

Homogenising SoHO/EIT and SDO/AIA 171Å Images: A Deep Learning Approach

Subhamoy Chatterjee¹, Andrés Muñoz-Jaramillo¹, Maher Dayeh^{2,3}, Hazel M. Bain^{4,5}, Kimberly Moreland^{2,3,4,5}

¹ Southwest Research Institute, Boulder, CO, USA.

² Southwest Research Institute, San Antonio, TX, USA.

³ The University of Texas at San Antonio, San Antonio, TX, USA.

⁴ Cooperative Institute for Research in Environmental Sciences, University of Boulder, CO, USA.

⁵ Space Weather Prediction Center, NOAA, Boulder, CO, USA.

e-mail: subhamoy.chatterjee@swri.org

Received _____; accepted _____

Abstract

Extreme Ultraviolet images of the Sun are becoming an integral part of space weather prediction tasks. However, having different surveys requires the development of instrument-specific prediction algorithms. As an alternative, it is possible to combine multiple surveys to create a homogeneous dataset. In this study, we utilize the temporal overlap of SoHO/EIT and SDO/AIA 171 Å surveys to train an ensemble of deep learning models for creating a single homogeneous survey of EUV images for 2 solar cycles. Prior applications of deep learning have focused on validating the homogeneity of the output while overlooking the systematic estimation of uncertainty. We use an approach called ‘Approximate Bayesian Ensembling’ to generate an ensemble of models whose uncertainty mimics that of a fully Bayesian neural network at a fraction of the cost. We find that ensemble uncertainty goes down as the training set size increases. Additionally, we show that the model ensemble adds immense value to the prediction by showing higher uncertainty in test data that are not well represented in the training data.

Subject headings: Sun: Solar extreme ultraviolet emission — Sun: Solar atmosphere — Sun: Solar corona — techniques: calibration — techniques: Convolutional neural networks

1. INTRODUCTION

Solar observations span across several past decades being recorded by multiple ground-based and space-based observatories. Due to changes in instrumentation, the datasets differ in resolution (spatial and temporal), field-of-view, dynamic range, and noise characteristics. While dealing with long-term solar data it requires building custom algorithms to detect solar features/events for each

instrument. Cross-calibrating those surveys and creating a single homogeneous dataset for the entire span of observation enables the scientific community to make long-term studies and discover underlying patterns without spending effort to deal with instrument differences.

There have been several efforts to homogenize solar images from different surveys using traditional approaches such as oversampling of pixels using interpolation, intensity rescaling etc. Recent studies find a significant improvement over those baseline approaches in homogenization tasks such as super-resolving solar magnetograms ([Jungbluth et al. 2019](#)) utilizing state-of-the-art machine learning (ML) approaches such as Convolutional Neural Networks.

Deep-learning approaches applied in scientific domains suffer from the shortcoming of limited availability of data. Limited data creates barriers for ML models to learn the diversity of natural phenomena. Thus, in such situations trusting a point estimate by an ML model on unseen data can result in underestimation and overestimation. Creating model-ensembles can help mitigate this problem by generating prediction uncertainty. There are different approaches to creating deep learning model-ensembles that range from multiple selections of data to random disconnections in neural networks ([Gal & Ghahramani 2015](#)) and approximating Bayesian Neural Networks by adding several anchors that drive the regularization term in the model loss ([Pearce et al. 2018](#)). Studies find the application of usage model-ensembles to be useful in overcoming the demerits of data scarcity in scientific domains such as heliophysics ([Chatterjee et al. 2022](#))

In this study, we use convolutional neural networks to homogenize and estimate the uncertainty of, full-disc solar extreme ultraviolet image in 171\AA from Extreme-ultraviolet Imaging Telescope (EIT;[Delaboudinière et al. \(1995\)](#)) onboard Solar and Heliospheric Observatory (SoHO) and Atmospheric Imaging Assembly (AIA;[Lemen et al. \(2012\)](#)) onboard Solar Dynamics Observatory

(SDO) using their overlapping period.

2. OBSERVATIONAL DATA

We use level-1 SoHO/EIT and SDO/AIA 171Å full-disc images (Figure 1) with a cadence of 1 day and over the period 2010-2020 for our study. We make sure that the EIT-AIA pairs are observed at the same time.




Fig. 1.— Level-1 representative images from SoHO/EIT and SDO/AIA from the same date. *Left panel* shows an EIT 171Å image at the native resolution and field-of-view. *right panel* shows an AIA 171Å image at native resolution and field-of-view (FoV). A comparison of the images clearly depicts the higher FoV of EIT.

