




AR-Assisted Craniotomy Planning for Tumour Resection

Joost Wooning¹ and Mohamed Benmahdjoub²  and Theo van Walsum²  and Ricardo Marroquim¹ 

¹Delft University of Technology, Netherlands
²Erasmus MC, Netherlands

Abstract

Craniotomy is a procedure where neurosurgeons open the patient's skull to gain direct access to the brain. The craniotomy's position defines the access path from the skull surface to the tumour and, consequently, the healthy brain tissue to be removed to reach the tumour. This is a complex procedure where a neurosurgeon is required to mentally reconstruct spatial relations of important brain structures to avoid removing them as much as possible.

We propose a visualisation method using Augmented Reality to assist in the planning of a craniotomy. The goal of this study is to visualise important brain structures aligned with the physical position of the patient and to allow a better perception of the spatial relations of the structures. Additionally, a heat map was developed that is projected on top of the skull to provide a quick overview of the structures between a chosen location on the skull and the tumour. In the experiments, tracking accuracy was assessed, and colour maps were assessed for use in an AR device. Additionally, we conducted a user study amongst neurosurgeons and surgeons from other fields to evaluate the proposed visualisation using a phantom head. Most participants indeed agree that the visualisation can assist in planning a craniotomy and feedback on future improvements towards the clinical scenario was collected.

(see <https://www.acm.org/publications/class-2012>)

CCS Concepts

• *Human-centered computing* → *Scientific visualization; Heat maps*; • *Applied computing* → *Life and medical sciences*;

1. Introduction



Figure 1: An example of our AR visualisation to plan a craniotomy. A phantom skull is used for our evaluations.

To remove a brain tumour, sometimes it is necessary for surgeons to gain direct access to the brain by making an opening in the skull.

The procedure of creating this opening in the skull is called a craniotomy. Often, most of the healthy tissue between the skull opening and the tumour is removed, hence, to minimise damaging important brain tissue, this access path needs to be carefully planned. For this task imaging techniques are frequently used and specific structures are segmented. A neurosurgeon then looks at important structures in the brain which should be avoided, such as vessels and certain brain areas and fibres linked to specific neurological functions.

Commonly, surgical navigation systems are used to align the image data to the patient, often using landmarks [CDS*19], which allow to locate specific positions in the MRI data on physical positions on the patient. Surgeons can then directly draw the craniotomy location on the head of the patient before performing the actual intervention.

A downside of current surgical navigation methods is that the surgeon must constantly look at a screen instead of the patient, reducing hand-eye coordination and requiring more experience [HTB*17, VMR*14]. Furthermore, most navigation systems use a regular display, which can make it difficult to convey all spatial information of the 3D structures.

Augmented reality (AR) has great potential in improving this procedure in several ways. AR allows to superpose information directly on the patient, giving greater context and avoiding deviating

the attention to external displays. Moreover, the 3D models of the brain structures may increase the perception of spatial relations between them. Finally, an AR device may potentially replace the neuronavigation system, for example, by using the built-in camera of the device the position of the patient can be tracked using fiducial markers.

This work offers two main contributions. First, the proposal of an AR-based visualisation to assist in the planning of a craniotomy for tumour resection. The method combines traditional 3D visualisation of the structures with planning maps to summarise the risks of a possible path. Second, we provide an evaluation of the use of colour maps in such scenario and the tracking accuracy using the HoloLens. We conducted a user study to evaluate the visualisation with surgeons from different fields, including neurosurgeons. Even though the user study was carried out with a phantom head, we were able to gather useful insights which can help further research towards clinical trials.

2. Related Work

In a previous study, Stadie et al. proposed the use of virtual reality to plan a craniotomy [SKS*11]. However, the results of the planning are hard to bring to the patient space during the procedure itself. On the other hand, the use of AR allows to integrate digital data in the same space with the real patient. This may reduce planning time since the procedure can be done entirely in place.

Common AR devices come in the form of head-mounted displays (HMDs) or handheld devices [SPR*18]. HMDs can be composed of single displays [MMBML*16] or stereoscopic displays. The latter provides an enhanced depth perception. Nevertheless, even with stereoscopic displays extra depth cues, such as realistic shading, are necessary to achieve good depth perception [KOCC14].

Two types of stereoscopic displays are available: video see-through displays (VST) and optical see-through devices (OST) [DDROL16]. One limitation to the use of OST displays is the impossibility of visualising dark colours. This is due to the light added to the view of the user from the real scene. As a consequence, it is for example not possible to trivially render realistic shadows [MIKS19].

Another limitation of some OST-HMDs is the small field of view (FoV) of the display. On the HoloLens [Hol], for instance, the display's FoV is $30 \times 17^\circ$, this corresponds to $32 \times 18\text{cm}$ at a distance of 60cm . This small FoV is not necessarily an issue in medical settings since the user tends to focus on a small part of their vision [KKJ*17].

2.1. Object tracking

To align the 3D model with the patient, the head must be constantly tracked. Even if the patient is in a stationary position, this step is necessary since the HMD moves. There are two categories of tracking, outside-in and inside-out.

With outside-in techniques, external tracking systems are used to track the position of the patient and the HMD, then the relative position of the patient to the HMD is calculated to display the

model. The benefit is that external tracking systems, with a known high accuracy, can be used. The tracking data can be sent to the HMD where the relative position of the head to the HMD can be calculated. A visualisation error of 1.5mm was reported using this method [MNS*19]. Another approach is to use markers which can be detected both by the HoloLens and an external tracking system [HSH19].

