Denoising of InSAR Phase Closure

Elizabeth Wig

Abstract—InSAR phase closure arises when the phases of three InSAR interferograms are combined circularly. Phase closure has been shown to correspond to soil moisture and has potential as a means to measure soil moisture. However, phase closure measurements are very noisy. Current practice uses only linear averaging to reduce noise, but this also reduces the resolution. This paper will explore the possibility of using other denoising techniques to improve the signal-to-noise ratio while maintaining a high resolution. Each method: the linear mean, median, bilateral filter, non-local means, and DnCNN network, will be tested and compared to the ground truth in situ soil moisture in a region of Oklahoma. The results will show whether image processing-based denoising techniques are able to outperform a basic linear average when used on phase closure images to estimate soil moisture.

Index Terms—Computational Imaging, Denoising, Radar, InSAR

1 Introduction

InSAR, or interferometric synthetic aperture radar, compares phases between two radar passes over the same area (usually from an aircraft or satellite) to see how the ground is changing. InSAR is frequently used to capture deformation, whether from earthquakes, volcanoes, fracking, or other sources. The phase differences between the two radar passes are used to form an interferogram, which maps how the area is changing in phase across the radar image. The nature of these interferograms is very noisy. To reduce noise (and reduce large file sizes), the standard practice is to do "multilooking" or spatial averaging. For example, taking 4 "looks" each in range and azimuth means averaging every 4 pixels in each direction, which reduces the file size by 16 times. Since original files can be quite large, this reduction does decrease the resolution but increases the SNR, and many applications do not require the high resolution – and in fact find the computational power required to store and process the high-resolution images to be too taxing.

For some applications, a higher resolution is desired. Denoising methods such as bilateral filters, non-local means, and neural networks, may be able to intelligently denoise radar images while preserving a high resolution and not losing data through averaging.

A more novel area within InSAR research is looking at phase closure. While conventional InSAR differences the phase between two radar acquisitions over the same area, phase closure compares three images in a cycle, creating three interferograms and differencing the phase between the three. Figure 1 shows a schematic of the process of calculating phase closure. As shown in the equations, phase closure appears to have some correspondence to soil moisture (as well as vegetation, and other properties on the ground) [1], [2], [3].

$$\phi_{closure} = \phi_{12} + \phi_{23} + \phi_{31} \neq 0 \tag{1}$$

$$\phi_{closure} \propto \text{soil moisture} \tag{2}$$

Phase closure only arises in images that have been multilooked, or spatially averaged. An image with no multilooking will have zero phase closure. For this reason, some level of basic spatial averaging is necessary to generate phase closure images. However, it is possible that, instead of a high amount of multilooking used to generate the phase closure images (e.g. 30 looks in each direction), a smaller number of looks could be used instead (e.g. 10 looks in each direction) and supplemented with sophisticated image processing techniques to achieve a higher resolution in the final data. This paper will explore that possibility.

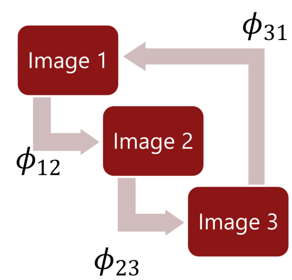

Fig. 1. Phase closure cycle around three InSAR images. The three interferograms have phases ϕ_{12}, ϕ_{23}, and ϕ_{31}.

2 Related Work

InSAR is a well-established field at this point, and some people have done research into, for example, using CNNs, nonlocal means, and other denoising methods to improve the quality of InSAR images [4], [5], [6].

The study of phase closure in InSAR is relatively new, and in particular, the method outlined in the following section was first presented at a conference in December [7]. This method is based on work by Michaelides, Ansari, and De Zan [1], [2], [3] showing that InSAR phase closure appears to change according to soil moisture. However, prior to this presentation, a method for estimating soil moisture from phase closure had not been put forth. Ansari and De Zan showed that there was a persistent phase bias which had some relation to soil moisture, and Michaelides et. al. showed that it was possible to correct this phase. This paper builds on that work.