3. DATA ALIGNMENT

3.1. Re-projection

The EIT-Sun-AIA angle, focal length and detector differences necessitate re-projecting EIT, AIA images to a common point-of-view and field-of-view. We use the ‘reproject’ module of Python to map AIA images to EIT detector with 4 times the original resolution of EIT (Figure 2). The re-projection of the corona is non-trivial. So, we make an assumption that coronal structures lie on a sphere centered on Earth, with a radius of Sun-Earth distance.




Fig. 2.— Outcome of re-projection. *Left panel* shows a EIT 171 Å image; *Right panel* shows a re-projected AIA image to match the FoV and point-of-view seen by EIT detector.

3.2. Template matching

Re-projection is insufficient when we want to perform one-to-one mapping of EIT to AIA patches. Several other factors such as resolution, and optical aberration result in pixel shifts and may affect our homogenization objective. To make an accurate alignment we perform template matching (Brunelli 2009) using the following steps-

1. Divide the EIT full disc images into contiguous non-overlapping 64×64 patches.
2. Blow up each of those patches by 4 times and map that location to the re-projected AIA.
3. Correlate that blown-up EIT patch within the corresponding AIA location having a margin of $(1/8)^{th}$ patch size on all 4 sides.
4. Find the 256×256 region from the AIA window providing the best correlation.
5. Follow the above steps after generating another set of non-overlapping patches bounded by the center pixels of the original patches.

The steps above produce 365 patch-pairs [64×64 EIT, 256×256 AIA] per full-disc EIT-AIA image-pair (Figure 3).

3.3. Data preparation: training, validation and test set

We sample decade-long data (05/2010- 12/2020) within the training, validation, and test set by adding the first 7 months of each year to the training set, August and September to the validation set, and, the rest to the test set. We also put the entire 2015 data aside in the test set and tune the model hyperparameters based on the validation set performance.




Fig. 3.— Benefit of patch alignment while creating input-target pair for EIT-to-AIA transformation. *Top row* shows the effect of patch alignment while the *bottom row* shows unaligned patches. While comparing the unaligned AIA patches with the aligned ones, the vertical and horizontal shift correction becomes apparent by looking at the bright region on top-right. Different columns for the AIA patch represent steps toward undersampling for visual comparison with EIT at native resolution.

4. DL-based IMAGE TRANSFORMATION

4.1. Deep learning model

We prepare the Deep learning (DL) model to transform 64×64 EIT patches to 256×256 AIA patches. The model is based on the super-resolution convolution neural net (CNN) presented in [Deudon et al. \(2020\)](#). However, our network is simple as we don't have multiple views and avoid shift-net by pre-aligning the EIT, and AIA patches. The model consists of an encoder and a decoder structure. The encoder consists of 2D convolution layers and residual blocks. The decoder consists of 4x upsampling and 2D convolution layers. In each residual block, two convolution layers are used and the convolution layer outcome is added to the input layer to generate the output. We use a kernel size of 5×5 , 32 feature layers and a parametric ReLU activation in all the convolution layers except for the final layer that has only one feature layer and a kernel size of 1×1 . We use reflection padding the deal bring edge continuity in the adjacent super-resolved tiles.

4.2. Loss function definition

To minimize the difference between the target (T) and DL predicted (P) image we consider the following as elements of loss function-

1. Mean squared error: $MSE = \frac{1}{N} \sum_i (T_i - P_i)^2$ where N stands for the number pixels in output patches.
2. Histogram ([Wang et al. 2018](#)) difference: $H = \frac{1}{K} \sum_i (HT_i - HP_i)^2$ where K stands for the number of bins in the histograms.
3. Negative of Structural Similarity Index Metric ([Zhou Wang et al. 2004](#)): $S = -\frac{(2\mu_T\mu_P+C_1)(2\sigma_{T,P}+C_2)}{(\mu_T^2+\mu_P^2+C_1)(\sigma_T^2+\sigma_P^2+C_2)}$ with $C_1 = (K_1L)^2, C_2 = (K_2L)^2$, L being dynamic range and $K_1, K_2 \ll 1$. $\mu_X, \sigma_X, \sigma_{X,Y}$ respectively represent mean of X , standard deviation of X and covariance of X, Y

4. Gradient difference: $G = \frac{1}{N} \sum_i (\nabla_x T_i - \nabla_x P_i)^2 + \frac{1}{N} \sum_i (\nabla_y T_i - \nabla_y P_i)^2$

Our objective is to perform the best performing weighted combination of MSE and one of the other losses. To identify the weights we combine MSE with each of the remaining components and calculate the best weight through a set of diagnostics as described in the following subsection.