Inside-out tracking on the other hand relies exclusively on the HMD's hardware, which performs the tracking and the visualisation. The setup time using inside-out tracking is shorter compared to the outside-in tracking. This makes it potentially advantageous for more time-critical situations.

Inside-out tracking can be achieved using various techniques such as using spatial maps or fiducial markers. The first technique uses spatial maps to track the position of the HMD in space, which does not track the patient, thus it can only be used if the patient is stationary. The second technique uses fiducial markers to superpose the 3D models on the correct location with respect to the patient. Since the markers are attached to the patient, any motion of the latter can be compensated. Consequently, the models are kept in the correct position and orientation as long as the fiducial markers are tracked. If the patient is stationary, it is also possible to place the markers in the background instead.

The HoloLens provides a spatial map of the surroundings which it continuously maintains. Several studies have been performed to find the accuracy of this spatial map and show that it is not stable. Disruptions in the spatial map can displace the models by around 5mm [VRCPI17]. Another study found that the HMD localisation has a positional error of around 10mm [LDZES18].

By using fiducial markers instead, a camera can keep track of predefined patterns on the markers. In this manner, objects with attached markers can be tracked and, consequently, the relative position between the HMD and the objects can be computed. Several studies investigated the creation of these markers, Garrido et al. [GJMSMCMJ14] list some of these and report that ArUco [GJMSMCMC16, RRMSMC18] provides a highly reliable method with a pixel error of about 0.16 with regards to the marker's corners location. Other proprietary solutions, such as Vuforia [Vuf], achieve a surface point localisation error of up to 1.92mm using fiducial markers [FJDV18].

Previous studies investigated the use of the HoloLens and fiducial markers as a navigation system [AQNK17, CDS*19]. The reported accuracy results ranged from 1.5mm to 5mm using various techniques. Unfortunately, these studies provide a single number accuracy, lacking insights with regards to error distribution over the tracking volume and the influence of the various parameters on the accuracy.

2.2. Visualisation for craniotomy planning

Previous work has been done in visualisations for craniotomy planning without AR. Many of these use MRI slices in combination with rendering surface models of anatomical structures [VGBD10]. Two studies included a view where the user can look down a proposed craniotomy position to see the structures between the created skull opening and the tumour [Die10, OHM10].

Kaniz et al. used a technique where the tumour emits rays from its surface towards the skull [Kai10]. These rays are attenuated by structures that should be avoided. The position on the skull with the higher number of rays represents the location with the least structures between the tumour and that part of the skull.

The opposite is also possible, tracing rays from the brain surface to the tumour. Here a single point in the brain is chosen as the destination and the amount of important structures between every point on the brain and the destination is calculated. This is then mapped to a colour resulting in a heat map on the brain surface [ZZM10].

Cutolo et al. used a VST-HMD visualisation to perform a craniotomy [CMC*17]. They mention that a downside of VST-HMDs is that the camera feed is not perfectly aligned with the eyes. In their work the craniotomy has been planned beforehand and the visualisation is used to show where to cut the patient's skin and skull. More recently, Incekara et al. [ISDV18] realised a pilot study for neurosurgery using the HoloLens. However, their work was limited to only viewing the tumour with the AR system and comparing the planning with standard neuronavigation.

3. Method

Below we describe our approach to AR for neurosurgical planning. We first discuss the tracking and coordinate systems alignment, subsequently we describe the registration method used for aligning the virtual model to the patient, the visualisation approach for the planning, and the rendering approach. Our method relies on the following imaging data: fMRI structures, vessels, DTI fibers, segmented skull and tumour. For a more comprehensive reading on fMRI and DTI for neurosurgery we refer the reader to the work from Ganslandt et al. [GNBG16] and Wolf et al. [WNM16].

3.1. Tracking

The tracking method in this work is based on markers and the ArUco library for detection. The frontal RGB HoloLens camera is used for tracking, and the mode with FoV of 48 degrees since it is larger than the HoloLens' FoV (30 degrees). Since the markers are close to the target subject, there is no need for wider FoVs during tracking. Moreover, a combination of markers was used to reduce the noise.

Apart from ArUco's built-in optimisations to detect markers along several frames, a Kalman filter [KYW18] was employed as a temporal filter in order to reduce the amount of noise during pose detection. The filter keeps track of previously detected poses, and provides an estimate of the next pose. When the next pose is retrieved from the frame it can be combined with the estimated pose to get the resulting filtered pose.

3.2. Registration

The planning approach projects image-based brain structures on the patients' head. This requires an alignment (registration) of the virtual patient model to the physical patient model. For this, we followed a landmark-based registration approach [AHB87], where

anatomical landmarks were pinpointed in the imaging data and with a tracked pointer on the physical model.

3.3. Planning map

The goal of the planning map is to give a quick overview of the structures between the skull and the tumour. Moreover, by mapping importance to the structures, the map also indicates which paths are potentially safer than others, that is, which paths avoids the removal of critical structures.




Figure 2: Illustration of the planning map. The skull is depicted in black, the tumour in pink and the structures to be avoided in red (for example important vessels). The planning map is constructed as a texture from the user's point of view. The resulting map is projected on the skull surface. Paths with less structures are drawn in green, with more structures in red, and orange in between. Weights can be assigned to certain structures to have more or less influence in the resulting texture.

To render planning maps in real-time an orthogonal projection camera is used at the position of the viewer pointing towards the tumour. The structures are then projected using this camera and the result is stored as a texture. Figure 2 demonstrates the concept in an example of a 2D cross section.

