

Slice-Aware Sampling in CS-based Deep Learning Brain MRI Reconstruction

Original

Slice-Aware Sampling in CS-based Deep Learning Brain MRI Reconstruction / Spinazzola, E., Prono, L., Pareschi, F., Rovatti, R., Setti, G.. - STAMPA. - (2026), pp. 609-613. (2026 IEEE International Symposium on Circuits and Systems (ISCAS) Shanghai (Chi) 24-28 May 2026) [10.1109/ISCAS66217.2026.11562767].

Availability:

This version is available at: 11583/3013871 since: 2026-09-21T09:01:29Z

Publisher:

IEEE

Published

DOI:10.1109/ISCAS66217.2026.11562767

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2026 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

Slice-Aware Sampling in CS-based Deep Learning Brain MRI Reconstruction

Elisabetta Spinazzola^{*}, Luciano Prono^{*}, Fabio Pareschi^{*,†}, Riccardo Rovatti^{†,‡}, and Gianluca Setti^{§,‡}

^{*}DET, Politecnico di Torino, Italy - Email: {elisabetta.spinazzola, luciano.prono, fabio.pareschi}@polito.it

[†]DEI, [‡]ARCES, University of Bologna, Italy - Email: riccardo.rovatti@unibo.it

[‡]CEMSE, King Abdullah University of Science and Technology (KAUST), Saudi Arabia - Email: gianluca.setti@kaust.edu.sa

Abstract—Deep Neural-based reconstruction for Compressed Sensing (CS) has emerged as a powerful approach for accelerating brain Magnetic Resonance Imaging (MRI), minimizing scan duration and patient discomfort, surpassing traditional methods. Among the numerous available architectures, ADMM-CSNet stands out as one of the most effective state-of-the-art solutions. However, this approach is affected by a high standard deviation in the reconstruction accuracy. In this work, we investigate the cause of this variability. Specifically, we show that brain MRI itself has a high variability in information content, meaning that the first and last sections are relatively simple to reconstruct, while central views are more complex.

To overcome this problem, we adopt the Weighted Random Sampling (WRS) strategy, enhancing both reconstruction quality and robustness. Although this is not typically applied to this kind of dataset, it allows a refinement of the ADMM-CSNet framework, directing the network's learning process to prioritize central slices over lateral slices. Experimental results demonstrate that the proposed approach consistently outperforms the baseline ADMM-CSNet and other state-of-the-art methods in terms of mean achieving a 22% lower standard deviation than the original model. Finally, the experimental results on pathological MRI images further validate the efficacy of the proposed method, showing an average PSNR gain of 3 dB over the original baseline.

I. INTRODUCTION

Magnetic Resonance Imaging (MRI) is a key tool in modern medical diagnostics due to its ability to provide high-resolution, non-invasive, and multi-contrast images of soft tissues. It is widely used for detecting and monitoring a variety of conditions, from neurological disorders and cardiovascular diseases to musculoskeletal injuries and cancers. Despite its clinical value, one of the major limitations of MRI is its inherently long acquisition time. MRI studies typically last between 20 and 60 minutes, depending on the anatomical region and imaging protocols. Prolonged scan times can lead to patient discomfort and motion artifacts, while also decreasing scanner efficiency and limiting MRI use in urgent contexts as emergency care or pediatric settings [1]. The long acquisition time in MRI comes from the need to sample a large amount of data in the k-space (frequency domain) to accurately reconstruct high-quality images. Traditional Nyquist-based sampling methods state that a certain number of measurements must be taken to ensure accurate image reconstruction, leading to prolonged scan durations. This requirement presents a significant bottleneck in clinical workflows and has driven the exploration of acceleration techniques. In this sense, as shown in Fig. 1, Compressed Sensing (CS) [2] emerges as a powerful solution to this challenge by enabling accurate image reconstruction




Fig. 1. CS framework in MRI scan: an encoder acquires undersampled data y , and a decoder reconstructs $\hat{x}$ close to the original image x .