4.3. Approximate Bayesian Ensembling

Having a small amount of training data it is important that we estimate the uncertainty in model inference. We use a state-of-the-art technique called Approximately Bayesian Ensembling (ABE) that draws a set of anchor weights from a prior distribution and factor them through an additional regularization term to the loss function defined by-

$$Loss_n = \frac{1}{N} \|\mathbf{I} - \hat{\mathbf{I}}_n\|_2^2 + \frac{1}{N} \|\mathbf{\Gamma}^{1/2} \cdot (\boldsymbol{\theta}_n - \boldsymbol{\theta}_{anchor,n})\|_2^2 \tag{1}$$

where $diag(\mathbf{\Gamma})_i = \frac{\sigma_\epsilon^2}{\sigma_{prior,i}^2}$ and N is the number of datapoints

During the training phase, the weights ($\boldsymbol{\theta}_n$) are optimized but the anchors ($\boldsymbol{\theta}_{anchor,n}$) are kept unchanged. We randomly select a set of four anchors and optimize a model for each giving rise to the ensemble of models. Here the index n runs over the ensemble members. We set σ_ϵ^2 to be the mean of the AIA patch histogram. We use anchored regularization in the last layer of our model to select both priors and anchors from a Glorot uniform (Glorot & Bengio 2010) distribution with $\sigma_{prior,i}^2 = \frac{6}{n_l + n_{l+1}}$ where n_l and n_{l+1} are the numbers of units in two consecutive layers. We rewrite $diag(\mathbf{\Gamma})_i$ as $\frac{k_1(n_l + n_{l+1})}{6}$. We use k_1 (with a default value of $\sigma_\epsilon^2 \approx 0.5$) to change the regularization weight and another parameter k_2 (with a default value of 1) to scale the standard deviation of the anchor weight distribution. We examine the sensitivity of the results to the anchored regularization with different values of k_1 (later referred to as ‘LAMBDA’) and k_2 (later referred to as ‘SD’).

5. Results

5.1. Model prediction, diagnostics and comparison with baseline approach

We develop a set of diagnostics to compare the model inference between the baseline and outcomes for different loss function combinations. For this purpose, we collect all the validation patch pairs and model inference on those. We generate 2-dimensional histograms through logarithmic binning (uniform bins in log space) from the tuple of pixel intensity from model inference and target. We calculate the median and standard deviation for each inferred intensity bin. This generates a curve representing the expected AIA target and uncertainty for each inferred value. From those, we estimate the deviation from the expected rule i.e. ‘target = expected’. We name that metric as ‘% relative residual’. We also add a metric called ‘% relative variance’ representing the standard deviation and median ratio.

From the visual inspection and metrics (Figures 4,5), we find that all the DL outcomes give rise to superior results as compared to those corresponding to simple Upsampling and scaling. As depicted in Figure 4, the DL outcomes get rid of the bright point-like artifact present in the EIT patch. Among the DL outcomes, we use MSE loss as a baseline and examine the improvement caused by additional terms in the loss function. We find that unlike MSE additional loss terms drive the ‘% relative residual’ to be within $\pm 5\%$ for an inferred intensity range of [0,6] (Figure 5). We also observe that MSE + 400*Grad performs the best for higher intensities (> 6). Also, by visual inspection, one can discern finer details in MSE+400*Grad more easily as compared to other DL outcomes.

The results show that (although better than the baseline) the deep-learning model cannot reach the same level of detail as the target AIA 171 Å images (Figure 4). This could be because of the ill-posed nature of the problem with AIA having higher diffraction-limited resolution and better pixel

pitch than EIT. It could also be the limited training data that is not able to constrain the outcome enough or that the small-scale information can't be uniquely recovered from the large scale. More experiments should be performed to understand the optimal number of loss function terms beyond which performance can't be improved given the size of the training data. Recent advancement in conditional GANs also inspires further experiments to bring back surface texture information but physical validation is needed for them to be fit for scientific applications.




Fig. 4.— Effect of loss function variations of inferred high-resolution image. The baseline outcome is depicted by the 3rd column while starting from the 4th column all the outcomes towards the right are from a DL model with loss function shown at the titles. The effect of loss function variation is highly apparent the model outcome. For example, MSE + 20*Hist does not show much micro-details as depicted by the MSE + 400*Grad. It can also be seen that the DL outcomes get rid of the apparent artifact (bright dot on patch bottom half) on EIT.