The current path is defined as the region between the tumour and the projected tumour on the skull in the direction of the viewer. The idea is then to accumulate on the 2D texture the structures along one ray, that is, along the path of a texel and the tumour. In case the structure is a DTI Fiber or vessel, we set constant weights since their actual sizes is not precisely defined, for example, the radius of the DTI fiber is arbitrarily defined in this system. For our tests, we used 0.5 and 0.2 as weights for vessels and DTI fibers, respectively.

For the fMRI, another solution is needed since they have arbitrary shapes. Projecting the front and back faces and computing the depth for each structure along each ray is a possibility, but since there is no guarantee that all shapes are convex, there might be multiple depth intervals for each ray. Therefore, ray tracing is used to compute the contributions of the fMRI structures to the planning map.

Finally, the options to turn the planning map off during the visualisation, keep the contour of the map only, as well as locking

the map for a given viewpoint are included in the system. The surgeon can then move around and further inspect the selected path. Figure 3 depicts some modalities of the visualisation.




Figure 3: Some modalities of the visualisation system. In this case, the planning map was locked in a position different from the current viewing angle. The images show: a) the 3D structures; b) adding the planning map; c) only 3D structures in the planning path (between the map and the tumour); d) contour only of the planning map. Note that the captured image shows the models more opaque than when viewed with the HoloLens.

3.4. Rendering

To shade the models a Blinn-Phong reflection model is used. In the case of the DTI fibers, plane primitives are shaded appropriately in a shader program to reassemble cylinders. Moreover, the fibers are shaded using the fractional anisotropy values, but other variables can be mapped to the colour if desired.

Finally, transparency was included in the renderer to be able to visualise structures behind others. Nevertheless, to implement alpha blending the order of the primitives, or fragments, need to be sorted, which is an expensive operation. Therefore, we resort to an approximate solution by sorting the structures based on prior-knowledge which avoids high computational costs in real-time.

Since it is our target, the last rendering layer is always the tumour, and the skull is the first layer. We pre-sort the other structures based on the following assumptions: most important vessels are near the skull; DTI Fibers are selected close to the tumour; and fMRI areas are also close to the tumour. Hence, the final order is skull, vessels, DTI fiber, fMRI areas, and tumour.

Naturally, this is not fail proof and may lead to fragments being wrongly composed, but only minor issues were noticed while avoiding a major performance hit in such limited computational scenario.

4. Experiments and results

To assess the method described, several experiments have been performed. The following subsections describe the data used, and the experiments on assessing the tracking accuracy, tests on the colour map selection for the HoloLens, and a user evaluation study.

4.1. Data

We developed a multi-modal renderer for the HoloLens. For the experiments with the renderer, the data from the 2010 IEEE Visualization Contest [IEE] was used. It contains both fMRI and CT scans. The isosurfaces for the skull, tumour, and vessels were extracted in addition to an fMRI area, by filtering voxels with significance less than 0.01%, and the DTI Fibers.

To perform our tests, a phantom skull was used and 3D the structures were manually transformed to fit within it because it did not match the shape of the phantom skull. Even though this might lead to implausible anatomical situations, upon consultation with the neurosurgeons during the design process, this was not viewed as an issue to evaluate the visualisation.

4.2. Tracking accuracy

The accuracy of the setup was assessed under various scenarios:

- *Camera resolution:* The tested resolutions are: 1408×792 , 896×504 , 704×396 , and 448×252 , the latter two are linearly interpolated downscaled frames created from the first two resolutions.
- *Marker size:* Four marker sizes were tested: 15, 30, 45 and 60 millimetres. The distance of the marker to the camera was kept relatively constant, at arm's length, about 50cm.
- *Incident angle:* The accuracy was tested at approximately 0° , 20° , 40° , and 60° . As the marker were hand-held, these angles were approximated as best as possible.
- *Marker motion:* No method was used to measure the accuracy at specific values. Instead, during the experiment the markers were moved randomly, then afterwards the motion was retrieved from the logged poses.

An optical tracking system was used as reference standard for the marker position. The optical tracking system used in the evaluation was the Polaris Vega VT [Pol]. It tracks IR-reflective spheres as markers with an RMS of 0.12mm. The coordinate system of the optical tracker N is used as the reference space. Markers tracked by the HoloLens are in the coordinate system of the HoloLens camera C . Both systems give a pose of the markers or spheres as a translation and rotation, we refer to these as a rigid transformation T_C^M from the HoloLens coordinate system to the marker M .

Spheres tracked by the optical system were attached to the HoloLens ($T_N^{H_{spheres}}$) and the markers ($T_N^{M_{spheres}}$). A pointer provided with the tracking system was used to retrieve the positions

of the marker corners relative to $M_{spheres}$, and then used to calculate $T_{M_{spheres}}^M$. The marker pose in N was calculated as follows:

$$T_N^M = T_N^{M_{spheres}} * T_{M_{spheres}}^M .$$

Figure 4 illustrates the setup.




Figure 4: Photos of the components used in the experiment. a) HoloLens with optical spheres attached. b) Marker with optical spheres attached. c) Pointer used to point out the corners of the marker.

Using $T_N^{H_{spheres}}$ and T_N^M the relative position of the marker to the spheres on the HoloLens is defined as follows:

$$T_H^M = (T_N^{H_{spheres}})^{-1} * T_N^M .$$

Subsequently, a calibration is performed to find the transformation between T_C^M and T_H^M . The optimal transformation is found using a least squares method [AHB87]. With this calibration, the marker's position in the HoloLens camera space as tracked by the optical tracker can be calculated as follows:

$$T_{C_N}^M = T_{calibration} * T_{H_{spheres}}^M .$$

Finally, the error represents the transformation between T_C^M and $T_{C_N}^M$:

$$T_{error} = T_C^M * (T_{C_N}^M)^{-1} .$$

The error can then be split into a translation and a rotation as Euler angles.

