

DRR-Guided Angulation

Leveraging Pre-operative CTA to Optimize Vessel Visualization in DSA-Assisted Thrombectomy

DSAIT5000:
Cristian Cutitei

Delft University of Technology

DRR-Guided Angulation

Leveraging Pre-operative CTA to Optimize
Vessel Visualization in DSA-Assisted
Thrombectomy

by

Cristian Cutitei


Supervisors: X. Zhang, F. Nijenhuis, T. van Walsum
Committee: J. Yang
Project Duration: September, 2025 - May, 2026
Presentation Date: 20th of May 2026
Faculty: Faculteit Elektrotechniek, Wiskunde en Informatica, Delft

Cover: doctor getting patient ready for ct scan - Freepik -
Magnific.com [https://www.magnific.com/free-photo/
doctor-getting-patient-ready-ct-scan_25053924.htm](https://www.magnific.com/free-photo/doctor-getting-patient-ready-ct-scan_25053924.htm)
Style: TU Delft Report Style, with modifications by Daan Zwaneveld

DRR-Guided Angulation: Leveraging Pre-operative CTA to Optimize Vessel Visualization in DSA-Assisted Thrombectomy

Cristian Cutitei

Pattern Recognition and Bioinformatics



I. Introduction

A. Background Information

Acute Ischemic Stroke (AIS) represents a critical challenge to public health across Europe, remaining a primary driver of long-term disability and mortality. Current epidemiological data indicate that ischemic events comprise approximately 85% of all stroke cases within the European region (Prendes et al., 2023). These events typically stem from an acute arterial occlusion, which precipitates a cascade of brain tissue damage due to the critical deprivation of oxygen and essential nutrients.

Current clinical guidelines establish Mechanical Thrombectomy as the gold standard for Large Vessel Occlusion treatment (Powers et al., 2019).

Fluoroscopy is the primary imaging modality used

during thrombectomy, the process through which an embolism is removed using a microwire to reach and retrieve or break the blockage. By creating an X-ray image of the patient, whilst they are injected with a contrast agent for better radiation absorption, a projection of the patient's vessels is obtained. By taking an X-ray image of the patient before any contrast agent is injected, a baseline is obtained. When subtracting the image taken with contrast from the baseline image, a map of only the vessels is obtained; this procedure is named Digitally Subtracted Angiography (DSA). This is very valuable when detecting obstructions. DSA provides a 2-D projection of complex, 3-D vascular anatomy.

Although mechanical thrombectomy is the standard of care for large vessel occlusions, the procedure relies entirely on fluoroscopic guidance where the

thrombus remains radiolucent (Munich et al., 2019). Interventionalists must navigate using radio-opaque markers and vessel "roadmaps." Consequently, precise visualization of the anatomy proximal and distal to the occlusion is crucial to prevent complications like perforation or distal embolization.

Both fluoroscopy and DSA are acquired in what is called a C-Arm machine. The machine features a table on which the patient lies horizontally, and around it a C-shaped guide rail, supporting an X-ray source at one end and a sensor on the opposite end. The C-Arm is capable of capturing images from multiple angles around the patient. The geometry of this machine is discussed in the appendix subsection VIII.B. In DSA, a common obstacle for interventionists conducting the procedure is vessel **foreshortening** and **overshadowing**, which is where the C-arm angle warps the real geometry of the vessel or obscures the target vessel, respectively. Traditionally, finding the "optimal acquisition angle", the angle that best preserves the geometry and does not obscure the vessel, involves trial-and-error, using extra contrast agent and ionizing radiation, which consumes precious time and increases the risk of contrast-induced nephrotoxicity and increased intake of ionizing radiation.

Despite the effectiveness of mechanical thrombectomy, the selection of an optimal C-arm acquisition angle remains a time-consuming and operator-dependent process. Current trial-and-error approaches increase procedure duration, radiation exposure, and contrast agent usage, thereby elevating patient risk. This work addresses the need for a systematic and patient-specific method to determine optimal imaging angles that minimize vessel foreshortening and overlap.

B. Related Work

Historically, automated view planning research has been dominated by cerebral aneurysm treatment. Methods such as those proposed (van der Weide et al., 1998) utilize 3D rotational angiography to identify the "perpendicular ring" of a vessel, ensuring the aneurysm neck is viewed in profile. However, translating these methods to ischemic stroke requires a shift from point-based targets (aneurysms) to segment-based targets (arterial trunks). Our approach adapts these geometric principles to the specific needs of thrombectomy roadmapping, focusing on the preser-

vation of the longitudinal axis of the target vessel.

Another, more recent approach, gave the user the flexibility of choosing the optimal angle. The advancement made to procedural planning introduced by (Panse et al., 2020) has focused on mitigating the limitations of 2D fluoroscopic guidance through computer-assisted visualization, by utilizing augmented reality to assist interventionists in identifying C-arm positions that minimize vessel foreshortening and overshadowing.

C. Contribution

The purpose of this paper is to report on the investigation of automatic approaches to finding optimal acquisition angles when conducting mechanical thrombectomy. The contributions made are:

- a mathematical formulation of the problem as well as a theoretical optimization function
- two automatic methods to find the near-optimal acquisition angle for mechanical thrombectomy in a target vessel, namely a grid search and a gradient descent approach
- assessment of the proposed solution on a large custom dataset derived from expertly annotated volumes

The remainder of this paper is organized as follows: Section II provides a detailed overview of the proposed methods. Section III describes the experimental setup, including the custom dataset creation, evaluation metrics and experiments. Section IV presents the results and discussions of the presented methods. Finally, section IV concludes the paper. There is additional information about the methods inside the appendix, mainly used for reproducibility.

II. Methodology

For Digitally Reconstructed Radiographs Guided Angulation, two methods were devised, each leveraging different assumptions about the search space.

The first method used grid search. This approach assumed that optimal acquisition angles are wide enough to be discovered by a sweep of the entire search space using a predetermined angle offset. This is essentially a brute-force approach, the optimal acquisition angle being the best-scoring one out of all the possible angles. The number of DRRs that need to

be generated is mostly dependent on the angle offset; therefore, the compute time is known beforehand.

