

Placental Vessel-guided Hybrid Framework for Fetoscopic Mosaicking

Sophia Bano^a, Francisco Vasconcelos^a, Anna L. David^b, Jan Deprest^c, Danail Stoyanov^a

^aWellcome/EPSRC Centre for Interventional and Surgical Sciences(WEISS) and Department of Computer Science, University College London, London, UK; ^bFetal Medicine Unit, University College London Hospital, London, UK; ^cDepartment of Development and Regeneration, University Hospital Leuven, Leuven, Belgium

ARTICLE HISTORY

Compiled September 13, 2022

ABSTRACT

Fetoscopic laser photocoagulation is used to treat Twin-to-twin transfusion syndrome, however, this procedure is hindered because of difficulty in visualizing the intraoperative surgical environment due to limited surgical field-of-view, unusual placenta position, limited manoeuvrability of the fetoscope and poor visibility due to fluid turbidity and occlusions. Fetoscopic video mosaicking can create an expanded field-of-view image of the fetoscopic intraoperative environment, which could support the surgeons in localizing the vascular anastomoses during the fetoscopic procedure. However, classical handcrafted feature matching methods fail on in vivo fetoscopic videos. An existing state-of-the-art method on fetoscopic mosaicking relies on vessel presence and fails when vessels are not present in the view. We propose a vessel-guided hybrid fetoscopic mosaicking framework that mutually benefits from a placental vessel-based registration and a deep learning-based dense matching method to optimize the overall performance. A selection mechanism is implemented based on vessels' appearance consistency and photometric error minimization for choosing the best pairwise transformation. Using the extended fetoscopy placenta dataset, we experimentally show the robustness of the proposed framework, over the state-of-the-art methods, even in vessel-free, low-textured, or low illumination non-planar fetoscopic views.

KEYWORDS

Fetoscopy; Vessel segmentation; Deep learning; Registration; Video mosaicking; Twin-to-twin transfusion syndrome

1. Introduction

Twin-to-twin transfusion syndrome (TTTS) is a rare fetal anomaly that affects the twins sharing a monochronic placenta. It is caused by abnormal placental vascular anastomoses on the placenta, leading to uneven flow of blood between the two fetuses (Baschat et al. (2011)). Fetoscopic Laser Photocoagulation (FLP) is used to treat TTTS, however, this procedure is hindered because of difficulty in visualizing the intraoperative surgical environment due to limited surgical field-of-view (FoV), unusual placenta position, limited manoeuvrability of the fetoscope and poor visibility

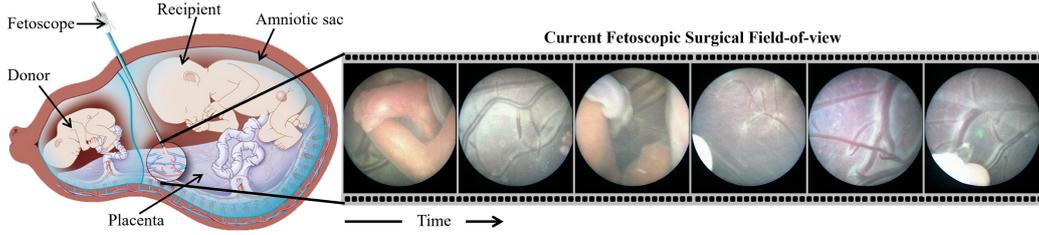


Figure 1.: Illustration of a fetoscopic laser photocoagulation surgery (left) in which a fetoscope is inserted into the amniotic cavity and is used to localize and ablate the vascular anastomoses sites. The current fetoscopic field-of-view (right) is limited and contains occlusion due to the presence of fetus and floating amniotic fluid particles.

due to fluid turbidity and occlusions (as shown in Fig. 1). This adds to the surgeon’s cognitive load and may result in increased procedural time and missed treatment, leading to persistent TTTS. Fetoscopic video mosaicking can create a virtual expanded FoV image of the fetoscopic intraoperative environment, which may provide computer-assisted interventions support in localizing the vascular anastomoses during the FLP procedures.