3 PROPOSED METHOD

This section will first discuss the general method for finding soil moisture values from phase closure, including the two stages where spatial averaging is currently performed. It will then move to discuss the method for this paper, and the potential of replacing spatial averaging, especially in the second stage, with an image processing technique. Instead of a simple linear denoising technique, the methods used will include a median, bilateral, non-local means, and pre-trained DnCNN denoising.

3.1 Data Sources

We tested the potential for integrating phase closure to measure soil moisture at an InSAR swath across the middle of Oklahoma, shown in Fig. 2. The Mesonet system of environmental monitoring systems [8] provided daily *in situ* soil moisture data at eleven sites across this swath at 5 cm depth. To account for variations in soil moisture, and because the phase triplets in InSAR were acquired over periods of at least 24 days (three passes with a 12-day repeat cycle), we temporally averaged the soil moisture measurements across 41 days surrounding the date of the middle image in the phase triplet.

We derived the phase triplets from raw L0 Sentinel-1 data ($\lambda = 5.5cm$), which we processed using 30 x 30 looks, with some additional spatial averaging after processing, for a final pixel size of 900x750 m [9]. This pixel size can be reduced if the number of looks is reduced.

To construct phase triplets, we used only adjacent triplet sets, so the temporal baselines were short - generally 12 days between acquisitions and 24 days across the entire triplet set. Shorter baselines have been generally shown to correspond to higher phase misclosure [2], which fades as temporal baseline increases.

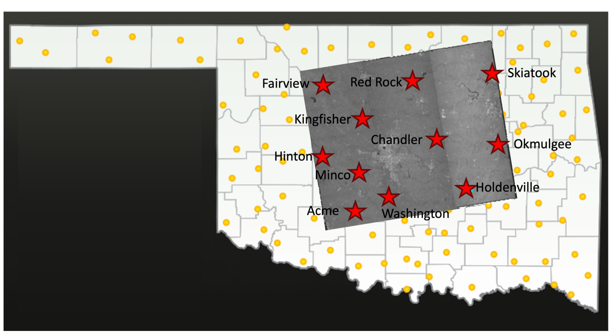

Fig. 2. Location of eleven Mesonet ground sites and InSAR swath in Oklahoma.

3.2 Original Data Reduction Approach

Our data reduction approach is based on the assumption that nonzero phase closure is caused by changing dielectric properties in the area being imaged, such as changes in bulk scattering processes due to changing soil moisture or vegetation. The key assumption is that a change in dielectric properties (such as a change in soil moisture) corresponds to phase closure. This suggests that integrating the phase closure over time could correspond to the soil moisture itself, rather than its derivative or rate of change.

We integrate the phase closure over time using a cumulative sum function, and find that a bias is present at all eleven sites (Fig. 3a). The integrated phase closure at all sites has a roughly monotonic increase over time. However, with the foreknowledge that soil moisture is very unlikely to monotonically increase over the course of years, this monotonic increase likely contains a bias–soil moisture should fluctuate within a fixed range from 0% to 100% saturated. We remove the bias by fitting an appropriately shaped signal to each cumulatively summed phase closure, and then subtracting that signal to force the integrated phase closure to be zero-mean.

We tested several shapes of fitting, including a linear fit, a polynomial fit, and a square root fit (whose shape imitates a random walk). We found that the random walk fit was the best match to the shape of the integrated phase closure. We fit this shape to the integrated phase closure signals at all 11 sites, and then subtract the fitted signal. Figure 3b shows the fitted random walk signal with the integrated phase closure at the Skiatook site.

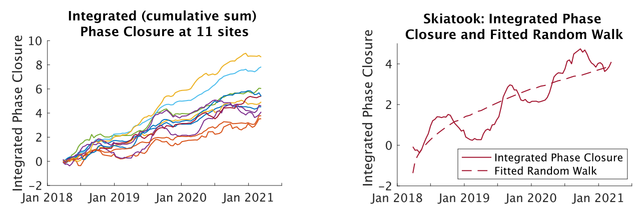

(a) (b)

Fig. 3. (a) At all sites, phase closure appears to near-monotonically increase, suggesting some type of bias in the signal. (b) We found that the best fit to this bias was a "random walk" or square-root shaped signal.

