



HAL
open science

A2V: A semi-supervised domain adaptation framework for brain vessel segmentation via two-phase training angiography-to-venography translation

Francesco Galati, Daniele Falcetta, Rosa Cortese, Barbara Casolla, Ferran Prados, Ninon Burgos, Maria A. Zuluaga

► To cite this version:

Francesco Galati, Daniele Falcetta, Rosa Cortese, Barbara Casolla, Ferran Prados, et al.. A2V: A semi-supervised domain adaptation framework for brain vessel segmentation via two-phase training angiography-to-venography translation. BMVC 2023, 34th British Machine Vision Conference, Nov 2023, Aberdeen, United Kingdom. hal-04195756v1

HAL Id: hal-04195756

<https://hal.science/hal-04195756v1>

Submitted on 7 Sep 2023 (v1), last revised 13 Sep 2023 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A2V: A Semi-Supervised Domain Adaptation Framework for Brain Vessel Segmentation via Two-Phase Training Angiography-to-Venography Translation

Francesco Galati¹
galati@eurecom.fr

Daniele Falchetta¹

Rosa Cortese²

Barbara Casolla³

Ferran Prados^{4,5}

Ninon Burgos⁶

Maria A. Zuluaga^{1,7}

¹ EURECOM, Biot, France

² University of Siena, Italy

³ Centre Hospitalier Universitaire de Nice, France

⁴ Centre for Medical Image Computing
University College London, UK

⁵ e-health Center
Universitat Oberta de Catalunya
Barcelona, Spain

⁶ Sorbonne Université, Paris Brain
Institute, CNRS, Inria, Inserm,
AP-HP, Paris, France

⁷ School of Biomedical Eng. & Life
Sciences, King's College London, UK

Abstract

We present a semi-supervised domain adaptation framework for brain vessel segmentation from different image modalities. Existing state-of-the-art methods focus on a single modality, despite the wide range of available cerebrovascular imaging techniques. This can lead to significant distribution shifts that negatively impact the generalization across modalities. By relying on annotated angiographies and a limited number of annotated venographies, our framework accomplishes image-to-image translation and semantic segmentation, leveraging a disentangled and semantically rich latent space to represent heterogeneous data and perform image-level adaptation from source to target domains. Moreover, we reduce the typical complexity of cycle-based architectures and minimize the use of adversarial training, which allows us to build an efficient and intuitive model with stable training. We evaluate our method on magnetic resonance angiographies and venographies. While achieving state-of-the-art performance in the source domain, our method attains a Dice score coefficient in the target domain that is only 8.9% lower, highlighting its promising potential for robust cerebrovascular image segmentation across different modalities.




Figure 1: Maximum intensity projection (MIP) of a magnetic resonance angiography (left) and minimum intensity projection (mIP) of a magnetic resonance venography (right) along the three spatial axes.

1 Introduction

The accurate segmentation of the cerebrovascular tree is critical to multiple diagnostic and therapeutic procedures and to the study of the brain’s health and associated pathologies. Due to the intricate morphology of the cerebrovascular tree [1], different clinical applications require different imaging techniques to visualize and analyze the brain vessels. Besides the differences in the image appearance and intensity distributions, image modalities may also vary in the type of vessels they target, i.e., either arteries or veins (see example in Figure 1). This diversity results in a large distribution shift that causes automatic segmentation models to fail when used across different image modalities. Given the variety of imaging techniques to assess the brain vasculature [2], the lack of cross-modality generalization is particularly limiting. Training, deploying, and maintaining a segmentation model per image modality is an expensive and highly demanding task.

Domain adaptation (DA) has been an active research field to tackle the problem of distribution shift, also known as domain shift, by finding a mapping from the source data distribution to the target distribution with little to no labels available in the target domain [3]. Recent studies have shown promising results for different organs [4], among which the brain and its main substructures [5, 6, 7]. Nevertheless, no current work has focused specifically on the brain vessels. This can be explained by two factors. First, brain vessels are relatively small objects within a large image volume [8] and the preservation of fine details is still a challenging task for state-of-the-art DA methods. As a consequence, small objects are often merged with the background, thus impeding their segmentation [9]. Second, while there is a correlation between the morphology of arteries and veins, they occupy different positions in the brain, and arteries generally have larger sizes. This results in a domain gap that encompasses variations in the overall image appearance, but also includes dissimilarities in the structure and spatial arrangement of cerebral vasculature between the source and target domains. As a result, there are currently no methods that can seamlessly be used to segment both the arteries and the veins.