TABLE I
COMPARISON OF THE MRI CS-DECODING STATE OF ART METHODS ON THE IXI BRAIN DATASET ON GAUSSIAN 30% UNDERSAMPLING MASK

Method	PSNR Mean [dB]	PSNR Std [dB]
ADMM [10]	21.15	4.72
Plain Unet [11]	37.87	5.33
KIGAN [15]	36.08	1.05
ADMM-CSNet [16]	50.27	8.10

(decoder stage) from significantly faster acquisitions based on fewer measurements than traditionally required (encoder stage). Based on the principles of sparsity and incoherence, CS takes advantage of the fact that many medical images can be sparsely represented in a suitable transform domain. Undersampling the original image into fewer measurements, it can be recovered by applying nonlinear reconstruction algorithms [2]–[10]. Recently, Deep Neural Networks (DNNs) have replaced traditional optimization algorithms, demonstrating superior performance in the decoder stage [11]–[14]. In particular, among the DNN-based reconstruction approaches, ADMM-CSNet [16] demonstrates a competitive performance in brain MRI reconstruction. Table I summarizes its performance compared to other approaches we have reproduced.

Although ADMM-CSNet overcomes many of the inherent limitations of traditional CS techniques, the reconstruction performance in brain MRI is often afflicted by the high standard deviation of the accuracy. This is not necessarily a consequence of the network architecture itself, but rather due to the intrinsic characteristics of the dataset and anatomical regions. Specifically, MRI datasets are composed of multiple acquisitions that include two-dimensional slices or three-dimensional volumes of the brain. During the scanning process, the brain is sampled across multiple planes or sections, each exhibiting variability due to anatomical, physiological, and acquisition-related factors. The slices acquired at the

beginning and at the end of the scan are generally simpler to reconstruct, since they contain less anatomical information and fewer structural variations. In contrast, the central slices of the acquisition are richer in energy and anatomical details, making them more complex and challenging to reconstruct.

To solve this, we present a refined version of ADMM-CSNet, aimed at reducing the standard deviation in image reconstruction. We optimize the training process by applying a Weighted Random Sampling (WRS) [17] strategy, so that the network can learn more effectively from those slices that are more difficult to reconstruct (the central slices). We validate the resulting model both on the IXI Dataset [18] and on pathological MRI for Brain Tumor with Bounding Boxes dataset [19], demonstrating superior reconstruction quality and generalization capability compared to the standard training approach.

The remainder of the paper is structured as follows. Section II outlines the fundamentals of the CS theory. Then, Section III presents the original ADMM-CSNet model and the proposed training approach. We report numerical results and evidence in Section IV. Finally, the conclusion is drawn.

II. CS FRAMEWORK

CS is a signal acquisition method that allows the retrieval of the information using fewer measurements compared to what Nyquist-Shannon sampling theorem dictates, for signals with certain properties, namely *sparsity* and *incoherence*. The former means that there exists a linear operator $\mathbf{D}$ that maps the signal $\mathbf{x} \in \mathbb{R}^n$, to a space where its representation $\boldsymbol{\xi} = \mathbf{D}^{-1}\mathbf{x}$ is sparse, i.e., it contains at most $k \ll n$ non-zero entries.

Conversely, an encoding operator $\mathcal{A}$ is *incoherent* if it spreads information in a way that retains the important details. In MRI images, the encoding of the original signal $\mathbf{x} \in \mathbb{R}^n$ into a measurement $\mathbf{y} = \mathcal{A}(\mathbf{x}) \in \mathbb{R}^m$, with $m < n$, is achieved by means of a sub-sampling mask $\mathbf{M} \in \{0, 1\}^n$ containing m ones. We can build $\mathbf{M}$ by using different approaches, among which the most employed are cartesian masks, radial masks, spiral masks and Gaussian probability masks.