Figure 4 shows the bias-corrected integrated phase closure plotted next to the soil moisture for the Skiatook site. The two quantities appear to be anticorrelated. We can thus investigate the possibility of using this parameter, which we derived from phase closure, as a means of tracking soil moisture. This bias-corrected integrated phase closure appears to track *in situ* data much more closely than the original, non-integrated phase closure.

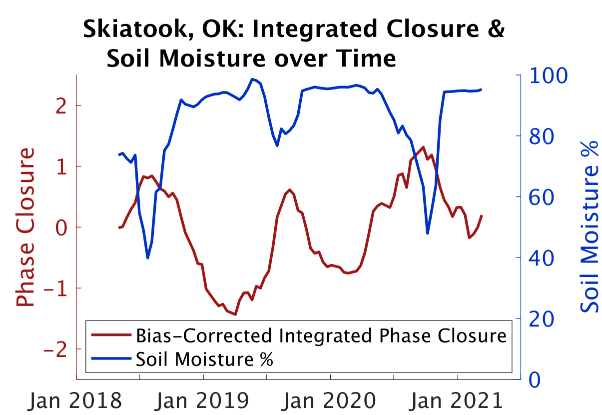

Fig. 4. Bias-corrected integrated phase closure appears to be strongly anticorrelated with soil moisture.

We create a scatterplot of the two parameters graphed in Fig. 4 against one another. Figure 5 plots the soil moisture against the bias-corrected integrated phase closure at the Skiatook site.

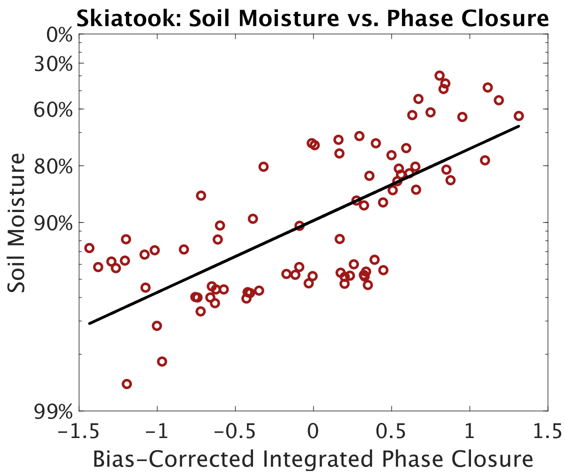

Fig. 5. The parameter of bias-corrected integrated phase closure appears to have a strong correlation with soil moisture. The correlation coefficient $|R| = 0.71$.

We can then create a fit line between these two quantities and use the fit line to make a prediction of soil moisture from the phase closure, shown in Fig. 6. This prediction appears to track quite closely, rising and falling in the same areas. While the amplitude does not track exactly, particularly at times with very low soil moisture, this parameter does appear to track when soil moisture is rising and falling accurately.

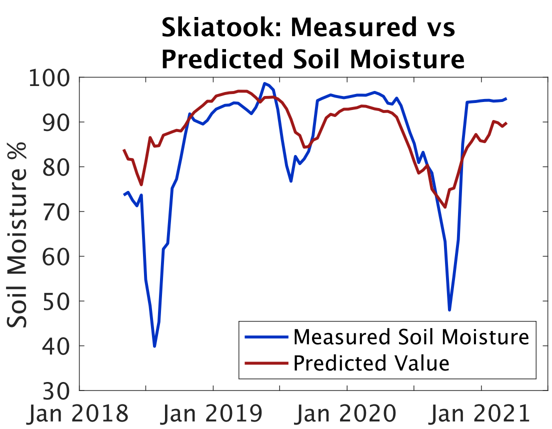

Fig. 6. The prediction of soil moisture based on phase closure at the Skiatook site appears to track quite closely, though it does not always capture areas of very low soil moisture.