Given that brain arterial segmentation is a simpler and more widely explored task [10], we aim at segmenting brain veins by relying primarily on annotated data from angiography images. To this end, we develop a semi-supervised domain adaptation framework that accomplishes both image-to-image translation and semantic segmentation of the cerebrovascular tree, relying on a disentangled and semantically rich latent space to represent heteroge-

neous volumetric data and to perform image-level adaptation from source to target domain, where vessels have different properties. The ability to disentangle volume-related and vessel-related image properties enables our method to generate a translation that transforms veins into arteries only in appearance, while preserving crucial spatial information such as shape and location. This bridges the gap between the two domains, which in turn enables easier generalization from labeled source data to unlabeled target data and thereby facilitates the segmentation of the vessels in both domains. Additionally, we reduce the typical complexity of cycle-based architectures [53] by decreasing the number of components from multiple to three: one generator, one discriminator, and one encoder. This reduction leads to more efficient utilization of computational resources and results in a simpler model, enabling faster experimentation at training and easier deployment of the model in real-world scenarios. We assess our method in the segmentation of vessels from magnetic resonance (MR) angiography and venography images, while mainly relying on annotated angiographies. Our results are particularly promising for the application of DA methodologies to the less-explored area of brain vein segmentation.

2 Related Work

Multi-modal brain vessel segmentation. Although 3D brain vessel segmentation has been extensively studied [46], few works have focused on the problem of cross-modality generalization. In [48], the authors use morphological operators on paired T1-weighted MR and time-of-flight magnetic resonance angiography (MRA) sequences, whereas [53] presents a tensor voting strategy using paired computed tomography angiography and 3D phase-contrast MR images. Although these works use the same algorithm to segment arteries from different modalities, each modality requires a different initialization process or modality-specific parameter tuning. Most recently, [9] proposed a neural architecture to segment arteries and veins from different MR sequences. Artery segmentations from time-of-flight MRA volumes resulted in a Dice score coefficient of 79.32%, whereas vein segmentations were only assessed through visual inspection. A further limitation is that the network requires separate training for each modality.

Domain Adaptation. Supervised DA methods [4, 52] simplify model training by assuming that a small number of labeled data in the target domain are available. These methods, however, require labeled target data, which is particularly costly to obtain for brain vessels. Unsupervised DA (UDA) techniques avoid the use of target domain labels, leveraging the potential information available in readily accessible unlabeled data. These techniques can perform adaptation at the input/image-[27], feature-[24] or output-[10] level; or a combination of two of the categories [2]. Among these, image-level alignment methods have gained a lot of traction due to the success of image-to-image translation methods [50], which convert source-domain images to target-domain images. However, many image-alignment approaches rely on adversarial training [2, 17, 50], which is known to be unstable. Lastly, only recently researchers have started to investigate semi-supervised domain adaptation for medical segmentation [5, 14].

Closer to our objective, two recent works rely on DA techniques to tackle domain shift for vessel segmentation. Peng et al. [19] use a disentangled representation and two segmentation models, each specialized to its own domain, in a collaborative UDA learning module for 2D retinal segmentation. Gu et al. [6] propose a semi-supervised method for cross-anatomy domain adaptation on 2D images that employs domain-specific batch normalization and cross-