5.2. Sensitivity of Uncertainty to Training set Size

We select four anchors and train a model with each of those without changing the training and validation set. Initially, we train the model with 20000 EIT-AIA patch pair and then increase

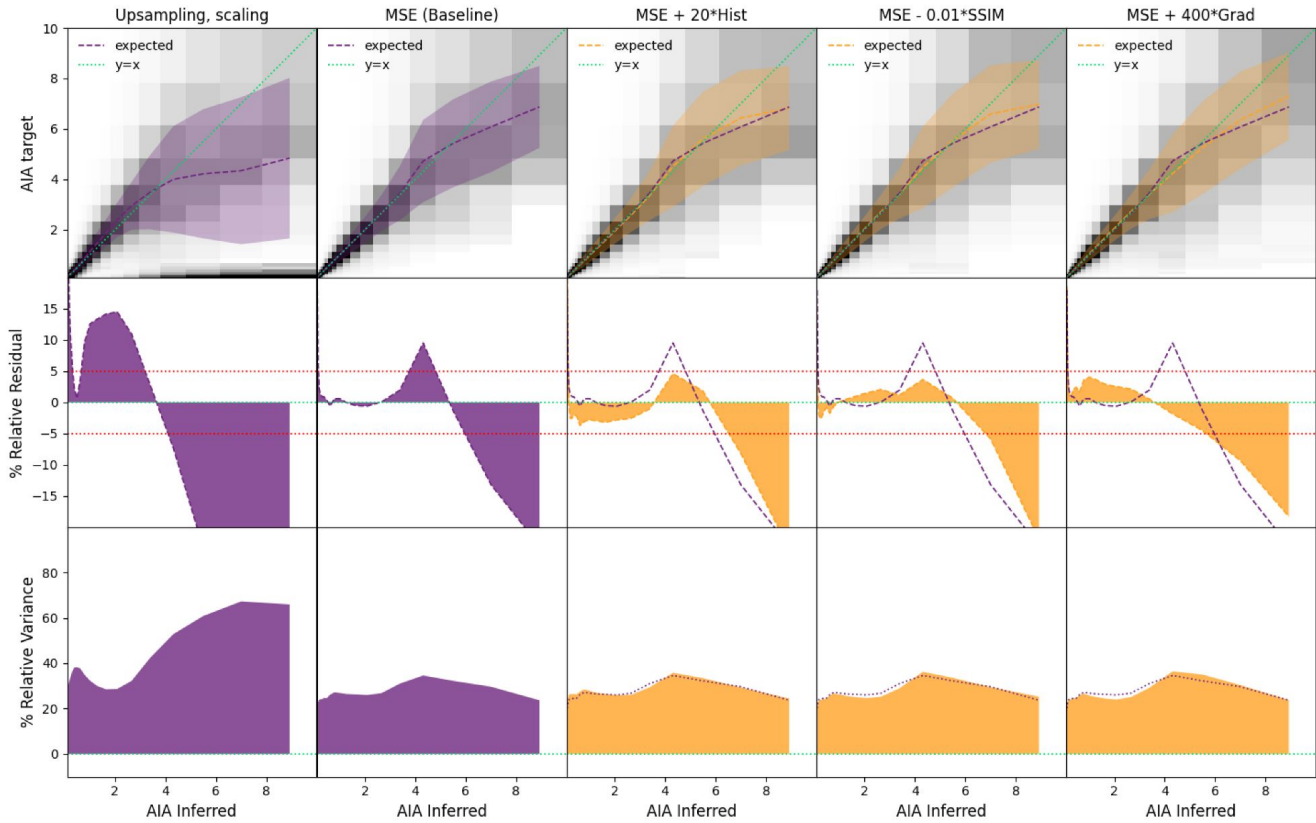


Fig. 5.— Change in metrics defining the quality of inferred images with change in the loss function. The leftmost column shows the outcome of upsampling and scaling. 2nd to 5th column starting from left show the outcomes of the DL model acquired through minimizing a loss printed as the title of each column. The top row shows a 2D histogram between the target and inferred AIA intensities across all the test patches. The dashed lines represent the median line calculated through each input inferred intensity bin and the shaded region represents the inter-quartile range. The green diagonal line represents ‘ $y=x$ ’. The middle row depicts the relative residual for meaning percentage deviation of the median line from ‘ $y=x$ ’ line. The horizontal dotted lines mark $\pm 5\%$ relative residual. The bottom row represents relative variance meaning the percentage uncertainty of outcome w.r.t. median. The outcome for MSE loss is overplotted as a baseline on all other loss function outcomes.

the number of training patches to 70000. To record the change in performance, we generate 2-dimensional histograms between ensemble standard deviation vs. input EIT intensity over all the patch pixels for training, validation, and test sets. The ensemble standard deviation ($\sigma_{ensemble}$) is calculated over 64 pixels (4×4 for 4 models) corresponding to each EIT pixel. From 2D histograms, we find a clear impact of training set size on the ensemble uncertainty. We observe that for each EIT intensity, the distribution of ensemble uncertainty when trained with 70000 patches shifts towards smaller values as compared to when trained with 20000 patches (Figure 6).