3.3 Denoising Methods

As previously mentioned, there are two denoising stages, and the current practice uses linear averaging for both. The first, the multilooking stage, is currently done with a speed-optimized Fortran code and is required to turn the very noisy initial radar measurements into something resembling an image. (For example, compression using traditional image compression methods on this initial image barely make a dent in its size). The second stage has more flexibility and a starting format that is more image-like, so this was chosen as the stage with greater potential for optimization of the denoising method. Pretrained methods trained on traditional images would have a greater chance of working on this stage-two image. Figure 7 shows the steps taken to find an estimate for soil moisture from InSAR phase closure. The denoising step is the part of the flow that is being optimized in this work.

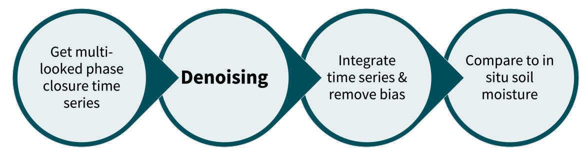

Fig. 7. High-level overview of processing steps to get soil moisture estimates from InSAR phase closure.

Five denoising methods were tested on the original images. The original processing flow used linear averaging in stage two on an image that had been multilooked 30 times in each direction (meaning the resolution before the second processing stage was 900x worse than the original radar image). The denoising methods were used on each radar image in the time series to see whether they could improve the resolution and accuracy over linear averaging.

Because a higher resolution was desired, the denoising methods were tested after stage-one multilooking with 10 looks in each direction (meaning the resolution of the initial image in stage two was 9 times higher than in previous work).

The types of denoising included a linear mean, used as the control and past standard for this experiment; a median, used with the same kernel size as the mean; a MATLAB built-in bilateral filter made using imbilatfilt and optimized over degree of smoothing; a MATLAB built-in non-local means filter using imnlmfilt and also optimized for degree of smoothing; and a pretrained built-in DnCNN neural net in MATLAB. The DnCNN had no controllable parameters, but a parameter sweep was run over the degree of smoothing for the bilateral and non-local means functions. The results of the parameter sweep are shown in Fig. 8.

A relatively high degree of smoothing was found to be optimal for both the bilateral and non-local means filters. The parameter sweep found that a degree of smoothing of 100000 for the bilateral filter and of 3 for the non-local means filter was ideal. Once the optimal filters were found, each image in the time series of phase closure images was filtered using the appropriate spatial filters. The filtered images were then processed using the integration and bias removal method outlined in the previous section. Then the bias-corrected integrated phase closures were compared to the soil moisture. This involved both calculating the correlation coefficient between the corrected phase closure for each filtering method, as well as running a regression of the soil moisture against the phase closure to generate an estimate

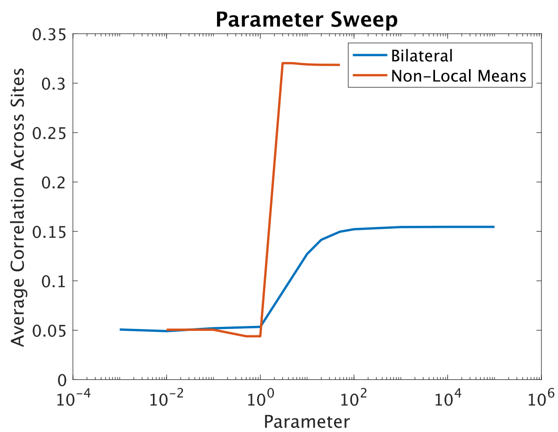

Fig. 8. Parameter sweep by degree of smoothing for the Bilateral and Non-Local Means built-in MATLAB filters. The parameters were chosen to maximize the correlation between the filtered image series and the soil moisture time series.

of what it might look like if only InSAR was used to estimate the soil moisture using each method.