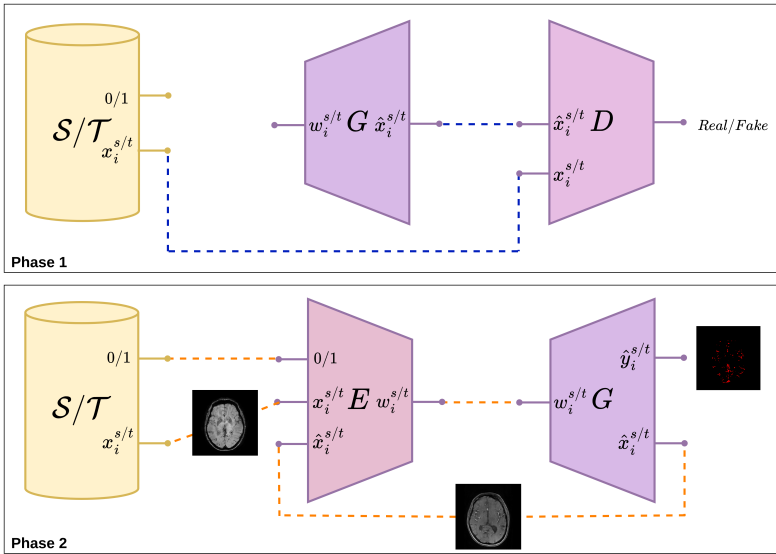


Figure 2: Two-phase training algorithm. Images x_i from domains \mathcal{S} and \mathcal{T} are fed into the model, composed of the generator G , the discriminator D and the encoder E . The modules are trained in two separate phases. Phase 1 (top) trains G in an adversarial fashion to learn a smooth and semantically rich latent space. Phase 2 (bottom) trains E to perform image-to-image translation and refines G to also generate segmentation masks \hat{y}_i^s and \hat{y}_i^t .

domain contrastive learning within a self-ensembling mean-teacher framework. Although these two approaches have reported promising results, we argue that the task of segmenting the cerebrovascular tree is more arduous, primarily due to the significant gap that exists between arteries and veins, as well as the intricate three-dimensional complexities inherent in their structure.

3 Method

Let us denote \mathcal{S} and \mathcal{T} as the source and target domains. We denote $S = \{x_i^s, y_i^s\}_{i=1}^N$ a set of N labeled images from \mathcal{S} , i.e., a set of labeled angiographies, $T_U = \{x_j^t\}_{j=1}^M$ a set of M unlabeled images from \mathcal{T} , i.e., a set of unlabeled venographies, and $T_L = \{x_k^t, y_k^t\}_{k=1}^m$ a target labeled dataset consisting of $m \ll M$ annotated images, also from \mathcal{T} . Figure 2 illustrates our proposed end-to-end framework to segment venographies using annotated angiographic data in a semi-supervised fashion by encapsulating conditional image generation, image-to-image translation, and image segmentation.

Our framework consists of three elements: 1) a generator G , which is responsible for constructing a smooth and semantically rich latent space that can represent a heterogeneous range of images from \mathcal{S} and \mathcal{T} and for producing the final segmentations in \mathcal{T} ; 2) a discriminator D that enables the training of G ; and 3) an encoder E that performs image-to-image translation across domains. The three modules are trained in two phases. The first phase trains G to build its smooth and semantically rich latent space \mathcal{W} . The second phase trains E to achieve image-to-image translation. By decoupling the training into two phases, we

guarantee that only the first phase involves an adversarial training routine, a highly desirable property to enforce stable and fast convergence.

Phase 1. G is trained to learn a function g that maps a latent vector $w_i \in \mathcal{W}$ into the corresponding image \hat{x}_i , in either \mathcal{S} or \mathcal{T} . The images generated by G are fed into the discriminator D , which is trained in an adversarial manner along with the generator G to distinguish between real and fake samples. In the meantime, the generator’s objective is to fool the discriminator by retrieving images of indistinguishable quality from the original images, for both \mathcal{S} and \mathcal{T} .

Without loss of generality, we adopt a StyleGAN2 architecture [10] for G and D , given its well-established ability to disentangle high-level features in the latent space \mathcal{W} during training. The property of disentanglement grants the generator the ability to control and vary specific features of the generated images independently of others. This provides control over the generated images. Additionally, manipulation of features such as intensities, textures, location and shape of vessels can be handled separately, allowing the next phase to establish mappings between arteries and veins at different semantic levels.

Phase 2. The encoder E is trained to learn a function e that inverts the generative function g , i.e., discover the corresponding latent representation of an image in both \mathcal{S} and \mathcal{T} . The main goal of E is to perform image-to-image translation under the assumption that all vessel properties are disentangled within the latent space \mathcal{W} . Due to this property, the encoder is able to establish correspondences between the two source and target domains at a high level of abstraction [25, 26], discovering mappings that differentiate between various vessel characteristics, such as altering vessel intensities, while preserving their spatial arrangement. To achieve this, E is trained in three possible configurations, which are alternated by providing the model with an interchangeable binary label. When the label aligns with the domain of the inputted images, the encoder retrieves their reconstructions within the same domain (Configuration 2.1). When the label is inverted, the encoder retrieves inter-domain translations in the opposite domain, i.e., angiogram-to-venogram (Configuration 2.2) or venogram-to-angiogram (Configuration 2.3) translations. By feeding the model with opposite labels in succession, the encoder gains a cyclical behavior. This enables the computation of cycle-consistency losses without the need for two encoder-decoder pairs, which represents a reduction in complexity and computational cost compared to traditional cycle-based approaches [31].