Several techniques have been proposed in the literature for fetoscopic mosaicking (Reeff et al. (2006); Daga et al. (2016); Gaisser et al. (2018); Tella-Amo et al. (2019); Bano et al. (2020b, 2021, 2019, 2020a, 2022); Alabi et al. (2022); Casella et al. (2022)), each having their own strength and weaknesses, however, the majority of them were unable to overcome the existing challenges in fetoscopic videos that hinders robust mosaicking. Classical video mosaicking methods (Reeff et al. (2006); Daga et al. (2016)) that used handcrafted features (e.g., SIFT, SURF) perform poorly on in vivo fetoscopic videos due to low resolution, poor visibility, honeycomb or blur effect due to fibre-based fetoscope, floating particles and texture paucity or repetitive texture challenges that inherently exists in fetoscopy. Hence, classic computer vision methods for video mosaicking are not suitable for fetoscopic video mosaicking. Fusion of visual tracking with electromagnetic pose sensing has also been studied, but only in ex vivo experiments (Tella-Amo et al. (2019, 2018)). A direct registration method for mosaicking has also presented which was only validated on a single in vivo fetoscopic video (Peter et al. (2018)). Recently, deep learning-based mosaicking methods have also been reported (Bano et al. (2019, 2020b); Alabi et al. (2022); Casella et al. (2022)) for fetoscopic video mosaicking. (Bano et al. (2019, 2020b)) approach restricted the model to estimate only euclidean transformation, thus bounding the drifting error that led to inaccuracies. A recent intensity-based image registration (Bano et al. (2020a)) method relies on placental vessel segmentation maps for registration. This method facilitated in overcoming some visibility challenges, but it failed when the predicted segmentation map was inaccurate or inconsistent across frames, or in views with thin or no vessels. Recent computer vision literature has introduced deep learning-based interest point descriptors (Sarlin et al. (2020)) and detector-free dense feature matching (Sun et al. (2021)) techniques, showing robustness in multiview feature matching. Such techniques can be explored for fetoscopic mosaicking to improve the performance of fetoscopic mosaicking.

To overcome the existing literature limitations, we propose a vessel-guided hybrid framework for creating robust and reliable mosaics. The framework optimizes the performance by fusing the state-of-the-art in fetoscopic mosaicking and computer vision, namely, vessel-based registration (Bano et al. (2020a)) and detector-free local feature

matcher (Sun et al. (2021)) for registration methods, respectively. Our framework introduces a selection mechanism based on appearance consistency of placental vessels and photometric error minimization for choosing the best pair-wise transformations. Through both qualitative and quantitative comparison performed using the extended fetoscopy placenta dataset (Bano et al. (2020a)), we show the robustness of the proposed framework over the existing methods. Our key contribution lies in proposing a mosaicking framework for computer-assisted intervention application which is robust even in the absence of vessels and presence of heavy floating particles, low illumination, non-planar views and spotlight light source. The existing fetoscopic mosaicking methods do not show robustness to all these challenging conditions in a single framework. The proposed hybrid framework brings us closer towards translating the fetoscopic mosaicking framework into clinical settings, which in turn could help in reducing surgeon’s cognitive load during fetoscopic procedures.

2. Method

The proposed framework consists of two parallel registration methods, namely, placental vessel-based direct registration (Bano et al. (2020a)) and detector-free dense (LoFTR) matcher (Sun et al. (2021)) as shown in Fig. 2. Each method performs matching of two consecutive frames I_t and I_{t+1} followed by their registration for estimating pairwise transformations $H_V^{t,t+1}$ and $H_L^{t,t+1}$ from the vessel-based and the dense matcher methods, respectively. A vessel-guided transformation selection strategy is then proposed that also minimizes the photometric errors, thus enabling robust mosaicking. The proposed framework allows generation of mosaics from long fetoscopic video sequences without drift accumulation.

2.1. Placental Vessels Registration

Placental vessel segmentation allows overcoming visibility related challenges such as moving occlusions (floating amniotic fluid particles) and specular view-dependent illumination that can result in inaccurate feature matches. We utilize the placental vessel segmentation and registration method from (Bano et al. (2020a)) as this is the state-of-the-art method in in vivo fetoscopic mosaicking.