Figure 5 shows the poses from the HoloLens and the optical tracker next to each other. By interpolating the results from the optical tracker to the time points that are recorded by the HoloLens, the error of the marker pose detection in the HoloLens can be computed.

Figure 6 represents the accuracy obtained by using different resolutions and marker sizes. It is notable that the accuracy increases with larger markers and larger resolutions. Consequently, the results with a lower camera resolution of 448×252 were poor. For the resolution 704×396 , the results obtained are similar to those from the resolution 1408×792 .

Figure 7 shows that the translational accuracy decreases when the incident angle increases. However, the rotational accuracy reaches optimum around 40° .

Markers can move by rotating or by changing their position. We tested both scenarios and noted that, regarding the translation, the accuracy does decrease when the markers move faster. On the other

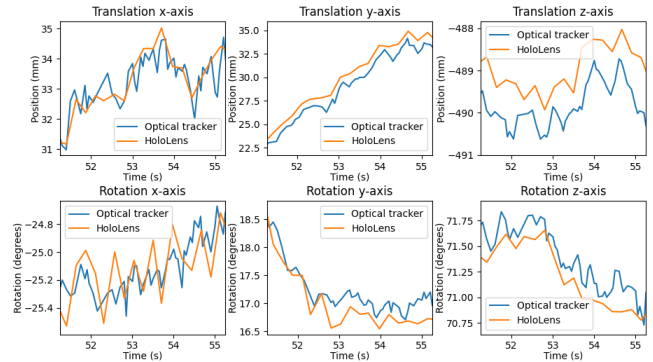


Figure 5: Marker pose in the coordinate system of the HoloLens camera. The optical tracker pose is aligned to this coordinate system using the performed calibration. A camera resolution of 1408×792 and a marker size of 60mm were used for this experiment.

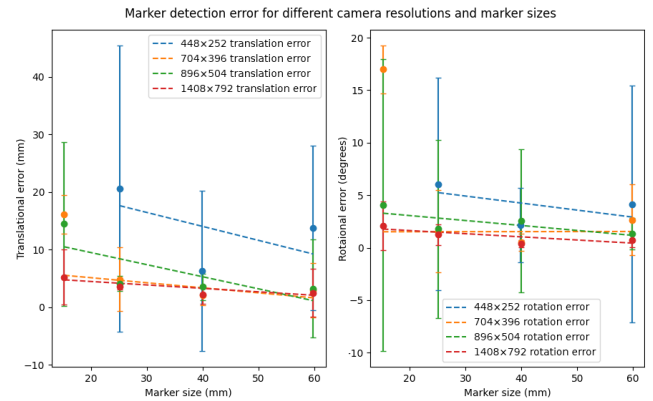


Figure 6: Error in marker poses from HoloLens for different marker sizes and resolutions. The dashed line represents the linear regression of the error. For resolution 448×252 with marker size 15mm the system was unable to detect the marker.

hand, we noted no strong correlation between the accuracy and rotation speed.

Another test using stationary markers was performed to assess the error introduced by noise in the pose detection. A 60mm marker and the HoloLens were fixed on a table at a distance of 70cm from each other. It was not possible to compare these results to the optical tracker, since the calibration only included a small depth range.

As noted from Figure 8, the noise is roughly 5 times less in the x/y directions than in z direction. Nevertheless, the error introduced by the noise is much smaller than the error found when comparing poses to the optical tracker.

To investigate the factors that influence the error introduced by the noise, the standard deviations of the results were analysed at different resolutions and incident angles (Figure 9).



Figure 7: The relation between marker detection error and incident angle. The dashed line represents the linear regression of the error. The incident angle is defined as the combined rotation along the x and z axes.




Figure 8: Tracking results from a stationary marker, both markers and HoloLens were placed on a table. Camera resolution is 1408×792 and marker size is 60mm. The distance from the camera to the markers was approximately 70cm. Note that the Kalman filter does not perform well with very small noise.

Even for the worst case the noise is quite low, with the maximum amount of noise at less than 1mm and 0.3° .

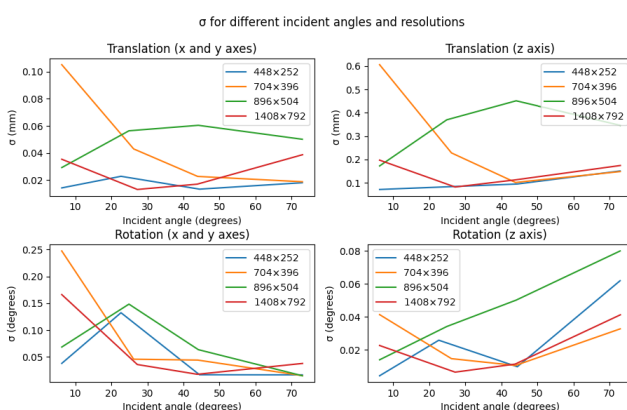


Figure 9: Standard deviations of stationary marker detection for different incident angles and resolutions. The marker used in this test were 60mm large. The distance from the camera to the markers was approximately 70cm.

4.3. Colour maps

Since in the HoloLens dark colours are mapped to transparency, colour maps in the HoloLens may not perform in the same manner as on regular displays (see Figure 10). Therefore, the choice of colour maps was investigated for use with the HoloLens and the planning map. The first step is to find the colour mapping used in the HoloLens which calculates the opacity.