In order to achieve segmentation, the generator is extended by adding a label-synthesis branch [30], consisting of three fully-connected layers attached to the feature vectors of G . With this, we aim to avoid the use of a separate segmentation module, thus decreasing computational requirements. By freezing all other parameters in the framework, the branch is optimized in isolation, with the assistance of labeled samples from \mathcal{S} and T_L , to return a semantic segmentation mask that aligns with a generated image. Differently from reconstruction and cycle-consistency, segmentation losses are computed only when labels are available, i.e., for \mathcal{S} and T_L . After one cycle, losses are propagated into E and G only once for the most recent pass. Figure 3 provides a detailed breakdown of Phase 2.

Inference. Given an unseen target image x^t , our model generates its reconstruction in the original target domain, i.e., \hat{x}^t , and its inter-domain translation, which accounts for translating x^t into \mathcal{S} , i.e., \hat{x}^s . Additionally, G generates the corresponding segmentation mask, \hat{y}^t and \hat{y}^s , from its label-synthesis branch for both \hat{x}^t and \hat{x}^s . Both masks contain valuable information pertaining to vessel segmentation within the target domain. For this reason, the final




Figure 3: Phase 2 of the training algorithm alternating intra-domain (2.1) and inter-domain (2.2 and 2.3) configurations. We compute the sum of mean squared error and LPIPS [29] for reconstruction and cycle-consistency losses \mathcal{L}_R , and the sum of Dice and cross-entropy for segmentation losses \mathcal{L}_S . The backpropagation of \mathcal{L}_R updates only the weights of E , while \mathcal{L}_S affects both E and G .

predicted segmentation mask is obtained as the average of \hat{y}^t and \hat{y}^s before the final argmax operation.

4 Experiments and Results

4.1 Experimental Setup

Datasets and Preprocessing. We use a subset of 49 time-of-flight (TOF) MRA volumes selected at random from the OASIS-3 database [12], with median dimension $576 \times 768 \times 232$ and voxel size $0.60 \times 0.30 \times 0.30$, as the source database \mathcal{S} . This selection includes 41 cognitively normal patients, as well as 16 patients at various stages of cognitive decline, all adults ranging in age from 42 to 95 years. For the target database \mathcal{T} , we use a set of 28 susceptibility weighted imaging (SWI) venographies, with median dimension $480 \times 480 \times 288$ and voxel size $0.50 \times 0.50 \times 0.50$. The database was obtained from retrospective studies previously conducted at <ANONYMIZED> and includes adult subjects with no lesions visible on SWI. The source database and 11 volumes in the target database have associated brain and vessel masks. Brain masks were obtained using SynthStrip [8] for the source database and HD-BET [9] for the target database, while vessel masks were manually annotated by an

expert.

Each volume was standardized based on its mean and standard deviation and rescaled to have a uniform spacing of $[0.5, 0.5, 0.5]$ mm between adjacent voxels. Next, 2D slices were extracted from the volumes, and each slice was clipped and normalized in the range $[-1, +1]$. The slices were also padded and cropped to a fixed size of 512×512 pixels. The segmentation masks were one-hot encoded into a single label of size $3 \times 512 \times 512$ (with an extra dimension for the background).

Setup. The training set consists of 45 TOF images (S), 17 SWI venographies without annotations (T_U) and 3 annotated SWI venographies (T_L). Four annotated SWIs are used for the validation set, whereas the remaining four annotated angiographies and four annotated venographies are kept for testing.

We compare our method against four DA state-of-the-art methods:

- 1) **Synergistic image and feature adaptation (SIFA)** [20], an UDA technique based on image-to-image translation for multi-class medical segmentation;
- 2) **SynthSeg** [21], an UDA 3D output-level alignment method based on synthetic data generation for brain synthesis and segmentation;
- 3) **Contrastive Semi-supervised learning for Cross Anatomy Domain Adaptation (CS-CADA)** [9], a semi-supervised DA method relying on feature alignment, which has been tested for 2D coronary artery segmentation; and
- 4) **DCDA** [19], an UDA technique based on style representation transfer, specifically designed for retinal vessel segmentation.