Given two consecutive frames I^t and I^{t+1} , a UNet (Ronneberger et al. (2015)) with ResNet50 (He et al. (2016)) backbone, pretrained on the fetoscopy placenta dataset¹, is used for obtaining the predicted vessel maps S^t and S^{t+1} . Similar to (Bano et al. (2020a); Peter et al. (2018)), the registration between I^t and I^{t+1} is approximated with an affine transformation, as the use of projective transformations in fetoscopy data has empirically been shown to lead to poor results. This is because fetoscopic scene is only piece-wise planar. Intensity based direct registration is applied using a pyramidal Lucas-Kanade framework that minimizes the photometric error between S^t and S^{t+1} through Levenberg-Marquardt optimization. A circular FoV mask of the fetoscopic image is used to perform registration while neglecting the black background regions. This gives the estimated affine transformation $H_V^{t,t+1}$ between I^t and I^{t+1} .

Since the placental vessel-based registration method is driven by predicted vessel maps, it tends to fail when the predicted maps are inaccurate or inconsistent across

¹Fetoscopy Placenta Dataset: <https://www.ucl.ac.uk/interventional-surgical-sciences/fetoscopy-placenta-data>

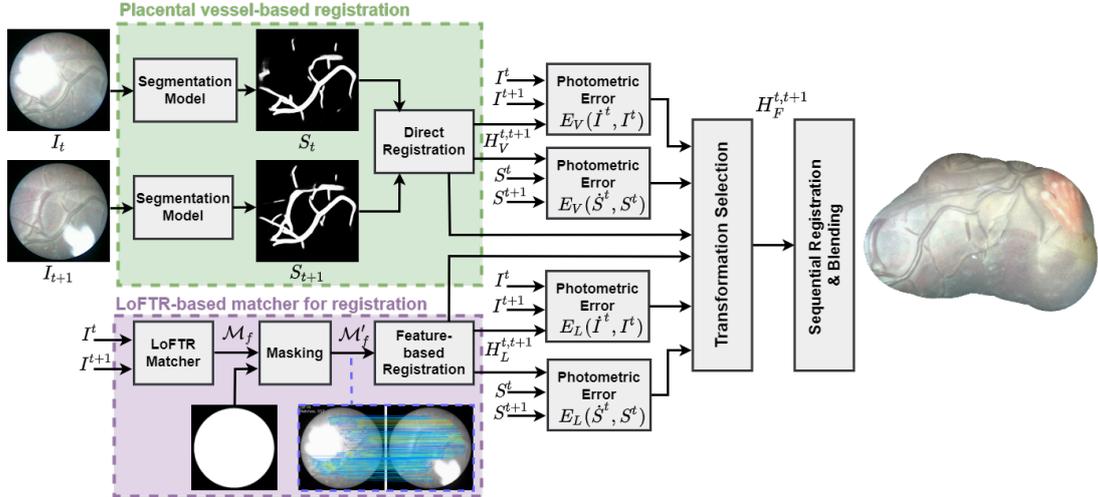


Figure 2.: The proposed framework first estimates the affine transformations $H_V^{t,t+1}$ and $H_L^{t,t+1}$ between RGB frames (I_t, I_{t+1}) using vessel-based registration and LoFTR-based dense feature matches' registration, respectively. Photometric errors ($E_V(I^t, I^t)$, $E_V(S^t, S^t)$) between RGB I_t and reprojected \hat{I}_t frames and between vessel map S_t and reprojected vessel map \hat{S}_t are computed, respectively. Likewise, $E_L(I^t, I^t)$, $E_L(S^t, S^t)$ are computed for the LoFTR-based transformation estimate $H_L^{t,t+1}$. The best estimate ($H_F^{t,t+1}$) is then selected based on vessel segmentation consistency and minimum photometric errors. Finally, pairwise transformations are sequentially registered to form an expanded FoV image.

frames or in views with thin or no vessels.

2.2. Detector-Free Dense Feature Matching for Registration

Unlike classical feature matching methods that perform feature detection and description followed by their matching, the recently proposed LoFTR (Sun et al. (2021)) method takes a hierarchical approach and first establishes pixel-wise dense matches at a coarse level and later refines the good matches at a fine level.

