Attention-Aware DAE for Automated Solar Coronal Loop Segmentation

Prabal Man Dhaubhadel[1], Jong Kwan Lee[1], Qing Tian[1,2]
[1]Dept. of Computer Science, Bowling Green State University, Bowling Green, OH 43403, USA
[2]Dept. of Computer Science, University of Alabama at Birmingham, Birmingham, AL 35294, USA
pdhaubh@bgsu.edu, leej@bgsu.edu, qtian@uab.edu

ABSTRACT

This paper introduces an enhanced Denosing Autoencoder (DAE) model, incorporating a novel attention mechanism, for the segmentation of solar coronal loops. This work is based on DAE framework to address the segmentation challenges posed by intricate structures of coronal loops which also appear with other solar features and image noises. Specifically, we introduce Encoding-Aware Decoding Attention (EADA) to all decoding stages of DAE, which resulted in improvement in coronal loop segmentation. Our models are validated through experiments on a synthetic image dataset of 11,000 images and a test dataset of 165 real coronal images of the NASA's Solar Dynamics Observatory (SDO) satellite mission. Compared to the state-of-the-art coronal loop segmentation baseline, our attention-enhanced model results in better loop gap-filling and higher segmentation metrics (i.e., 3.6% increase in accuracy, 11.4% better recall and 5.6% higher precision).

Keywords

Denoising Autoencoders (DAE), Attention Block, Solar Coronal Loop Segmentation, Loop Gap-Filling

1 INTRODUCTION

In recent years, deep learning algorithms have emerged as powerful tools in image processing and computer vision, offering significant improvements in handling complex patterns and structures in various application domains. In this paper, we introduce an attention-enhanced deep learning model for accurate solar feature segmentation.

Solar coronal loops, which are highly structured and dynamic features in the solar corona, play a crucial role in understanding solar physics and the mechanisms behind solar activities. These solar activities can impact our daily life (e.g., disrupting electrical grids, damaging communication satellites and navigation systems). The coronal loops, visible in extreme ultraviolet (EUV) and X-ray wavelengths, are fundamentals in studying the Sun's magnetic field structure and its evolution. However, the accurate segmentation of these loops from solar images is a very challenging task due to their intricate structures, varying brightness, and the presence of other solar features and image noises. Fig. 1 is an example of coronal loop sub-image from NASA's Solar Dynamics Observatory (SDO) satellite mission. As shown in the image, the coronal loops (i.e., the bright arcing structures) have complex shapes with varying intensities and blurry boundaries. There are also non-loop features and noises on the image.

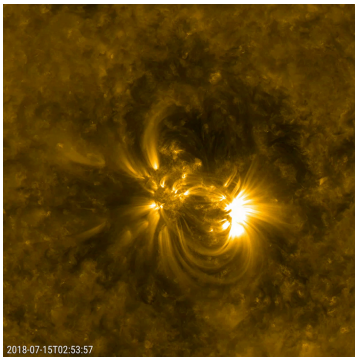

Figure 1: A sub-image of SDO coronal image (image credit: NASA's SDO Mission).

In this paper, we present a Denoising Autoencoder (DAE) enhanced by a new attention module, aimed at advancing coronal loop segmentation. To our knowledge, our work is the first attempt to consider attention-based DAE for coronal loop segmentation. The paper is organized as follows. In Section 2, we discuss related works. Section 3 introduces our new method and Section 4 outlines the experimental setup followed by experimental results. Section 5 concludes the paper.

2 RELATED WORK

2.1 U-Net and DAE

Machine learning/deep learning algorithms have been used in many image segmentation applications (e.g., [**gon16**, **lk16**, **tme22**]). Among them, U-Net [**rfb15**] and Denoising Autoendcoder (DAE) [**vlb08**] have widely been used recently.

U-Net (e.g., [**rfb15**, **bkc17**]) is based on the encoder-decoder architecture that follows a symmetric contracting and expansive path. U-Net's unique feature lies in its use of skip connections, facilitating the direct information propagation from contracting layers to corresponding layers in the expansive path. These horizontal propagation structures, known as skip connections, allow the concatenation of contractive path information and expansive path information before up-sampling.

DAE [**vlb08**] is a class of autoencoders that normally operate on an unsupervised learning paradigm, where the network is trained to reconstruct clean data from corrupted input during the training phase. The training process involves introducing corruption to the input data and optimizing the model to minimize the reconstruction error, encouraging the network to capture meaningful features while filtering out irrelevant noise. Feature extraction and dimensionality reduction tasks are efficiently handled by the denoising autoencoder (e.g., [**fbh17**, **vll10**, **xmy16**]). DAEs showed promising results in removing noise while completing the encoding-decoding process. However, not as many DAEs have not been studied for segmentation.

2.2 Solar Coronal Loop Segmentation

Several methods have been introduced for solar coronal loop segmentation. Oriented Connectivity-based method (OCM) [**lng06a**] and Dynamic Aperture-based method (DAM) [**lng06b**] employed image processing and constructive feature segmentation by exploring physical constraints and coronal loop feature's intensity profiles, respectively. An enhanced version of the OCM has also been introduced [**asc10**]. A solar loop mining system which includes a block-by-block loop segmentation for retrieving coronal loop images has been introduced by Durak *et al.* [**dnc09**, **dnc10**]. McAteer *et al.* [**mka10**] utilized a 2D Wavelet-based smoothing function as an edge detector for segmenting coronal loops. Lee and Tang [**let11**] presented an active contour-based model where the minimization of the spline energy constrained by the physical shape of the coronal loop's intensity profiles allowed more robust loop segmentation. Zhiming *et al.* [**zxz19**] presented a coronal loop segmentation method that exploits a clustering algorithm based on approximated local directionality determined by a match and image enhancing filters. To our best knowledge, a recent work by Moradi *et al.* [**mlt21**] was the only work that explored a deep learning algorithm (i.e., a U-Net) for segmenting the coronal loops.

2.3 Attention Module for Enhanced Image Segmentation

Self-Attention [**gsy21**], Bottleneck Attention Module (BAM) [**pwl18**], Convolutional Block Attention Module (CBAM) [**wpl18**] that focus on channel and spatial attention, and criss-cross attention [**hwh19**] are all attention-equipped approaches to reinforce their respective neural networks. Using denoising autoencoders, missing data imputation is achieved by implementing mask attention in DAEMA [**tfj21**]. Real Image Denoising [**anb19**] with feature attention uses attention in a modular manner and achieves superior performance. Motivated by the use of attention in neural networks, we propose our feature attention block for the denoising autoencoder that directs the model's attention to foreground pixels in segmenting coronal loops. As we will show later, this novel approach leads to improved performance metrics. It can also enhance the segmentation of loops previously categorized as "disconnected". We refer to this segmentation of disconnected loops into one cohesive loop as the *gap-filling effect* in this paper.