The second method used a gradient descent approach. To enable this, it utilized a differentiable framework for generating DRRs. While traditional DRR generation is computationally expensive and non-differentiable, frameworks like DiffDRR (Gopalakrishnan and Golland, 2022) introduce GPU-accelerated, differentiable rendering. By making the projection operator differentiable with respect to the C-arm geometry, it became possible to use gradient-based optimization to rapidly simulate and find the optimal projection. This method assumed that optimal acquisition angles would be missed by a large offset in the aforementioned grid and that they could be found by simply rotating the volume around incrementally. To avoid local minima, multiple starts were created, but considerably fewer than the grid-search generates, 45 to be exact. An additional benefit of this method is that it provided an additional quality measure of optimal acquisition angles, namely, the optimal angle width by studying how steep or shallow the gradient is in that region.

A. Input

Both of the methods mentioned use the same input. These inputs include a CTA volume of the subject's brain, a segmentation of the vessels in the subject's brain, a specific segmentation of the target vessel (the one to compute the best view for), and a pivot point marking the center of the C-Arm rotation, which is ideally the occlusion point but can be any point inside the target vessel segmentation. All of the input can be extracted automatically from the CTA volume. To obtain an accurate 3D representation of the vascular tree and overcome the limitations of conventional thresholding (which often fails due to noise and overlapping structures), deep learning architectures are utilized. Specifically, DTU-Net (Liu et al., 2025) is used as the backbone for generating the real-world labelmaps, leveraging its dual-path features to capture both fine vessel details and global spatial context. Additionally, foundation models like VesselFM (Ahmad et al., 2024), trained via a masked autoencoder framework, can be leveraged to provide robust latent representations of the vascular tree's inherent topology. The pivot point can be found from

automatic occlusion detection (Kumar et al., 2023).

Additional assumptions regarding the input include specific search space angle bounds of -180° to 0° for the azimuth coordinate and -45° to 45° for the elevation coordinate. Furthermore, the system assumes a Field of View of 256 mm, a source-to-detector distance of 1045 mm, a detector pixel spacing of 1 mm, and a detector radius of 810 mm around the pivot. These parameters were derived from the specifications of C-Arm machine models produced by Philips (Philips Healthcare, n.d.), as well as through consultations with interventional radiologists.

B. Optimization function

To determine whether one angle is preferred over the other, an objective function whose values are aligned with the interventionists is needed. The most common sources of errors in the 2D projection of the 3D vessel structure are overshadowing and foreshortening.

The ideal viewing angle has the maximum intensity of projected pixels associated with the target vessel, especially its proximal regions, with the least target pixels overlap with other vessel structures or when their boundaries collide in the image plane.

Let the imaging volume be $\Omega \subset \mathbb{I}^3$. Let $\xi \in \mathbb{I}^2$, denoting the coordinate of a pixel in an image. $\mathbb{I}$ symbolize intergers.

Let the two (disjoint) labeled subsets be:

$$\begin{array}{lll} S_1 \subset \Omega & S_2 \subset \Omega & S_1 \cap S_2 = \emptyset \\ \text{(target vessel)} & \text{(rest of)} & \\ \text{segment)} & \text{vessel tree)} & \end{array}$$

While S_1 and S_2 represent the precise anatomical boundaries, the gradient-based optimization is performed over a modified volume Ω_{grad} . This ensures a wider capture range for the cost function, allowing the optimizer to "sense" and move away from potential overlaps. Ω_{grad} defines this smooth occupancy via a Signed Distance Transform applied to the subsets S_k , followed by a sigmoidal transfer function.

Let $d_k : \Omega \rightarrow \mathbb{R}$ be the signed distance to the boundary of S_k :

$$d_k(X) = \begin{cases} -\text{dist}(X, \partial S_k), & X \in S_k, \\ +\text{dist}(X, \partial S_k), & X \notin S_k. \end{cases} \quad (1)$$

The gradient-ready occupancy $f_k(X)$ is then obtained via the logistic function, where $\tau > 0$ controls the

sharpness of the boundary transition:

$$f_k(X) = \frac{1}{1 + \exp(d_k(X)/\tau)}. \quad (2)$$

Using this formulation ensures the objective is differentiable over the entire domain.

To prioritize the proximal regions of the target vessel over the distal branches, the weight w is defined as a spatially varying function $w(X)$ based on the geodesic distance.

Let $X_0 \in S_1$ be the reference coordinate representing the proximal origin of the target vessel. The geodesic distance $D_{\text{geo}}(X, X_0)$ between any point $X \in S_1$ and the root X_0 is defined as the length of the shortest path connecting them.

To ensure stability across different vessel lengths, we compute a normalized distance $s(X) \in [0, 1]$:

$$s(X) = \frac{D_{\text{geo}}(X, X_0)}{\max_{Y \in S_1} D_{\text{geo}}(Y, X_0)} \quad (3)$$

The final geodesic weight $w(X)$ is then assigned using an exponential decay function, which heavily penalizes distal regions while preserving proximal importance:

$$w(X) = w_{\min} + (w_{\max} - w_{\min}) \exp(-\beta \cdot s(X)) \quad (4)$$

where β controls the rate of exponential decay, and $w_{\min}$ and $w_{\max}$ define the lower and upper bounds of the assigned weights.

We define the weighted (geodesic) importance function for the target vessel $\tilde{\ell}_1$ and the binary occupancy for the background tree ℓ_2 :

$$\tilde{\ell}_1(X) = \begin{cases} w(X), & X \in S_1 \\ 0, & \text{otherwise} \end{cases}$$

$$\ell_2(X) = \begin{cases} 1, & X \in S_2 \\ 0, & \text{otherwise} \end{cases}$$

The line-integrated importance and occupancy are:

$$\begin{aligned} \tilde{I}_1(\xi) &= \int_0^{+\infty} \tilde{\ell}_1(r(\xi, t)) dt \\ I_2(\xi) &= \int_0^{+\infty} \ell_2(r(\xi, t)) dt \end{aligned}$$

The weighted projection objective is:

$$\begin{aligned} J(p, u) &= \underbrace{\int_{\mathbb{R}^2} \tilde{I}_1(\xi) d\xi}_{\text{visible importance}} \\ &\quad - \lambda_1 \underbrace{\int_{\mathbb{R}^2} \tilde{I}_1(\xi) I_2(\xi) d\xi}_{\text{weighted overlap penalty}} \\ &\quad - \lambda_2 \underbrace{\int_{\mathbb{R}^2} |\nabla \tilde{I}_1(\xi)| |\nabla I_2(\xi)| d\xi}_{\text{weighted silhouette penalty}} \end{aligned} \quad (5)$$

where $\lambda_1, \lambda_2 > 0$ control the strength of the penalties. The first term rewards visibility of the target vessel, while the second and third terms penalize rays whose projections of S_1 and S_2 overlap in depth or touch in the image plane. The default values for λ_1 is 1 and for λ_2 is 0.5, these weights were used throughout the experiment, unless stated otherwise.