Given I^t and I^{t+1} , a standard convolutional neural network architecture is used to extract dense features at coarse and fine levels from both frames. Coarse local features are fed into the LoFTR module, which uses a transformer with positional encoding, and self and cross-attention layers to transform coarse features into position and context dependent local feature descriptors. A confidence matrix is obtained by matching these descriptors using a differentiable matching layer. Matches in the confidence matrix that are higher than a predefined threshold and that satisfy the mutual nearest neighbour criteria are selected as coarse-level matches. Coarse to fine feature matches \mathcal{M}_f are then obtained by taking a local window size from fine-level features at each coarse match positions and applying the LoFTR module to it. For more detail, refer to (Sun et al. (2021)), in which it is shown that LoFTR produces high-quality matches even in regions with low-textures, motion blur or repetitive patterns; making it an ideal matching module for fetoscopic mosaicking.

For registration, a circular mask covering only the fetoscopic FoV is first used to obtain matches \mathcal{M}'_f only in the visible fetoscope region. Registration is then

approximated as an affine transformation using the RANdom SAmple Consensus (RANSAC) method. The obtained transformation is refined by using only the inliers with Levenberg-Marquardt optimization that further reduces the transformation error. This gives the affine transformation estimate $H_L^{t,t+1}$ that defines the alignment between I^t and I^{t+1} through the LoFTR-based matching.

We note through empirical experimentation that LoFTR matching is affected by the light source intensity and the resulting view-dependent reflectance in the surgical scene that can result in drift error during sequence registration.

2.3. Vessel-guided Transformation Selection

Vessel-guided transformation selection aims at finding the best affine transformation from $H_V^{t,t+1}$ and $H_L^{t,t+1}$ based on the vessels' appearance consistency, percentage of vessels with respect to the fetoscopic FoV and minimum reprojection error. Let $\dot{I}^t = H^{t,t+1}I^{t+1}$ be the reprojected frame obtained by warping I^{t+1} using the estimated transformation $H^{t,t+1}$. The photometric error between \dot{I}^t and I^t is obtained using,

$$E(\dot{I}^t, I^t) = \frac{1}{n} \sum_{i=1}^n |\dot{I}_i^t - I_i^t|, \quad (1)$$

where n is the total number of pixels in a frame. Four photometric errors are computed using the input frames, segmentation maps and two estimated transformations. $E_V(\dot{I}^t, I^t)$ measures the photometric error between I^t and reprojected \dot{I}^t obtained using $H_V^{t,t+1}$. Likewise, $E_V(\dot{S}^t, S^t)$, $E_L(\dot{I}^t, I^t)$ and $E_L(\dot{S}^t, S^t)$ are computed (as shown in Fig. 2).

A rule-based strategy is defined for transformation selection based on the qualitative observations made from the vessel-based and LoFTR matcher-based registration methods. Let x^{t+1} denotes the percentage of vessel class pixels with respect to the total number of pixels in the FoV mask in S^{t+1} . And y^{t+1} denotes the percentage difference between vessel class pixels in S_t and S^{t+1} . We empirically found that vessel consistency can be guaranteed by ensuring $x^{t+1} > 15\%$ and $y^{t+1} < 25\%$ of x^{t+1} . In pairs of frames where vessels are consistent across frames, the affine transformation estimate from the method that gives the lowest errors between $E_V(\dot{S}^t, S^t)$ and $E_L(\dot{S}^t, S^t)$ is selected as the final transformation $H_F^{t,t+1}$. When the vessel consistency conditions are not satisfied, error measurements based on vessel maps become inaccurate. In this case, we select the transformation estimate of the method that reports lower among $E_V(\dot{I}^t, I^t)$ and $E_L(\dot{I}^t, I^t)$ errors.

2.4. Sequential Registration

Once the pairwise transformations are obtained, next step is to compute the relative transformations with respect to a reference frame for mosaic generation. The relative transformation of I_l with respect to a reference frame I_k is computed by applying left-hand matrix multiplication,

$$\mathbf{H}_F^{k,k+l} = \prod_{i=k}^{k+l-1} \mathbf{H}_F^{i,i+1}. \quad (2)$$

where $l > k$ and l is the length of the fetoscopy sequence. This gives an expanded FoV image of the placental surface. To create seamfree mosaics, blending is applied using the Enblend² software.

3. Dataset and Experimental Setup

For experimental analysis, we use an extended version of the publicly available fetoscopy placenta dataset that was introduced in (Bano et al. (2020a)). In the extended version, each video sequence contains an additional 100 frames. The dataset contains 6 in vivo fetoscopic video sequences from six different FLP procedures. The addition of 100 extra frames in each sequence resulted in frames having either weak or no vessels. The number of frames in each video are reported in Fig. 3. We note that there are large inter and intra-case variabilities in the fetoscopic videos. These videos are of varying visual quality having low resolution, poor visibility due to floating amniotic fluid particles and artefacts due to spotlight source, texture sparsity and non-planar views (Video 3 and 5) due to anterior placenta imaging.

