
VConstruct: Filling Gaps in Chl-a Data Using a Variational Autoencoder

Matthew Ehrler

Department of Computer Science
University of Victoria
Victoria B.C
mehrler@uvic.ca

Neil Ernst

Department of Computer Science
University of Victoria
Victoria B.C
nernst@uvic.ca

Abstract

Remote sensing of Chlorophyll-a is vital in monitoring climate change. Chlorophyll-a measurements give us an idea of the algae concentrations in the ocean, which lets us monitor ocean health. However, a common problem is that the satellites used to gather the data are commonly obstructed by clouds and other artifacts. This means that time series data from satellites can suffer from spatial data loss. There are a number of algorithms that are able to reconstruct the missing parts of these images to varying degrees of accuracy, with Data INterpolating Empirical Orthogonal Functions (DINEOF) being the current standard. However, DINEOF is slow, suffers from accuracy loss in temporally homogenous waters, reliant on temporal data, and only able to generate a single potential reconstruction. We propose a machine learning approach to reconstruction of Chlorophyll-a data using a Variational Autoencoder (VAE). Our accuracy results to date are competitive with but slightly less accurate than DINEOF. We show the benefits of our method including vastly decreased computation time and ability to generate multiple potential reconstructions. Lastly, we outline our planned improvements and future work.

1 Introduction

Phytoplankton and ocean colour are considered “Essential Climate Variables” for measuring and predicting climate systems and ocean health [3]. Measuring phytoplankton and ocean colour is cost effective on a global scale as well as relevant to climate models. Chlorophyll-a (Chl-a) is a commonly used metric to estimate phytoplankton levels (measured in units of mg/m^3) and can be derived from ocean colour [1]. Additionally, Chl-a can also be used to detect harmful algae blooms which can be fatal to marine life [11]. As climate change progresses harmful algae blooms will increase in frequency. An increase of $2^{\circ}C$ in sea temperature will double the window of opportunity for Harmful Algae Blooms in the Puget Sound [9]. Several different satellites provide Chl-a measurements but the Sentinel-3 mission¹ will be the focus of this paper.

One of the biggest problems faced when using these measurements is the loss of spatial data due to clouds, sunglint or various other factors which can affect the atmospheric correction process [11]. Various algorithms exist to reconstruct the missing data with the most effective being those based off of extracting Empirical Orthographic Functions (EOF) from the data [12]. The most accurate and commonly used of these algorithms is Data INterpolating Empirical Orthogonal Functions (DINEOF) which iteratively calculate EOFs based on the input data [5, 12]. DINEOF is fairly slow [12] and performs poorly in more temporally homogenous waters such as a river mouth [5].

¹<https://sentinel.esa.int/web/sentinel/missions/sentinel-3>

Machine learning has also been successful in reconstructing Chl-a data. Park et al. use a tree based model to reconstruct algae in polar regions [10]. This method is effective but requires knowledge of the domain to properly tune it. This makes it much less generalizable and therefore less effective than DINEOF as DINEOF works with no a priori knowledge. DINCAE is a very new machine learning approach to reconstructing data [2], which has also been shown to work on Chl-a [4]. DINCAE is accurate, but shares a drawback with DINEOF in that it can only generate a single possible reconstruction. Being able to see multiple potential reconstructions and potentially select a better one based on data that may not be able to be easily or quickly incorporated into the model. For example if we had Chl-a concentrations manually measured from missing areas, we could then generate reconstructions until we find one that better matches the measured values. This would be much faster than changing DINCAE or DINEOF to incorporate the new data.

The approach we outline in this paper is based on the Variational Autoencoder (VAE) from Kingma et al. [8] as well as Attribute2Image's improvements in making generated images less random [13]. The dimensionality reduction in a VAE is somewhat similar to the Singular Value Decomposition (SVD) used in DINEOF. The potential to leverage performance improvements using quantum annealing machines with VAEs was another motivation. [7].

In this paper we apply a model similar to Attribute2Image as well as Ivanov et al's inpainting model to Chl-a data from the Salish Sea area surrounding Vancouver Island [6, 13]. We compare it to the industry standard DINEOF using experiments modeled after Hilborn et al.'s experiments [5]. This area was chosen as it contains both areas of high and low temporal homogeneity in terms of algae concentrations, which was determined by Hilborn et al to be something DINEOF is sensitive to [5].