3 OUR METHODOLOGY

In this section, we introduce our new attention-aware denoising autoencoder (AtnDAE) for coronal loop segmentation. Particularly, we propose the *Encoding-Aware Decoding Attention (EADA)*, a new attention mechanism that enhances the segmentation of coronal loops by integrating attention at each stage of the decoding path, taking into account the corresponding layer from the encoding path. Fig. 2 illustrates the

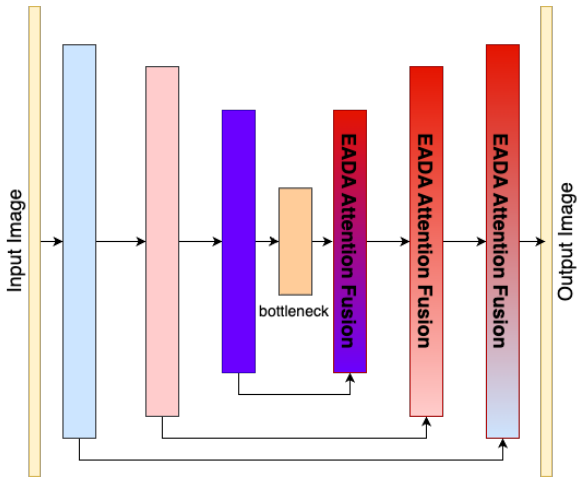

Figure 2: Illustration of our Attention-aware DAE architecture. Details of the Encoding-Aware Decoding Attention (EADA) block are shown in Fig. 3.

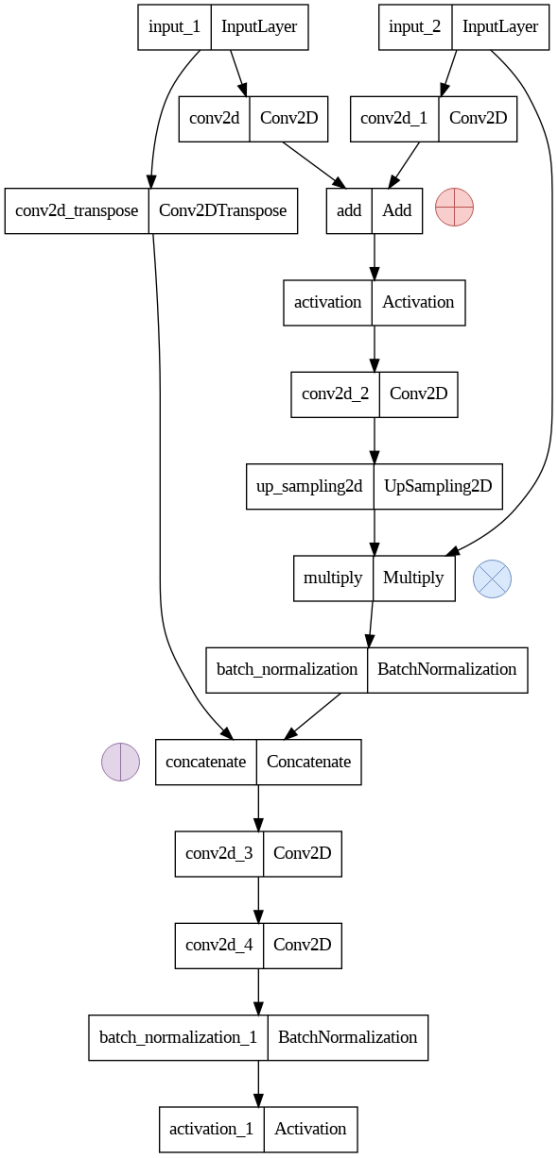

Figure 3: Encoding-Aware Decoding Attention (EADA) block. This block is integrated at each stage of our attention-enhanced DAE (Fig. 2). It utilizes information from both the decoding input (input_1) and its corresponding encoding features (input_2).

overall architecture of our AtnDAE, with the details of our EADA attention fusion block depicted in Fig. 3. Our novel attention mechanism is particularly beneficial for addressing imbalanced data like corona images, where the background pixels substantially outnumber the foreground pixels. Each EADA block selectively emphasizes the foreground loop structures, which are often overshadowed by the complex background of solar images.

As shown in Fig. 2, our EADA attention fusion block is integrated at every stage of the decoding path, leveraging information from both the decoding input (i.e., input_1 in Fig. 3) and the corresponding encoding layer (i.e., input_2 in Fig. 3) at each stage. This setup ensures that the EADA mechanism has access to both more location-accurate features in the encoder and highly semantic features in the decoder. The information from the encoding phase directs more accurate attention toward the foreground loop structures, facilitating the refinement and amplification of the distinctive curvilinear features found in coronal loops. Through dot multiplication, such attention is applied to a transformed combination of encoding and decoding features. The above attention-enhanced results are then concatenated with the regular decoding output, followed by further convolutions before arriving at the output for the EADA block.

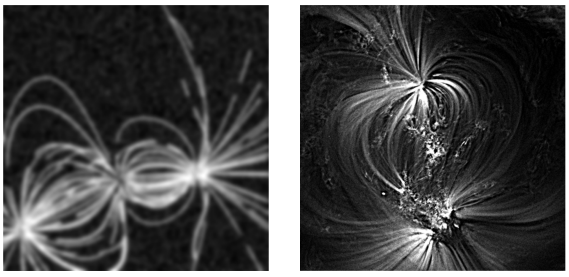

(a) synthetic image (b) real image

Figure 4: A sample of synthetic and real images.

In our attention-aware DAE, batch normalization is used to reduce the chance of overfitting and improve the training stability. ReLU is used as the activation function.

4 EXPERIMENTS AND RESULTS

This section first outlines the dataset and experimental setup utilized in our study. Then, we provide a comparative assessment of our novel attention-aware denoising autoencoder (AtnDAE), including both quantitative and qualitative analyses. Specifically, we compare the performance of the U-Net and traditional DAE models in the context of coronal loop segmentation and present the results of our improved attention-aware DAE.