Since the ground-truth transformations are not available for in vivo fetoscopy, we use the quantitative metric, referred as N -frame SSIM, proposed by (Bano et al. (2020a)). The N -frame SSIM quantifies the accumulated drift error in N frames by computing the structural similarity index measure (SSIM) between the current frame and warped frame at N frame distance. N can take values from 2 to 5 frames. For comparison, we report the 1 to 5-frame distance SSIM for the state-of-the-art vessel-based (Bano et al. (2020a)), LoFTR (Sun et al. (2021)) matcher-based and the proposed hybrid methods in Fig. 4. Additionally, we present the qualitative results in Fig. 3. These results are discussed in detail in Sec. 4.

The segmentation model is implemented in PyTorch and trained on the vessel segmentation dataset using the same hyperparameters reported in (Bano et al. (2020a)) and on a single Tesla V100-DGXS-32GB GPU of an NVIDIA DGX-station. For LoFTR matcher, we use the pretrained model (trained on the ScanNet dataset (Dai et al. (2017))) for obtaining the fine-level matches between two consecutive frames. We observe that matches returned through this network on fetoscopy data are already robust and repeatable. Retraining or fine-tuning of this network on fetoscopy data was not possible because of the lack of ground-truth data.

4. Results and Discussion

Qualitative and quantitative comparison are presented in Fig. 3 (refer to the supplementary video) and Fig. 4, respectively, using the 6 in vivo fetoscopy videos from the extended fetoscopy placenta dataset. Note that classical feature matching based (Reeff et al. (2006); Daga et al. (2016)) and RGB intensity based (Bano et al. (2020a)) techniques do not work on fetoscopic videos (as mentioned in Sec. 1). Hence, comparison is mainly performed with the existing state-of-the-art method (Bano et al. (2020a)) of fetoscopic mosaicking.

From Fig. 3, we observe that the vessel-based method outperformed in Video 1 and 2 due to strong vessel appearance in these videos. LoFTR matcher-based also performed well, but introduced some registration errors (marked with red circle in