2 Method

2.1 Dataset and Preprocessing

The dataset we use comes from the Algae Explorer Project². This project used 1566 images taken daily from 2016-04-25 to 2020-09-30. For our experiments we use a 250x250 pixel slice from each day to create a 1566x250x250 dataset.

We then preprocess the data for DINEOF using a similar process to Hilborn et al. [5]. The data for VConstruct uses a similar process to Han et al. [4].

We select five days for testing at random from all days that have very low cloud cover. This allows us to add artificial clouds and measure accuracy by comparing to the original complete image.

2.2 DINEOF Testing

As we are using different satellite data than Hilborn et al., we cannot compare directly to their results and need to devise a similar experiment [5]. Since DINEOF is not "trained" like ML models, we cannot do conventional testing with a withheld dataset. For our experiment we use the five testing images selected in preprocessing, and then overlay artificial clouds on these images to create our testing set, which is then inserted back into the full set of images. Samples are shown in the Appendix, Figs. 2 and 3. After running DINEOF we then compare these reconstructions with the known full image and report accuracy.

This scheme slightly biases the experiment towards DINEOF as DINEOF has access to the cloudy testing images when generating EOFs where VConstruct does not. This is unfortunately unavoidable but the effect seems minimal.

2.3 VConstruct Model

The VConstruct model is based on the Variational Autoencoder [8], it consists of an encoder, decoder and attribute network. All network layers are fully connected layers with ReLU activation functions.

The encoder and decoder layers function exactly like they do in a conventional VAE. The encoder network compresses an image down to a lower dimensional latent space and learns a distribution it can later sample from during testing when the complete image is unknown. The decoder takes

²<https://algaeexplorer.ca/>

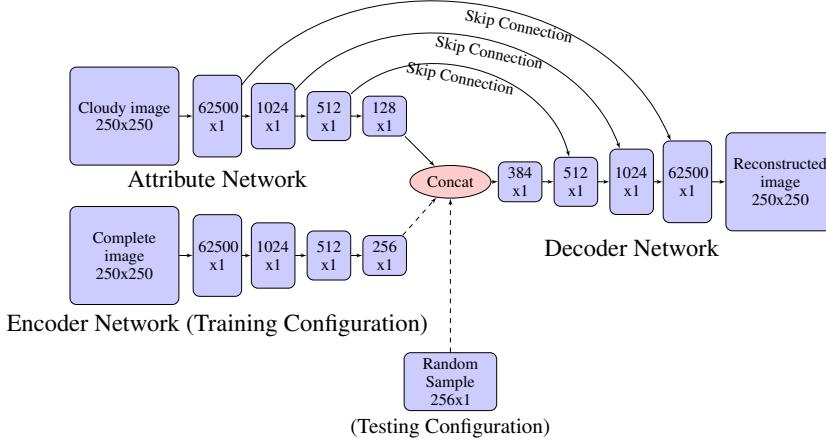


Figure 1: Training Configuration for VConstruct

the output of the encoder network, or random sample from the learnt distribution, and attempts to reconstruct the original image. We use Kullback–Leibler divergence and Reconstruction loss for our loss function.

The attribute network is based off the work of Yan et al. and Ivanov et al. [6, 13]. The network extracts an attribute vector from a cloudy image which represents what is “known” about the cloudy image. This attribute vector then influences the previously random image generation of the decoder network so that it generates a potential reconstruction.

These three networks make up the training configuration of VConstruct and can be seen in Fig. 1. When testing we cannot use the Encoder network as we do not know the complete image, so the network is replaced with a random sample from the distribution learnt in training. The parts that switch out are indicated by the dashed lines.

2.4 VConstruct Testing

We train VConstruct by using all of the complete images marked in preprocessing (minus the five testing images which are withheld) with artificial clouds overlaid. The model is trained for 150 epochs. After training we use the five testing images, randomly selected in preprocessing, with the same artificial cloud mask as DINEOF and calculate the same metrics.