4 Experimental Results and Discussion

Figure 9 shows the estimated soil moisture using each method using a linear regression against the ground truth from the in situ stations - at a sample site in Skiatook. From this plot of estimated (and true) soil moisture over time, it can be seen that some denoising methods result in better estimates than others. None fully capture the troughs of very low soil moisture in the summers of 2018 and 2020, but the nonlocal means method does a very good job capturing the summer 2019 dip in soil moisture. The bilateral-filtered time series does worse and does not even seem to vary along with the soil moisture. The DnCNN-derived estimate also performs relatively poorly, while the mean and median filters match the ground truth moderately well.

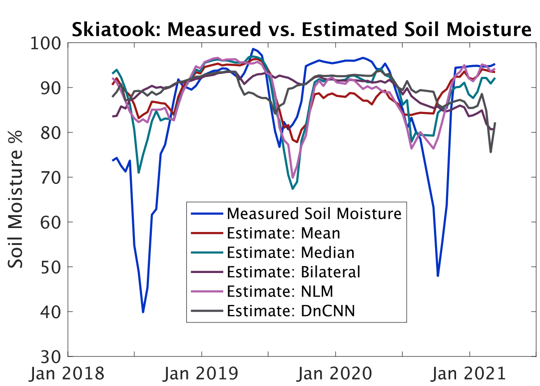

Fig. 9. Prediction of soil moisture at Skiatook site using image time series from each denoising method. The ground-truth measured soil moisture is in bright blue. In aggregate, none of the denoising methods perfectly captures the soil moisture, but some do better than others.

Quantitatively, the bilateral filter also has worse performance. Table 1 shows the correlation between the InSAR-derived parameter and the soil moisture, averaged across the 11 sites, and its standard deviation. This was calculated for each method.

TABLE 1
Mean and Standard Deviation of Correlation (between InSAR phase closure and soil moisture) for each Denoising Method

Method	Average Correlation	Standard Deviation
Linear Mean	0.2914	0.1569
Median	0.3109	0.1483
Bilateral	0.1546	0.1839
Non-Local Means	0.3203	0.1763
DnCNN	0.0693	0.3094

It is also of interest to see how each filter affects an individual InSAR phase closure image, to better understand how the filtering may change the output. Figures 10 through 14 show the first image in each time series for a qualitative analysis of how the filters affect the radar image. Overlaid with the image is a dot representing the correlation between the soil moisture and the phase closure-derived parameter at each site. Redder dots indicate higher correlation and a better match to soil moisture. There is a significant amount of variation among sites, but some methods have overall higher correlation across sites than others. The biggest difference in the images is the blur level. The mean, median, and non-local means filtered images look significantly blurrier than the bilateral and DnCNN images - and they have higher performance. It seems that preservation of edges or the features that the DnCNN network was trained on does not improve performance for the phase closure images. The median and non-local means both perform slightly better than the mean filter. This may be because the median is more robust to outliers than a simple mean, while the non-local means can pull data from disparate areas with similar shapes or statistical properties, such as agricultural fields.

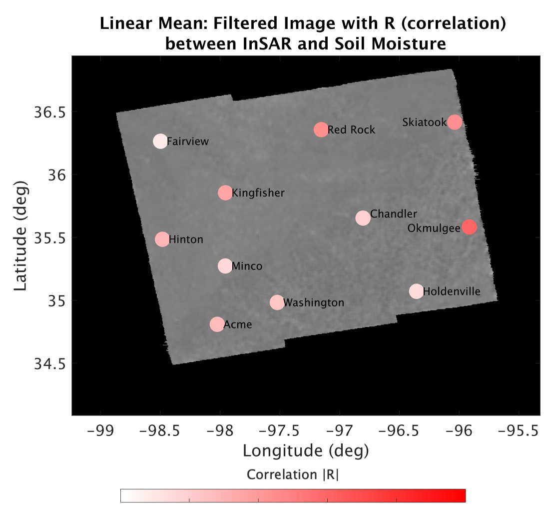

Fig. 10. Sample phase closure image filtered with linear mean of 11x11 pixels (chosen based on desired resolution). This image has a fairly high level of blur.