5.3. Sensitivity of Uncertainty to ensembling parameters

We also record changes in the distribution of ensemble standard deviation ($\sigma_{ensemble}$) for values of input intensity and parameters defining anchor distribution (k_2 or ‘SD’) and weightage of the regularization term (k_1 or ‘LAMBDA’) in the loss function over training, validation, and test set. It appears that for all the input intensities and anchor parameters, the $\sigma_{ensemble}$ distributions have a longer tail in the validation and test set as compared to the training set. The difference between validation/test and training set distribution diminishes for higher input intensities. However, the distributions get wider for higher input intensities. We don’t find as pronounced an effect by changing the anchor parameters. However, a close inspection reveals that the distribution gets wider with a reduction in regularization weightage for $I=0.64$. Also, for $I = 0.76$ the second peak in the distribution gets less pronounced for lower weightage (Figure 7).

5.4. Sensitivity of uncertainty to observational artifacts

On a representative solar image quadrant, we find that regions that are not well represented in the training data (e.g. coronal regions) produce high disagreement among the ensembles. Also,

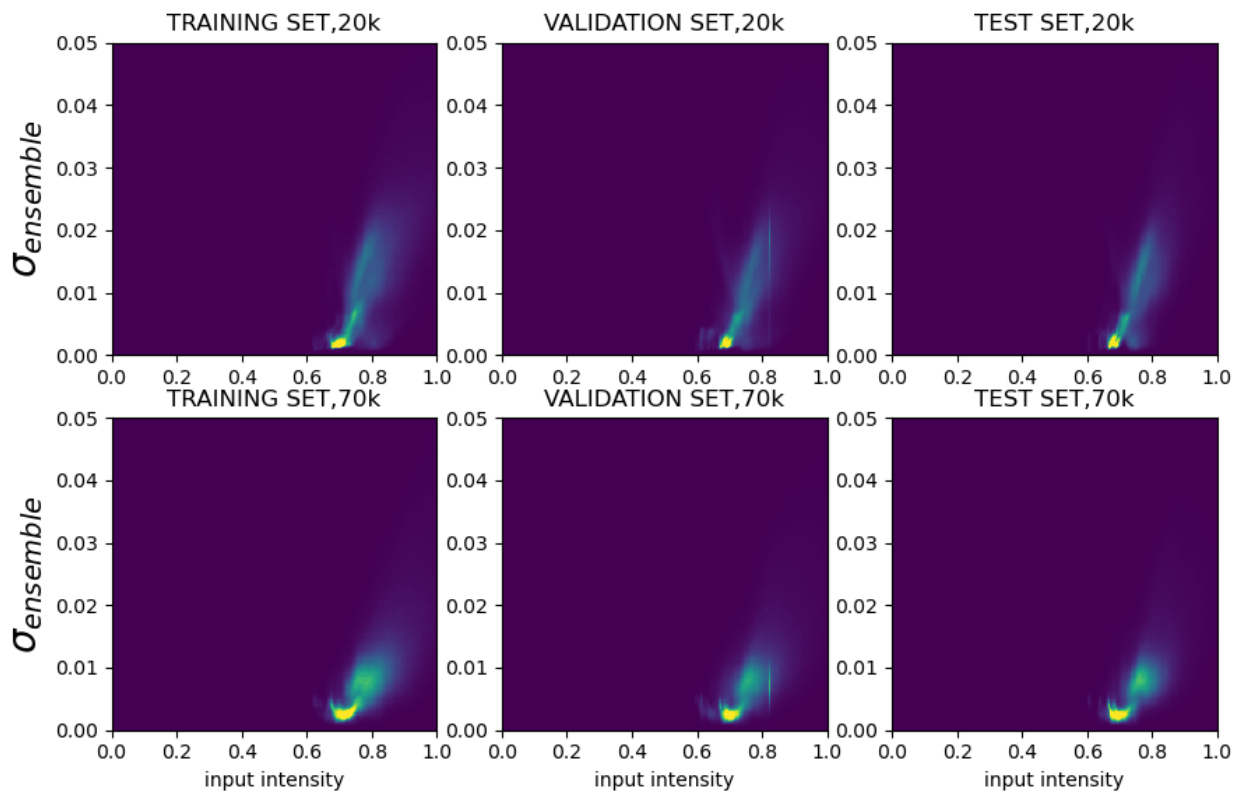


Fig. 6.— Change in inference uncertainty with training set size. Each panel represents a 2D histogram between input pixel intensity and standard deviation of ensemble inference ($\sigma_{ensemble}$) across 4 super-resolved pixels. The top row represents the 2D histograms made from training, validation, and test patches when the models were trained with 20000 patch pairs. The bottom row represents the 2D histograms made from training, validation, and test patches when the models were trained with 70000 patch pairs. A clear reduction in ensemble SD can be observed especially for higher-intensity pixels when the training set gets bigger.