C. Heuristic

The objective function mentioned Equation 5 does not address the second source of major errors, namely foreshortening. To minimize this error, the camera viewing direction $u^\top$ must be as perpendicular as possible to the normal vector $\hat{v}_s$, which is the vector formed by the two ends of the vessel label.

$$u^\top \hat{v}_s = \cos(\theta)$$

where θ is the angle between the two vectors. Perfect orthogonality corresponds to $\theta = \pi/2$, or equivalently $u^\top \hat{v}_s = 0$.

The second objective function is therefore defined as

$$J_\perp(u, a, b) = (u^\top \hat{v}_s)^2 \quad (6)$$

The first term of the first optimization function Equation 5 automatically captures this property, given that it tries to maximize the total area projected of the target vessel.


III. Experimental Setup

The role of the experiments is to test the validity of the assumptions made, find hyperparameters, the equivalence of the presented methods, and, finally, real-world effectiveness.

A. Dataset

For this study, we utilize the TOP-BRAIN 2025 dataset, released as part of the TopBrain Segmentation Challenge for Whole Brain Vessel Anatomy held in conjunction with The Medical Image Computing and Computer-Assisted Intervention Society (MICCAI) 2025 (Yang et al., 2025).

The dataset consists of high-resolution cerebral angiographic scans from both MRI and CTA modalities, from which we selected only the CTA volumes and labels. The multiclass label map consists of 35 labels arteries, a table of which is present in the appendix subsection VIII.C.



The figure illustrates example labels of the TopBrain Dataset. Each color represents a different vessel class (i.e. L-M1, BA)

Figure 1. TopBrain Labels


This dataset was chosen due to the high-quality, hand-annotated, multiclass labels and because of the CTA modality.

From this TopBrain dataset, a custom one, that suits the needs for testing the aforementioned methods was created using the following algorithm:

The labels containing veins were removed, and for every volume, two vessel volumes were created by combining vessels belonging to the left-mid and right-mid arteries, mimicking contrast being injected in only one brain hemisphere. Every arterial label was paired with one of the aforementioned label groups (e.g. R-MCA paired with right-mid label). The pivot point was chosen as the median point in the vessel segment, while the vessel segment ends were determined using PCA on the volume, taking the largest dimension as length and selecting the two voxels that were furthest apart. The proximal end of the vessel segment was chosen based on which point was closer along the axis of blood flow for that particular vessel label, according to an atlas. Finally, the individual vessel segment

volume was exponentially weighed using ($w_{\min} = 0.1$, $w_{\max} = 1.0$, and $\beta = 1.5$) based on the geodesic distance from the proximal end Equation 4, and the target vessel segments along with the remaining vessel tree can be observed Figure 2.

The resulting dataset comprised 729 combinations between distinct vessel labels and the vessel tree they belong to.



■ represents the entire volume belonging to the vessel tree. ■ represents the target vessel segment. The target segment has a gradient from ■ to ■, representing the transition from the proximal to distal end respectively.


Figure 2. Example of Vessel Labels

Additionally, to test real-world scenarios, where vessel segmentations are not precisely manually annotated, another dataset was derived from the custom one.

The dataset construction maintained the original CTA volumes from TopBrain, while vessel labels were generated using the DTU-Net automatic extraction method. To ensure realistic labels for the target vessel label, the weighted label from the previous dataset was masked with the DTU-Net vessel prediction, as illustrated in Figure 3. To keep consistent with the prior iterations, the pivot point remained unchanged. Finally, both vessel segment ends and the proximal point were selected following the same steps as in the previous custom dataset.

B. Evaluation Metrics

The objective function Equation 5 is parametrized by the sphere coordinates (azimuth and elevation). For every pair, the function will output a score, which can be interpreted as the quality of that angle. The larger the number the better, and it is all relative, the units used do not belong on any scale. Still, the azimuth



The ■ segmentation is larger and noisier as opposed to segmentations obtained from the TopBrain dataset.

Figure 3. Example of Realistic Vessel Labels


and elevation pair is most important for consistency, therefore to be able to compare methods with one another, a baseline was created.

Since the base grid-search algorithm computes the objective for the entire search space, we can safely assume that if an optimal angle pair exists, the grid-search algorithm can find it. Setting the angle offset to a small offset will ensure the angle is not missed by sparse sampling.

C. Experiments

To test whether the first term of the first optimization function Equation 5 automatically minimizes foreshortening, the search space was reduced to only capture the angles that are situated proximal to the ring perpendicular to the vessel segment $\hat{v}_s$.

To test whether the proposed solution has real-world viability, a user study was conducted. The experiment presented two interventionists with 40 CTA volumes of patients' brains. Each volume had its own respective arterial vessels labeled, and the target vessel was labeled individually, akin to Figure 2. The target labels were R-M1, L-M1, R-M2, L-M2, R-M3, and L-M3, chosen due to common ischemic stroke in this region, the size of the vessels, and the variability of visualization difficulty. The interventionists were then asked to move the volume in 3D space such that the target vessel label is projected onto the viewport with the least overshadowing and foreshortening. For each of these volumes, the angle chosen by the interventionist, along with the time it took to complete the experiment, was saved. These angles were then compared with



The camera icons represent the C-Arm sensor. In the left image, the cyan ■ dots represent all the positions from which an angle quality computation will be computed. In the right image, the red ■ segment represents the target vessel, the blue line is perpendicular to the vessel segment, and the yellow ■ lines represent the boundaries of this reduced search space, 30° from the blue ■ line.

Figure 4. Comparison of the size of the search space

each other to get an insight into inter-rater variability for the same volumes, comparing both time Δ_{Time} and distance. $\Delta\theta = \arccos(|\vec{v}_1 \cdot \vec{v}_2|)$ where

$$\vec{v} = \begin{bmatrix} \cos(el) \cos(az) \\ \cos(el) \sin(az) \\ \sin(el) \end{bmatrix}$$

Finally, the angles output by the proposed solution were compared against the interventionists' manually discovered angles.

