

Tumour-focused 3D-2D Registration in Liver Laparoscopy

Mohammad Zohaib⁴, Erol Ozgur^{1*}, Mohammad Alkhatib¹,
Emmanuel Buc², Bertrand Le Roy³, Youcef Mezouar¹,
Adrien Bartoli²

^{1*}Institut Pascal, Clermont Auvergne INP, Clermont-Ferrand, France.

²University Hospital, Clermont-Ferrand, France.

³University Hospital, Saint-Etienne, France.

⁴MaLGa-DIBRIS, University of Genoa, Genoa, Italy.

*Corresponding author(s). E-mail(s): erol.ozgur@sigma-clermont.fr;

Abstract

Purpose. Augmented reality (AR) is a promising tumour-resection guidance tool in liver laparoscopy. Existing methods register a preoperative liver 3D model (reconstructed from CT or MRI) to 2D laparoscopic images. They use surface landmarks to solve for registration, then transfer and overlay the internal tumour. This has two main weaknesses, hindering clinical adoption. First, computing the registration is highly challenging due to *(i)* multimodal and multidimensional correspondences (colourless 3D model to 2D RGB images), *(ii)* restricted field of view (proximity of laparoscope to parenchyma), and *(iii)* liver deformations (breathing and instrument interactions). Second, extrapolating from the surface to deeper tumours increases uncertainty and may result in inaccurate AR.

Methods. We propose Tumour-focused 3D-2D registration (**Tf32**), a method that registers the tumour, not the parenchyma. **Tf32** exploits laparoscopic ultrasonography (LUS), which provides intraoperative cross-sections of the tumour and can be reliably localised with respect to the laparoscope. **Tf32** matches the LUS tumour profile to a set of tumour profiles simulated preoperatively from the preoperative tumour 3D model. This initial registration is then refined using Iterative Closest Point (ICP). A multiple-hypothesis implementation handles ambiguities and local minima arising from the finite profile sample set.

Results. **Tf32** achieved **91.07%** success rate in meeting the **10 mm** oncological margin on semi-synthetic data and established a new state-of-the-art on a public benchmark dataset.

Conclusion. Compared with competitors, Tf32 does not require manual initialisation nor the full liver 3D model, and is notably more accurate, making it well-adapted to clinical use.

Keywords: 3D-2D registration, tumour, liver laparoscopy.

1 Introduction

The accurate localisation of liver tumours is critical for precise resection and preservation of healthy tissue in laparoscopic surgery [1]. Existing guidance methods register a preoperative 3D model of the parenchyma and tumours to intraoperative 2D laparoscopic images (LAP). This 3D model is reconstructed from CT or MRI data, which are modalities generally unavailable during surgery. The accuracy of 3D-2D registration is critical to realise a precise tumour overlay on the surgical images, but remains a major challenge [2]. Most registration methods perform a complete liver registration [3–6] from anatomical landmarks defined at the liver surface. The landmark visibility and localisation may be uncertain; given the high flexibility of the liver, their predictive power for the intra-parenchymal volume registration, including the tumours, may be low, causing inaccurate guidance. Low accuracy and the requirement for the surgeon to edit the landmarks hinder clinical adoption, as they may cause misguidance, interrupt the surgical workflow, and prevent real-time video processing.

We propose Tf32, a fast and clinically-feasible guidance method. Akin to previous work, Tf32 registers a preoperative 3D model to a 2D LAP image. In contrast, Tf32 directly registers the tumour, without relying on the rest of the liver parenchyma. The key idea is to exploit LUS, an intraoperative modality used in laparoscopic and robot-assisted hepatectomy. The LUS probe enters the abdominal cavity through the same keyholes as the other instruments and is laid directly in contact with the liver parenchyma during surgery. Being a 2D modality, LUS allows the surgeon to localise the tumour, without however providing the tumour’s spatial extent¹. Tf32 fixes this core limitation of LUS by revealing the full 3D tumour directly on the laparoscopic image. Registration remains valuable even when the tumour is visible on LUS, for three reasons: *(i)* electrocautery, acoustic shadowing, and liver mobilisation intermittently blind or degrade LUS; *(ii)* LUS provides only 2D cross-sections, making the 3D tumour extent difficult to perceive and compromising margin control; and *(iii)* trocar-constrained probe motion can cause subtle undetected rotations that displace the perceived tumour location increasingly with depth. AR addresses all three challenges by providing a stable, operator-independent 3D spatial reference. Technically, Tf32 matches the automatically-segmented 2D LUS tumour profile to a set of profiles simulated prior to surgery from the 3D tumour model, with an advanced ambiguity-aware mechanism. This provides a fast and robust initialisation, which is then refined by means of ICP. Unlike existing methods, Tf32 does thus not require an external

¹Existing methods for reconstructing 3D volumes from 2D free-hand US images are not directly applicable to LUS, as probe pose in a fixed coordinate frame cannot be obtained without integrating a tracking system — which existing LUS probes lack.

or manual initialisation, markers, or trackers. Being training-free, it is patient-generic and can be used off-the-shelf in any procedure. The strongest requirements are that the LUS probe’s pose be estimated from the LAP image [7] and that the tumour 2D profile be segmented in the LUS image [8–10]; both tasks are solved by existing methods, without hindering Tf32’s advantages. We show that Tf32 performs accurately on a semi-synthetic dataset and outperforms existing methods on a public benchmark.

Section 2 reviews related work. Section 3 describes Tf32. Section 4 presents experimental results. Section 5 concludes.

2 Related work

The data available in liver laparoscopy are the preoperative CT² and intraoperative LAP and LUS. We review registration methods for CT-LAP, which form the majority of guidance methods, for CT-LUS-LAP (to which Tf32 belongs), and for LUS-CT. All reviewed methods in the first two categories register the complete preoperative 3D liver model to LAP images and thus critically depend on the liver’s visibility in LAP.