The second step is to use this knowledge to composite the partially transparent colour map over a background and assess the performance.

A collection of 228 colour maps consisting of colour maps from Matplotlib as well as a few external ones were considered. Diverging or categorical maps were discarded leaving 58 perceptually linear maps. The lightness L^* value in the CIE Lab colour space is used as a metric for how perceptually uniform a given colour map is. Linear change of the other variables in this colour space should also provide perceptually linear colour maps, but these do not perform well when the size of the subject is small [Kov15]. Therefore, only the L^* parameter was used for evaluation. For easier analysis CAM02-UCS colour space was used, which has the same relevant properties and variables as CIE Lab [LCL06].

For a good colour map, the L^* value is expected to linearly increase or decrease, meaning that the colour map is perceptually uniform. Furthermore, the linear change should be over a large range of L^* values, making it easier to distinguish between the different values. Finally, these properties should also hold when the colour map is used in the HoloLens and composited over a non-uniform background. Because all dark colours are mapped to transparent, the colour map value is composited over whatever the user is looking at. For this experiment two fixed backgrounds were used, a white background and a sine wave varying from black to white with five periods over the whole length of the colour map. Note that the result for a black background would simply be the original colour map.

Figure 11 is an example of the *Reds* colour map composited over different backgrounds. The performance of the colour map can be analysed using a plot of the lightness values (see Figure 12).

Figure 13 illustrates how a perceptually linear colour map fails in different manners when used with the HoloLens. For most cases, the linear scale is broken with the composition over a background, which can be noticed as a zig-zagging pattern in the image. Moreover, many maps become too transparent for some ranges. 44 colour maps that presented similar behaviour were excluded.

From the remaining 14 candidates, shown in Figure 14, colour maps with very dark colours (which become transparent in the HoloLens) were excluded. These were colour maps such as the *cubehelix*, *cividis*, *viridis*, and *Purples* colour map. Even though the *summer* colour map is not very dark, it is very flat with a white background. Moreover, maps that are non-linear over at least one of the backgrounds were removed, thus colour maps such as *plasma*, *CET_L19*, *CET_L12*, *winter*, and *CET_L18*.

The remaining four maps preserve the original design when composited over a background. From these maps, *Wistia* was selected, since it is the only one that is close to being linear in a perceptual




Figure 10: Visualisation of some randomly selected colour maps and how they look when rendered in the HoloLens, with dark colours mapped to transparency. Every colour map is divided by a green line. For each one, the original colour map is drawn on top and the version with dark colours mapped to transparency on bottom. It is notable that for most maps there is a visible difference between both versions. Many of these colour maps are by design perceptually uniform, but are not when this mapping is applied.




Figure 11: Reds colour map. Left is the original map (or mapped over black background). Centre and right are the map composited over a white and sine wave backgrounds, respectively. Since this colour map is lighter for lower values, the background is less visible at this range.

scale. This colour map has a small lightness range by design, but it also has very little dark sections, with the lowest opacity at 0.95, meaning that the colour map will be very similar to the original no matter the background.

4.4. User evaluation

During the development of our approach, we had informal discussions with neurosurgeons. To provide a more structured evaluation




Figure 12: Reds colour map lightness plots, one plot for every background. On the x-axis the colour map input value linearly increases. The colour is then mapped using this value and resulting lightness values is plotted on the y axis.

of the visualisations, we conducted a user study. The goal was to evaluate if domain experts agree that it offers the correct insights in the data. Two main aspects are important: if the spatial perception of the structures is adequate; and if the planning map allows the user to make a more informed decision of where a craniotomy should be performed.

First, the participants had some time to get comfortable with the prototype and how to use it. Then they were asked to perform the following task: *Mark where you think is the best position and orientation to perform the craniotomy.*

Three different scenarios were created, and each participant was




Figure 13: Plots for some colour maps which were not used because of transparency issues. For every plot, on the x-axis the colour map input value linearly increases. Overlap in the plots is due to illustrative purposes. The colours in the plots show the mapped colour value of the input. For all colour maps, there are three sub-plots: original with no transparency; composited over a white background; and sine wave background.




Figure 14: Plots of the final 14 candidate colour maps. The layout of the plots is the same as in Figure 13

asked to perform the task for all scenarios in a randomised order. Prior to the first evaluation, a neurosurgeon was consulted over the scenarios, to make sure they were realistic.

The three scenarios are:

- Small tumour of 2.4cm diameter, at the surface of the skull.
- Large tumour of 8.4cm diameter, at the surface of the skull.
- Deep tumour, same scenario as the small tumour, but the tumour was artificially displaced due to lack of an appropriate dataset.

Following the tasks the participants were asked to answer a questionnaire (Table 1). Additionally, notes were taken during the sessions to perform some qualitative assessment.

Apart from the complete user study, shorter evaluations with other users were performed. In this case, users were asked to try out the visualisation with only the large tumour scenario and no possibility to turn parts of the visualisation on or off. Afterwards they were given the same questionnaire, except for the questions regarding the different scenarios.

A total of ten surgeons from different fields participated in this evaluation, the participants had an average of 11 years of experience. We categorised the participants into three groups: neurosurgeons, long evaluations and all the participants. Two neurosurgeons participated in the evaluation, with an average of 20 years of experience.

Six people participated in the long evaluation, including the neurosurgeons, with an average of 17 years of experience.

Since limited time was available for each participant, the visualisation options were controlled by the experimenters whenever the participant asked to turn on or off a feature.