A target vessel may have many near-optimal acquisition angles, and it is up to the interventionist to choose which one they prefer. The proposed solution can output the top n viewing angles according to the objective function Equation 5. The angles output by the proposed solution are different from the manually annotated ones, therefore, an additional user experiment was devised. Two weeks after the first experiment, the same interventionists were shown four images of the original volumes to reduce bias. These images featured optimal acquisition angles from four sources: the top two automated solutions, the interventionist's own previous choice, and the choice made by their colleague. The scope of this experiment is to allow the specialists to judge the images output by the automatic solution, especially in comparison with manually discovered angles.

The user study was conducted to evaluate the clinical preference for four different DRR projections across 40 distinct cases. A interventional-radiologist ranked the projections from 1 (Best Visibility) to 4 (Worst Visibility). Visual inspection of the rank distribution demonstrates a clear preference for the optimized projections over the baseline directions. No pre-existing distribution for the projection images was presumed. The null hypothesis is that all 4 projections are equally optimal. To get a conclusion for this assumption, the Friedman test was conducted.

Since the possible ranks are limited to 1 through 4, the sum of ranks assigned within each individual case is constant and equal to 10. Under the null hypothesis that all projections are of equivalent quality, each projection would be expected to receive a uniform distribution of ranks across all cases. Consequently, the mean rank for each of the four projections should converge to 2.5. The Friedman test evaluates deviations from this expected mean by aggregating the differences between the observed average ranks and 2.5. Substantial deviations result in a larger test statistic and, correspondingly, a p-value below the significance threshold of 0.05, thereby providing statistical evidence that the interventionist exhibits a consistent preference among the projections.


Finally, the role of each term of the Equation 5 was inspected. An ablation study was conducted to determine whether the effects on the results match assumptions.

IV. Results

The plots Figure 5 and Figure 6 show that annotator agreement depends on the specific anatomy being viewed. Certain vessels (R-M1 and R-M2) appear to have consistency, while others (L-M1 and R-M3) present discrepancies in both the chosen camera angles and the time taken to annotate.

Figure 5 highlights how different the two annotators were from one another in finding an optimal camera view. The annotators achieved the highest levels of inter-rater reliability regarding the R-M1 and R-M2 categories, with median angular differences around 18° to 20° and narrow interquartile ranges.

In contrast, L-M1 presents a skewed distribution. While its median error remains relatively low at approximately 22° , the interquartile range stretches up




Boxplot of mean, median and variance of the angular difference between the two interventional radiologists' manual angle choices on the y-axis, mapped across different vessel labels on the x-axis.

Figure 5. Inter-annotator Angle Difference per Anatomy

to 65° , with the upper bound exceeding 80° . L-M2 and L-M3 present moderate medians but wide overall spreads.

R-M3 shows a distinct pattern. The median angular difference is large, approximately 82° , accompanied by a wide interquartile range (spanning from roughly 48° to 85°).




Plot of mean, median and variance of the time difference between interventional radiologists manual angle on the y-axis, for different vessel labels, x-axis.

Figure 6. Time Difference per Anatomy

Figure 6 shows the difference in how long it took each annotator to complete the task. Annotators were consistent in their timing for R-M1 and L-M2, with differences under 20 seconds.

The M3 segments exhibited the greatest variance in annotation duration. Specifically, L-M3 demonstrated a broad distribution of completion times, with a median of 95 seconds and an upper bound reaching 200 seconds.

An analysis of the combined angular differences



Boxplot of mean, median and variance of the angular difference between the interventional radiologists' manual angle choices and the proposed automated solution on the y-axis, across different vessel labels on the x-axis. I_1 and I_2 represent the first and second interventionists.

Figure 7. Manual vs Automated Angle Differences


between manually selected viewing angles and those generated by the automated system (Figure 7) reveals a consistent, systematic divergence in optimal view selection. However, the maximum discrepancies rarely exceed 90° . Across the majority of the evaluated vascular anatomies, the median angular discrepancy hovers between 50° and 70° for both annotators when compared to the automated baseline.

Interestingly, the comparative performance between the individual interventionists and the algorithm varies heavily by segment. For instance, in the L-M2 anatomy, Interventionist 1 (I_1) deviates from the automated solution by a median of nearly 68° , whereas Interventionist 2 (I_2) aligns more closely, with a median difference of approximately 44° . Conversely, in the L-M1 segment, I_1 demonstrates closer alignment with the algorithm (median 35°) than I_2 (median 57°).

As shown in the rank frequency distribution Figure 8, Optimizer_Top_1 performed the best nominally, receiving the 1st-place ranking in 18 out of 40 cases (45%) and placing either 1st or 2nd in over 70% of the evaluations. Similarly, Optimizer_Top_2 showed strong performance, predominantly capturing 2nd-place rankings (19 cases). Conversely, the baseline projections (Interventionist_1 and Interventionist_2) performed worse, with both being ranked as the absolute worst projection (4th place) in 18 out of 40 cases each.

Following a Friedman test, a Nemenyi post-hoc analysis was conducted to determine the statistical significance of the observed differences. The resulting


Frequency of Ranks per Method



Distribution of ranks for both the manually annotated and the top_2 optimizer picks. O_1 and O_2 represent top_1 and top_2 optimizer picks, respectively. I_1 and I_2 represent the first and second interventionists.

Figure 8. Stacked Bar Frequencies

Nemenyi Post-Hoc Test P-Values



Heatmap of p-values, denoting statistical difference between solutions. O_1 and O_2 represent top_1 and top_2 optimizer picks. I_1 and I_2 represent the first and second interventionists.

Figure 9. Nemenyi Post-Hoc Test

pairwise p-value matrix highlights three primary findings. First, both optimized projections demonstrated statistical preference over the baseline projections; specifically, Optimizer_Top_1 significantly outperformed Interventionist_1 ($p < 0.001$) and Interventionist_2 ($p < 0.001$), while Optimizer_Top_2 similarly surpassed Baseline_Dir_1 ($p = 0.003$) and Baseline_Dir_2 ($p = 0.003$). Second, no statistically significant difference in clinical preference was observed between the two optimizers ($p = 0.803$); although Optimizer_Top_1 secured a higher number of absolute first-place votes, the overall rank distributions remained statistically

comparable. Finally, the analysis revealed no statistical distinction between the two baseline directions ($p=1.000$).