2.1 CT to Laparoscopic Image Registration

We first review optimisation-based methods. These methods are not automatic, hence incompatible with real-time video processing, but are historically important. Koo et al. [4] proposed the first deformable method, which uses a Neo-Hookean model. The 3D model is registered using the liver shading, silhouette and anatomical landmarks. The method requires the surgeon to mark the anatomical landmarks on the LAP image, and requires 2-3 minutes, typically. Adagolodjo et al. [11] concurrently proposed deformable registration with a biomechanical model. The image constraints are the liver contours, manually marked in the LAP image by the surgeon. The method requires the surgeon to provide an initial pose. Espinel et al. [3] later proposed a hybrid method combining [11] with interactive deformation. Finally, Robu et al. [12] rigidly register the 3D model to a reconstructed intraoperative surface. Similarly to prior work, the method requires the surgeon to mark the liver contours in the LAP image.

The methods which more recently followed include learnt modules and are automatic. Labrunie et al. [13] perform rigid registration using coarse-to-fine pose estimation. The method relies on anatomical landmarks and the silhouette found by a U-Net. Although automatic, the landmarks suffer from under-detection, which hinders registration accuracy. Labrunie et al. [6] then proposed Liver Mesh Recovery (LMR), which performs the first learning-based registration, by regressing pose and deformation parameters from image primitives. The method is fast, as the computation cost is largely transferred to preoperative training. However, LMR is patient-specific, requiring retraining or fine-tuning for each patient. Finally, Mhiri et al. [14] proposed a neural method also exploiting image primitives via a lightweight Multi-Layer Perceptron (MLP) architecture. Training uses synthetic data generated via an as-rigid-as-possible method from the preoperative 3D model. The method remains patient-specific and does not generalise to deformations absent from the training data.

²We use ‘CT’ to mean the liver 3D model reconstructed from either CT or MRI, as the former is more commonly used.

2.2 CT, LUS and Laparoscopic Image Co-registration

Combining the three available modalities is challenging. Montaña-Brown et al. [15] proposed a learning-based method for pose estimation relying on patient-specific synthetic rendering and requiring retraining for each patient. Registration remains unstable due to liver deformation. More recently, Kalantari et al. [16] have proposed a geometric method for deformable registration. It exploits the cylindrical shape of the LUS probe to compute its pose, disambiguating the rotational ambiguity about its axis using a contact constraint. Although automatic, the method’s computational cost related to its complex optimisation process, typically requires 4 minutes per frame.

2.3 LUS to CT Registration

Existing methods in hepatectomy use Content-Based Image Retrieval (CBIR) for global initialisation. Ramalhinho et al. [17] match LUS vessel centroids and areas against a preoperative database. The method was then improved using Deep Hashing to learn descriptors for vessel lumen variability [18]. These vessel-based methods have four fundamental limitations. 1) *Geometric ambiguity*. Vessels are repetitive throughout the liver and matching a subset to the full 3D tree is thus inherently ambiguous. **Tf32** uses tumour geometry. 2) *Search space efficiency*. Methods [17, 18] search the entire liver parenchyma extent to resolve vascular branch ambiguities. **Tf32** restricts the search domain to a surgeon-marked surface patch at the probe position, drastically improving speed and accuracy. 3) *Training and generalisation*. Method [18] requires Siamese Network training with patient-specific fine-tuning. **Tf32** is training-free and patient-generic. 4) *Robustness*. Vessel methods are sensitive to segmentation errors: missing a vessel can compromise retrieval. **Tf32**’s tumour profile approach is robust to such errors.

3 Methodology

Tf32 requires the preoperative 3D liver model and seeks to register its 3D tumour model to a live intraoperative video stream from the laparoscope by exploiting the LUS image stream. An overview of **Tf32** is shown in figure 1. The algorithm with all steps and parameters is:

- Process the preoperative inputs and store the results: $P1(N, \theta)$, $P2$, $P3(t)$, $P4$, $P5$
- Repeat:
 - Process the intraoperative inputs and transmit the results: $I1$, $I2$, $I3(n)$
 - Perform registration: $R1(k, \ell)$, $R2$, $R3$
 - Perform augmentation: A

The ‘repeat’ loop processes new incoming frames. We next explain each step and parameter.

3.1 Preoperative Inputs and Processing

3.1.1 Inputs

We start from the patient’s preoperative 3D liver model, conveying the liver surface and a target tumour, in a preoperative coordinate frame. We request the surgeon to mark a free-shape patch on the liver surface, indicating the locations where they are likely to position the LUS probe during surgery.

3.1.2 LUS Contact Constraints

The LUS probe’s transducer must be in contact with the parenchyma to deliver an image. Therefore, the LUS probe is allowed three degrees of freedom: moving along two tangential directions on the surface patch and rotating about the surface normal. The LUS imaging plane is aligned with the surface normal.

3.1.3 Processing

We prepare LUS tumour profile matching by synthesising candidate profiles in 5 steps:

- (P1) We generate sample LUS probe contact poses on an $N \times N$ grid covering the liver surface patch and rotate the LUS probe at each grid node between 0° and 360° with a step of θ degrees. We experimentally found that $N \geq 20$ and $\theta \leq 6^\circ$ yield satisfying results – hence a total number of $\geq 20 \times 20 \times 360/6 = 24000$ samples.

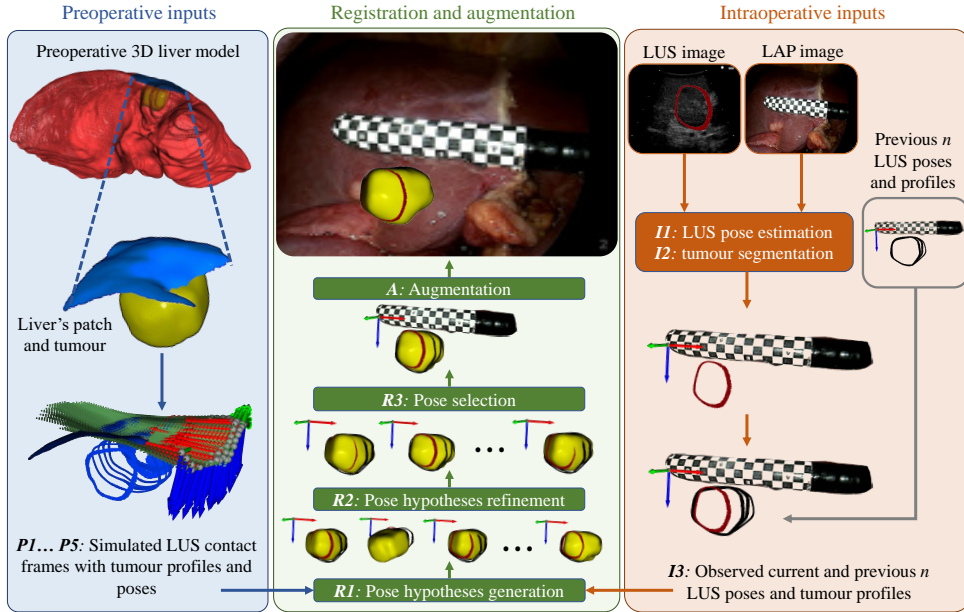


Fig. 1: Overview of the proposed Tf32 method, with steps Px , Ix , Rx and A ($x = 1, 2, \dots$) indicated as in the detailed description given in the main text.