A compressed signal is reconstructed by means of the optimization of the sparse vector $\boldsymbol{\xi}$ by solving the following convex optimization problem

$$\min_{\mathbf{x}} \frac{1}{2} \|\mathcal{A}(\mathbf{x}) - \mathbf{y}\|_2^2 + \lambda \|\boldsymbol{\xi}\|_1, \quad (1)$$

where ℓ_1 -norm promotes sparsity in the solution, and $\lambda > 0$ is a regularization parameter. Among different approaches, Alternating Direction Method of Multipliers (ADMM) [10] solves the optimization problem by introducing auxiliary variables $\{z_1, z_2, \dots, z_L\}$ in

$$\begin{aligned} \hat{\mathbf{x}} = \min_{\mathbf{x}} \frac{1}{2} \|\mathcal{A}(\mathbf{x}) - \mathbf{y}\|_2^2 + \sum_{l=1}^L \lambda_l \|\mathbf{z}_l\|_1 \\ \text{s.t. } \mathbf{z}_l = \mathbf{D}_l \mathbf{x}, \forall l \in \{1, 2, \dots, L\} \end{aligned} \quad (2)$$

where $\hat{\mathbf{x}}$ is the reconstructed signal and $\mathbf{D}_l$ is each one of L filters. ADMM solves this constrained problem via iterative updates of $\mathbf{x}$, $\{z_1, \dots, z_L\}$ and a scaled dual variable $\mathbf{u}$, as described in [16].




Fig. 2. ADMM-CSNet architecture. Multiple Net-blocks perform a full reconstruction step. The second block processes data iteratively 8 times, emulating the ADMM algorithm. Inside a Net-block, data passes through a first reconstruction layer R , then into two convolutional layers which perform the ADMM auxiliary variable update (Z blocks), followed by a multiple update module MU .

III. PROPOSED APPROACH

A. ADMM-CSNet

In this section, we briefly describe the ADMM-CSNet architecture [16]. The method derives from the aforementioned ADMM solver, where each iteration of the optimization algorithm is mapped into a *Net-block*, as depicted in Fig. 2. In this framework, the solution is obtained by iteratively updating the variables elaborated by the *Net-block*, through a looped structure in which the input data pass eight times through the same block. Each *Net-block* is composed of a reconstruction layer R , which corresponds to the $\mathbf{x}$ update in the ADMM algorithm. Then, the auxiliary variable $\mathbf{z}$ update is performed by the Z block, which in turn is composed of a convolutional layer and a non-linear transform layer, optimized to enforce sparsity from data. Finally, the scaled dual variable $\mathbf{u}$ update is computed by the multiple update layer MU , allowing the network to adapt how the constraint enforcement evolves during training.

B. Weighted training approach

While ADMM-CSNet has demonstrated better performance compared to other state-of-the-art methods in CS MRI reconstruction, as shown in Table I, one observed limitation is the large standard deviation of the reconstruction quality across different slices. This variability is mainly due to the intrinsic characteristics of MRI datasets [20]–[22]. Specifically, an MRI volume is composed of multiple slices acquired along different planes of the brain. On the one hand, we can distinguish peripheral slices, which are acquired at the beginning and end of the scan and typically contain less anatomical information, which are therefore relatively easy to reconstruct. On the other hand, central slices are more challenging to reconstruct reliably, as they are richer in structural detail. As a result, the reconstruction error distribution across slices is imbalanced, with some slices being consistently harder to reconstruct.

To address this issue, we introduce a WRS [17] strategy during training. This method is commonly employed in machine learning to mitigate class imbalance by oversampling

underrepresented categories. In our case, we adapt this concept to account for the imbalance between easy-to-reconstruct and difficult-to-reconstruct slices. We detail the approach in Algorithm 1.

Algorithm 1 Weighted Random Sampling (WRS)

Require: A population V of a weighted items $\{v_1, v_2, \dots, v_a\}$ with weights $\{w_1, w_2, \dots, w_a\}$, and a desired sample size b

Ensure: A set S of size b with a weighted random sample

- 1: Initialize $S \leftarrow \emptyset$
- 2: **for** $k = 1$ to b **do**
- 3: Compute probabilities for each item $v_i \in V \setminus S$:

$$p_i(k) = \frac{w_i}{\sum_{j \in V \setminus S} w_j}$$

- 4: Randomly select one item $v_i \in V \setminus S$ with probability $p_i(k)$
 - 5: Insert v_i into S
 - 6: **end for**
 - 7: **return** S
-

Concretely, we assign a reference group to each training slice, either central or lateral. In a MRI acquisition, of 150 slices, those from index 40 to 110 are labeled as central, while the remaining ones are labeled as lateral. Then, we give a sampling weight w according to its group within the MRI volume: we use $w = 2$ for central slices, which contain a higher anatomical complexity and $w = 1$ for peripheral slices. During training, we sample slices proportionally to these weights. This means that central slices have a higher probability of being selected and presented to the network compared to peripheral ones.