²Enblend: <http://enblend.sourceforge.net/>

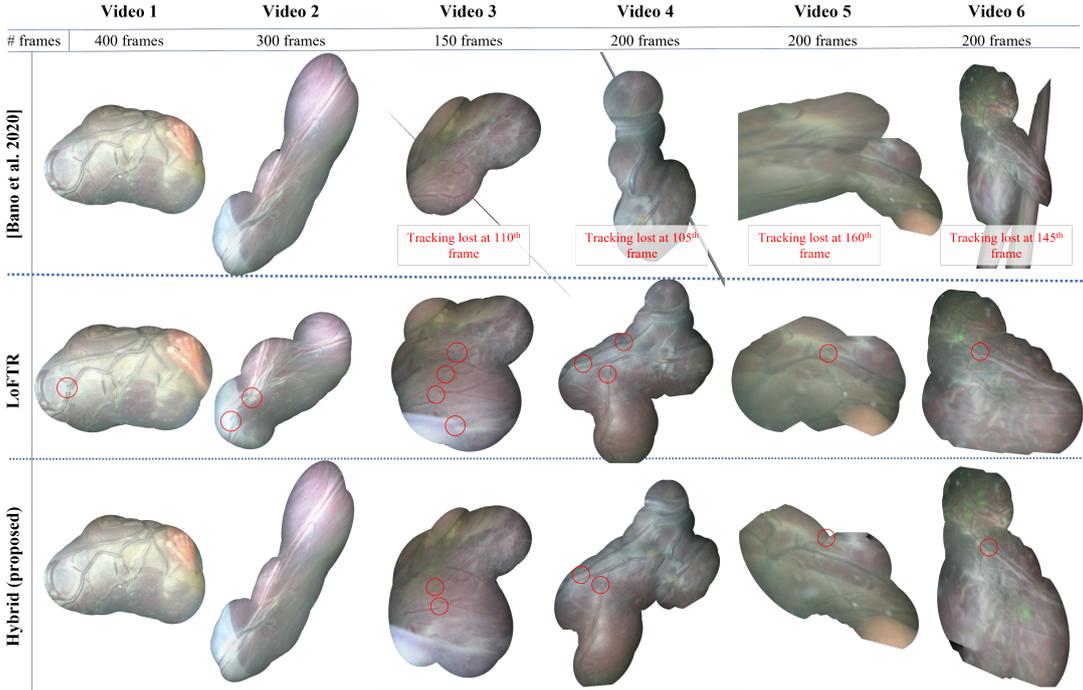


Figure 3.: Qualitative comparison of the proposed hybrid framework (third row) with the vessel-based (Bano et al. (2020a)) (first row) and LoFTR (Sun et al. (2021)) (second row) matcher-based methods using the extended fetoscopy placenta dataset. We can visually observe that the hybrid framework is robust even in the absence of vessels and non-planar views. Refer to the supplementary material for video-based comparison.

Fig. 3). Our proposed hybrid method converged towards selecting the transformations from the vessel-based method because of the consistent appearance and dominant presence of vessels throughout these videos. Video 3 to Video 6 show more challenging scenarios where the vessels are either very thin (video 3 and 6) or thick (close-up view in Video 4) or are not present in some frames. Moreover, Video 3 and 5 shows an anterior placenta, making these views highly non-planar. This negatively influenced the vessel-based method, resulting in increase drifting errors and tracking failures in Video 3, 4, 5 and 6 at frame 110^{th} , 105^{th} , 160^{th} and 145^{th} , respectively. The errors are mostly because of insufficient vessels present in the scene or false negative in predicted vessel maps where the segmentation network failed to properly segment thin vessels. On the other hand, LoFTR matcher-based provided stable mosaics, but resulted in some inconsistencies in registration (marked with red circle in Fig. 3). Our hybrid approach optimize itself to select the best from the two methods, hence the registration error are visibility reduced, giving reliable mosaics.

The quantitative results presented in Fig. 4 shows the 1 to 5-frame SSIM measurements for the three methods under comparison on the 6 in vivo video clips. We make similar observations from here as that of Fig. 3, where performance of all the methods is comparable in Video 1 and Video 2. Vessel-based registration failed for some frames in Video 3 to Video 6, hence we can observe low SSIM values with increasing frame distance. This also shows that the drifting error is large in Video 3 to Video 6 for the vessel-based methods. In these videos, we observe significantly low interquartile range