Each sample contact pose is expressed in the preoperative coordinate frame and represents a contact coordinate frame.

- (P2) We build an LUS imaging plane at each LUS probe contact frame.
- (P3) We intersect the LUS imaging planes with the preoperative 3D tumour model to extract cross-sectional tumour profiles bounded by the transducer length t . This may yield partial profiles.
- (P4) We retain only the LUS probe contact frames whose tumour profiles cover at least half of their full profiles simulated by the unbounded LUS imaging planes.
- (P5) We form preoperative tuples, where each tuple contains: a sample LUS probe contact frame, the corresponding cross-sectional tumour profile, and the preoperative 3D tumour model’s pose, both expressed in that contact frame.

3.2 Intraoperative Inputs and Processing

3.2.1 Inputs

At surgery time, we receive (i) the current 2D LUS and LAP images, from the LUS probe and laparoscope respectively, and (ii) n previous LUS probe poses with the corresponding observed tumour profiles, used for temporal smoothing.

3.2.2 Processing

We prepare the profiles for matching:

- (I1) We estimate the current LUS probe pose from the current LAP image (*e.g.*, using the markerless method [7]).
- (I2) We segment the current LUS image to obtain the observed tumour profile (*e.g.*, using a neural method [10]).
- (I3) We gather the current LUS probe pose and tumour profile with the previous n ones, forming a set of $n + 1$ profiles coherently positioned in space and expressed in the current LUS probe frame.

3.2.3 Assumptions

We assume that the $n + 1$ LUS images show profiles from the same tumour and that the imaging planes are reasonably spaced, typically 1 to 5 *mm* apart. We assume that the LUS probe is not draped. These are mild assumptions, widely satisfied in practice.

3.3 Registration and Augmentation

3.3.1 Registration

We use the preoperative and intraoperative data and proceed in 3 steps:

- (R1) *Pose hypotheses generation:*
 - (R1.1) We generate multiple registration hypotheses by comparing the current observed LUS tumour profile to the preoperatively-generated profiles. We use the Hausdorff distance after pose alignment and keep the k best matches. Due to tumour

symmetry, some of these poses may be incorrectly mirrored, which is mitigated at step *R1.3*.

(*R1.2*) We retrieve the associated tumour poses of the best k matches from their preoperative tuples and instantiate the k corresponding preoperative 3D tumour models at these poses in the current LUS probe frame.

(*R1.3*) We then find the best ℓ matches by comparing each preoperative 3D tumour model instance to the n previous observed LUS tumour profiles using the Hausdorff distance.

(*R2*) *Pose hypotheses refinement.* We align each of the best ℓ instances of the preoperative 3D tumour model to the current observed LUS tumour profile with at most 10 ICP iterations. This yields ℓ refined instances of the preoperative 3D tumour model.

(*R3*) *Pose selection:*

(*R3.1*) We compute the Hausdorff distance between each refined preoperative 3D tumour model instance and the n previous observed LUS tumour profiles. We select the refined instance with the smallest Hausdorff distance.

(*R3.2*) Finally, we transform the selected instance into the laparoscope’s coordinate frame and consider it as the registered tumour.

3.3.2 Augmentation

(*A*) Given the laparoscope’s intrinsic calibration parameters, we augment the registered tumour on the current LAP image following [19].

4 Experimental Results

We conducted semi-synthetic and clinical data experiments. Following [20], our success criterion for registration is that the maximal point distance of the registered tumour volume to the ground-truth, as measured by the Hausdorff distance, be lower than the 10 *mm* oncological margin. This 10 *mm* threshold is the established oncological margin for liver resection [21, 22]. All experiments use Fujifilm’s L44LA probe. We set the Tf32 parameters as $N = 20$, $\theta = 6^\circ$, $t = 44$ *mm* for the L44LA probe, $n \in \{1, 2, 3\}$, $k = 15$ and $\ell = 5$.

4.1 Semi-synthetic Data Experiments

The dataset was collected under IRB00008526-2019-CE58. The dataset includes 21 distinct preoperative 3D tumour models, reconstructed from CT scans of 16 anonymous patients, with some patients contributing multiple tumours. The cases span a broad range of sizes and shapes to ensure comprehensive clinical evaluation. We first present the synthesis of realistic intraoperative scenarios and then the evaluation of Tf32.