3 Results and Discussion

Table 1 presents the results of reconstructing the five randomly selected testing days. We show results for an area off the coast of Victoria and an area by the mouth of the Fraser River. RMSE (Root Mean Squared Error) and R^2 (Correlation Coefficient) are reported. The Fraser River mouth is an area of high temporal homogeneity, which is identified by Hilborn et al. as a problem area for DINEOF [5]. The actual reconstructed images can be found in the appendix.

For the Victoria Coast VConstruct matches DINEOF’s RMSE and R^2 in 3/5 days but DINEOF has a better average score. For the Fraser River Mouth we see VConstruct outperforms DINEOF on 2/5 tests and nearly matches its average score, particularly in R^2 .

3.1 Other Benefits of VConstruct

VConstruct also provides a few benefits unrelated to accuracy, the first being computation time. VConstruct is parallelized and runs on a GPU. Once trained VConstruct is able to reconstruct in roughly 10 milliseconds as opposed to the 10 minutes it took for DINEOF on the testing computer.

Table 1: Testing Results. Last row reflects overall mean performance.

RMSE				R^2			
Victoria Coast		Fraser River Mouth		Victoria Coast		Fraser River Mouth	
DINEOF	VConstruct	DINEOF	VConstruct	DINEOF	VConstruct	DINEOF	VConstruct
.104	.125	.183	.152	.247	-.089	.759	.834
.093	.096	.209	.234	.667	.646	.788	.736
.078	.08	.131	.119	.569	.552	.797	.833
.071	.086	.154	.193	.736	.614	.789	.688
.067	.068	.164	.176	.499	.472	.898	.883
.0826	.091	.1684	.1748	.544	.439	.806	.791

This decrease in computation time allows researchers to reconstruct much larger datasets, which was an important concern raised by the oceanographer we consulted for this project.

VConstruct also has a few advantages that apply to DINCAE (the recent Chl-a approach from [2]) as well as DINEOF. Currently VConstruct is fully atemporal, meaning that we do not need data from a previous time period to perform reconstructions. This is significant as it allows us to reconstruct data even if nothing is known about previous time periods.

Since VConstruct is based off of a VAE we can resample the random distribution to provide different possible images. From an oceanographic perspective, this allows us to generate new possible reconstructions. This is useful when subsequently collected field-truthed data was from a missing area that invalidated the initial reconstruction. For example, the dataset we are using is field-truthed using HPLC derived Chl-a measurements from provincial ferries. Since reconstruction only takes a few milliseconds we could generate and test 1000s of possible images in the same time it takes for DINEOF to run.

3.2 Future Work

We evaluated the approach using two specific test areas. Expanding the training set by using data from other areas in the Salish Sea is important, because different oceanographic areas have different factors affecting Chl-a concentrations. The Salish Sea describes waters including Puget Sound, Strait of Georgia, and the Strait of Juan de Fuca in the US Pacific Northwest/Western Canada. We plan on making the accuracy testing more rigorous in the next iteration. We also plan on testing the effects of adding temporality to the input data. We initially chose to pursue atemporality as data is very commonly missing. However, temporal data is likely to improve accuracy when available. Lastly, VConstruct uses fully connected layers for simplicity but DINCAE has shown success using convolutional layers so this will be tested in the future.

4 Conclusion

We have shown that VConstruct and machine learning in general can be used to reconstruct remotely sensed measurements of Chl-a, which is important in oceanographic climate change research. Even though VConstruct does not match or beat DINEOF in every accuracy test, we feel we have shown its potential for highly accurate reconstructions, particularly in areas of high homogeneity where DINEOF performs poorly. We also show VConstruct's other potential benefits, including better computation time as well as its ability to generate a high number of different potential reconstructions. Remote sensing is an important part of monitoring the climate and climate change, but is limited by cloud cover and other factors which result in data loss. These factors make data reconstruction an important part of climate change research.

5 Acknowledgements

Special thanks to Yvonne Coady, Maycira Costa, Derek Jacoby, and Christian Marchese for their input and feedback.