Figure 10 illustrates the relative impact of each weighted term within the optimization function across all vessel labels. The visualizations demonstrate that the second term, which penalizes vessel overshadowing, contributes minimally to the cumulative score across the entire dataset.

A similar pattern can be observed in the magnitude of the gradients Figure 11, showing low magnitudes for the second term t_2 as compared to the others.

An ablation study was conducted where the weight of the second term was reduced to 0, effectively removing the effect it had on the optimization. The same equivalency test to the baseline resulted in approximately 90% of the original angles being retained.

Because the Top-Brain dataset contains labels for all the major arteries inside the human brain, we computationally found how easy it is to find an optimal view for each label Figure 12 by creating a mean and variance from each score for every label.

Figure 13 it can be observed that the discrepancy in quality between the angles found varies little by score. The same pattern is present regarding the perpendicular heuristic; only some labels lose notable quality. The largest discrepancy in scores is observed in the realistic label scenario.

V. Discussion

Regarding the inter-annotator agreement on optimal camera views, there is no uniform error rate across different vessels. Notably, the time spent does not correlate with angular agreement. While annotators were consistent in timing for the L-M2 segment (Figure 6), the angle graph (Figure 5) shows they strongly disagreed on the final result for this same anatomy.

For L-M1, the wide interquartile range up to 65° and beyond (Figure 5) indicates that while the radiologists agreed on an optimal angle for some L-M1 cases, there were many instances where their chosen camera angles were vastly different. L-M2 and L-M3 share a similar issue. For R-M3, the large median angular difference (82°) and wide interquartile range suggest a systematic discrepancy or a divergence in protocol application between the two annotators for this specific vessel, often resulting in selected views that are nearly

orthogonal to one another.


Comparing the manual and automated angles (Figure 7) reveals that the median angular discrepancy hovers between 50° and 70° . Despite intra-anatomy variations, the overarching trend suggests that the algorithmic selections generally does not fully align with the clinical intuition or habits of the human experts.

To identify why interventionists prefer the views generated by the proposed solution, it is necessary to examine the specific mechanisms within the optimization objective function Equation 5 that drive the selection of these optimal viewpoints. As shown in the results (Figure 10), the second term, which penalizes overshadowing, contributes minimally to the score. Nevertheless, the optimal viewpoints selected by the algorithm consistently achieve effective, visual, overshadowing mitigation, indicating that the other terms already optimize for that. Further analysis of why the property is captured by the other terms not specifically designed for that, can yield a more robustly formulated optimization function.

The ablation study indicates that exclusion of the second term would not affect the quality of the views and would simplify the formulation. No effect would be seen in the computation time of the optimization since the bulk of the operations is taken up by the DiffDRR renderer. A more in depth ablation study can find whether other properties can emerge from the current optimization formulation Equation 5 or one that includes the second optimization function Equation 6 as a separate weighted term.


When analyzing the optimal view scores per label (Figure 12), one has to take into account the assumptions made when designing the dataset. Score differences can be attributed to both the size of the vessel and the location.

Finally, analyzing the label sensitivity (Figure 13) reveals that the largest discrepancy occurs in the realistic label scenario. Due to a large number of voxels belonging to non-target vessels, the target label gets overshadowed from every angle. A solution to clean the disconnected voxels from the volume, along with classification of the voxels, which would allow the same steps used when working with the TopBrain dataset, would greatly improve the results. Additional experiments would need to be conducted with interventionists to establish whether the score difference can indeed be used as a quality measure.



Plot of mean and variance of the scores of the terms that compute the optimization function Equation 5. The first term seeks to maximize the visibility of the target vessel in the 2D projection. The second term minimizes target vessel overshadowing in the 2D projection. Finally, the third term seeks to minimize the amount of false vessel tangents in the 2D projection. This is done for every artery class present in our modified dataset subsection III.A.

Figure 10. Separate Terms of Optimization function




Plot of mean and variance of the scores the terms that compute the optimization function Equation 5. The plot is over 80 iterations of the gradient-descent method of 20 cases.

Figure 11. Median Gradient Magnitude

VI. Visualization

After getting the idea that the proposed solution is worth pursuing further and will be a useful tool to the future interventionist, there was a shift in focus. The additional meetings with the medical staff moved from testing the efficacy of our solution to improving the user experience.

The implementation at the time could be viewed as a simple function that, given a CTA volume, can automatically output a list of all possible views scored based on the optimization function. The output was not designed to be human readable, but to be highly




Plot of mean and variance of the scores computed by the optimization function Equation 5 of every class present in custom dataset subsection III.A.

Figure 12. Label Score Analysis

modular and easily coupled to additional components. The first step was querying the output of the function for the top n angles. This, while being the easiest way to find the best angles for a certain case, has drawbacks. Namely, the reference point for the computed angles are the pre-operative CTA and the computed angles are displayed in sphere coordinates in the coordinate space of the DiffDRR rendering engine. The assumptions are that the patient has not changed position from the pre-operative CTA and that the interventionist is familiar with the coordinate space of DiffDRR, both of which are unrealistic assumptions.

This has led to the search of an intuitive way of dis-




Heatmap, for every class present in our modified dataset subsection III.A, the average score gap between methods is plotted. The larger the difference, the worse the proposed solution by that method is. NP represents no perpendicular, meaning the use of the entire search space. P represents the perpendicular heuristic applied. R represents realistic labels. Ablation represents the removal of the second term of the Equation 5.

Figure 13. Label Sensitivity Heatmap

playing the results of the optimization. Brainstorming sessions with the interventionist and residents yielded in solving both main issues with the approach at the time. By further processing the the pre-operative CTA, a face mask of the surface of the head can be obtained by first thresholding the volume based on the air and skin values and applying a cube-marching algorithm on the filtered volume to transform the voxel based representation to a surface, vertex based representa-

tion. This new mesh will, based on the resolution of the CT scan, accurately represent the skin surface of the patient. In the next step, the list of scores provided by the optimization, are normalized then arranged into a a sphere (the hemisphere is mirrored) surrounding the patients head. Finally, the face mesh is scored (colored), by projecting the sphere onto it. This process is displayed in this figure Figure 14.



