5%)

Table 2: Influence of specularities in the matching results (in EndoMapper *test set*). all feat: total number; w/o S: number without features that fall into specularities.

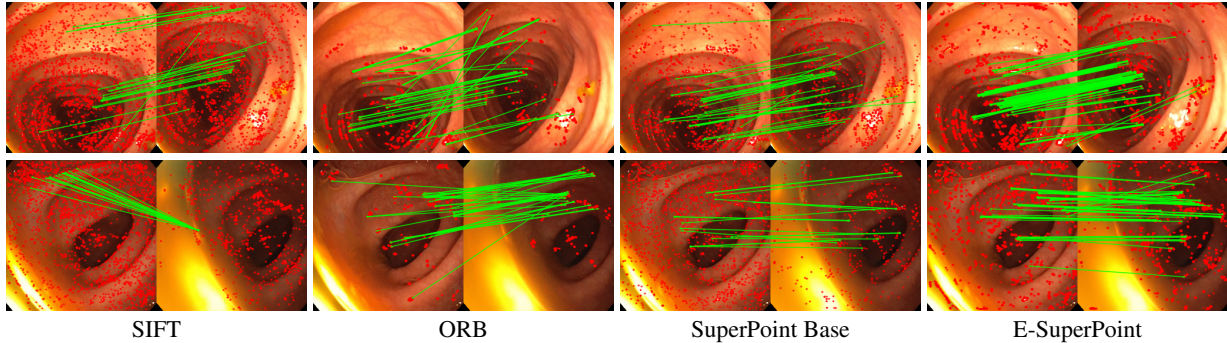


Figure 2: Features (red circles) and inlier matches after RANSAC (green lines) obtained for two pairs of 1 second apart frames using different features.

Specularities. E-SuperPoint is designed to encourage feature extraction avoiding specularities. This experiment evaluates this with the number of features and inliers when features located in specularity pixels are discarded. We consider a pixel part of a specularity if the intensity value is over 0.7. Table 2 summarizes these results, showing that the baseline models lose a significant amount of features and inliers if we ignore specularity features (*w/o S*), confirming the suspicion that they fire too much on specularities in this environment. In contrast, the proposed E-SuperPoint effectively removes reliance in specularities and allows the detector to focus on other image patterns which have higher chances of being stable. Fig. 2 shows several matching examples. Higher resolution version of them and additional examples can be found in the supplementary material.

Rotation estimation from matches. We compare the RANSAC-estimated essential matrix and the pseudo-ground truth essential matrix to compute the rotation estimation error for every pair of images. These values are summarized in Tab. 3, where we show the values of relevant percentiles of the errors obtained. E-SuperPoint achieves lower rotation error than all the other methods. Additionally, when counting the percentage of the estimations that obtain an error lower than 30 degrees, E-SuperPoint succeeds **71.6%** of the time, while the second best, SuperPoint Base, only 61.3%. Figure 2 shows some examples where we can see that E-SuperPoint is more robust than other methods: Top row example shows how E-SuperPoint ignores the specularities in the image completely (there are no features extracted on top of it), and the bottom row example shows that our model finds more matches and better spread over the image than the other methods. The improvement is largely due to the adaptation we made to better deal with specularities for feature extraction.

Percentile	10 th	20 th	30 th	40 th	median
SIFT	3.2	5.5	8.4	12.4	20.1
ORB	3.7	7.3	13.4	27.4	45.8
SP Base	3.2	5.8	9.2	13.4	19.8
E-SP	2.6	4.8	7.5	10.5	14.5

Table 3: Rotation estimation error for pairs of frames 1 second apart. Percentile values of the errors obtained by each method.

5 Conclusions

This work² studies the performance of local features in *in the wild* endoscopic environments, using data captured during routine medical practice. We compare the effectiveness of hand-crafted local features against deep learning-based ones, in particular SIFT, ORB and SuperPoint. Although hand-crafted features are still a dominant choice in this field, we show how deep learning based features can surpass them in the considered challenging environments. Besides, we have trained and adapted the general-purpose SuperPoint to better fit the challenges of endoscopic imagery. Our evaluation, on endoscopies of different patients, is focused on the quality of the recovered 3D camera motion. Our results show that SuperPoint adaptation provides more numerous and non-specular features, and more disperse correspondences, essential for accurate and robust 3D geometry estimations.