In this strategy, the number of samples b drawn per epoch equals the total number a of training images. However, the use of a WRS modifies the selection probabilities, favoring central slices over peripheral ones. This strategy directly alters the training distribution: instead of the model being exposed equally to all slices (which would bias learning toward the easier peripheral slices), it is guided to encounter more frequently the difficult central slices that dominate reconstruction error. As a result, the network is encouraged to refine its feature extraction and reconstruction capability specifically in the regions where it previously underperformed, while still maintaining full dataset coverage.

As training progresses, the network thus learns to balance performance across the dataset. Peripheral slices remain well reconstructed—since their structure is simpler—while central slices benefit from the increased training frequency. The net effect is a reduction of the standard deviation of the accuracy across slices and a measurable improvement in average reconstruction quality.

IV. NUMERICAL RESULTS

A. Datasets and Hyperparameters

The high variability in MRI acquisition processes results in different datasets with very different characteristics. We

used two of the most well-known datasets used for training DNN-based CS-MRI decoders. For the method tuning and evaluation, we considered the T1-weighted 3D healthy brain data from the IXI Dataset [18]. Data from this collection have been recorded using a Philips 3T and 1.5T system. Specifically, the dataset selected for our analysis includes 581 scans from approximately 573 subjects. Each scan is made of around 150 2D slices uniformly cropped to 256×256 pixels. We split the dataset in a training, validation and testing set composed of 373, 92 and 116 3D scans, respectively.

To further validate the generalization capabilities of the method, we employed pathological MRI slices from the MRI for Brain Tumor with Bounding Boxes dataset [19], containing high-quality MRI images of brain tumors with detailed annotations. The dataset includes a total of 5249 MRI images divided into training and validation sets. Each image is annotated with bounding boxes in YOLO format [23] and labels corresponding to one of the three classes of brain tumors, i.e., Glioma, Meningioma and Pituitary. We extracted from this dataset a validation set of 512 images, including MRI scans acquired from sagittal, axial, and coronal views. By using this dataset, we assess how well the model generalizes to images with real pathological features, complementing standard validation on healthy MRI data.

We train the model end-to-end by using MRI image slices from IXI on k -space data with Adam optimizer with a learning rate of 0.0002. The loss is based on the normalized reconstruction error (NRE), formulated as the Frobenius norm of the difference between the original and reconstructed image. Training was carried out with a batch size of 1, as in the original work, for 10 epochs.

B. Metrics

We evaluate the performance of the method with Peak Signal-to-Noise Ratio (PSNR)

$$\text{PSNR} = 10 \log \frac{\|I\|_{\max}^2}{\frac{1}{n} \|I - I_r\|_F^2} \quad (3)$$

where I is the original image, I_r is the reconstructed image, $\|\cdot\|_F$ denotes the Frobenius norm and $\|\cdot\|_{\max}$ denotes the maximum absolute value. PSNR is typically used as an estimation of the difference between two images and is commonly employed to assess MRI CS-reconstruction performance.

C. Method evaluation

To evaluate the performance of the proposed method and its improvements in reconstruction quality, we compute the mean of the PSNR on the entire IXI test set and compare it with other state-of-art-approaches, using a Gaussian undersampling mask with 30% of the total samples. The refined method achieves a PSNR of (51.37 ± 6.35) dB, demonstrating a higher PSNR in mean compared to the original ADMM-CSNet model and also among state-of-the-art MRI reconstruction DNNs, as shown in Table I. At the same time, tests performed on a cartesian sampling mask show an improvement from (37.11 ± 8.82) dB to (39.00 ± 7.97) dB. In terms of reconstruction consistency, the proposed method achieves a 22% reduction




Fig. 3. Comparison of PSNR distributions on the validation set: the proposed method's results (blue) show a more stable and improved reconstruction quality compared to the original model (red).