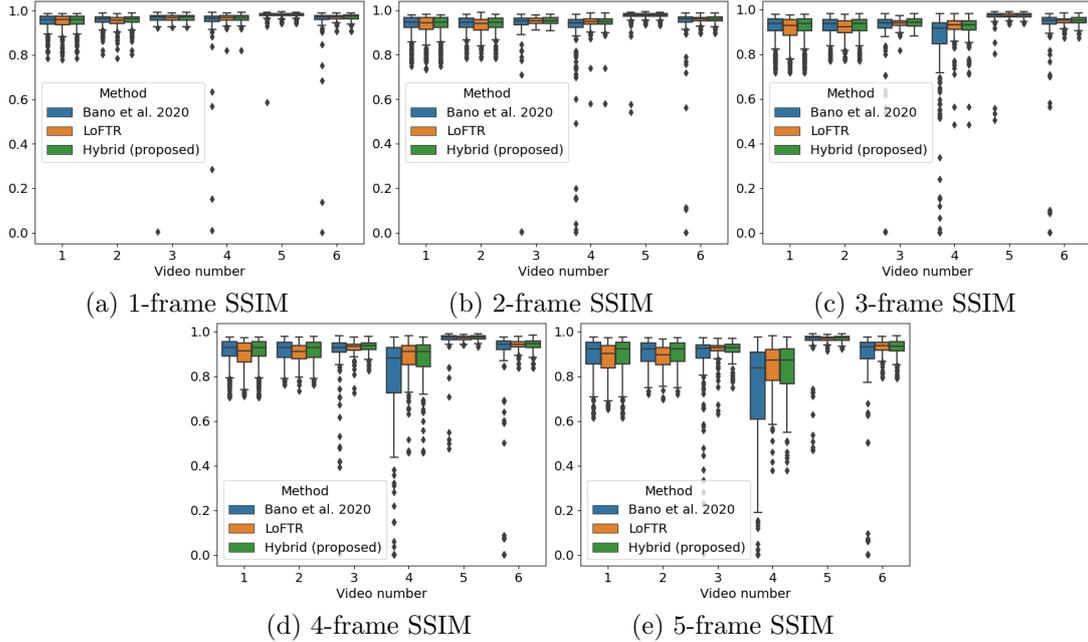


Figure 4.: Quantitative comparison of the proposed hybrid framework with the vessel-based (Bano et al. (2020a)) and LoFTR (Sun et al. (2021)) matcher-based methods using the N -frame SSIM metric. Note that hybrid selects the best affine transformation from both vessel-based and LoFTR matcher-based methods, and does not accumulate drifting error.

and high median 5-frame SSIM for LoFTR matcher-based and hybrid methods compared to the vessel-based method. Hybrid method performance is significantly better in Video 3 than the LoFTR matcher-based method, and is comparable to LoFTR matcher-based in Video 4 to Video 6. This is because the majority of the pair-wise transformations from LoFTR-based are better than the vessel-based in these videos.

The experimental results show that the proposed vessel-guided hybrid framework is optimized to select the best pair-wise transformations from the vessel-based and LoFTR matcher-based methods, overcoming the limitations of these methods. As a result, the proposed hybrid framework is robust even in the absence of vessels and presence of heavy floating particles, low illumination, non-planar views and spotlight light source. Our method significantly advances the literature of fetoscopic mosaicking, and paves the way towards translating such a framework into clinical settings for assisting surgeons during fetoscopic procedures. **Future work** involves further reducing the registration error through loop closure and bundle adjustment (Li et al. (2021)), designing a real-time application based on the hybrid approach and testing its usability through in-lab and clinical trails.

5. Conclusions

We propose a vessel-guided hybrid fetoscopic video mosaicking framework for generating reliable virtual expanded field-of-view image of the intraoperative fetoscopic environment. The proposed framework benefited from both placental vessel-based registration (Bano et al. (2020a)) method and detector-free feature matching with trans-

formers (LoFTR) (Sun et al. (2021)) method used as a robust matcher for registration, resulting in overcoming the limitations of individual methods. Using an extended version of the publicly available fetoscopy placenta dataset (Bano et al. (2020a)), we experimentally showed that the proposed hybrid framework optimized itself to select the best pair-wise transformations from the two methods, hence showing significant performance improvement over the existing state-of-the-art (Bano et al. (2020a)) on fetoscopic mosaicking. The proposed framework is robust even in vessel-free, low-textured or low illumination non-planar views, which shows its potential towards clinical translation for assisting the surgeons during the TTTS procedure.

Acknowledgments

This research was supported by the Wellcome/EPSRC Centre for Interventional and Surgical Sciences (WEISS) [203145/Z/16/Z]; the Engineering and Physical Sciences Research Council (EPSRC) [EP/P027938/1, EP/R004080/1, EP/P012841/1]; the Royal Academy of Engineering Chair in Emerging Technologies Scheme, and Horizon 2020 FET Open (863146). For the purpose of open access, the author has applied a CC BY public copyright licence to any author accepted manuscript version arising from this submission.