The largest drawback reported by the participants was marker tracking, specifically, jitter and latency. Some participants also highlighted difficulties using with the device, while others experienced problems with the small FoV of the display.

The most informative part of the visualisation was the general 3D visualisation. Participants appreciated the fact that with AR the spatial relations between the tumour and other brain structures are clearly visible. According to the participants, this system has multiple advantages compared to other systems: it shows a 3D visualisation which can be viewed from different angles; it might reduce the planning time while still being accurate; it can lessen the damage to healthy tissue; and, finally, the craniotomy location can be easily drawn on the head using this AR-based method.

The two neurosurgeons suggested that some structures could be added: the ventricles, cranial nerves, corpus callosum and smaller vessels. However, they mentioned that the neurosurgeons themselves would prefer to do segmentation of the structures since they know what is important.

During the evaluation most clinicians asked to turn the planning map off at some point and only used the regular rendering of the structures to plan the craniotomy. One surgeon used the map to confirm his own planning.

Even though the planning map was perceived to be useful through the questionnaire with closed-ended questions, the open-ended questions and the discussions with the surgeons suggest that there are more factors that can influence the craniotomy planning aside from the amount of structure between the skull and the tumour. Practical problems such as how the skull can be cut open also play a role in the decision, and there are structures such as the cerebral falx which cannot be operated through. Nevertheless, it is possible to add them to the planning map.

Another issue was the shape of the planning map which is not always a simple outline of the tumour. For example, for a small tumour a small craniotomy is not necessarily the right solution, as it might not be possible to operate through it. Moreover, for deep large tumours a smaller craniotomy might still be large enough to operate through [SKS*11]. The planning map feature should then incorporate ways to change the shape of the craniotomy interactively.

5. Discussion

In this work a AR visualisation system is proposed for aiding surgeons in planning craniotomies. It involves rendering 3D structures overlaid on top of the skull as well as generating a dynamic planning map that summarises a possible craniotomy path. Moreover, an investigation of colour maps that are suitable for usage with the HoloLens and an analysis of the tracking accuracy were provided.

Our experiments on the tracking accuracy demonstrated that the tracking accuracy depends mostly on the resolutions of the camera, and less on the marker size. The average accuracy as measured

#	Question	Neurosurgeons	Long	All
1	It was clear what the different structures are	4.0	4.2	4.1
2	I was able to distinguish structures by their colour	4.5	4.3	4.1
3	The shape and size of the structures was clear	4.0	4.2	4.1
4	The spatial position of the structures was clear	4.0	4.0	3.5
5	The distance between the different structures was clear	4.0	4.0	3.4
6	It was clear if structures were behind or in front of other structures	4.0	4.3	4.0
7	The depth of the tumour in the head was clear to see	4.0	4.0	3.9
8	I did not perceive the latency of the system as an issue	3.0	3.2	2.6
9	I did not perceive jittery movement of the models as an issue	3.0	3.2	2.6
10	It felt like the alignment of the models to the head was accuracy enough	3.5	3.7	2.9
11	The visualisation did make clear how many structures lay between the craniotomy surface and the tumour	4.0	4.3	3.8
12	The visualisation did make clear what specific structures lay between the craniotomy surface and the tumour	3.5	4.1	3.7
13	I was able to easily move the planning map on the skull surface	3.5	3.5	2.8
14	The planning map helped me in finding the right position for the craniotomy	4.5	4.2	3.5
15	Structures hiding behind each other was not a problem	4.0	3.8	3.4
16	The visualisation was useful in the small superficial tumour scenario	4.0	3.8	N.A.
17	The visualisation was useful in the small deep tumour scenario	3.5	4.0	N.A.
18	The visualisation was useful in the large superficial tumour scenario	4.0	3.8	N.A.

Table 1: Closed questions average answers. Each question was answered with a Likert scale from (1) totally disagree to (5) totally agree.

in our experiments under optimal conditions is slightly less than 5 mm, which is considerably worse compared to common navigation systems, and likely insufficient for accurate planning. Even though the accuracy may not yet be enough to allow use in practice, it is likely to improve with newer devices, such as the HoloLens 2. Other investigative directions may also lead to improvements regarding the tracking accuracy, such as combining data from the depth sensor or the grey scale cameras provided with the HoloLens. Using other types of markers might also aid in this regard.

In addition, we assess an appropriate colour map for use in HoloLens, using criteria related to robust to projection over various backgrounds, and also a linear appearance. Based on the 'simulated' plots of the colour maps, only a few preserved the original design when projected over backgrounds. We did not explicitly evaluate the user perception of this colour map, and this would be an interesting topic for further investigations as well.

A user study confirmed the possible benefits of the proposed planning strategy and also points to future necessary steps towards clinical application. Users noted that spatial relations between structures are clearer than with 2D displays and provide more context of the operating site. The planning maps are promising but should incorporate more flexibility, such as interactively changing the shape and size. This also requires more control options for the user that can work in a seamless manner in the operating room.

We did not assess metrics such as planning time, or to what extent the resulting planning is of good quality. Such questions are relevant and interesting, and given the results of our current study, a follow-up study assessing the added value of the approach for clinical practice would be a logical next step. Another natural follow-up step is to test the visualisation with real human heads, instead of phantoms. Here, further investigation is necessary to handle issues related to hair and head movement.