4.1 Dataset

In our experiments, 11,000 synthetic coronal images and 165 real coronal images were used. The synthetic coronal images were created using the scheme, including the image noises, used in [**lng06a**, **lng06b**, **let11**] as it was one of the most popularly used synthetic images of the solar corona. The synthetic images were of 256×256 and the coronal loops were represented as the magnetic field lines derived from a physical magnetic model. (Solar physicists consider the corona loops as the traces of the magnetic fields around the Sun.) The primary benefit of using the synthetic images based on this scheme is that the ground truth masks (i.e., labels)

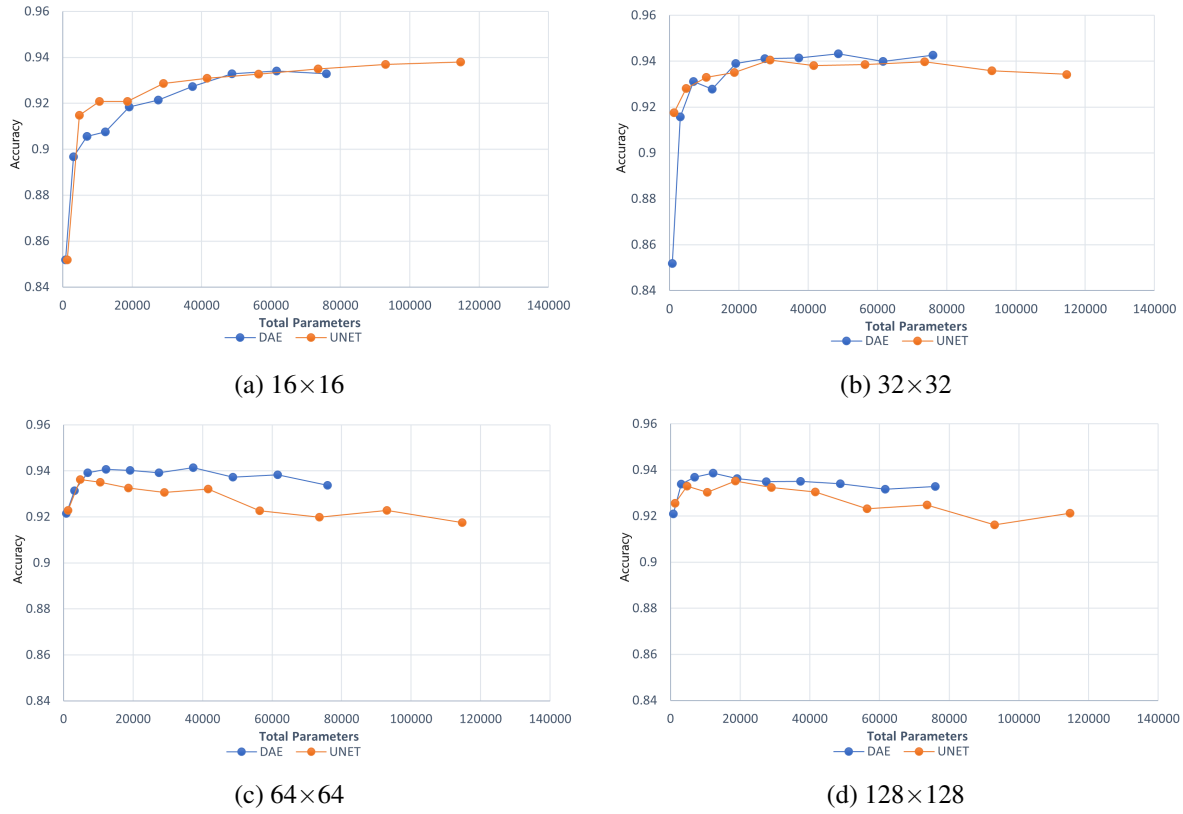

(a) 16×16 (b) 32×32

(c) 64×64 (d) 128×128

Figure 5: Accuracy over range of parameters for different bottleneck sizes-16×16, 32×32, 64×64, 128×128.

are available for training. In our synthetic image set, we have included loops that are broken (e.g., "gaps" within a loop, abruptly ending loops) which are often present in the real coronal loops due to 3D-to-2D projection, viewpoint occlusion, and noise and non-loop features in the image. No other prior study have considered this in coronal loop segmentation. 10,000 synthetic images were used for 80-20 training-validation split. The rest of 1,000 synthetic images were used as the test set. The real coronal images were obtained from the NASA's SDO satellite mission. The ground truth masks (i.e., labels) were determined from manual segmentation. Fig. 4 shows a sample of synthetic and real corona images.

4.2 Experimental Setup

All executions utilize a conda environment with TensorFlow version 2.8, cuda version 11.2.2, and Python version 3.8.12. A test on a range of 10 models is performed to select the best-performing model on each of the bottleneck sizes. Each model is trained for 200 epochs. The training of AtnDAE is optimized using the Adam optimizer [**kib14**] with a learning rate carefully selected to balance fast convergence and training stability. A range of learning rates between 0.1 to 0.00001 was tested. For DAE, a learning rate of 0.01 is selected while a learning rate of 0.002 is used for U-Net. To make a fair comparison, all the key hyperparameters are kept the same for both types of architectures, including bottleneck size, filter count, layer count, input and output dimensions, and batch size.

4.3 Quantitative Comparison

In this and the next subsection, we conduct quantitative and qualitative comparative analyses between U-Net and DAE models in the segmentation of coronal loops, providing rationale for selecting the DAE model as the basis for our approach. The improved results of our attention-aware DAE will be shown in Sec. 4.5.

4.3.1 Accuracy, Recall, Precision

Fig. 5 illustrates the accuracy over a range of different numbers of parameters for four different bottleneck sizes. The top sub-figures demonstrate that accuracy tends to plateau at approximately 94% for bottleneck sizes of 16×16 and 32×32. For larger bottleneck sizes of 64×64 and 128×128, the accuracy of the DAE saturated at around 94%, whereas the U-Net's accuracy decreases, leveling off at about 93%. The DAE reached its highest accuracy with a bottleneck size of 32×32. The most significant discrepancy in accuracy between DAE and U-Net occurred at bottleneck sizes larger than 64×64. Among all 40 models tested, the DAE reached

Table 1: Performance metrics across bottleneck sizes using the synthetic dataset.