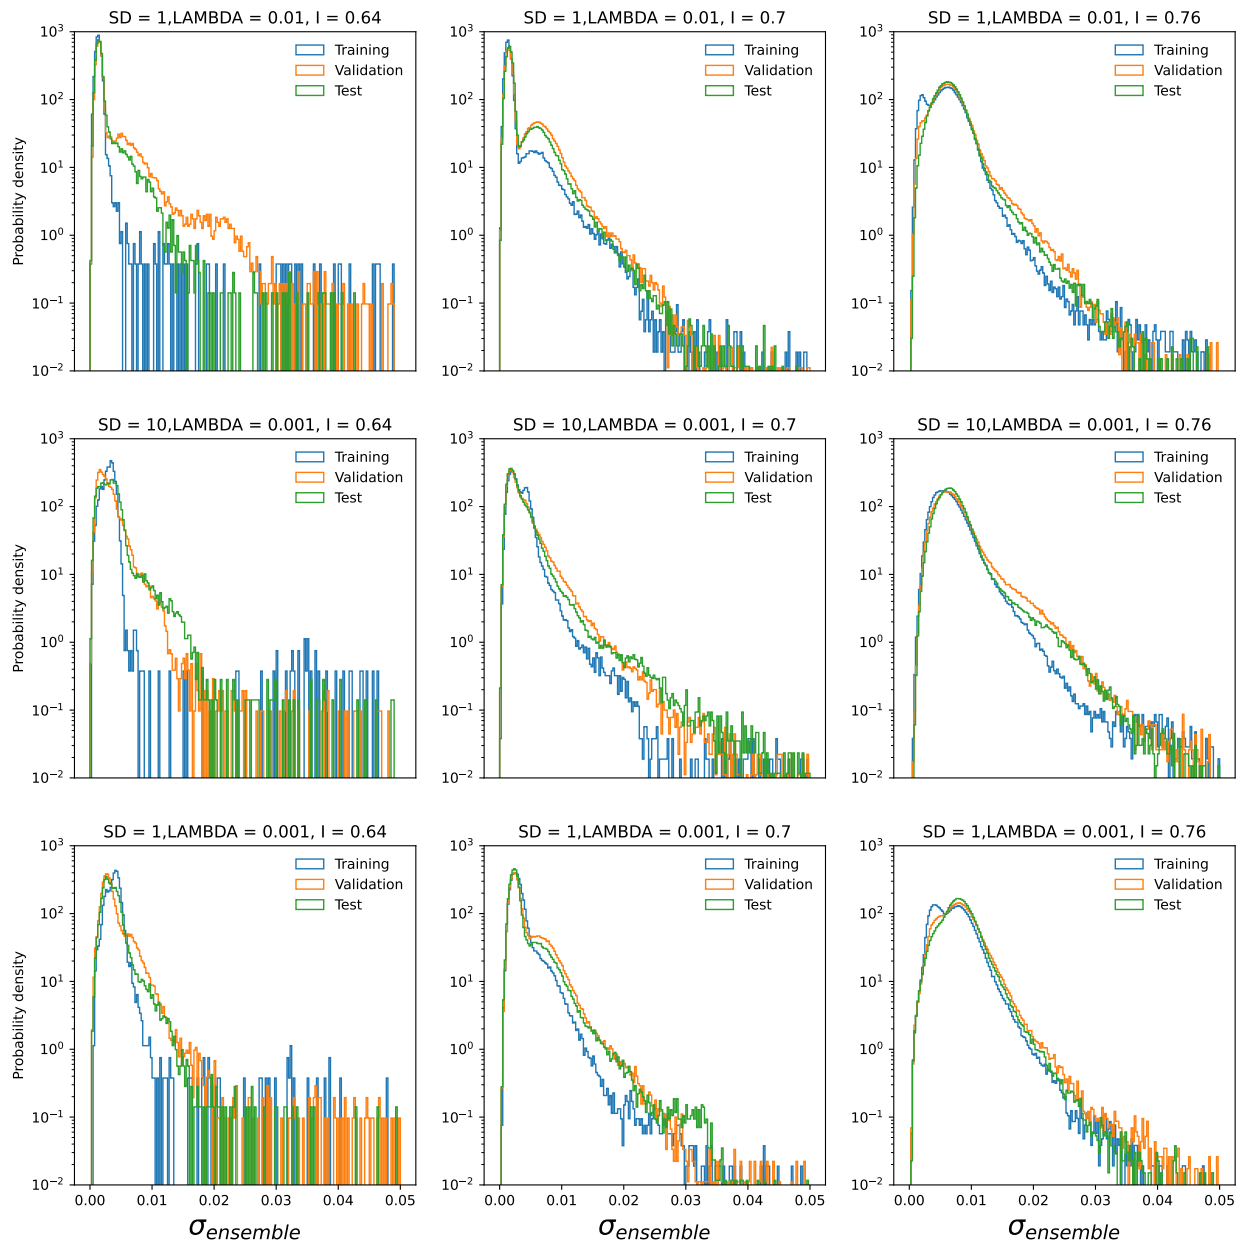


Fig. 7.— Change in ensemble uncertainty ($\sigma_{ensemble}$) with the intensity of input pixels. Each panel shows uncertainty histograms on training, validation, and test sets for a particular input intensity and setting of the ensemble run. Each row represents the effect of a particular standard deviation scaling factor (SD) of the regularization term’s anchor distribution and weightage (LAMBDA). Each column represents the outcome for a particular input intensity (I). The stronger difference in histogram tails can be observed for the lowest input intensity i.e. I = 0.64 with validation and test set outcomes being more extended towards higher uncertainty values.