6. Conclusion

We presented a study towards using AR for planning of craniotomies in neurosurgical planning, and performed experiments to assess the tracking accuracy of our prototype setup and appropriate colour maps. In addition, a user evaluation was performed. Whereas the accuracy of the tracking implemented in our approach currently is too low to be of use for surgical navigation, the user study revealed that the integrated visualisation of brain structures inside the skull, and projected on the skull surface, may be of benefit for planning in clinical practice. Further developments should focus on more accurate tracking and alignment, and a more flexible and intuitive interaction.

7. Acknowledgements

The dataset form the IEEE Visualization contest is courtesy of Prof. B. Terwey, Klinikum Mitte, Bremen, Germany.

References

- [AHB87] ARUN K. S., HUANG T. S., BLOSTEIN S. D.: Least-Squares Fitting of Two 3-D Point Sets. *IEEE Transactions on Pattern Analysis and Machine Intelligence PAMI-9*, 5 (1987), 698–700. doi:10.1109/TPAMI.1987.4767965. 3, 5
- [AQNK17] AZIMI E., QIAN L., NAVAB N., KAZANZIDES P.: Alignment of the Virtual Scene to the Tracking Space of a Mixed Reality Head-Mounted Display. *arXiv e-prints* (3 2017), arXiv:1703.05834. URL: <http://arxiv.org/abs/1703.05834>. 2
- [CDS*19] CAO A., DHANALIWALA A., SHI J., GADE T., PARK B.: Image-based marker tracking and registration for intraoperative 3D image-guided interventions using augmented reality. *arXiv e-prints* (8 2019), arXiv:1908.03237. 1, 2
- [CMC*17] CUTOLO F., MEOLA A., CARBONE M., SINCERI S., CAGNAZZO F., DENARO E., ESPOSITO N., FERRARI M., FERRARI V.: A new head-mounted display-based augmented reality system in neurosurgical oncology: a study on phantom. *Computer Assisted Surgery* 22, 1 (2017), 39–53. doi:10.1080/24699322.2017.1358400. 3