Category	Accu.	Recall	Prec.	#Params
16×16_DAE	0.921	0.649	0.783	27,457
16×16_U-Net	**0.929**	**0.699**	**0.795**	28,986
32×32_DAE	**0.941**	0.766	**0.825**	27,457
32×32_U-Net	0.940	**0.774**	0.816	28,986
64×64_DAE	**0.939**	**0.767**	**0.812**	19,101
64×64_U-Net	0.931	0.722	0.792	18,653
128×128_DAE	**0.939**	0.749	**0.821**	12,257
128×128_U-Net	0.930	**0.794**	0.750	10,588

a peak accuracy of 94.1%, slightly higher than the U-Net counterpart.

Next, the accuracy, recall, and prevision for the models of similar sizes are compared and shown in Table 1. Each model had 4 bottlenecks. The bold values indicate the best performance within each bottleneck size for the respective metric. As the bottleneck size increases, the best performance shifts between the models for different metrics. Overall, DAE outperforms U-Nets, particularly in terms of accuracy and precision at large bottleneck sizes.

4.3.2 *ROC Curve and AUC*

The Receiver Operator Characteristic (ROC) curve plots the true positive rate against the false positive rate, and the Area Under the Curve (AUC) summarizes the ROC Curve. Across the different bottleneck sizes, the ROC Curves were very similar, converging nearly to 1.0 before false positive rate reaching 0.2. (The ROC Curve plots are omitted due the space limit.) The AUCs for U-Net and DAE were also similar, approximately around 0.973. (Here, we note that the AUC of our new attention-aware DAE was 0.992.)

4.3.3 *Confusion Matrix*

Fig. 6 and Fig. 7 illustrate the confusion matrices for DAE and U-Net, respectively. (For an imbalanced dataset like the coronal images, the number of 0s, i.e. background, is much higher than the number of 1s, i.e., loop pixels.) According to the results, the percentage of predicted true positives is higher for DAE for the bottleneck sizes of 32×32, 64×64, and 128×128. (This matches the accuracy shown in Table 1.)

4.3.4 *Performance and Training Data Size*

For the next analysis, we explore how the size of the training data impacts model performance. This investigation aids in identifying which model type is more data efficient and which is more susceptible to overfitting. Fig. 8 shows the accuracy of the U-Net and DAE for training data sizes of 500, 1,000, 5,000, and 10,000. According to the results, DAE performed better across

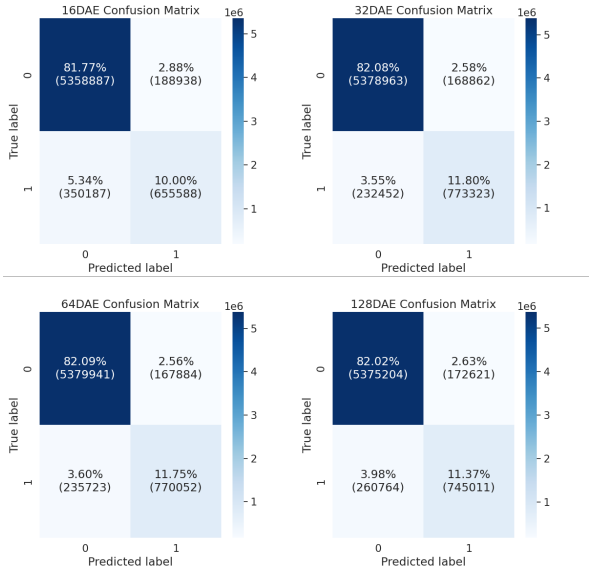

Figure 6: Confusion Matrices for DAE on our dataset across each of the tested bottleneck sizes - 16×16, 32×32, 64×64 and 128×128.

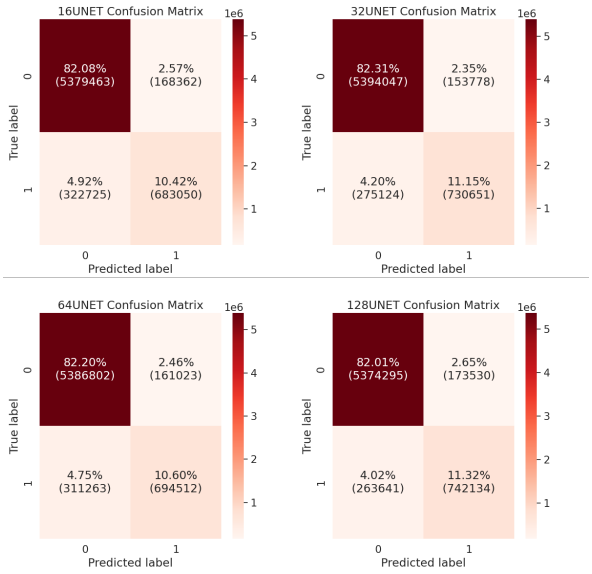

Figure 7: Confusion Matrices for U-Net on our dataset across each of the tested bottleneck sizes - 16×16, 32×32, 64×64 and 128×128.

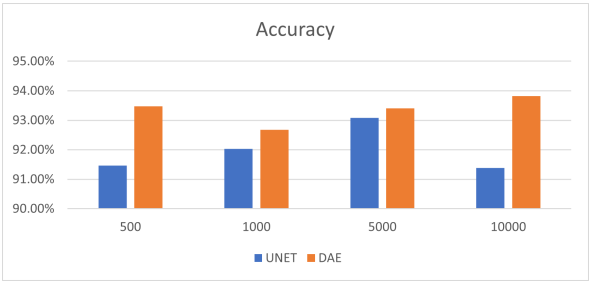

Figure 8: Accuracy vs. training data size (e.g., 500, 1,000, 5,000, and 10,000). Model bottleneck size: 16×16.

all training data sizes. Table 2 shows different model performance metrics across various training data sizes. As we can see, U-Net is less data efficient as the accuracy is much lower when the training data is limited (e.g., 500). This lower accuracy also demonstrates U-Net's higher vulnerability to overfitting. As the training data amount increases, the accuracy of U-Net improves. However, it decreases again when the data size reaches 10,000. On the other hand, DAE consistently maintains higher accuracy across varying sizes of training data.

4.4 Qualitative Comparison

In this subsection, we offer a qualitative assessment of U-Net and DAE by discussing visual observations of model predictions and evaluating the models' effectiveness in gap-filling.

4.4.1 Visual Observations

Fig. 9 shows the predictions of U-Net and DAE for both a synthetic coronal image sample (in (a)) and a real image sample (in (b)). The test image and its ground truth (i.e., labels) are shown on the left-most column of each sub-figure. The prediction results of the models using four different bottleneck sizes are shown on the four right columns, labeled by the bottleneck sizes.