References

- Alabi O, Bano S, Vasconcelos F, David AL, Deprest J, Stoyanov D. 2022. Robust fetoscopic mosaicking from deep learned flow fields. *International Journal of Computer Assisted Radiology and Surgery*. 17(6):1125–1134.
- Bano S, Casella A, Vasconcelos F, Moccia S, Attilakos G, Wimalasundera R, David AL, Paladini D, Deprest J, De Momi E, et al. 2021. Fetreg: Placental vessel segmentation and registration in fetoscopy challenge dataset. arXiv preprint arXiv:210605923.
- Bano S, Casella A, Vasconcelos F, Qayyum A, Benzinou A, Mazher M, Meriaudeau F, Lena C, Cintorrino IA, De Paolis GR, et al. 2022. Fetreg2021: A challenge on placental vessel segmentation and registration in fetoscopy. arXiv preprint arXiv:220612512.
- Bano S, Vasconcelos F, Amo MT, Dwyer G, Gruijthuijsen C, Deprest J, Ourselin S, Vander Poorten E, Vercauteren T, Stoyanov D. 2019. Deep sequential mosaicking of fetoscopic videos. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer. p. 311–319.
- Bano S, Vasconcelos F, Shepherd LM, Vander Poorten E, Vercauteren T, Ourselin S, David AL, Deprest J, Stoyanov D. 2020a. Deep placental vessel segmentation for fetoscopic mosaicking. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer. p. 763–773.
- Bano S, Vasconcelos F, Tella-Amo M, Dwyer G, Gruijthuijsen C, Vander Poorten E, Vercauteren T, Ourselin S, Deprest J, Stoyanov D. 2020b. Deep learning-based fetoscopic mosaicking for field-of-view expansion. *International journal of computer assisted radiology and surgery*. 15(11):1807–1816.
- Baschat A, Chmait RH, Deprest J, Gratacós E, Hecher K, Kontopoulos E, Quintero R, Skupski DW, Valsky DV, Ville Y, et al. 2011. Twin-to-twin transfusion syndrome (TTTS). *Journal of Perinatal Medicine*. 39(2):107–112.
- Casella A, Bano S, Vasconcelos F, David AL, Paladini D, Deprest J, De Momi E, Mattos LS, Moccia S, Stoyanov D. 2022. Learning-based keypoint registration for fetoscopic mosaicking. arXiv preprint arXiv:220713185.
- Daga P, Chadebecq F, Shakir DI, Herrera LCGP, Tella M, Dwyer G, David AL, Deprest J,

- Stoyanov D, Vercauteren T, et al. 2016. Real-time mosaicing of fetoscopic videos using SIFT. In: *Medical Imaging 2016: Image-Guided Procedures, Robotic Interventions, and Modeling*; vol. 9786. International Society for Optics and Photonics. p. 97861R.
- Dai A, Chang AX, Savva M, Halber M, Funkhouser T, Nießner M. 2017. ScanNet: Richly-annotated 3D reconstructions of indoor scenes. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. p. 5828–5839.
- Gaisser F, Peeters SH, Lenseigne BA, Jonker PP, Oepkes D. 2018. Stable image registration for in-vivo fetoscopic panorama reconstruction. *Journal of Imaging*. 4(1):24.
- He K, Zhang X, Ren S, Sun J. 2016. Deep residual learning for image recognition. In: *IEEE Conference on Computer Vision and Pattern Recognition*. p. 770–778.
- Li L, Bano S, Deprest J, David AL, Stoyanov D, Vasconcelos F. 2021. Globally optimal fetoscopic mosaicking based on pose graph optimisation with affine constraints. *IEEE Robotics and Automation Letters*. 6(4):7831–7838.
- Peter L, Tella-Amo M, Shakir DI, Attilakos G, Wimalasundera R, Deprest J, Ourselin S, Vercauteren T. 2018. Retrieval and registration of long-range overlapping frames for scalable mosaicking of in vivo fetoscopy. *International journal of computer assisted radiology and surgery*. 13(5):713–720.
- Reeff M, Gerhard F, Cattin P, Gábor S. 2006. Mosaicing of endoscopic placenta images. *INFORMATIK 2006–Informatik für Menschen, Band 1*.
- Ronneberger O, Fischer P, Brox T. 2015. U-Net: Convolutional networks for biomedical image segmentation. In: *International Conference on Medical image computing and computer-assisted intervention*. Springer. p. 234–241.
- Sarlin PE, DeTone D, Malisiewicz T, Rabinovich A. 2020. SuperGlue: Learning feature matching with graph neural networks. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. p. 4938–4947.
- Sun J, Shen Z, Wang Y, Bao H, Zhou X. 2021. Loftr: Detector-free local feature matching with transformers. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. p. 8922–8931.
- Tella-Amo M, Peter L, Shakir DI, Deprest J, Stoyanov D, Iglesias JE, Vercauteren T, Ourselin S. 2018. Probabilistic visual and electromagnetic data fusion for robust drift-free sequential mosaicking: application to fetoscopy. *Journal of Medical Imaging*. 5(2):021217.
- Tella-Amo M, Peter L, Shakir DI, Deprest J, Stoyanov D, Vercauteren T, Ourselin S. 2019. Pruning strategies for efficient online globally consistent mosaicking in fetoscopy. *Journal of Medical Imaging*. 6(3):035001.