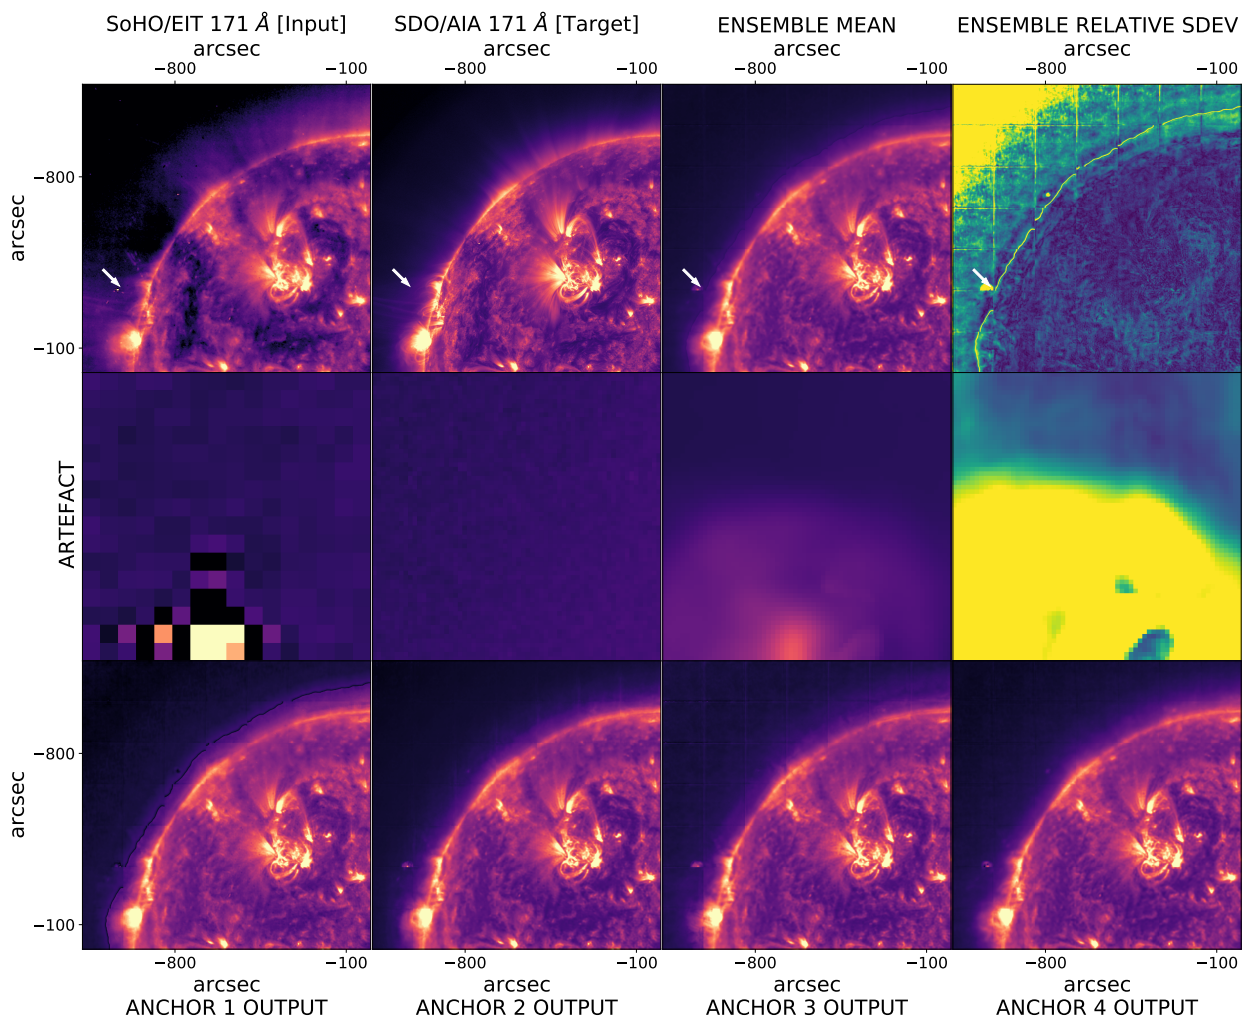


Fig. 8.— Ensemble inference on a single image quadrant of a representative EUV 171 Å image. The top row starting from left depicts the model input, target, ensemble inference mean, and ensemble relative standard deviation (σ/μ). The bottom row shows the outcomes of different ensemble members. It can be observed those bright point-like artifacts (marked by white arrows in the top row and zoomed in the middle row of images) on the input images show high uncertainty and the across ensemble members. Also, the image for relative ensemble standard deviation depicts high uncertainty for the coronal region in the top left corner.