In the predictions on synthetic images, both U-Net and DAE produced very comparable results, with the predicted curves appearing fuzzier when smaller bottleneck sizes were utilized in both models. In predictions on the real image, U-Net demonstrated a great ability to accurately capture curvilinear structures when the bottleneck size is small (16×16). This is likely attributed to the skip connections in U-Net. However, as the bottleneck size increased, DAE's predictions improved and seemed more accurate, whereas U-Net produced significantly more false positives than DAE.

4.4.2 Gap-Filling Effects

As described in Section 4.1, certain loops may exhibit discontinuities. Fully segmenting such loops would help solar physicists in understanding the underlying properties of the coronal loops. Fig. 10 demonstrates the gap-filling performance by the U-Net and DAE (we select the bottleneck size of 32×32 according to Table 1). Sub-figure (a) shows the ground truth magnetic field lines. Sub-figure (b) is the synthetic corona image generated from the magnetic field lines shown in sub-figure (a). As shown in the figure, there are disconnected loops, such as those in the top left and bottom left corners. Sub-figures (c) and (d) show the predictions of U-Net and DAE. As shown in the figures, both prediction models achieved reasonably well accurate segmentation. However, the DAE demonstrated a more noticeable gap-filling effect than the U-Net.

Table 2: Results of 16×16×256 model on various training data sizes. # of parameters were 784,385 and 1,179,121 for DAE and U-Net, respectively

Models	Size	Accu.	Recall	Prec.
DAE	500	**93.48%**	**78.67%**	**78.77%**
U-NET	500	91.47%	78.67%	68.47%
DAE	1,000	**92.68%**	66.94%	**80.35%**
U-NET	1,000	92.03%	**79.70%**	70.40%
DAE	5,000	**93.41%**	**86.00%**	73.82%
U-NET	5,000	93.09%	80.28%	**74.88%**
DAE	10,000	**93.82%**	80.30%	**78.51%**
U-NET	10,000	91.38%	**80.52%**	67.56%

4.5 Results of New Attention-aware DAE

Next, we will analyze the performance of our attention-aware DAE (i.e., AtnDAE) based on the proposed Encoding-Aware Decoding Attention (EADA) mechanism. Compared with competing baselines, our novel AtnDAE exhibited superior segmentation performance, encompassing the gap-filling effect, in both quantitative and qualitative dimensions. As shown in Table 3, our AtnDAE outperformed DAE significantly in accuracy, recall, and precision, underscoring the effectiveness of our attention mechanism.

Fig. 11 shows the prediction of our AtnDAE on a synthetic image example. In the figure, regions of the image where the gap-filling effect is noticeable are highlighted with red circles. As shown in the figure, all loops were not only distinctly segmented (with no blurring or missing loops), but all gaps were also effectively filled. Fig. 12 demonstrates the comparison of the gap-filling effect among U-Net, DAE, and our new AtnDAE. Our AtnDAE not only achieved superior clarity in segmenting the loops but also effectively addressed gaps within them, areas overlooked by both the U-Net and DAE models. Fig. 13 shows the confusion matrix for our AtnDAE model. Compared to the DAE and U-Net models, there was an increase in both true positives and true negatives, accompanied by a decrease in both false positives and false negatives. For example, the false positives and false negatives were 3.55% and 2.58% for the DAE, 4.20% and 2.35% for the U-Net, and 1.54% and 1.81% for the AtnDAE. (Refer to the top-right sub-images in Fig. 6 for the DAE and in Fig. 7 for the U-Net, respectively.)

More segmentation results obtained with our AtnDAE are displayed in Fig. 14.

Table 3: Performance of DAE and AtnDAE.

Model	Accu.	Recall	Prec.
32x32-DAE	0.939	0.741	0.829
32x32-AtnDAE	**0.967**	**0.888**	**0.872**