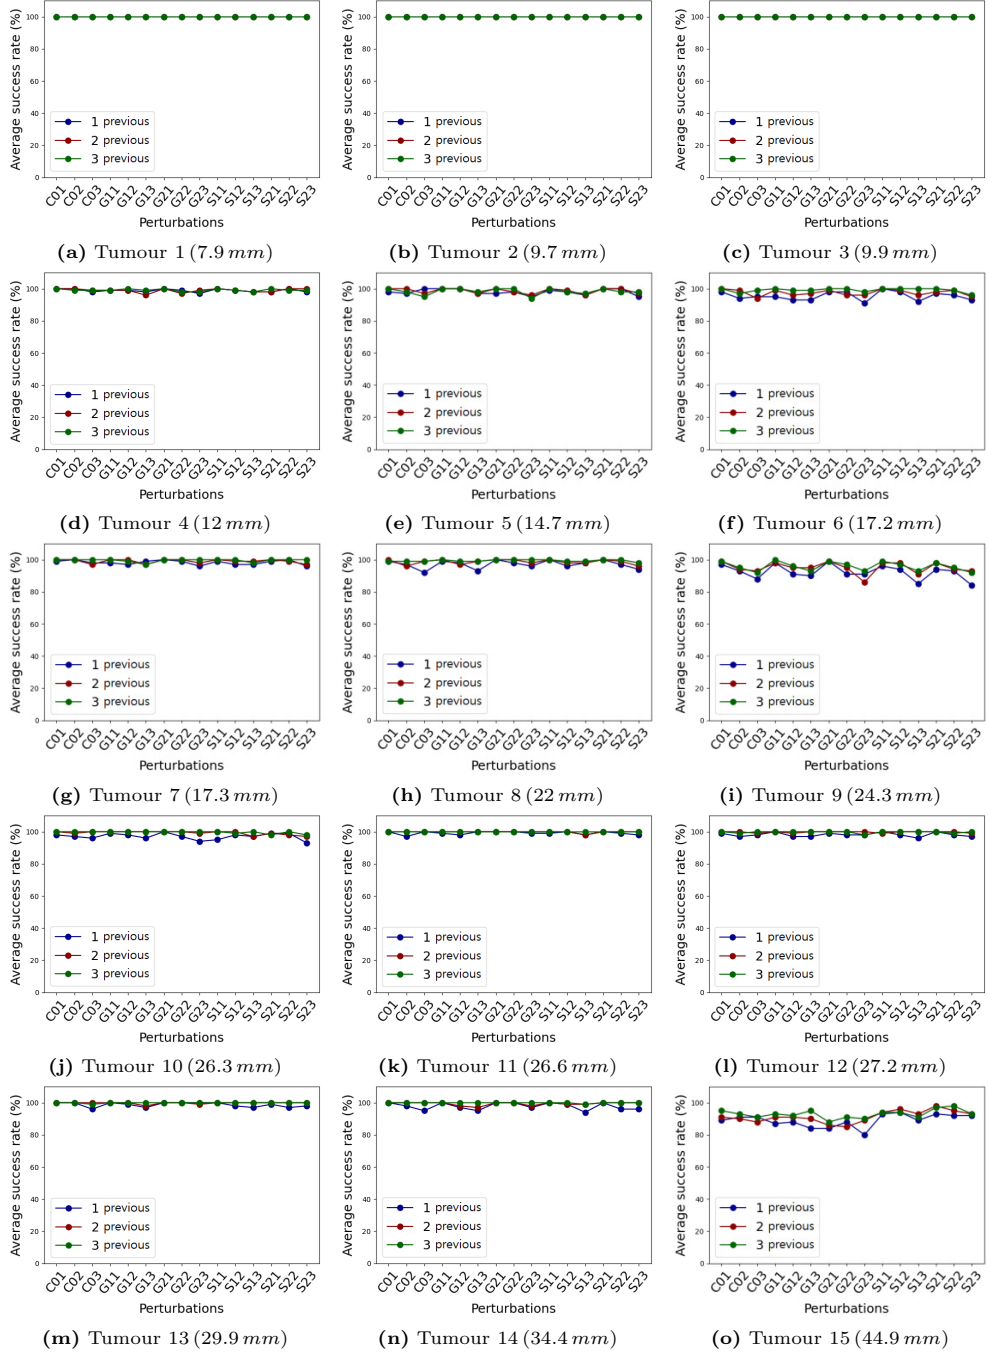


Fig. 2: Success rates for 15 tumours, sorted from smallest (a) to largest (o). The largest tumour roughly matches the transducer length ($t = 44 \text{ mm}$).

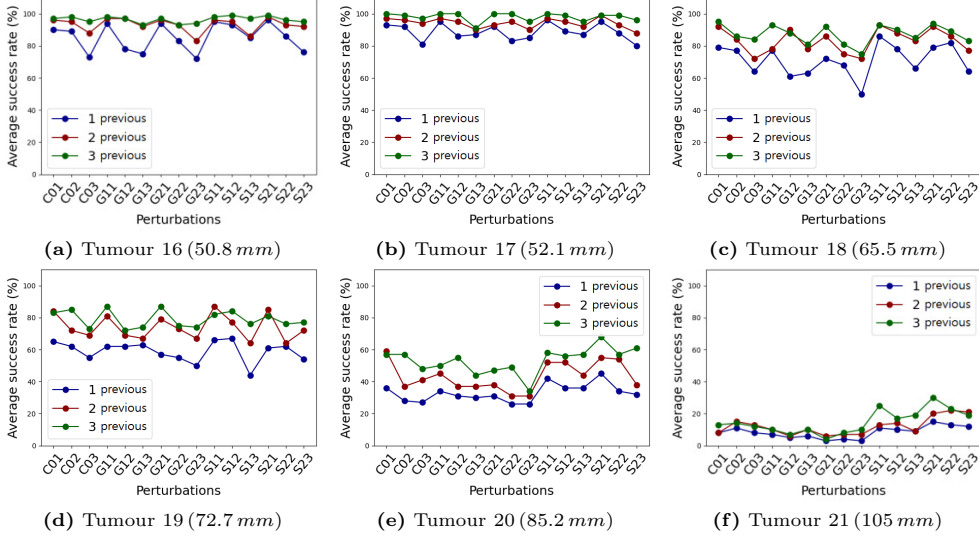


Fig. 3: Success rates for 6 tumours, sorted from smallest (a) to largest (f). The tumours are larger than the transducer length ($t = 44 \text{ mm}$).

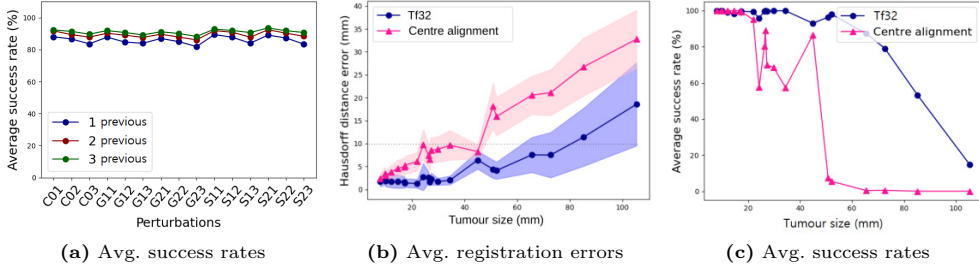


Fig. 4: Average success rates and registration errors for 21 tumours. (a) Success rates versus perturbations. (b, c) Registration errors and success rates versus tumour size using $n = 3$.