In Figure 15 section, we can see quantitatively that the average correlation is best for the mean, median, and non-local-means denoising methods. The bilateral and DnCNN have worse performance overall, although all have a high

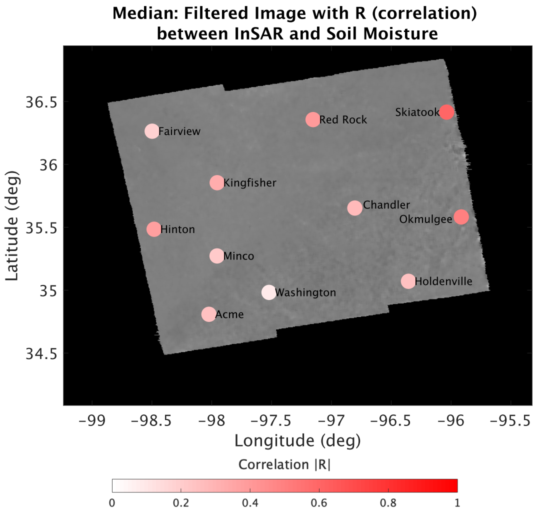

Fig. 11. Image filtered with a median fitler over 11x11 pixels (chosen based on desired resolution). This image, like the linear one, has high levels of blur. However, it looks slightly sharper, probably because the median is better about preserving edges rather than smoothing over them.

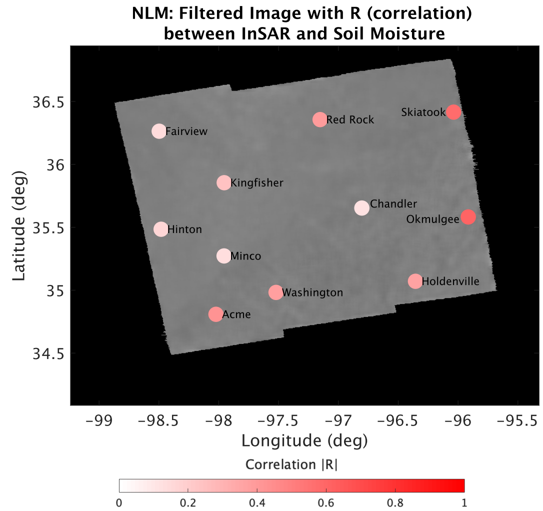

Fig. 13. Sample image filtered with a non-local means filter with an optimized degree-of-smoothing parameter of 3. This image has a higher level of blur, more similar to the linear mean and median filters. There is a wide spread of goodness-of-fit (redder dots = better fit), but in general the fit is relatively good with this method.

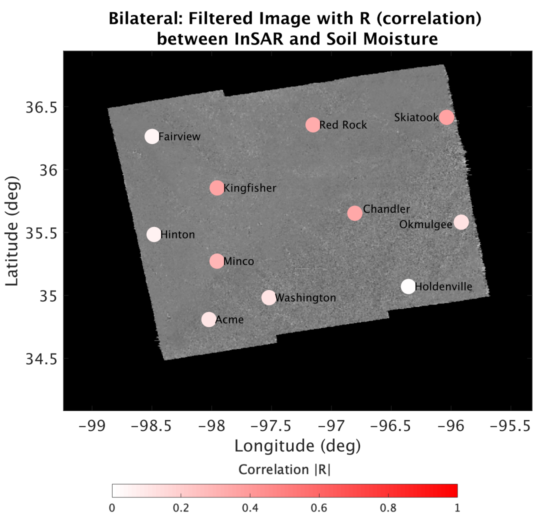

Fig. 12. Sample image filtered with a bilateral filter with the optimized smoothing parameter of 100000. Even with the "high" level of smoothing for this pre-built filter, it looks noticeably sharper and has more fine grains than the mean or median filtered images. The lighter colored dots at the locations of the soil moisture measurement sites indicate that the correlation between this phase closure and soil moisture is worse.