Fig. 4. PSNR distributions on the validation set with respect to the indexed slices. The proposed method (blue) obtains more uniform PSNR values across all the slices and higher mean performance compared to the original one (red).

of the standard deviation compared to the original ADMM-CSNet. While other approaches may exhibit a lower standard deviation, they do not attain comparably high mean PSNR values. This limitation might arise because, during training, these networks tend to specialize on easier slices, dedicating less learning capacity to the more challenging central slices, which are critical for achieving high-fidelity reconstructions across the entire MRI volume.

To better illustrate the consistency of the results, Fig. 3 shows the Probability Density Function (PDF) of PSNR across the test set. Compared to the baseline (original ADMM-CSNet), the proposed method exhibits a narrower distribution, indicating reduced variability, and a pronounced peak at higher PSNR values, reflecting improved reconstruction quality across the dataset.




Fig. 5. Visual comparison of a reconstructed slice of the MRI Brain Tumor with Bounding Boxes dataset, obtained with the proposed method and the original model, with respect to the fully sampled image. The image shows the tumor region in detail with the PSNR values evaluated on the tumor box.

On the other hand, Fig. 4 illustrates the PSNR values with respect to the slice indices (ranging from 1 to 150, corresponding to the beginning and end of the acquisition). From this, it is evident that the central slices achieve higher PSNR values compared to the baseline, which is also where the retrieval of information is most important from the clinical point of view. This results in a reduced standard deviation and overall improved reconstruction performance across the total IXI validation dataset.

Finally, we validated the proposed approach on pathological data using the MRI for Brain Tumor with Bounding Boxes dataset. With the baseline ADMM-CSNet, we achieve a mean PSNR of 44.92 dB. In contrast, the proposed Weighted training approach attains a mean PSNR of 47.92 dB, which is 3 dB higher compared to vanilla training. We evaluated the PSNR values specific to the tumor bounding boxes, featuring higher scores for the proposed method (31.11 dB) compared to the original model (29.20 dB). We observe that the reduction of the standard deviation of the model accuracy also contributes to improved generalization capabilities, resulting in enhanced performance on distinct datasets. Fig. 5 shows a slice case of the achieved performance.

V. CONCLUSION

In this work, we proposed a refinement of the ADMM-CSNet framework which leverages a WRS training strategy, addressing the variability of the reconstruction performance across scan MRI slices. By prioritizing the more challenging central slices during training, the model achieves higher consistency and improved reconstruction quality. Experimental evaluations confirm that our approach not only reduces the standard deviation by 22% compared to the original ADMM-CSNet, but also delivers superior results on pathological MRI images, introducing an improvement of 3 dB. This demonstrates that this training approach improves the robustness of DNN-based reconstruction models for clinical applications in accelerated MRI reconstruction.