Additionally, we include a Sato filter [22] for vessel enhancement as a baseline model. The segmentation results obtained on the test set are quantitatively assessed using the Dice coefficient (Dice), the centerlineDice (cDice) [23], precision and recall.

Implementation Details. We implemented our framework in PyTorch 1.9.1. Both Phase 1 and Phase 2 were executed with a batch size of 4 images and trained respectively for 250k and 50k iterations. We retained the model with the best validation performance for the final evaluation. The generator G and discriminator D are based on the architectures of StyleGAN2 [10], while the label-synthesis branch is adapted from DatasetGAN [60]. The input images are encoded into the extended latent space $\mathcal{W}+$ of StyleGAN2, as formulated in [20]. The architecture of E is inspired by [28], which utilizes a ResNet backbone and multiple output branches: one for latent code prediction, the others for feature tensor prediction. These are filtered and passed to G through a dynamic skip connection module [26], which establishes fine-level content correspondences.

For SIFA, SynthSeg, CS-CADA and DCDA, we used the implementations proposed in the original publications, following the respective training schedules. The Sato filter was implemented in C++ using the ITK library¹. We trained, validated and tested our proposed method as well as the state-of-the-art methods on two NVIDIA GeForce RTX 2080 Ti GPUs.

4.2 Results

We first assess the intra-domain vessel segmentation performance of our method using the angiography images within the test set. Our model achieves Dice of $79.3 \pm 4.4\%$, a cDice of $78.7 \pm 5.3\%$, precision of $83.1 \pm 1.4\%$ and a recall of $76.2 \pm 7.4\%$. The obtained results are comparable with state-of-the-art methods for artery segmentation [9, 15].

¹<anonymized>

Table 1: Performance comparison of different DA methods in the target domain, i.e., four testing venographies. We report mean Dice, Precision, Recall and cIDice (in %) with standard deviations.

		SIFA	SynthSeg	CS-CADA	DCDA	Sato	Ours
Dice	Vessels	0.8 ± 0.2	37.3 ± 4.4	51.4 ± 1.7	4.5 ± 0.4	44.2 ± 7.2	70.4 ± 2.4
	Brain	91.5 ± 0.4	79.6 ± 3.8	91.5 ± 0.8	-	-	97.5 ± 0.2
Precision	Vessels	11.6 ± 1.2	42.3 ± 9.2	58.6 ± 6.7	14.8 ± 3.5	42.7 ± 6.4	66.8 ± 5.2
	Brain	84.8 ± 0.7	69.4 ± 5.5	89.6 ± 0.8	-	-	97.6 ± 0.3
Recall	Vessels	0.4 ± 0.1	33.9 ± 1.6	46.2 ± 2.2	2.7 ± 0.2	46.1 ± 9.3	74.9 ± 3.0
	Brain	99.3 ± 0.1	93.6 ± 0.5	93.5 ± 1.1	-	-	97.4 ± 0.5
cIDice	Vessels	0.8 ± 0.2	48.2 ± 4.7	58.0 ± 2.8	3.9 ± 0.2	50.0 ± 6.7	74.8 ± 2.4

We evaluate our method’s performance in the target domain, i.e., venographies, and we compare it against state-of-the-art DA methods and the Sato filter. For the sake of fairness, in addition to vessel segmentation, we report the results obtained for cross-modality brain segmentation since most methods (i.e., SIFA, SynthSeg and CS-CADA) have been originally conceived to segment large objects, such as the brain. Table 1 reports the obtained results.

While the brain segmentation results are generally satisfactory, most of the methods show poor performance at segmenting vessels, being surpassed by the simpler Sato filter. Despite being provided with only 3 target annotations, the proposed method succeeds in segmenting vessels by linking arteries and veins across the two modalities, achieving high performance in both domains. In particular, our proposed method bridges the large domain gap that encompasses not only low-level features, such as intensities and textures, but also high-level features, such as variations in the locations and shapes of arteries and veins.