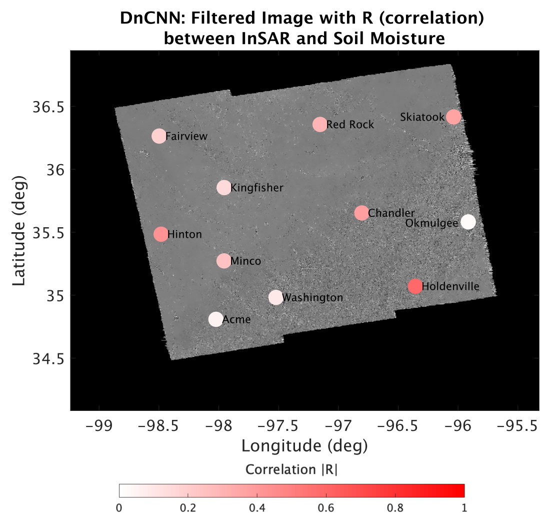

Fig. 14. Sample image filtered with MATLAB's pre-trained DnCNN denoising neural network. This image has the finest grains and highest resolution of the set, and also the lowest correlations with the in situ soil moisture data.

standard deviation. Moreover, even the best-performing sophisticated denoising method does not significantly outperform the basic mean and median filters.

One constraint to weigh against better performance is processing time. In this area, the mean and median filters have the clear advantage (each taking less than a few seconds), while the bilateral and non-local means filters each take a few minutes on a laptop CPU, and applying the DnCNN takes over half an hour. If processing time is at all a constraint, mean or median have the clear advantage, especially given the marginal advantage that the non-local means has over either.

This project showed the limitations of image processing networks from the EE367 class. While many methods work well for traditional images, they falter for radar images, where the statistics are different and there is less ability to smooth or compress easily. The method that was the most highly trained to conventional images, the pre-trained DnCNN neural net, performed the worse on the phase closure images. The bilateral filter, which retained sharp lines for the sample images from class, did not seem to retain the right details to give good quality soil moisture predictions. Of the more sophisticated image processing techniques, the non-local means performed the best. Intuitively, the strength of the non-local means method–leveraging areas with similar statistics from different parts of the image–seems to be better adapted to radar phase closure images. The algorithm may be finding and combining data from different regions with similar statistical properties, such as wheat fields.

A clear need if this work were to continue would be images trained on radar, or specifically on phase closure. There is not much phase closure data processed or available for use currently, so without a strong training set, the use of these methods is perhaps premature. Where out-of-the-box image processing methods falter, more tailored methods can step in, but the area of study of InSAR phase closure is relatively new and small, so no methods have been developed. This leaves lots of room for future work in the area as the properties of InSAR phase closure become more well-understood.

Figure 15 shows a final comparison of the methods' accuracy to the soil moisture ground truth, as well as the standard deviation. For all methods, the standard deviation from one site to another was high–in part due to a low number of sites used, but also because the sites did vary in quality.

This variation in accuracy of InSAR measurements from one site to another is an active area of interest in my current research. Characterization by land cover type and vegetation level (using measurements such as the Normalized Difference Vegetation Index) is ongoing work that will hopefully yield more accurate estimates. For example, soil moisture may be better estimated in areas with less vegetation cover or only certain types of vegetation, or there may be some compensation for a thicker tree canopy. More advanced image processing techniques may be able to categorize estimates based on land cover type and other properties. Future study will explore these possibilities more.