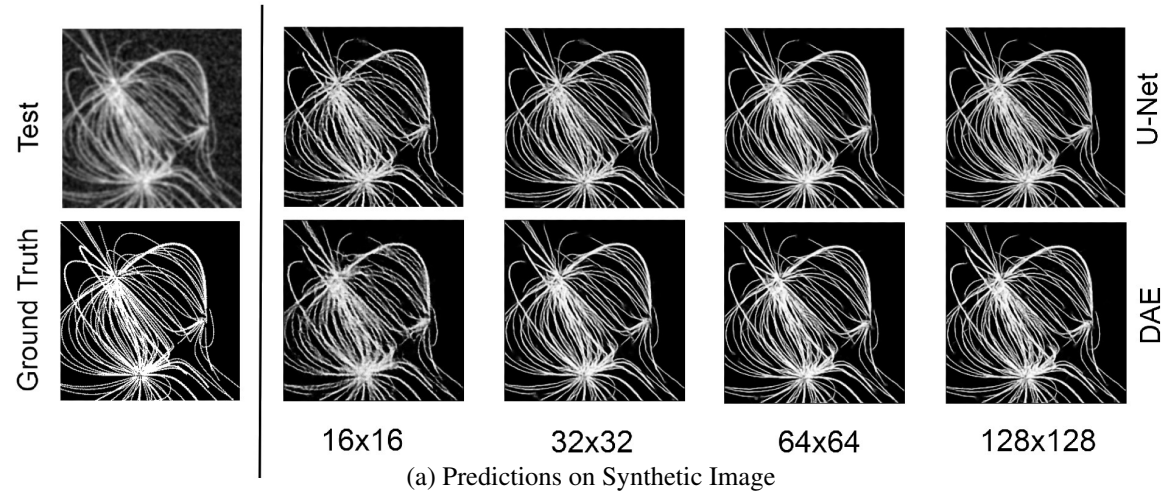

(a) Predictions on Synthetic Image

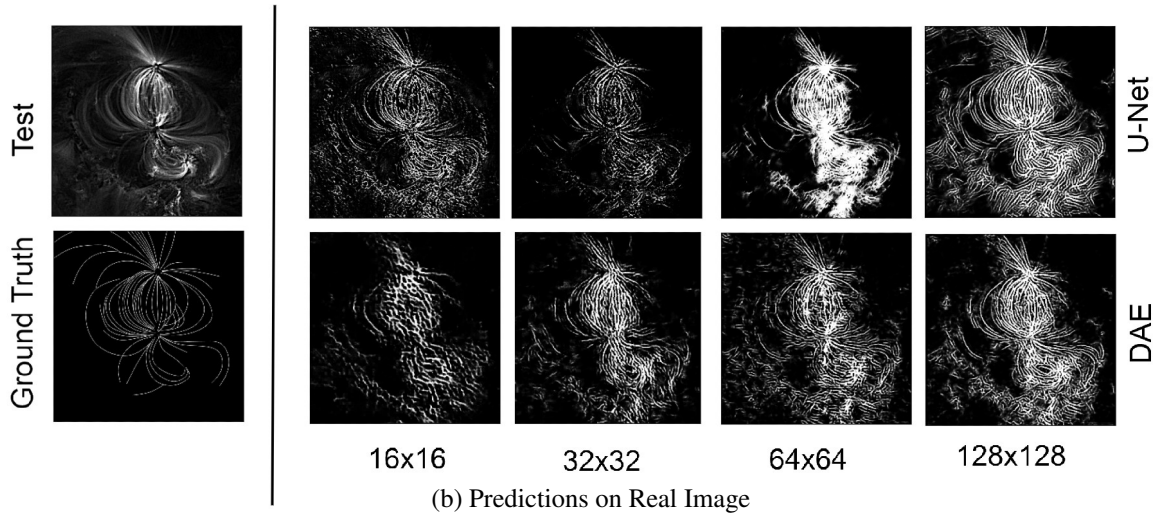

(b) Predictions on Real Image

Figure 9: Predictions of U-Net and DAE using different bottleneck sizes.

4.6 Ablation Study: Influence of Loss Functions

We have investigated nine loss functions for coronal loop segmentation. Table 4 presents the selected loss functions with their corresponding performance metrics. In the table, the top three values for each metric are highlighted. The binary cross-entropy, Binary Dice, and Dice Loss had the highest accuracy. Combo Loss, Focal Tversky, and Tversky Loss had the highest recall. The binary cross-entropy, Focal Loss, and Binary Tversky had the highest precision. Given the results, we use binary cross entropy for all previous experiments.

Table 4: Performance of different loss functions. Base model: 32x32 DAE. Accu.: Accuracy, Prec.: Precision.

Loss Function	Accu.	Recall	Prec.
BinaryCrossEntropy	**0.9390**	0.7670	**0.8120**
Binary Dice	**0.9388**	0.8511	0.7633
Dice Loss	**0.9378**	0.8446	0.7615
Focal Loss	0.9375	0.7426	**0.8189**
HED Loss	0.9242	0.9325	0.6775
Combo Loss	0.9216	**0.9387**	0.6673
Binary Tversky	0.9190	0.5160	**0.8913**
Focal Tversky	0.9121	**0.9514**	0.6359
Tversky Loss	0.9117	**0.9517**	0.6347

5 CONCLUSION

In conclusion, we present a new approach for automated segmentation of solar coronal loops through the development of the attention-aware Denoising Autoencoder (AtnDAE). This novel approach, integrating an Encoding-Aware Decoding Attention (EADA) mechanism, has significantly improved the segmentation accuracy, demonstrating superior performance over competing models such as the U-Net and DAE. The quantitative and qualitative assessments highlight the AtnDAE's improved capacity for achieving higher levels of accuracy, precision, and recall. Additionally, they underscore its effectiveness in addressing the inherent challenges of gap-filling in coronal loop imagery. These findings underscore the significance of integrat-

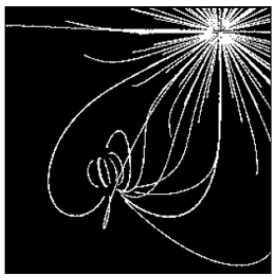

(a) ground truth mask

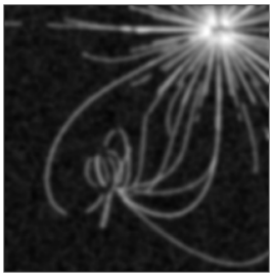

(b) test image with gap

(c) U-Net prediction

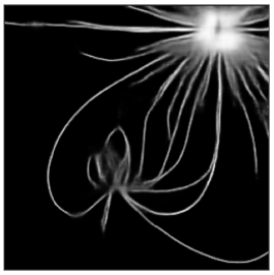

(d) DAE prediction

Figure 10: Gap-filling by U-Net and DAE. Number of layers: 6. Bottleneck size: 32×32.

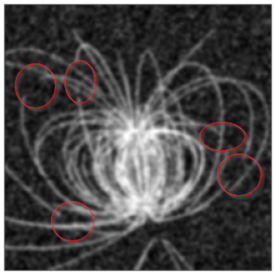

(a) Synthetic Image

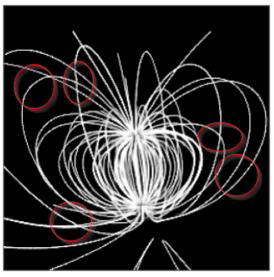

(b) AtnDAE Prediction

Figure 11: AtnDAE prediction on a synthetic corona image with gap-filling effects marked. Number of layers: 6. Bottleneck size: 32×32.

ing attention mechanisms into deep learning models for intricate pattern recognition tasks. Moreover, they emphasize the potential of the AtnDAE to make substantial contributions to our understanding of solar physics, offering a more reliable tool for analyzing coronal loops.

The comparative analyses showed the robustness of DAE models against varying conditions, including different bottleneck sizes and training data volumes, affirming its resilience to overfitting. The exploration of multiple loss functions has provided additional insights, guiding the selection of the most effective loss function for this specific task.

For future work, our aims include devising a neighborhood-aware loss function capable of addressing more intricate features and a systematic evaluation of synthetic data. Additionally, we seek to investigate the integration of our approach with

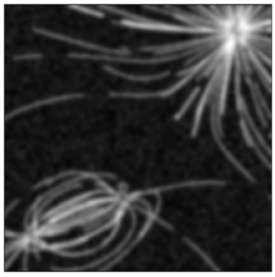

(a) test image

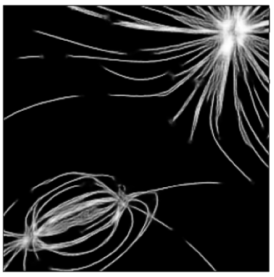

(b) U-Net prediction

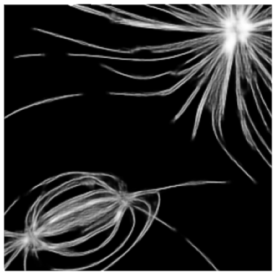

(c) DAE prediction

(d) AtnDAE prediction

Figure 12: Gap-filling by U-Net, DAE, and AtnDAE. Number of layers: 6. Bottleneck size: 32×32.