4.1.1 Generation of Realistic Intraoperative Scenarios

LUS Poses

For each preoperative 3D tumour model, we simulated 100 random LUS contact ‘current’ poses on the associated liver surface patch \mathcal{P} . For each current pose, we generated n ‘previous’ LUS poses with imaging planes parallel to the current plane. These previous poses are randomly placed at distances of 1–5 mm, 6–10 mm, and 11–15 mm from the current pose, along a direction orthogonal to the transducer and tangential to the patch surface. This yields 100 configurations per tumour model \mathcal{T} , each containing one current and n previous LUS poses.

Perturbations

We present three perturbation types to emulate the intraoperative laparoscopic surgery conditions:

- *Per1. Alteration of preoperative tumour shape.* This emulates LUS segmentation errors, tumour evolution, and deformation, by generating three intraoperative target shapes: (i) C , the convex hull of the preoperative tumour; (ii) $G\alpha$, the convex hull uniformly grown by $\alpha\%$; (iii) $S\alpha$, the convex hull uniformly shrunk by $\alpha\%$.
- *Per2. Alteration of target tumour pose with respect to current LUS pose.* This emulates liver deformation from surgical manipulation, heartbeat, and breathing. We applied a random rigid-body transformation to the target tumour (rotation by θ around a random axis through the centroid and translation of d along another random direction). Let γ be the distance from the tumour centroid to the LUS transducer. For perturbation amplitude β , we set $\theta = \beta^\circ$ and $d = (\beta/100)\gamma mm$, so translation scales with depth. Together, *Per1* and *Per2* emulate both liver and tumour deformation under LUS probe contact. Since tumours are smaller and stiffer than the liver, they primarily shift rather than deform as the probe pressure increases.
- *Per3. Alteration of previous LUS pose and tumour profile pair.* This emulates hand-held laparoscope movement. We apply a random rigid-body transformation to a previous LUS-profile pair: rotation by θ around a random axis through the transducer centroid, followed by translation of d along another random direction. For perturbation amplitude β , we set $\theta = \beta^\circ$ and $d = \beta mm$.

Scenarios

A scenario contains source and target data. The source data are a preoperative 3D tumour model \mathcal{T} and its associated liver surface patch \mathcal{P} . The target data is formed in four steps:

1. Altering a preoperative 3D tumour shape using perturbation *Per1* with an amplitude α (e.g., $C0$, $G1$, $S2$).
2. Displacing the shape using perturbation *Per2* with amplitude β .
3. Generating the cross-sectional profiles using the imaging planes of the current and n previous LUS poses generated earlier. The LUS imaging planes are bounded by the LUS probe’s transducer length, here $44 mm$.
4. Displacing each previous LUS pose and cross-sectional tumour profile pair using perturbation *Per3* with an amplitude β .

We encode the target data as $\mathcal{X}\alpha\beta$ where $\mathcal{X} \in \{C, G, S\}$. Finally, we generate 15 target data scenarios per preoperative 3D tumour model: $C01$, $C02$, $C03$, $G11$, $G12$, $G13$, $G21$, $G22$, $G23$, $S11$, $S12$, $S13$, $S21$, $S22$, and $S23$. Each scenario is played 100 times with each current and corresponding n previous LUS data. We repeat these with $n = 1, 2, 3$.

4.1.2 Results and Analysis

We compare the registered preoperative 3D tumour shape \mathcal{T} with the ground-truth intraoperative 3D tumour shape \mathcal{X} from a scenario $\mathcal{X}\alpha\beta$ using the Hausdorff distance. Figures 2 and 3 show success rates across perturbation scenarios, averaged over 100 configurations for all 21 tumour cases. Performance is high for tumours within the 44 mm transducer length, and decreases beyond this physical limit, characterizing Tf32’s robustness across tumour sizes and perturbation scenarios. We show in figure 4 the average success rates and registration errors across all tumour registrations. We observe that registration improves with number of previous LUS data. The current LUS tumour profile alone yields multiple solutions; including previous profiles eliminates rotationally symmetric and mirrored solutions. Across all scenarios and LUS pose configurations with 21 tumour cases, Tf32 achieves average success rates of 91.07%, 89.4%, and 85.95% in meeting the 10 mm oncological margin criterion with three, two and one previous profile, respectively. In figure 4 (b) and (c), we show the average registration errors and success rates versus tumour size using 3 previous profiles. We observe in figure 4 (b) that Tf32 meets the 10 mm oncological margin criterion for tumours smaller than 6 cm. We also provide for comparison in figure 4 (b) a simple method aligning the preoperative 3D tumour centre to the current LUS profile centre. This method meets the oncological margin criterion for tumours smaller than 2 cm, as expected.

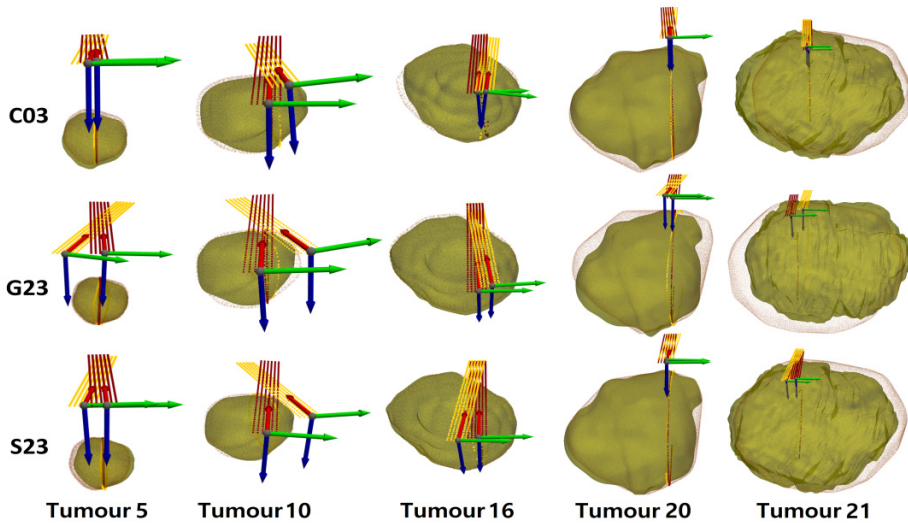


Fig. 5: Registration results for tumour cases 5, 10, 16, 20 and 21 from scenarios $C03$, $G23$ and $S23$, showing LUS contact surfaces (registered: yellow; target: red) and tumours (registered: green; ground truth: semi-transparent red).