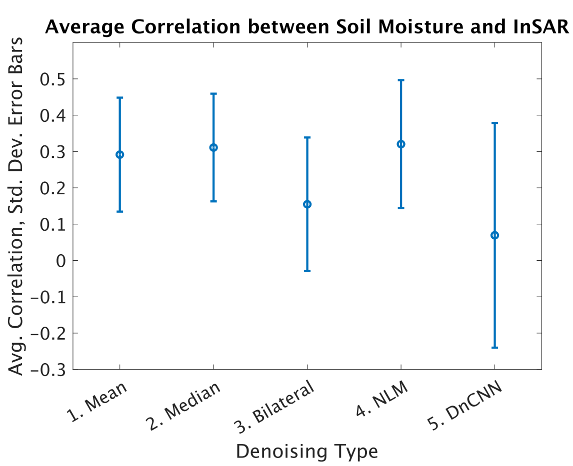

Fig. 15. Comparison of methods (a graphical presentation of the data in Table 1). The mean accuracy, again represented by correlation between the InSAR time series and the ground-truth soil moisture, is shown with a dot, and the error bars represent the standard deviation in each direction.

5 Conclusion

In conclusion, it's clear that the conventional image filters are not optimized for radar. Data trained specifically on radar or on phase closure would be necessary for better-correlated results. Non-local means performed best: it may have been able to synthesize areas with similar traits, such as wheat fields. However, even the best sophisticated denoising method did not significantly outperform the basic mean and median filters. Especially considering processing time, the mean and median may be the best filters for now. This demonstrates the limitations of the computational imaging methods from the EE367 class, and shows that optimization for optical images does not necessarily translate to radar.

Future work will involve more closely associating land-cover type and vegetation to characterize soil moisture matching in different areas. While these results show decent correlation, ultimately more averaging trumps all other methods here. In other work with 30x30 looks instead of 10x10 (as presented at [7]), the results are better-correlated: closer to 0.4 or 0.5 rather than the 0.3 that is the best here. None of these methods is a magic bullet for this problem, but it was interesting to explore!

Acknowledgments

The author would like to thank and credit Howard Zebker and Roger Michaelides for their collaboration on the original work that this denoising class project is based on.

References

[1] F. D. Zan, A. Parizzi, P. Prats-Iraola, and P. López-Dekker, "A sar interferometric model for soil moisture," *IEEE Trans. Geosci. Remote Sens.*, vol. 52, no. 1, pp. 418–425, January 2014.

[2] H. Ansari, F. D. Zan, and A. Parizzi, "Study of systematic bias in measuring surface deformation with sar interferometry," *IEEE Trans. Geosci. Remote Sens.*, vol. 59, no. 2, pp. 1285–1301, February 2021.

[3] R. Michaelides, H. Zebker, and Y. Zheng, "An algorithm for estimating and correcting decorrelation phase from insar data using closure phase triplets," *IEEE Trans. Geosci. Remote Sens.*, vol. 57, no. 12, pp. 10 390–10 397, 2019.

[4] S. Mukherjee, A. Zimmer, N. K. Kottayil, X. Sun, P. Ghuman, and I. Cheng, "Cnn-based insar denoising and coherence metric," in *2018 IEEE SENSORS*, 2018, pp. 1–4.

[5] W. Rui, Y. Yanan, and Z. wenli, "Interferometric phase stack denoiseing via nonlocal higher order robust pca method," in *IGARSS 2019 - 2019 IEEE International Geoscience and Remote Sensing Symposium*, 2019, pp. 1685–1688.

[6] G. Baier, C. Rossi, M. Lachaise, X. X. Zhu, and R. Bamler, "A nonlocal insar filter for high-resolution dem generation from tandem-x interferograms," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 56, no. 11, pp. 6469–6483, 2018.

[7] E. Wig, R. Michaelides, and H. Zebker, "Fine-resolution measurement of soil moisture from insar phase closure," in *American Geophysical Union Conference*, December 2021.

[8] F. B. et. al., "The oklahoma mesonet: A technical overview," *J. Atmos. Oceanic Technol.*, vol. 12, 1995.

[9] "Copernicus sentinel data 2018-2021," retrieved from ASF DAAC, some processing by ESA.

Elizabeth Wig Elizabeth Wig is a PhD candidate at Stanford University in Howard Zebker's group, where she researches InSAR. She is particularly interested in InSAR phase closure, permafrost, and scattering from soil moisture and vegetation.