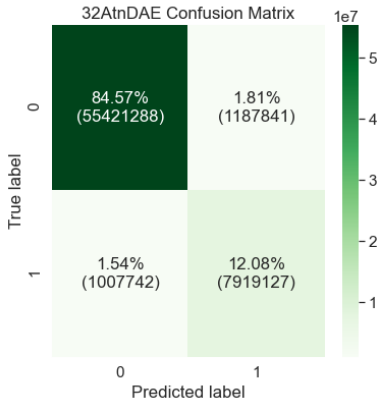

Figure 13: Confusion Matrices for the new AtnDAE on our dataset with the bottleneck size of 32×32.

other deep learning architectures, such as transformers. Furthermore, we aspire to expand the application of this methodology to tackle challenging segmentation tasks in various other fields of scientific interest.

6 ACKNOWLEDGMENTS

This research was supported in part by the David and Amy Fulton Endowed Professorship in Computer Science at Bowling Green State University. This work would not have been possible without the computing resources provided by the Ohio Supercomputer Center. We also thank the reviewers for their valuable comments which improved our paper.

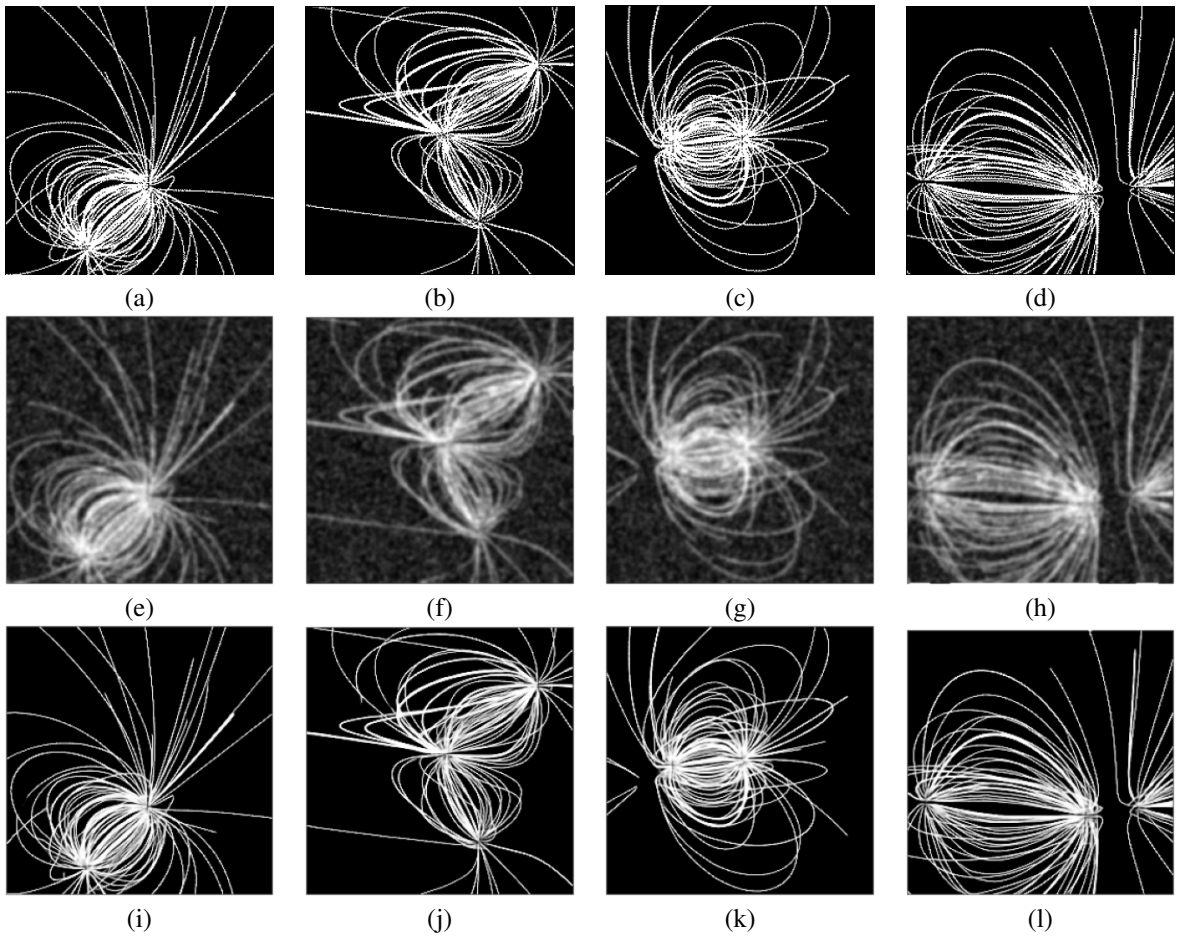

Figure 14: Segmentation results by AtnDAE: (a-d) ground truth, (e-h) test image, (i-l) AtnDAE predictions. Number of layers: 6. Bottleneck size: 32×32.

7 REFERENCES

[AnB19] Anwar, S., and Barnes, N., Real Image Denoising With Feature Attention, on Proc., IEEE/CVF Int'l Conf. on Computer Vision, October 2019, pp. 3155–3164.

[Asc10] Aschwanden, M.J., A Code for Automated Tracing of Coronal Loops Approaching Visual Perception, Solar Physics, Vol. 262, 2010, pp. 399–423

[BJP17] Benson, B., Jiang, Z., Pan, W.D., Gary, G.A., and Hu, Q., Determination of Linear Force-Free Magnetic Field Constant Alpha using Deep Learning, in Proc., 2017 Int'l Conf. on Comp. Science and Comp. Intelligence, 2017, pp. 760–765.

[BKC17] Badrinarayanan, V., Kendall, A., and Cipolla, R., Segnet: A Deep Convolutional Encoder-Decoder Architecture for Image Segmentation, IEEE Trans. on Pattern Analysis and Machine Intelligence, Vol. 39 (12). 2017, pp. 2481–2495.

[DNC09] Durak, N., Nasraoui, O., and Schmelz J., Coronal Loop Detection from Solar Images, Pattern Recognition, Vol. 42 (11), 2009, pp. 2481–2491.

[DNC10] Durak, N., Nasraoui, O., and Schmelz J., Automated Coronal-Loop Detection based on Cuntour Extraction and Contour Classification from the SOHO/EIT Images, Solar Physics, Vol. 264, 2010, pp. 383–402.