References

- [1] Azagra, P., Sostres, C., Ferrandez, Á., Riazuelo, L., Tomasini, C., Barbed, O.L., Morlana, J., Recasens, D., Batlle, V.M., Gómez-Rodríguez, J.J., et al.: Endomapper dataset of complete calibrated endoscopy procedures. arXiv preprint arXiv:2204.14240 (2022)
- [2] Borgli, H., Thambawita, V., Smedsrud, P.H., Hicks, S., Jha, D., et al.: Hyperkvasir, a comprehensive multi-class image and video dataset for gastrointestinal endoscopy. *Scientific data* **7**(1), 1–14 (2020)
- [3] Chadebecq, F., Vasconcelos, F., Mazomenos, E., Stoyanov, D.: Computer vision in the surgical operating room. *Visceral Medicine* **36**(6), 456–462 (2020)
- [4] DeTone, D., Malisiewicz, T., Rabinovich, A.: Superpoint: Self-supervised interest point detection and description. In: *Conference on Computer Vision and Pattern Recognition Workshops*. IEEE (2018)
- [5] Di Febbo, P., Dal Mutto, C., Tieu, K., Mattocchia, S.: Kcnn: Extremely-efficient hardware keypoint detection with a compact convolutional neural network. In: *CVPR Workshops*. IEEE (2018)
- [6] Espinel, Y., Calvet, L., Botros, K., Buc, E., Tilmant, C., Bartoli, A.: Using multiple images and contours for deformable 3d-2d registration of a preoperative ct in laparoscopic liver surgery. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer (2021)
- [7] Gómez-Rodríguez, J.J., Lamarca, J., Morlana, J., Tardós, J.D., Montiel, J.M.: SD-DefSLAM: Semi-direct monocular SLAM for deformable and intracorporeal scenes. In: *International Conference on Robotics and Automation*. IEEE (2021)
- [8] Jau, Y.Y., Zhu, R., Su, H., Chandraker, M.: Deep keypoint-based camera pose estimation with geometric constraints. In: *International Conference on Intelligent Robots and Systems*. IEEE (2020), <https://github.com/eric-yyjau/pytorch-superpoint>
- [9] Jiang, W., Trulls, E., Hosang, J., Tagliasacchi, A., Yi, K.M.: Cotr: Correspondence transformer for matching across images. arXiv preprint arXiv:2103.14167 (2021)
- [10] Jin, Y., Mishkin, D., Mishchuk, A., Matas, J., Fua, P., Yi, K.M., Trulls, E.: Image matching across wide baselines: From paper to practice. *International Journal of Computer Vision* **129**(2), 517–547 (2021)
- [11] Laguna, A.B., Riba, E., Ponsa, D., Mikolajczyk, K.: Key.Net: Keypoint detection by handcrafted and learned cnn filters. In: *ICCV*. IEEE (2019)
- [12] Liao, C., Wang, C., Bai, J., Lan, L., Wu, X.: Deep learning for registration of region of interest in consecutive wireless capsule endoscopy frames. *Computer Methods and Programs in Biomedicine* **208**, 106189 (2021)

²This project has been funded by the European Union’s Horizon 2020 research and innovation programme under grant agreement No 863146 and Aragon Government FSE-T45_20R.