We visualise cases of two small tumours (5th and 10th), one medium tumour (16th) and two large tumours (20th and 21st) in figure 5 from scenarios $C03$, $G23$, and

S23. The registered LUS contact surface (yellow) and the intraoperative target surface (ground truth, red) are shown with their coordinate frames, together with the registered preoperative tumour (green) and the intraoperative target tumour (ground truth, semi-transparent red). Alignment of the LUS contact surfaces indicates successful registration. Although rotational misalignment is present in cases 5 and 10, the registered tumours still overlap well with the ground truth due to rotational symmetry. This owes to the general property that a solid’s spatial extent is invariant to symmetries. Thus, **Tf32** improves on rotational-symmetric tumours.

4.2 Clinical Data Experiments

We compared **Tf32** against existing methods [4, 6, 11, 13, 16] solving for the 3D-2D registration of the complete liver model. Only [16] uses the intraoperative LUS data. Code and data were not publicly released for [17, 18]; as they follow complex pipelines, they could not be reproduced.

We evaluate quantitatively the methods using a public clinical benchmark [20], including four patients. We excluded the first 3 frames from each patient since **Tf32** requires up to 3 previous frames. For Patient 2, we additionally excluded 7 frames where the probe observed the liver’s posterior surface, as **Tf32**’s patch was located on the anterior surface. **Tf32** works with any segmentation method; for fair comparison, we used the pre-segmented tumour profiles from the dataset. We further evaluated **Tf32**’s robustness to LUS probe pose estimation errors, simulating using *Per3* on all intraoperative probe poses with displacement amplitudes $\beta \in \{0, 3, 5, 7, 10\}$. These values span the typical accuracy of method [7], which has $\beta \approx 5$. Method [7] is markerless and trackingless, operating across varied camera angles and probe orientations without additional trocars or unnatural manoeuvres, maximising clinical usability and robustness.

A low 2D contour error is required but insufficient for a successful result; in contrast, the benchmark TRE [20] accounts for the entire tumour shape. We therefore use two evaluation metrics: the TRE from [20] and the minimal and maximal Hausdorff distance (HD) for each of the n registration profiles. As accurate tumour excision is the primary clinical goal, tumour registration error — directly reflecting oncological margin safety — is a highly relevant metric. We report TRE results in table 1 and HD results in table 2.

Tf32 outperforms existing methods. Indeed, existing methods fail to satisfy the oncological margin criterion of laparoscopic liver tumour resection in all frames of the four patients. In contrast, **Tf32** meets this criterion across all patients and frames (*i.e.*, SR=100%) under LUS pose perturbations $\beta \leq 5$ with adaptive selection of $n \in \{1, 2, 3\}$. Incorporating one, two, or three previous LUS profiles in **Tf32** yields similar accuracy. Slightly better results may in some cases be obtained with a single profile (*e.g.*, 2.81 mm vs 2.87 mm in Patient 3, table 1 for **Tf32** with $\beta=3$). This is because earlier profiles, while resolving rotational symmetry and mirroring ambiguities, may also be visible under slightly different laparoscope poses. This indicates that **Tf32** is already stable with a single previous LUS profile. We show two augmentations from the registered tumours of patients 1 and 3 in figure 6.

Table 1: Results on clinical benchmark data [20]. MA is manual rigid registration, M1 is [11], M2 is [4], A-Pose is [13], Opt-B is [13], LMR is [6], and ST is [16]. TRE is in *mm*. SR is the success rate (%): the proportion of frames meeting the 10 *mm* oncological margin. × indicates failure. * indicates GPU implementation.

	Patient1		Patient2		Patient3		Patient4		
Number of frames →	5		11		6		5		
Tumour size (mm) →	17.4		11.5		13		30		
Method ↓	TRE	SR	TRE	SR	TRE	SR	TRE	SR	Time (s)
MA	16.82	0	24.42	0	33.39	0	15.37	20	-
M1	11.11	20	27.59	0	29.65	0	15.50	0	53
M2	12.68	0	25.35	0	26.70	0	18.72	0	150
A-Pose	7.38	80	52.96	0	8.54	83	18.64	0	5±3*
Opt-B	6.91	80	52.85	0	8.66	83	15.69	0	43±47*
LMR-OLA	20.26	40	46.54	0	17.26	0	16.72	0	0.013*
LMR-OHA	24.38	0	54.02	0	19.36	0	14.98	0	0.013*
ST	4.21	100	×	×	10.25	66.7	14.22	0	300
Tf32 ($n=1, \beta=0$)	2.34	100	1.44	100	0.79	100	4.11	100	4.68
Tf32 ($n=2, \beta=0$)	2.80	100	1.40	100	0.71	100	4.03	100	4.68
Tf32 ($n=3, \beta=0$)	2.91	100	1.08	100	0.85	100	5.00	100	4.68
Tf32 ($n=1, \beta=3$)	2.54	100	2.04	100	2.81	100	4.25	80	4.68
Tf32 ($n=2, \beta=3$)	2.53	100	1.90	100	2.03	100	4.62	80	4.68
Tf32 ($n=3, \beta=3$)	3.29	80	2.83	91	2.87	100	4.58	100	4.68
Tf32 ($n=1, \beta=5$)	3.82	100	3.63	100	3.75	100	6.09	60	4.68
Tf32 ($n=2, \beta=5$)	4.39	80	3.45	91	4.04	83,3	5.30	60	4.68
Tf32 ($n=3, \beta=5$)	3.76	100	3.19	91	5.25	100	5.64	100	4.68
Tf32 ($n=1, \beta=7$)	5.81	100	4.69	91	4.15	100	6.84	20	4.68
Tf32 ($n=2, \beta=7$)	5.91	100	5.76	91	6.20	100	6.06	20	4.68
Tf32 ($n=3, \beta=7$)	5.14	100	5.86	82	4.47	100	6.36	60	4.68
Tf32 ($n=1, \beta=10$)	7.98	33.3	7.12	45.5	9.811	50	6.57	20	4.68
Tf32 ($n=2, \beta=10$)	6.68	80	8.03	45.5	9.47	50	11.39	0	4.68
Tf32 ($n=3, \beta=10$)	8.46	80	7.66	45.5	9.48	50	6.16	80	4.68