The figures depict an abbreviated pipeline of the entire optimization process. First, the cyan points in the first two images represent where the C-Arm sensor will capture the images from. The third image represents a heatmap of the scores computed by the optimization function Equation 5. The fourth and fifth images represent the quality of the possible angles overlaid on the patient's skin surface.

Figure 14. Pipeline of optimization process

VII. Conclusion

This study demonstrates the clinical viability of automating C-arm view planning for DSA-assisted mechanical thrombectomy using pre-operative CTA. By formalizing an optimization function that maximizes target vessel visibility while penalizing overshadowing and foreshortening, the proposed automated methods successfully replace the guesswork inherent in traditional trial-and-error angulation or standard views. Experimental results, validated through a user study with interventional radiologists, confirm that the automated DRR-guided algorithm statistically outperforms manual view selection. Furthermore, translating these complex algorithmic outputs into an intuitive 3D face-mesh visualization bridges the gap between mathematical optimization and practical clinical usability. Integrating this automated view-planning framework into the standard clinical workflow holds significant promise for reducing procedure durations, minimizing patient exposure to ionizing radiation and contrast agents. Future work should focus on refining the algorithm to better handle noisy, real-world automatic vessel segmentations and conducting expanded in-vivo clinical trials to quantify the exact reduction in operative delays and patient outcomes.

VIII. Reflections

Data Science and AI Technologies (DSAIT) are rapidly transforming the medical field, replacing manual, time-consuming tasks with automated workflows. Historically, tasks like medical image segmentation required meticulous effort from trained professionals. Today, AI is successfully automating these processes.

In my research, automatically finding the optimal angle for performing mechanical thrombectomy would not be practically feasible without a robust, modular, and fast automatic solution for vessel extraction, vessel classification, and generating DRRs. By accurately classifying complex 3D vascular anatomy using deep learning architectures, namely DTU-Net, AI removes the preprocessing burden of mapping every vessel inside the human brain.

From the CTA volume, 3D convolutional networks exist to pinpoint the exact location of occlusions.

From the vessel segmentation, 3D convolutional networks exist to classify the extracted vessel tree, resulting in 3D volumes of different vessel segments.

This evolution in DSAIT ensures that vital anatomical data is (almost) instantly available, bridging the gap between raw diagnostic imaging and immediate, actionable procedural planning for acute ischemic stroke interventions.

Time is important because, like in every other domain, time is money. When a procedure is shorter and requires fewer resources, it becomes more affordable, and many more patients will benefit from the treatment.

Time is also brain, as people working in stroke treatment portray the two as being equivalent. In case of stroke, every second counts, as more and more tissue begins to asphyxiate and die. The faster the procedure is conducted successfully, the more chances of recovery the patient has.

A step in creating a successful pipeline is to ensure it is not static and can be improved when one of its working parts receives an upgrade (i.e., more accurate

vessel segmentation) or is adapted for a different purpose (i.e., optimal angle for cardiac infarcts). The automated angle selection technology developed is designed to be modular. The module behaves like a function with intuitive input and detailed output, which can be used to display the results in however the radiology interventionist sees fit (i. e., 3D overlay, angles, heatmaps).

The responsible implementation of this technology yields benefits across a diverse set of clinical stakeholders, directly mitigating the risks of the current standard of care.

Patients are the primary beneficiaries. The traditional trial-and-error method of finding the correct angle requires repeated injections of contrast dye and continuous X-ray imaging. By predicting the optimal angle beforehand, this technology drastically reduces procedure duration, minimizes exposure to ionizing radiation, and lowers contrast agent use, directly decreasing the risks of acute radiation sickness and contrast-induced nephrotoxicity.

Furthering these points, patients' families will experience less burden both financially and in terms of time and resources, taking care of the person who suffered the stroke.

For **interventional radiologist**, this tool significantly reduces the cognitive load and fatigue for the procedure. By automatically providing an optimal view that minimizes vessel foreshortening and overshadowing, it gives radiologists a clear "roadmap." This helps them navigate microcatheters more safely, preventing severe complications such as vessel perforation or distal embolization. Automating the angle selection standardizes the quality of the procedure, ensuring optimal outcomes regardless of the individual operator's prior experience level, thereby improving the hospital's overall efficiency and standard of stroke care.

This is also a downside. By removing the need to conduct the procedure manually, in time, the essential practice to conduct this operation without the automatic pipeline will be impossible, leaving specialists with an experience gap. One example of such a case is in the 1970s and 80s, pediatricians and ENT specialists were experts at managing acute epiglottitis, frequently performing emergency intubations to save children from sudden airway closure. However, the introduction of the Hib vaccine caused cases to plum-

met by over 99%, effectively turning a routine clinical crisis into a medical rarity. Consequently, entire generations of modern residents may never encounter a "classic" case, leading to a significant decline in the specialized "muscle memory" once required to handle this specific life-threatening emergency.

For **hospitals and healthcare systems**, reducing the time required for each mechanical thrombectomy, hospitals can optimize the use of their operating rooms and expensive C-Arm machines. These cost savings can cascade down to insurance prices, since strokes are becoming more and more common due to aging populations. This may, in turn, make health insurance more affordable and may lead to more people getting the treatment they need.

Medical Device Manufacturers (OEMs) will also be affected by the general adoption of my proposed solution. Companies that design and manufacture imaging equipment (such as Siemens, Philips, or GE) are critical stakeholders in the applicability of this research. The modular nature of my angulation technology means it can be directly integrated into their next-generation C-Arm systems as a software upgrade, if it meets the system requirements, or requires separate revisions of the existing hardware that include the required acceleration hardware. By offering automated, DRR-guided angulation, manufacturers can provide immense value to hospitals, effectively bridging the gap between their physical hardware and the latest advancements in AI-driven surgical planning.

The broader scientific community benefits from the modularity of the pipeline I have developed. Because the technology separates the vessel segmentation phase from the mathematical optimization and angulation phase, future researchers do not need to rebuild the entire system from scratch. As new, more accurate automatic segmentation models or occlusion-detection algorithms are developed, they can be seamlessly "plugged into" this pipeline. This creates a sustainable, open-ended architecture that accelerates future innovations in medical data science. Moreover, this area of research is not explored; this publication may spark interest in developing it further, with my work as the starting point.