- [DDROL16] D'ANGELO T., DELABRIDA S. E., RABELO OLIVEIRA R. A., LOUREIRO A. A.: Towards a low-cost augmented reality head-mounted display with real-time eye center location capability. In *2016 VI Brazilian Symposium on Computing Systems Engineering (SBESC)* (2016), pp. 24–31. doi:10.1109/SBESC.2016.013. 2
- [Die10] DIEPENBROCK S.: Pre-Operative Planning of Brain Tumor Resections. *IEEE Visualization Contest* (2010). 2
- [FJDV18] FRANTZ T., JANSEN B., DUERINCK J., VANDEMEULEBROUCKE J.: Augmenting Microsoft's HoloLens with vuforia tracking for neuronavigation. *Healthcare Technology Letters* 5, 5 (10 2018), 221–225. doi:10.1049/htl.2018.5079. 2
- [GJMSMCMC16] GARRIDO-JURADO S., MUÑOZ-SALINAS R., MADRID-CUEVAS F. J., MEDINA-CARNICER R.: Generation of fiducial marker dictionaries using Mixed Integer Linear Programming. *Pattern Recognition* 51 (2016), 481–491. doi:https://doi.org/10.1016/j.patcog.2015.09.023. 2
- [GJMSMCMJ14] GARRIDO-JURADO S., MUÑOZ-SALINAS R., MADRID-CUEVAS F., MARÍN-JIMÉNEZ M.: Automatic generation and detection of highly reliable fiducial markers under occlusion. *Pattern Recognition* 47, 6 (6 2014), 2280–2292. doi:10.1016/J.PATCOG.2014.01.005. 2
- [GNBG16] GANSLANDT O., NIMSKY C., BUCHFELDER M., GRUMMICH P.: fmri in neurosurgery. In *fMRI Techniques and Protocols*. Springer, 2016, pp. 801–815. doi:10.1007/978-1-4939-5611-1_25. 3
- [Hol] HoloLens (1st gen) hardware. <https://docs.microsoft.com/en-us/holoLens/holoLens1-hardware>. Accessed: 2021-06-30. 2
- [HSH19] HEINRICH F., SCHMIDT G., HANSEN C.: A novel Registration Method for Optical See-Through Augmented Realty Devices and Optical Tracking Data. In *GI VR/AR Workshop* (2019). 2
- [HTB*17] HERRLICH M., TAVAKOL P., BLACK D., WENIG D., RIEDER C., MALAKA R., KIKINIS R.: Instrument-mounted displays for reducing cognitive load during surgical navigation. *International Journal of Computer Assisted Radiology and Surgery* 12, 9 (2017), 1599–1605. doi:10.1007/s11548-017-1540-6. 1
- [IEE] Ieee visualization contest data. <http://sciviscontest.ieeevis.org/2010/data.html>. Accessed: 2021-06-30. 4
- [ISDV18] INCEKARA F., SMITS M., DIRVEN C., VINCENT A.: Clinical Feasibility of a Wearable Mixed-Reality Device in Neurosurgery. *World Neurosurgery* 118 (2018), e422–e427. doi:10.1016/j.wneu.2018.06.208. 3
- [Kai10] KAINZ B.: Multivariate Beam Ray Obstacle Visualization for Brain Tumor Resection. *IEEE Visualization Contest* (2010). 3
- [KKJ*17] KUHLEMANN I., KLEEMANN M., JAUER P., SCHWEIKARD A., ERNST F.: Towards X-ray free endovascular interventions – using HoloLens for on-line holographic visualisation. *Healthcare Technology Letters* 4, 5 (10 2017), 184–187. doi:10.1049/htl.2017.0061. 2
- [KOCC14] KERSTEN-OERTEL M., CHEN S. J. S., COLLINS D. L.: An evaluation of depth enhancing perceptual cues for vascular volume visualization in neurosurgery. *IEEE Transactions on Visualization and Computer Graphics* 20, 3 (2014), 391–403. doi:10.1109/TVCG.2013.240. 2
- [Kov15] KOVESI P.: Good Colour Maps: How to Design Them. 2015. URL: <http://arxiv.org/abs/1509.03700>. 6
- [KYW18] KAM H. C., YU Y. K., WONG K. H.: An Improvement on ArUco Marker for Pose Tracking Using Kalman Filter. In *2018 19th IEEE/ACIS International Conference on Software Engineering, Artificial Intelligence, Networking and Parallel/Distributed Computing (SNPD)* (2018), pp. 65–69. doi:10.1109/SNPD.2018.8441049. 3
- [LCL06] LUO M. R., CUI G., LI C.: Uniform colour spaces based on CIECAM02 colour appearance model. *Color Research & Application* 31, 4 (8 2006), 320–330. URL: <https://doi.org/10.1002/col.20227>, doi:10.1002/col.20227. 6
- [LDZES18] LIU Y., DONG H., ZHANG L., EL SADDIK A.: Technical evaluation of HoloLens for multimedia: A first look. *IEEE Multimedia* 25, 4 (10 2018), 8–18. doi:10.1109/MMUL.2018.2873473. 2
- [MIKS19] MANABE S., IKEDA S., KIMURA A., SHIBATA F.: Shadow Inducers: Inconspicuous Highlights for Casting Virtual Shadows on OST-HMDs. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)* (2019), pp. 1331–1332. doi:10.1109/VR.2019.8798049. 2
- [MMBML*16] MARTINEZ-MILLANA A., BAYO-MONTON J.-L., LIZONDO A., FERNANDEZ-LLATAS C., TRAVER V.: Evaluation of Google Glass Technical Limitations on Their Integration in Medical Systems. *Sensors* 16, 12 (2016). doi:10.3390/s16122142. 2
- [MNS*19] MEULSTEE J. W., NIJSINK J., SCHREURS R., VERHAMME L. M., XI T., DELYE H. H. K., BORSTLAP W. A., MAAL T. J. J.: Toward Holographic-Guided Surgery. *Surgical Innovation* 26, 1 (2 2019), 86–94. doi:10.1177/1553350618799552. 2
- [OHM10] OSPELT D., HAUSD A., M S.: An Exploration and Planning Tool for Neurosurgical Interventions. *IEEE Visualization Contest* (2010). 2
- [Pol] Polaris vega. <https://www.ndigital.com/medical/products/polaris-vega/>. Accessed: 2021-06-30. 4
- [RRMSMC18] ROMERO-RAMIREZ F. J., MUÑOZ-SALINAS R., MEDINA-CARNICER R.: Speeded up detection of squared fiducial markers. *Image and Vision Computing* 76 (2018), 38–47. doi:https://doi.org/10.1016/j.imavis.2018.05.004. 2
- [SKS*11] STADIE A. T., KOCKRO R. A., SERRA L., FISCHER G., SCHWANDT E., GRUNERT P., REISCH R.: Neurosurgical craniotomy localization using a virtual reality planning system versus intraoperative image-guided navigation. *International Journal of Computer Assisted Radiology and Surgery* 6, 5 (2011), 565–572. doi:10.1007/s11548-010-0529-1. 2, 8
- [SPR*18] SOLBIATI M., PASSERA K. M., ROTILIO A., OLIVA F., MARRE I., GOLDBERG S. N., IERACE T., SOLBIATI L.: Augmented reality for interventional oncology: proof-of-concept study of a novel high-end guidance system platform. *European radiology experimental* 2 (12 2018), 18. doi:10.1186/s41747-018-0054-5. 2
- [VGBD10] VAILLANCOURT O., GIRARD G., BORÉ A., DESCOTEAUX M.: A Fiber Navigator for Neurosurgical Planning (NeuroPlanningNavigator). *IEEE Visualization Contest* (2010). 2
- [VMR*14] VIGH B., MÜLLER S., RISTOW O., DEPPE H., HOLDSTOCK S., DEN HOLLANDER J., NAVAB N., STEINER T., HOHLWEG-MAJERT B.: The use of a head-mounted display in oral implantology: A feasibility study. *International Journal of Computer Assisted Radiology and Surgery* 9, 1 (1 2014), 71–78. doi:10.1007/s11548-013-0912-9. 1
- [VRCP17] VASSALLO R., RANKIN A., CHEN E. C. S., PETERS T. M.: Hologram stability evaluation for Microsoft HoloLens. In *Medical Imaging 2017: Image Perception, Observer Performance, and Technology Assessment* (3 2017), Kupinski M. A., Nishikawa R. M., (Eds.), vol. 10136, International Society for Optics and Photonics, p. 1013614. doi:10.1117/12.2255831. 2
- [Vuf] Vuforia. <https://vuforia.com>. Accessed: 2021-06-30. 2
- [WNM16] WOLF R. L., NUCIFORA P. G., MELHEM E. R.: Dti in neurosurgical planning. In *Diffusion Tensor Imaging: A Practical Handbook*, Van Hecke W., Emsell L., Sunaert S., (Eds.). Springer New York, New York, NY, 2016, pp. 291–308. doi:10.1007/978-1-4939-3118-7_14. 3
- [ZZM10] ZHENG L., ZHANG Y., MA K.-L.: Quantitative Visualization of Access Paths for Preoperative Planning in Neurosurgery. *IEEE Visualization Contest* (2010). 3