structures of non-solar origin generate high uncertainty (Figure 8).

It is unknown at this moment how many ensemble members are optimal beyond which the uncertainty in reconstruction converges. Further investigations are needed for that.

6. DISCUSSION AND CONCLUSION

We thus trained an ensemble of convolutional neural networks to super-resolve SoHO/EIT 171 Å images to SDO/AIA 171 Å images using the overlapping time period. The salient features of this work are listed below:

1. We prepared ML-ready data consisting of aligned EIT 171 Å 64×64 patches and AIA 171 Å 256×256 patches as input and output images.
2. We trained a CNN based on residual blocks and upsampling layers to translate an EIT patch to an AIA patch
3. We found DL outcome is always superior to a simple upsampling based on intensity scaling and bi-cubic interpolation.
4. We added different terms in the loss function and found having SSIM or Image Gradient as additional terms helps improve the quality of the super-resolved images both visually and quantitatively as compared to MSE baseline.
5. We also estimated the uncertainty of the prediction by training an ensemble of models through Approximate Bayesian Ensembling (ABE).
6. We found that uncertainty improves while making the size of the training set larger.
7. We also found that the model ensemble generated high uncertainty for regions that are not well represented in the training set such as coronal regions, and regions of non-solar origin

such as artifacts.

We would like to highlight that ABE has been demonstrated to behave like a Bayesian Neural Network for toy applications for which there is ground truth. However, there is no Bayesian Neural Network ground truth for homogenization and super-resolution. Our results are encouraging in that ABE uncertainty has desirable and verifiable properties such as sensitivity to input data, robustness to changes in ensembling parameters, and sensible behavior to unexpected features.

As a future work, we would like to investigate the point of convergence of ensemble uncertainty. Also, translation between unpaired image domains ([Zhu et al. 2017](#); [Jarolim 2022](#)) can help make use of the large database beyond overlapping time-period and better constrain high-resolution textural information. We also plan to use our homogenized EUV data for scientific investigations in heliophysics toward independent verification of the data quality (or its scientific readiness) indicated by the metrics.

7. ACKNOWLEDGEMENTS

This work is mainly supported by SwRI internal research grant R6134 and NASA HGIO grant 80NSSC23K0416. The authors also acknowledge partial support by NASA SWO2R grant 80NSSC20K0290.

REFERENCES

- Brunelli, R. 2009, *Template Matching Techniques in Computer Vision: Theory and Practice* (Wiley Publishing)
- Chatterjee, S., Muñoz-Jaramillo, A., & Lamb, D. A. 2022, *Nature Astronomy*, 6, 796
- Delaboudinière, J. P., et al. 1995, *Sol. Phys.*, 162, 291
- Deudon, M., et al. 2020, *HighRes-net: Recursive Fusion for Multi-Frame Super-Resolution of Satellite Imagery*
- Gal, Y., & Ghahramani, Z. 2015, arXiv e-prints, arXiv:1506.02142
- Glorot, X., & Bengio, Y. 2010, in *Proceedings of Machine Learning Research*, Vol. 9, *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, ed. Y. W. Teh & M. Titterington (Chia Laguna Resort, Sardinia, Italy: PMLR), 249–256
- Jarolim, R. 2022, in *Proceedings of the 2nd Machine Learning in Heliophysics*, 39
- Jungbluth, A., et al. 2019, arXiv e-prints, arXiv:1911.01490
- Lemen, J. R., et al. 2012, *Sol. Phys.*, 275, 17
- Pearce, T., Leibfried, F., Brintrup, A., Zaki, M., & Neely, A. 2018, arXiv e-prints, arXiv:1810.05546
- Wang, Z., Li, H., Ouyang, W., & Wang, X. 2018, *CoRR*, abs/1804.09398
- Zhou Wang, Bovik, A. C., Sheikh, H. R., & Simoncelli, E. P. 2004, *IEEE Transactions on Image Processing*, 13, 600
- Zhu, J.-Y., Park, T., Isola, P., & Efros, A. A. 2017, in *2017 IEEE International Conference on Computer Vision (ICCV)*, 2242–2251