Regulatory agencies are highly invested in the transition from manual medical procedures to AI-assisted ones. For regulators, the primary concern is safety, reproducibility, and explainability. They will be im-

pacted by my approach as they need to test and validate my approach, as well as create regulations for how the technology can be used in the case of adoption. In the same category fall ethics committee members, who are stakeholders concerned with whether the training data for the automated segmentation models is representative of the entire demographic. Since my solution requires the use of other automated ones, they will be subject to the same rigors.

The reason why no neural networks were used in developing this solution is that using one would decrease almost every metric of quality. It would decrease explainability, robustness, and accuracy. The only theoretical advantage would be inference time, due to a reduction in redundant calculations (DRRs). A neural approach could also run on lower-power devices specifically meant to run neural networks.

References

- Prendes, C. F., Fuentes, B., Diez-Tejedor, E., Grefkes, C., Halvorsen, S., James, S., Katus, H. A., Landmesser, U., Lewis, B. S., Manolis, A. S., et al., "The Burden of Stroke in Europe: An Analysis of the Global Burden of Disease Study Findings," *Stroke*, Vol. 54, No. 6, 2023, pp. 1522–1531. doi: 10.1161/STROKEAHA.122.042022.
- Powers, W. J., Rabinstein, A. A., Ackerson, T., Adeoye, O. M., Bambakidis, N. C., Becker, K., Biller, J., Brown, M., Demaerschalk, B. M., Hoh, B., Jauch, E. C., Kidwell, C. S., Leslie-Mazwi, T. M., Ovbiagele, B., Scott, P. A., Sheth, K. N., Southerland, A. M., Summers, D. V., Tirschwell, D. L., and on behalf of the American Heart Association Stroke Council, "Guidelines for the Early Management of Patients With Acute Ischemic Stroke: 2019 Update to the 2018 Guidelines for the Early Management of Acute Ischemic Stroke: A Guideline for Healthcare Professionals From the American Heart Association/American Stroke Association," *Stroke*, Vol. 50, No. 12, 2019, pp. e344–e418. doi: 10.1161/STR.0000000000000211, URL <https://www.ahajournals.org/doi/abs/10.1161/STR.0000000000000211>.
- Munich, S. A., Vakharia, K., and Levy, E. I., "Overview of mechanical thrombectomy techniques," *Journal of NeuroInterventional Surgery*, , No. Neurosurgery 85 Suppl 1, 2019, pp. S60–S67. doi: 10.1093/neuros/nyz071.
- van der Weide, R., Zuiderveld, K., Mali, W., and Viergever, M., "CTA-based angle selection for diagnostic and interventional angiography of saccular intracranial aneurysms," *IEEE Transactions on Medical Imaging*, Vol. 17, No. 5, 1998, pp. 831–841. doi: 10.1109/42.736054.
- Panase, A., Flexman, M., Mory, B., Webb, P., and Keenan, P., "E-162 Optimal projection angle planning tool for C-arm in augmented reality," *Journal of NeuroInterventional Surgery*, Vol. 12, No. Suppl 1, 2020, p. A117. doi: 10.1136/neurintsurg-2020-snis.194, URL https://jn.is.bmj.com/content/12/Suppl_1/A117.1.
- Gopalakrishnan, V., and Golland, P., "Fast Auto-Differentiable Digitally Reconstructed Radiographs for Solving Inverse Problems in Intraoperative Imaging," *arXiv preprint arXiv:2208.12737*, 2022. doi: 10.48550/arXiv.2208.12737.
- Liu, S., Liu, R., Meijs, M., van der Lugt, A., Niessen, W. J., and van Walsum, T., "Segmentation-assisted vessel centerline extraction from cerebral CT Angiography," *Medical Physics*, Vol. 52, No. 7, 2025, p. e17855. doi: 10.1002/mp.17855, URL <https://doi.org/10.1002/mp.17855>.
- Ahmad, T., Kaul, C., Asad, M., Murray-Smith, R., and Yuan, K., "VesselFM: A Foundation Model for Human Vascular Tree Analysis," 2024. URL <https://arxiv.org/abs/2411.17386>.
- Kumar, S., Agarwal, A., Golla, S., Tanamala, S., Upadhyay, U., Chattoraj, S., Putha, P., and Chilamkurthy, S., "Mind the Clot: Automated LVO Detection on CTA using Deep Learning," *2023 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, 2023, pp. 2495–2504. doi: 10.1109/iccvw60793.2023.00264.
- Philips Healthcare, *Azurion 7 C20 with FlexArm: Specifications*, Philips Healthcare, n.d. URL <https://www.philips.ro/healthcare/product/HCNCVD207/azurion-7-m20-with-flexarm-image-guided-therapy-system>, accessed: April 3, 2026.
- Yang, K., Musio, F., Ma, Y., Juchler, N., Paetzold, J. C., Al-Maskari, R., Höher, L., Li, H. B., Hamamci, I. E., Sekuboyina, A., Shit, S., Huang, H., Prabhakar, C., de la Rosa, E., Wittmann, B., Waldmannstetter, D., Kofler, F., Navarro, F., Menten, M., Ezhov, I., Rueckert, D., Vos, I. N., Ruigrok, Y. M., Velthuis, B. K., Kuijff, H. J., Shi, P., Liu, W., Ma, T., Rokuss, M. R., Kirchhoff, Y., Isensee, F., Maier-Hein, K., Zhu, C., Zhao, H., Bijlenga, P., Hämmerli, J., Wurster, C., Westphal, L., Bisschop, J., Colombo, E., Baazaoui, H., Handelsmann, H.-L., Makmur, A., Hallinan, J., Soundararajan, A., Wiestler, B., Kirschke, J. S., Wiest, R., Montagnon, E., Letourneau-Guillon, L., Oh, K., Lee, D., Hilbert, A., Aydin, O. U., Rallios, D., Rieger, J., Tanioka, S., Koch, A., Frey, D., Qayyum, A., Mazher, M., Niederer, S., Disch, N., Holzschuh, J., LaBella, D., Galati, F., Falcetta, D., Zuluaga, M. A., Lin, C., Zhao, H., Zhang, Z., Zhang, M., You, X., Zhang, H., Yang, G.-Z., Gu, Y., Ra, S., Hwang, J., Park, H., Chen, J., Wodzinski, M., Müller, H., Mansouri, N., Autrusseau, F., Yalçın, C.,

Hamadache, R. E., Lisazo, C., Salvi, J., Casamitjana, A., Lladó, X., Estrada, U. M. L.-T., Abramova, V., Giancardo, L., Oliver, A., Casademunt, P., Galdran, A., Delucchi, M., Liu, J., Huang, H., Cui, Y., Lin, Z., Liu, Y., Zhu, S., Patel, T. R., Siddiqui, A. H., Tutino, V. M., Orouskhani, M., Wang, H., Mossa-Basha, M., Sato, Y., Hirsch, S., Wegener, S., and Menze, B., “Benchmarking the CoW with the TopCoW Challenge: Topology-Aware Anatomical Segmentation of the Circle of Willis for CTA and MRA,” , 2025. URL <https://arxiv.org/abs/2312.17670>.