- [13] Liu, X., Stiber, M., Huang, J., Ishii, M., Hager, G.D., Taylor, R.H., Unberath, M.: Reconstructing sinus anatomy from endoscopic video—towards a radiation-free approach for quantitative longitudinal assessment. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer (2020)
- [14] Liu, X., Zheng, Y., Killeen, B., Ishii, M., Hager, G.D., Taylor, R.H., Unberath, M.: Extremely dense point correspondences using a learned feature descriptor. In: Conference on Computer Vision and Pattern Recognition. IEEE (2020)
- [15] Ma, J., Jiang, X., Fan, A., Jiang, J., Yan, J.: Image matching from handcrafted to deep features: A survey. *International Journal of Computer Vision* pp. 1–57 (2020)
- [16] Ma, R., Wang, R., Pizer, S., Rosenman, J., McGill, S.K., Frahm, J.M.: Real-time 3D reconstruction of colonoscopic surfaces for determining missing regions. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer (2019)
- [17] Mahmoud, N., Collins, T., Hostettler, A., Soler, L., Doignon, C., Montiel, J.M.M.: Live tracking and dense reconstruction for handheld monocular endoscopy. *IEEE transactions on medical imaging* **38**(1), 79–89 (2018)
- [18] Mishchuk, A., Mishkin, D., Radenović, F., Matas, J.: Working hard to know your neighbor’s margins: local descriptor learning loss. In: International Conference on Neural Information Processing Systems (2017)
- [19] Mishkin, D., Radenovic, F., Matas, J.: Repeatability is not enough: Learning affine regions via discriminability. In: ECCV (2018)
- [20] Ono, Y., Trulls, E., Fua, P., Yi, K.M.: LF-Net: learning local features from images. In: International Conference on Neural Information Processing Systems (2018)
- [21] Ozyoruk, K.B., Gokceler, G.I., Bobrow, T.L., Coskun, G., Inctan, K., Almalioglu, Y., Mahmood, F., Curto, E., Perdigoto, L., Oliveira, M., et al.: EndoSLAM dataset and an unsupervised monocular visual odometry and depth estimation approach for endoscopic videos. *Medical image analysis* **71**, 102058 (2021)
- [22] Revaud, J., Weinzaepfel, P., de Souza, C.R., Humenberger, M.: R2D2: repeatable and reliable detector and descriptor. In: International Conference on Neural Information Processing Systems (2019)
- [23] Sarlin, P.E., DeTone, D., Malisiewicz, T., Rabinovich, A.: Superglue: Learning feature matching with graph neural networks. In: Conference on Computer Vision and Pattern Recognition. IEEE (2020)
- [24] Savinov, N., Seki, A., Ladický, L., Sattler, T., Pollefeys, M.: Quad-networks: Unsupervised learning to rank for interest point detection. In: Conference on Computer Vision and Pattern Recognition. IEEE (2017)
- [25] Schönberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: CVPR. IEEE (2016)
- [26] Schönberger, J.L., Zheng, E., Pollefeys, M., Frahm, J.M.: Pixelwise view selection for unstructured multi-view stereo. In: European Conference on Computer Vision (2016)
- [27] Stoyanov, D., Yang, G.Z.: Removing specular reflection components for robotic assisted laparoscopic surgery. In: International Conference on Image Processing. IEEE (2005)
- [28] Sun, J., Shen, Z., Wang, Y., Bao, H., Zhou, X.: Loftr: Detector-free local feature matching with transformers. In: CVPR. IEEE (2021)
- [29] Tian, Y., Fan, B., Wu, F.: L2-Net: Deep learning of discriminative patch descriptor in euclidean space. In: Conference on Computer Vision and Pattern Recognition. IEEE (2017)
- [30] Tian, Y., Balntas, V., Ng, T., Barroso-Laguna, A., Demiris, Y., Mikolajczyk, K.: D2d: Keypoint extraction with describe to detect approach. In: ACCV (2020)
- [31] Yi, K.M., Trulls, E., Lepetit, V., Fua, P.: Lift: Learned invariant feature transform. In: European Conference on Computer Vision. Springer (2016)
- [32] Zhang, L., Rusinkiewicz, S.: Learning to detect features in texture images. In: Conference on Computer Vision and Pattern Recognition. IEEE (2018)
- [33] Zhang, Z., Xie, Y., Xing, F., McGough, M., Yang, L.: Mdnet: A semantically and visually interpretable medical image diagnosis network. In: CVPR. IEEE (2017)
- [34] Zhou, Q., Sattler, T., Leal-Taixe, L.: Patch2pix: Epipolar-guided pixel-level correspondences. In: CVPR. IEEE (2021)

A Supplementary material




Figure 3: **Sample frames from *Train set***. They often present poor quality frames and many artifacts. We apply a simple mask (*Mask*) to ignore black corners, which are a constant artifact due to the endoscope camera type.




Figure 4: **Sample frames from Hyper-Kvasir [2]**, the labeled videos in “lower-gi-tract/quality-of-mucosal-view/BBPS-2-3”. They are significantly different from *Train set* and EndoMapper datasets, since they include additional augmented information superimposed to the left of the frames. The *Mask* allows us to ignore these parts and the black corners.

	Impl.	Max. Feat	Matching	Distance	Other parameters
SIFT	OpenCV	10000	Ratio (0.8)	L2 Norm	Contrast threshold= 0.1
ORB	OpenCV	10000	BFMatcher	Hamming	-
SuperPoint	[8]	10000	BFMatcher	L2 Norm	-

Table 4: **Implementation details**. For all our experiments we use OpenCV 3.4.2. RANSAC algorithm is set with confidence threshold of 0.9999, and all matches within a distance of 3 pixels of their matched pair when reprojecting are considered inliers.




Figure 5: **Additional examples in higher resolution** of features (red circles) and matches (green lines) obtained in pairs of 1 second apart frames from **EndoMapper** [1] dataset. As discussed in the main paper, our E-Superpoint achieves more reliable correspondences, more spread than the rest, and is able to mitigate the amount of features located in specularities.