Figure 4 presents a visual comparison of the results at different locations of the brain. We observe that, in some cases, spatial correspondence is lost during translation, e.g., DCDA mixes top and bottom slices. Methods like SIFA focus on translating the overall appearance of the image, perhaps because it has been designed to deal with the adaptation of objects at a larger scale. Although CS-CADA achieves the second highest Dice on the brain, its performance drops when detecting veins in our testing set. We hypothesize that, despite being designed to address significant domain gaps, CS-CADA may not be well-suited for highly complex 3D problems. Our method is able to disentangle the volume-related image properties, such as spatial information and appearance, and the vessel-related properties, such as intensities, textures, shapes, and locations. As a consequence, our method generates a translation that resembles a TOF image in terms of style while retaining the SWI’s content, thus leading to a fair vessel segmentation result.

5 Discussion and Conclusion

In this paper, we proposed an end-to-end semi-supervised domain adaptation framework for accurately segmenting 3D brain vessels using primarily annotations from arterial images, which are easier to obtain. We addressed the challenge of domain shift caused by the differences in imaging modalities, not only concerning the general appearance but also the different vessels they target (arteries or veins). Our approach relies on the StyleGAN2 architecture, which allows to represent heterogeneous volumetric data and bridge the large domain gap between angiography and venography brain images.




Figure 4: Visual comparison of the results produced by different methods for brain and vein segmentation from SWI images. Brain masks are indicated in red, vessels in green. From the 1st to the 3rd row we display in order a top-level, a middle-level and a bottom-level slice.

In accordance with the goal of reducing complexity in cycle-based architectures, our network was optimized and constructed from a single generator, discriminator, and encoder. We highlight how this design brings an advantage in terms of memory requirements and model comprehensibility, facilitating its debugging and understanding during experimentation and its deployment in real-world scenarios. Our method performs conventional image-level alignment with high performance, despite other techniques, such as SIFA or DCDA, build on more complex architectures. In particular, SIFA counts three distinct discriminators, two encoder-decoder networks and one segmentation module, while DCDA includes four encoder-generator pairs and two segmentation networks. These differences demonstrate the efficacy of our approach in building an efficient and intuitive model without sacrificing performance. Moreover, its training was stabilized and accelerated by designing a two-phase learning algorithm that minimizes the use of adversarial training.

Although our approach outperforms DA state-of-the-art methods, we acknowledge that there is still room for improvement, and this problem remains a challenging task. To the best of our knowledge, this is the first attempt to achieve artery-to-vein translation. Nevertheless, our results show a promising performance at segmenting the brain veins in a semi-supervised domain adaptation scenario, despite the inherent difficulty of the problem posed by the significant gap between veins and arteries and the intricate morphology of the cerebrovascular tree. Overall, our work highlights the potential of domain adaptation methodologies for improving brain vessel segmentation, and we hope it can inspire further research in this area.

Acknowledgments

FG, NB and MAZ are partially funded by the French government via Investments in the Future project (3IA Côte d’Azur, PRAIRIE 3IA Institute) managed by the ANR (ANR-19-P3IA-000, ANR-19-P3IA-0001). DF and MAZ are funded by the ANR JCJC project I-VESSEG (22-CE45-0015-01). NB has received funding from the ANR-10-IA Institut Hospitalo-Universitaire-6 (ANR-10-IAIHU-06).