5 Conclusion

We have presented a new method, Tf32, to solve the registration problem occurring in augmented reality guidance for liver surgery. In contrast to existing methods, we do not consider the entire liver parenchyma, but only the target tumour volume. Tf32 thus escapes the complex requirement of visible and reliable anatomical landmarks that existing methods rely on, and exploits LUS, which strongly boosts registration accuracy. Concretely, Tf32 meets the success criterion of a maximal error below the 10 *mm* oncological margin, which was never met by existing methods. Tf32 is currently limited to single-tumour registration. Future work will extend the method to multiple tumours, enable real-time implementation, expand the clinical study, refine preoperative pose sampling using surgical constraints (port locations, vascular proximity), and use tumours and vessels as intraparenchymal registration landmarks.

Table 2: Results on clinical benchmark data [20]. MA is manual rigid registration, M1 is [11], M2 is [4], A-Pose is [13], Opt-B is [13], LMR is [6], and ST is [16]. Minimum and maximum HDs are in *mm*, computed over $n = 2$ previous profiles.

	Patient1		Patient2		Patient3		Patient4	
Number of frames \rightarrow	5		11		6		5	
Tumour size (mm) \rightarrow	17.4		11.5		13		30	
Method \downarrow	Min	Max	Min	Max	Min	Max	Min	Max
MA	15.98	18.83	26.63	31.12	24.65	32.17	15.62	18.43
M1	12.61	30.75	19.57	29.09	28.18	57.31	16.56	50.10
M2	14.32	30.99	26.46	26.48	22.87	24.36	19.40	27.40
A-Pose	7.91	8.44	27.14	53.09	8.63	60.06	8.09	17.82
Opt-B	8.46	18.04	43.33	53.00	8.75	18.80	15.34	18.73
LMR-OLA	17.77	34.47	45.51	46.42	14.72	17.08	8.46	17.89
LMR-OHA	13.50	20.57	16.55	54.14	17.13	18.92	8.02	16.44
ST	6.55	8.35	\times	\times	11.43	14.77	15.42	31.90
Tf32 ($n=1, \beta=0$)	3.15	4.53	7.74	12.58	2.76	36.24	3.50	8.26
Tf32 ($n=2, \beta=0$)	5.20	11.41	12.84	15.26	2.97	6.84	8.30	19.06
Tf32 ($n=3, \beta=0$)	4.94	8.94	12.97	15.41	3.14	6.72	8.73	20.54
Tf32 ($n=1, \beta=3$)	5.51	7.49	14.19	24.11	4.46	5.38	7.87	10.44
Tf32 ($n=2, \beta=3$)	5.84	6.55	13.15	23.18	2.84	3.34	9.19	12.40
Tf32 ($n=3, \beta=3$)	6.34	8.83	13.10	22.81	4.12	4.94	9.37	14.11
Tf32 ($n=1, \beta=5$)	7.06	9.92	14.68	24.06	4.94	5.08	10.13	13.12
Tf32 ($n=2, \beta=5$)	6.44	7.96	14.42	22.89	4.92	6.54	10.57	13.14
Tf32 ($n=3, \beta=5$)	6.99	7.97	14.70	23.48	5.19	5.25	9.17	12.39
Tf32 ($n=1, \beta=7$)	11.21	15.71	16.64	26.40	4.58	5.03	12.48	15.59
Tf32 ($n=2, \beta=7$)	9.65	12.51	15.35	25.01	8.68	9.08	10.79	12.87
Tf32 ($n=3, \beta=7$)	9.41	11.59	15.65	24.10	6.68	7.00	9.74	10.30
Tf32 ($n=1, \beta=10$)	12.69	14.14	18.73	26.16	10.20	10.44	10.43	11.52
Tf32 ($n=2, \beta=10$)	10.44	12.46	17.29	24.69	10.11	10.84	14.94	15.59
Tf32 ($n=3, \beta=10$)	12.75	14.72	18.71	27.51	9.00	10.02	9.53	11.00

6 Compliance with Ethical Standards

Mohammad Zohaib declares to have no potential conflicts of interest. Erol Ozgur declares to have no potential conflicts of interest. Mohammad Alkhatib declares to have no potential conflicts of interest. Emmanuel Buc declares to have no potential conflicts of interest. Bertrand Le Roy declares to have no potential conflicts of interest. Youcef Mezouar declares to have no potential conflicts of interest. Adrien Bartoli declares to have no potential conflicts of interest.

All procedures involving human participants were in accordance with the ethical standards of the institutional and/or national research committee and with the 1964 Helsinki Declaration and its later amendments or comparable ethical standards. This study is also supported by ethical approval with ID IRB00008526-2019-CE58 issued by CPP Sud-Est VI in Clermont-Ferrand, France.

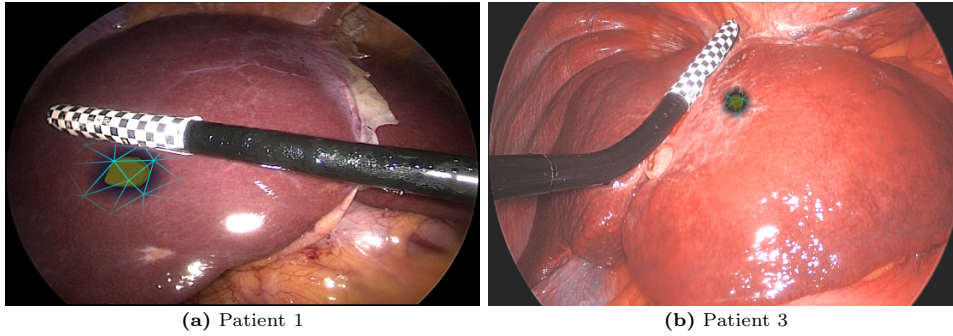


Fig. 6: Two tumour augmentations from patients 1 and 3 from the clinical data.