[FBH17] Fan, Z., Bi, D., He, L., Shiping, M., Gao, S., and Li, C., Low-level Structure Feature Extraction for Image Processing via Stacked Sparse Denoising Autoencoder, Neurocomputing, Vol. 243, 2017, pp. 12–20.

[Gon16] Gondara, L., Medical Image Denoising using Convolutional Denoising Autoencoders, in Proc., 2016 IEEE 16th Int'l Conf. on Data Mining Workshops, 2016, pp. 241–246.

[GSY21] Guo, C., Szemenyei, M., Yi, Y., Wang, W., Chen, B., and Fan, C., Sa-UNet: Spatial Attention U-Net for Retinal Vessel Segmentation, in Proc., 25th Int'l Conf. on Pattern Recognition, 2021, pp. 1236–1242.

[HWF21] Hu, X., Wang, Y., Fuxin, L., Samaras, D., and Chen, C., Topology-Aware Segmentation using Discrete Morse Theory, arXiv preprint arXiv:2103.09992, 2021.

[HWH19] Huang, Z., Wang, X., Huang, L., Huang, C., Wei, Y., and Liu, W., CCNet: Criss-Cross Attention for Semantic Segmentation, in Proc., IEEE/CVF Int'l Conf. on Computer Vision, 2019, pp. 603–612.

[Jad20] Jadon, S., A Survey of Loss Functions for Semantic Segmentation, in Proc., IEEE Conf. on Computational Intelligence in Bioinformatics and Computational Biology, 2020, pp. 1–7.

[KiB14] Kingma, D.P., and Ba, J., Adam: A Method for Stochastic Optimization, in Proc., 3rd Int'l Conf. for Learning Representations, 2014.

[LeT11] Lee, J.K., and Tang, W.K., Snake-based Technique for Automated Coronal Loop Segmentation, in Proc., 19th Int'l Conf. in Central Europe on Computer Graphics, Visualization, and Computer Vision, 2011, pp. 33–40

[LK16] Liskowski, P. and Krawiec, K., Segmenting Retinal Blood Vessels with Deep Neural Networks, IEEE Trans. on Medical Imaging, Vol. 35 (11), 2016, pp. 236–2380.

[LNG06a] Lee, J.K., Newman, T.S., and Gary, G.A., Oriented Connectivity-based Method for Segmenting Solar Loops, Pattern Recognition, Vol. 39 (2), 2006, pp. 246–259.

[LNG06b] Lee, J.K., Newman, T.S., and Gary, G.A., Dynamic Aperture-based Solar Loop Segmentation, in Proc., 2006 IEEE Southwest Symp. on Image Analy. & Interpret., Mar. 2006, pp. 91–94.

[MLT21] Moradi, S., Lee, J.K., and Tian, Q., Exploration of U-Net in Automated Solar Coronal Loop Segmentation, in Proc., 29th Int'l Conf. in Central Europe on Computer Graphics, Visualization, and Comp. Vision, 2021, pp. 227–235.

[MKA10] McAteer, R.T.J., Kestener, P., Arneodo, A., Khalil, A., Automated Dectecion of Coronal Loops Using a Wavelet Transformation Modulus Maxima Method, Solar Physics, Vol. 262, 2010, pp. 387–397.

[MZH20] Meng, Q., Zhang, J., Hu, Q., He, X., and Yu, J., LGNN: A Context-aware Line Segment Detector, in Proc., the 28th ACM Int'l Conf. on Multimedia, 2020, pp. 4364–4372.

[PWL18] Park, J., Woo, S., Lee, J.-Y., and Kweon, I, BAM: Bottleneck Attention Module, British Machine Vision Conf., Sept. 2018, pp. 1–14.

[RFB15] Ronneberger, O., Fischer, P., and Brox, T., U-Net: Convolutional Networks for Biomedical Image Segmentation, in Proc., Int'l Conf. on Medical Image Computing and Computer-assisted Intervention, 2015, pp. 234–241.

[SRS12] Singh, T.R., Roy, S., Singh, O.I., Sinam, T., Singh, K., and Singh, M., A New Local Adaptive Thresholding Technique in Binarization, arXiv preprint arXiv:1201.5227, 2012.

[TJF21] Tihon, S., Javaid, M.U., Fourure, D., Posocco, N., and Peel, T., DAEMA: Denoising Autoencoder with Mask Attention, in Proc., Int'l Conf. on Artificial Neural Networks, 2021, pp. 229–240.

[TME22] Theiner, J., Müller-Budack, E., and Ewerth, R., Interpretable Semantic Photo Geolocation, in Proc., IEEE/CVF Winter Conf. on Appl. of Comp. Vision, 2022, pp. 750–760.

[VLB08] Vincent, P., Larochelle, H., Bengio, Y., and Manzagol, P.-A., Extracting and Composing Robust Features with Denoising Autoencoders, in Proc., the 25th Int'l Conf. on Machine Learning, 2008, pp. 1096–1103.

[VLL10] Vincent, P., Larochelle, H., Lajoie, I., Bengio, Y., and Manzagol, P.-A., Stacked Denoising Autoencoders: Learning Useful Representations in a Deep Network with a Local Denoising Criterion, J. of Machine Learning Research, Vol. 11 (12), 2010, pp. 3371–3408.

[WPL18] Woo, S., Park, J., Lee, J.-Y., and Kweon, I.S., Cbam: Convolutional Block Attention Module, in Proc., European Conf. on Computer Vision, 2018, pp. 3–19.

[XiT15] Xie, S., and Tu, Z., Holistically-Nested Edge Detection, in Proc., IEEE Int'l Conf. on Computer Vision, 2015, pp. 1395–1403.

[XMY16] Xing, C., Ma, L., and Yang, X., Stacked Denoise Autoencoder based Feature Extraction and Classification for Hyperspectral Images, Journal of Sensors, 2016, pp. 1–10.

[YTS22] Yan, X., Tang, H., Sun, S., Ma, H., Kong, D., and Xie, X., After-UNet: Axial Fusion Transformer UNet for Medical Image Segmentation, in Proc., IEEE/CVF Winter Conf. on Appl. of Computer Vision, 2022, pp. 3971–3981.

[ZXZ19] Zhiming, S., Xiaoli, Y., Zhongquan, Q., and Hong-Bo, L., Automatic Detection and Extraction Algorithm of Coronal Loops Based on Match Filter and Oriented Directivity, Monthly Notices of the Royal Astronomical Society, Vol. 490 (4), 2019, pp. 5567–5584.