References

- [1] Benjamin Billot, Douglas N Greve, Oula Puonti, Axel Thielscher, Koen Van Leemput, Bruce Fischl, Adrian V Dalca, Juan Eugenio Iglesias, et al. Synthseg: Segmentation of brain MRI scans of any contrast and resolution without retraining. *Medical Image Analysis*, page 102789, 2023.
- [2] Cheng Chen, Qi Dou, Hao Chen, Jing Qin, and Pheng-Ann Heng. Synergistic image and feature adaptation: Towards cross-modality domain adaptation for medical image segmentation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(1): 865–872, 2019.
- [3] Vien Ngoc Dang, Francesco Galati, Rosa Cortese, Giuseppe Di Giacomo, Viola Marconetto, Prateek Mathur, Karim Lekadir, Marco Lorenzi, Ferran Prados, and Maria A. Zuluaga. Vessel-CAPTCHA: An efficient learning framework for vessel annotation and segmentation. *Medical Image Analysis*, 75:102263, 2022.
- [4] Mohsen Ghafoorian, Alireza Mehrtash, Tina Kapur, Nico Karssemeijer, Elena Marchiori, Mehran Pesteie, Charles R. G. Guttmann, Frank-Erik de Leeuw, Clare M. Tempny, Bram van Ginneken, Andriy Fedorov, Purang Abolmaesumi, Bram Platel, and William M. Wells. Transfer learning for domain adaptation in mri: Application in brain lesion segmentation. In *Medical Image Computing and Computer Assisted Intervention - MICCAI 2017*, pages 516–524, 2017.
- [5] Ran Gu, Jingyang Zhang, Guotai Wang, Wenhui Lei, Tao Song, Xiaofan Zhang, Kang Li, and Shaoting Zhang. Contrastive semi-supervised learning for domain adaptive segmentation across similar anatomical structures. *IEEE Transactions on Medical Imaging*, 42(1):245–256, 2023.
- [6] Hao Guan and Mingxia Liu. Domain adaptation for medical image analysis: A survey. *IEEE Transactions on Biomedical Engineering*, 69(3):1173–1185, 2022.
- [7] Jasper HG Helthuis, Tristan PC Van Doormaal, Berend Hillen, Ronald LAW Bleys, Anita A Harteveld, Jeroen Hendrikse, Annette Van der Toorn, Mariana Brozici, Jaco JM Zwanenburg, and Albert Van der Zwan. Branching pattern of the cerebral arterial tree. *The Anatomical Record*, 302(8):1434–1446, 2019.
- [8] Andrew Hoopes, Jocelyn S. Mora, Adrian V. Dalca, Bruce Fischl, and Malte Hoffmann. SynthStrip: skull-stripping for any brain image. *NeuroImage*, 260:119474, 2022. ISSN 1053-8119.
- [9] Fabian Isensee, Marianne Schell, Irada Pflueger, Gianluca Brugnara, David Bonekamp, Ulf Neuberger, Antje Wick, Heinz-Peter Schlemmer, Sabine Heiland, Wolfgang Wick, Martin Bendszus, Klaus H. Maier-Hein, and Philipp Kickingereder. Automated brain extraction of multisequence mri using artificial neural networks. *Human Brain Mapping*, 40(17):4952–4964, 2019.
- [10] Konstantinos Kamnitsas, Christian Baumgartner, Christian Ledig, Virginia Newcombe, Joanna Simpson, Andrew Kane, David Menon, Aditya Nori, Antonio Criminisi, Daniel Rueckert, et al. Unsupervised domain adaptation in brain lesion segmentation with adversarial networks. In *Information Processing in Medical Imaging: 25th International*

Conference, IPMI 2017, Boone, NC, USA, June 25-30, 2017, Proceedings 25, pages 597–609, 2017.

- [11] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [12] Pamela J. LaMontagne, Tammie LS. Benzinger, John C. Morris, Sarah Keefe, Russ Hornbeck, Chengjie Xiong, Elizabeth Grant, Jason Hassenstab, Krista Moulder, Andrei G. Vlassenko, Marcus E. Raichle, Carlos Cruchaga, and Daniel Marcus. Oasis-3: Longitudinal neuroimaging, clinical, and cognitive dataset for normal aging and alzheimer disease. *medRxiv*, 2019.
- [13] Amy Lin, Sapna Rawal, Ronit Agid, and Daniel M Mandell. Cerebrovascular imaging: which test is best? *Neurosurgery*, 83(1):5–18, 2018.
- [14] Xiaofeng Liu, Fangxu Xing, Nadya Shusharina, Ruth Lim, C-C Jay Kuo, Georges El Fakhri, and Jonghye Woo. Act: Semi-supervised domain-adaptive medical image segmentation with asymmetric co-training. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2022: 25th International Conference*, pages 66–76, 2022.
- [15] Michelle Livne, Jana Rieger, Orhun Utku Aydin, Abdel Aziz Taha, Ela Marie Akay, Tabea Kossen, Jan Sobesky, John D. Kelleher, Kristian Hildebrand, Dietmar Frey, and Vince I. Madai. A u-net deep learning framework for high performance vessel segmentation in patients with cerebrovascular disease. *Frontiers in Neuroscience*, 13, 2019.
- [16] Sara Moccia, Elena De Momi, Sara El Hadji, and Leonardo S Mattos. Blood vessel segmentation algorithms—review of methods, datasets and evaluation metrics. *Computer methods and programs in biomedicine*, 158:71–91, 2018.
- [17] Munan Ning, Cheng Bian, Dong Wei, Shuang Yu, Chenglang Yuan, Yaohua Wang, Yang Guo, Kai Ma, and Yefeng Zheng. A new bidirectional unsupervised domain adaptation segmentation framework. In Aasa Feragen, Stefan Sommer, Julia Schnabel, and Mads Nielsen, editors, *Information Processing in Medical Imaging*, pages 492–503, Cham, 2021. Springer International Publishing.
- [18] Nicolas Passat, Christian Ronse, Joseph Baruthio, J-P Armspach, and Jack Foucher. Watershed and multimodal data for brain vessel segmentation: Application to the superior sagittal sinus. *Image and Vision Computing*, 25(4):512–521, 2007.
- [19] Linkai Peng, Li Lin, Pujin Cheng, Ziqi Huang, and Xiaoying Tang. Unsupervised domain adaptation for cross-modality retinal vessel segmentation via disentangling representation style transfer and collaborative consistency learning. In *2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI)*, pages 1–5, 2022.
- [20] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding in style: A stylegan encoder for image-to-image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2287–2296, 2021.