REFERENCES

- [1] F. GharehMohammadi and R. A. Sebro, "Efficient health care: decreasing mri scan time," *Radiology: Artificial Intelligence*, vol. 6, no. 3, p. e240174, 2024. doi:10.1148/ryai.240174
- [2] D. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, 2006. doi:10.1109/TIT.2006.871582
- [3] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D: Nonlinear Phenomena*, vol. 60, no. 1, pp. 259–268, Nov. 1992. doi:10.1016/0167-2789(92)90242-F
- [4] M. Lustig, D. Donoho, and J. M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–1195, 2007. doi:10.1002/mrm.21391
- [5] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic Decomposition by Basis Pursuit," *SIAM Journal on Scientific Computing*, vol. 20, no. 1, pp. 33–61, Jan. 1998. doi:10.1137/S1064827596304010
- [6] E. Candes and T. Tao, "Decoding by linear programming," *IEEE Transactions on Information Theory*, vol. 51, no. 12, pp. 4203–4215, Dec. 2005. doi:10.1109/TIT.2005.858979
- [7] R. Tibshirani, "Regression shrinkage and selection via the lasso," *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 58, no. 1, pp. 267–288, 1996. doi:10.1111/j.2517-6161.1996.tb02080.x
- [8] I. Daubechies, M. Debrise, and C. De Mol, "An iterative thresholding algorithm for linear inverse problems with a sparsity constraint," *Communications on Pure and Applied Mathematics*, vol. 57, no. 11, pp. 1413–1457, 2004. doi:10.1002/cpa.20042
- [9] J. A. Tropp and A. C. Gilbert, "Signal Recovery From Random Measurements Via Orthogonal Matching Pursuit," *IEEE Transactions on Information Theory*, vol. 53, no. 12, pp. 4655–4666, Dec. 2007. doi:10.1109/TIT.2007.909108
- [10] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers," *Foundations and Trends in Machine Learning*, vol. 3, no. 1, pp. 1–122, Jul. 2011. doi:10.1561/22000000016
- [11] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, Eds. Cham: Springer International Publishing, 2015, pp. 234–241. doi:10.1007/978-3-319-24574-4_28
- [12] H. Cao, Y. Wang, J. Chen, D. Jiang, X. Zhang, Q. Tian, and M. Wang, "Swin-Unet: Unet-Like Pure Transformer for Medical Image Segmentation," in *Computer Vision – ECCV 2022 Workshops*, L. Karlinsky, T. Michaeli, and K. Nishino, Eds. Cham: Springer Nature Switzerland, 2023, pp. 205–218. doi:10.1007/978-3-031-25066-8_9
- [13] I. J. Goodfellow et al., "Generative Adversarial Nets," in *Advances in Neural Information Processing Systems*, vol. 27. Curran Associates, Inc., 2014. doi:10.48550/arXiv.1406.2661
- [14] Y. Song and S. Ermon, "Generative Modeling by Estimating Gradients of the Data Distribution," in *Advances in Neural Information Processing Systems*, vol. 32. Curran Associates, Inc., 2019. doi:10.48550/arXiv.1907.05600
- [15] R. Shaul, I. David, O. Shitrit, and T. Riklin Raviv, "Subsampled brain MRI reconstruction by generative adversarial neural networks," *Medical Image Analysis*, vol. 65, p. 101747, Oct. 2020. doi:10.1016/j.media.2020.101747
- [16] Y. Yang, J. Sun, H. Li, and Z. Xu, "ADMM-CSNet: A Deep Learning Approach for Image Compressive Sensing," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 42, no. 3, pp. 521–538, Mar. 2020. doi:10.1109/TPAMI.2018.2883941
- [17] P. Efraimidis and P. Spirakis, *Weighted Random Sampling*. Boston, MA: Springer US, 2008, pp. 1024–1027.
- [18] "IXI Dataset – Brain Development," <https://brain-development.org/ixi-dataset/>.
- [19] A. Sorour, "Mri for brain tumor with bounding boxes," <https://www.kaggle.com/datasets/ahmedsorour1/mri-for-brain-tumor-with-bounding-boxes/data>, kaggle dataset.
- [20] L. Oakden-Rayner, "Exploring large-scale public medical image datasets," *Academic radiology*, vol. 27, no. 1, pp. 106–112, 2020. doi:10.1016/j.acra.2019.10.006
- [21] K. A. Dishner et al., "A survey of publicly available mri datasets for potential use in artificial intelligence research," *Journal of Magnetic Resonance Imaging*, vol. 59, no. 2, pp. 450–480, 2024. doi:10.1002/jmri.29101
- [22] S. Ogawa, R. Menon, S.-G. Kim, and K. Ugurbil, "On the characteristics of functional magnetic resonance imaging of the brain," *Annual review of biophysics and biomolecular structure*, vol. 27, no. 1, pp. 447–474, 1998. doi:10.1146/annurev.biophys.27.1.447
- [23] A. Bochkovskiy, C.-Y. Wang, and H.-Y. M. Liao, "Yolov4: Optimal speed and accuracy of object detection," 2020.