Wang, X., Yang, J., Chen, Y., Ai, D., Hu, Y., and Wang, Y., “Optimal Viewing Angle Determination for Multiple Vessel Segments in Coronary Angiographic Image,” *Nuclear Science, IEEE Transactions on*, Vol. 61, 2014, pp. 1290–1303. doi: 10.1109/TNS.2014.2319336.

Appendix

A. Hyperparameter Tuning

Since the baseline was computed with an arbitrary angle size of 2.5° , the natural follow-up question is: **what angle offset performs best?** A large offset will decrease computation time, but can miss the optimal view. Additional offset angles of 5° and 10° were tested in comparison with the baseline 2.5° version.

Table 1. Similarity Ratios across different Experiment Configurations. Values are scaled to fit a single column width.

Experiment Configuration	Top 1	Top 3	Top 5
Ablation Study vs Baseline (2.5° step)			
ablation_NoPerpendicular	0.900	0.910	0.920
Perpendicular Search Space Heuristic			
Perpendicular_10°	0.080	0.100	0.110
Perpendicular_15°	0.100	0.140	0.150
Perpendicular_20°	0.120	0.170	0.180
Perpendicular_25°	0.150	0.200	0.220
Perpendicular_30°	0.180	0.240	0.260
Gridsearch Comparison			
NoPerpendicular_5deg	0.310	0.275	0.276
NoPerpendicular_10deg	0.101	0.092	0.089
Realistic Volume Configuration			
Real_NoPerpendicular	0.007	0.015	0.021

According to Table 1, 70% of the optimal acquisition angles are lost by doubling the offset to 5° . Doubling again loses almost 90%.

The search space itself can be viewed as a hyperparameter, its size and shape being dictated by both the angle offset and the bounding limits assumed. Observing the second optimization function, Equation 6, and the first term of the main optimization function, one can assume they are very similar in their objective. Therefore, one can argue, the best place to look for the optimum is the region closest to the perpendicular ring of the target vessel segment. To test this, an area bounded by 10° , 15° , 20° , 25° , and 30° around the perpendicular ring was evaluated and tested against the baseline Table 1.

The heuristic fails to find the same angles as the

baseline; therefore, the other terms of the optimization count for more than what the first term does towards reducing foreshortening.

The second proposed method is functionally different from the grid search approach, but the output is the same: a set of best-performing angles; therefore, comparisons can be made with the baseline. Since this approach is based on gradient descent, it makes use of a learning rate (LR) and number of iterations (NI) as its main hyperparameters. To find the best value for each, a subset of the dataset was created. The best LR was found using binary search. The binary search for the LR was 5 steps, between 0.05 and 2.0, resulting in a step size of 1.87. The best NI was found using linear search starting from 50 to 250 iterations, in increments of 50. The best value was 150 iterations.

B. C-Arm System Geometry




Figure 15. C-Arm Geomtry (Wang et al., 2014)

A typical angiography system, as illustrated in Figure 15, functions by rotating about two or three principal axes, defined by a primary angle, α , and a secondary angle, β . The system is further characterized by intrinsic geometric parameters, such as the **Source-to-Object Distance (SOD)** and the **Source-to-Detector Distance (SID)**, which govern image magnification, spatial resolution, and the position of the isocenter. The isocenter, O , for us the coordiantes of the occlusion, represents the mechanical center of rotation for the C-arm and is typically distinct from the midpoint of the SID.

To define the C-arm geometry, two coordinate systems are utilized. The World Coordinate System (x_w, y_w, z_w) is centered at the isocenter and establishes the spatial reference for rotations α (Left/Right

Anterior Oblique) and β (Cranial/Caudal); in DICOM standards, these correspond to the primary and secondary angles, respectively. They are the main parameters we are trying to optimize. The Source Coordinate System (x_c, y_c, z_c) is centered at the X-ray source, S , where the beam originates, effectively serving as the camera.

C. TopBrain Labels for CTA data

ID	Label Name	Label
1	Basilar artery (BA)	BA
2	Right P1/P2 of posterior cerebral artery (PCA)	R-P1P2
3	Left P1 and P2 of PCA	L-P1P2
4	Right internal carotid artery (ICA)	R-ICA
5	Right M1 of middle cerebral artery (MCA)	R-M1
6	Left ICA	L-ICA
7	Left M1 of MCA	L-M1
8	Right posterior communicating artery (Pcom)	R-Pcom
9	Left Pcom	L-Pcom
10	Anterior communicating artery (Acom)	Acom
11	Right A1/A2 of anterior cerebral artery (ACA)	R-A1A2
12	Left A1 and A2 of ACA	L-A1A2
13	Right A3 of ACA	R-A3
14	Left A3 of ACA	L-A3
15	Third A2 of ACA	3rd-A2
16	Third A3 of ACA	3rd-A3
17	Right M2 of MCA	R-M2
18	Right M3 of MCA	R-M3
19	Left M2 of MCA	L-M2
20	Left M3 of MCA	L-M3
21	Right P3 and P4 of posterior cerebral artery (PCA)	R-P3P4
22	Left P3 and P4 of PCA	L-P3P4
23	Right vertebral artery (VA)	R-VA
24	Left VA	L-VA
25	Right superior cerebellar artery (SCA)	R-SCA
26	Left SCA	L-SCA
27	Right anterior inferior cerebellar artery (AICA)	R-AICA
28	Left AICA	L-AICA
29	Right posterior inferior cerebellar artery (PICA)	R-PICA
30	Left PICA	L-PICA
31	Right anterior choroidal artery (AChA)	R-AChA
32	Left AChA	L-AChA
33	Right ophthalmic artery (OA)	R-OA
34	Left OA	L-OA