- [21] Yoshinobu Sato, Shin Nakajima, Hideki Atsumi, Thomas Koller, Guido Gerig, Shigeyuki Yoshida, and Ron Kikinis. 3D multi-scale line filter for segmentation and visualization of curvilinear structures in medical images. In *CVRMed-MRCAS'97: First Joint Conference Computer Vision, Virtual Reality and Robotics in Medicine and Medical Robotics and Computer-Assisted Surgery Grenoble*, pages 213–222, 1997.
- [22] Suprosanna Shit, Johannes C. Paetzold, Anjany Sekuboyina, Ivan Ezhov, Alexander Unger, Andrey Zhylka, Josien P. W. Pluim, Ulrich Bauer, and Bjoern H. Menze. clDice - a novel topology-preserving loss function for tubular structure segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16560–16569, 2021.
- [23] Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schulter, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. Learning to adapt structured output space for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [24] Jianghao Wu, Ran Gu, Guiming Dong, Guotai Wang, and Shaoting Zhang. Fpluda: Filtered pseudo label-based unsupervised cross-modality adaptation for vestibular schwannoma segmentation. In *2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI)*, pages 1–5, 2022.
- [25] Wayne Wu, Kaidi Cao, Cheng Li, Chen Qian, and Chen Change Loy. Transgaga: Geometry-aware unsupervised image-to-image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [26] Shuai Yang, Liming Jiang, Ziwei Liu, and Chen Change Loy. Unsupervised image-to-image translation with generative prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18332–18341, June 2022.
- [27] Kai Yao, Zixian Su, Kaizhu Huang, Xi Yang, Jie Sun, Amir Hussain, and Frans Coenen. A novel 3d unsupervised domain adaptation framework for cross-modality medical image segmentation. *IEEE Journal of Biomedical and Health Informatics*, 26(10): 4976–4986, 2022.
- [28] Xu Yao, Alasdair Newson, Yann Gousseau, and Pierre Hellier. A style-based gan encoder for high fidelity reconstruction of images and videos. In Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner, editors, *Computer Vision – ECCV 2022*, pages 581–597, 2022.
- [29] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [30] Yuxuan Zhang, Huan Ling, Jun Gao, Kangxue Yin, Jean-Francois Lafleche, Adela Barriuso, Antonio Torralba, and Sanja Fidler. Datasetgan: Efficient labeled data factory with minimal human effort. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10145–10155, 2021.

- [31] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [32] Qikai Zhu, Bo Du, and Pingkun Yan. Boundary-weighted domain adaptive neural network for prostate mr image segmentation. *IEEE Transactions on Medical Imaging*, 39(3):753–763, 2020.
- [33] Maria A Zuluaga, Roman Rodionov, Mark Nowell, Sufyan Achhala, Gergely Zom-bori, Alex F Mendelson, M Jorge Cardoso, Anna Miserocchi, Andrew W McEvoy, John S Duncan, et al. Stability, structure and scale: improvements in multi-modal ves-sel extraction for seeg trajectory planning. *International journal of computer assisted radiology and surgery*, 10:1227–1237, 2015.