Informed consent was obtained from the patients included in the study.

Acknowledgements. This work was funded by ANR JCJC - IMMORTALLS.

References

- [1] FP Zhong, YJ Zhang, Y Liu, SB Zou. Prognostic impact of surgical margin in patients with hepatocellular carcinoma: a meta-analysis, *Medicine*, **96**, 37, (2017).
- [2] Ali S, Y Espinel, Y Jin, P Liu, B Güttner, X Zhang, L Zhang, T Dowrick, MJ Clarkson, S Xiao, Y Wu, Y Yang, L Zhu, D Sun, L Li, M Pfeiffer, S Farid, L Maier-Hein, E Buc, A Bartoli. An objective comparison of methods for augmented reality in laparoscopic liver resection by preoperative-to-intraoperative image fusion from the MICCAI2022 challenge, *MIA*, **99**, 103371, (2025).
- [3] Y Espinel, E Özgür, L Calvet, B Le Roy, E Buc, A Bartoli. Combining visual cues with interactions for 3D–2D registration in liver laparoscopy, *ABE*, **48**(6), 1712–1727, (2020).
- [4] B Koo, E Özgür, B Le Roy, E Buc, A Bartoli. Deformable registration of a preoperative 3D liver volume to a laparoscopy image using contour and shading cues, *MICCAI*, (2017).
- [5] Y Espinel, N Rabbani, TB Bui, M Ribeiro, E Buc, A Bartoli. Keyhole-aware laparoscopic augmented reality, *MIA*, **94**, 103161, (2024).
- [6] M Labrunie, D Pizarro, C Tilmant, A Bartoli. Automatic 3D/2D deformable registration in minimally invasive liver resection using a mesh recovery network, *MIDL*, (2023).
- [7] MM Kalantari, E Ozgur, M Alkhatib, E Buc, B Le Roy, R Modrzejewski, Y Mezouar, A Bartoli. LARLUS: laparoscopic augmented reality from laparoscopic ultrasound, *IJCARS*, **19**(7), 1285–1290, (2024).
- [8] A Zifan, K Zhao, M Lee, Z Peng, LJ Roney, S Pai, JT Weeks, MS Middleton, AE Kaffas, JB Schwimmer, CB Sirlin. Adaptive evolutionary optimization of deep learning architectures for focused liver ultrasound image segmentation, *Diagnostics*, **15**(2), 117, (2025).

- [9] H Ryu, SY Shin, JY Lee, KM Lee, H Kang, J Yi. Joint segmentation and classification of hepatic lesions in ultrasound images using deep learning, *European Radiology*, **31**(11), 8733–8742, (2021).
- [10] KS Sheela, V Justus, RR Asaad, RL Kumar. Enhancing liver tumour segmentation with UNet-ResNet: Leveraging ResNet’s power, *THC*, **33**(1), 1–15, (2025).
- [11] Y Adagolodjo, R Trivisonne, N Haouchine, S Cotin, H Courtecuisse. Silhouette-based pose estimation for deformable organs application to surgical augmented reality, *IEEE/RSJ IROS*, (2017).
- [12] MR Robu, J Ramalhinho, S Thompson, K Gurusamy, B Davidson, D Hawkes, D Stoyanov, MJ Clarkson. Global rigid registration of CT to video in laparoscopic liver surgery, *IJCARS*, **13**(6), 947–956, (2018).
- [13] M Labrunie, M Ribeiro, F Mourthadhoi, C Tilmant, B Le Roy, E Buc, A Bartoli. Automatic preoperative 3D model registration in laparoscopic liver resection, *IJCARS*, **17**(8), 1429–1436, (2022).
- [14] I Mhiri, D Pizarro, A Bartoli. Neural patient-specific 3D–2D registration in laparoscopic liver resection, *IJCARS*, **20**(1), 57–64, (2025).
- [15] N Montaña-Brown, J Ramalhinho, B Koo, M Allam, B Davidson, K Gurusamy, Y Hu, MJ Clarkson. Towards multi-modal self-supervised video and ultrasound pose estimation for laparoscopic liver surgery, *ASMUS Workshop, MICCAI*, (2022).
- [16] MM Kalantari, E Ozgur, M Alkhatib, N Rabbani, Y Espinel, R Modrzejewski, B Le Roy, E Buc, Y Mezouar, A Bartoli. Stronger together: Registering preoperative imagery, LUS, and MIS liver images, *MICCAI*, (2025).
- [17] J Ramalhinho, HFJ Tregidgo, K Gurusamy, DJ Hawkes, B Davidson, MJ Clarkson. Registration of untracked 2D laparoscopic ultrasound to CT images of the liver using multi-labelled content-based image retrieval, *TMI*, **40**(3), 1042–1054, (2020).
- [18] J Ramalhinho, B Koo, N Montaña-Brown, SU Saeed, E Bonmati, K Gurusamy, SP Pereira, B Davidson, Y Hu, MJ Clarkson. Deep hashing for global registration of untracked 2D laparoscopic ultrasound to CT, *IJCARS*, **17**(8), 1461–1468, (2022).
- [19] K Hanifati, M Alkhatib, E Ozgur, E Buc, B Le Roy, H Rante, Y Mezouar, A Bartoli. Hidden tumour visualization in augmented monocular liver laparoscopy, *Healthcare Technology Letters*, special issue: selected papers from *AE-CAI Workshop at MICCAI*, (2025).
- [20] N Rabbani, L Calvet, Y Espinel, B Le Roy, M Ribeiro, E Buc, A Bartoli. A methodology and clinical dataset with ground-truth to evaluate registration accuracy quantitatively in computer-assisted laparoscopic liver resection, *CMBBE Imaging and Visualization*, **10**(4), 441–450, (2022).
- [21] D Vandeweyer, EL Neo, JWC Chen, GJ Maddern, TG Wilson, RTA Padbury. Influence of resection margin on survival in hepatic resections for colorectal liver metastases, *HBP*, **11**(6), 499–504, (2009).
- [22] S Thompson, J Tottz, Y Song, S Johnsen, D Stoyanov, S Ourselin, K Gurusamy, C Schneider, B Davidson, D Hawkes, MJ Clarkson. Accuracy validation of an image guided laparoscopy system for liver resection, *Medical Imaging*, (2015).