References

- [1] S. Alvain, C. Moulin, Y. Dandonneau, and F. M. Bréon, “Remote sensing of phytoplankton groups in case 1 waters from global SeaWiFS imagery,” *Deep Sea Research Part I: Oceanographic Research Papers*, vol. 52, no. 11, pp. 1989–2004, Nov. 2005. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0967063705001536>
- [2] A. Barth, A. Alvera-Azcárate, M. Licer, and J.-M. Beckers, “DINCAE 1.0: a convolutional neural network with error estimates to reconstruct sea surface temperature satellite observations,” *Geoscientific Model Development*, vol. 13, no. 3, pp. 1609–1622, Mar. 2020. [Online]. Available: <https://gmd.copernicus.org/articles/13/1609/2020/>
- [3] S. Bojinski and M. M. Verstraete, “(PDF) The Concept of Essential Climate Variables in Support of Climate Research, Applications, and Policy,” *ResearchGate*, 2014. [Online]. Available: https://www.researchgate.net/publication/271271716_The_Concept_of_Essential_Climate_Variables_in_Support_of_Climate_Research_Applications_and_Policy
- [4] Z. Han, Y. He, G. Liu, and W. Perrie, “Application of DINCAE to Reconstruct the Gaps in Chlorophyll-a Satellite Observations in the South China Sea and West Philippine Sea,” *Remote Sensing*, vol. 12, no. 3, p. 480, Jan. 2020. [Online]. Available: <https://www.mdpi.com/2072-4292/12/3/480>
- [5] A. Hilborn and M. Costa, “Applications of DINEOF to Satellite-Derived Chlorophyll-a from a Productive Coastal Region,” *Remote Sensing*, vol. 10, no. 9, p. 1449, Sep. 2018. [Online]. Available: <https://www.mdpi.com/2072-4292/10/9/1449>
- [6] O. Ivanov, M. Figurnov, and D. Vetrov, “Variational autoencoder with arbitrary conditioning.”
- [7] A. Khoshaman, W. Vinci, B. Denis, E. Andriyash, H. Sadeghi, and M. H. Amin, “Quantum variational autoencoder.”
- [8] D. P. Kingma and M. Welling, “Auto-Encoding Variational Bayes,” *arXiv:1312.6114 [cs, stat]*, May 2014, arXiv: 1312.6114. [Online]. Available: <http://arxiv.org/abs/1312.6114>
- [9] S. K. Moore, V. L. Trainer, N. J. Mantua, M. S. Parker, E. A. Laws, L. C. Backer, and L. E. Fleming, “Impacts of climate variability and future climate change on harmful algal blooms and human health,” in *Environmental Health*, vol. 7, no. S2. Springer, 2008, p. S4.
- [10] J. Park, J.-H. Kim, H.-c. Kim, B.-K. Kim, D. Bae, Y.-H. Jo, N. Jo, and S. H. Lee, “Reconstruction of Ocean Color Data Using Machine Learning Techniques in Polar Regions: Focusing on Off Cape Hallett, Ross Sea,” *Remote Sensing*, vol. 11, no. 11, p. 1366, Jan. 2019. [Online]. Available: <https://www.mdpi.com/2072-4292/11/11/1366>
- [11] D. Sirjacobs, A. Alvera-Azcárate, A. Barth, G. Lacroix, Y. Park, B. Nechad, K. Ruddick, and J.-M. Beckers, “Cloud filling of ocean colour and sea surface temperature remote sensing products over the Southern North Sea by the Data Interpolating Empirical Orthogonal Functions methodology,” *Journal of Sea Research*, vol. 65, no. 1, pp. 114–130, Jan. 2011. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1385110110001036>
- [12] M. H. Taylor, M. Losch, M. Wenzel, and J. Schröter, “On the Sensitivity of Field Reconstruction and Prediction Using Empirical Orthogonal Functions Derived from Gappy Data,” *Journal of Climate*, vol. 26, no. 22, pp. 9194–9205, Nov. 2013. [Online]. Available: <https://journals.ametsoc.org/jcli/article/26/22/9194/34073/On-the-Sensitivity-of-Field-Reconstruction-and>
- [13] X. Yan, J. Yang, K. Sohn, and H. Lee, “Attribute2Image: Conditional Image Generation from Visual Attributes,” *arXiv:1512.00570 [cs]*, Oct. 2016, arXiv: 1512.00570. [Online]. Available: <http://arxiv.org/abs/1512.00570>

A Reconstruction Test Images

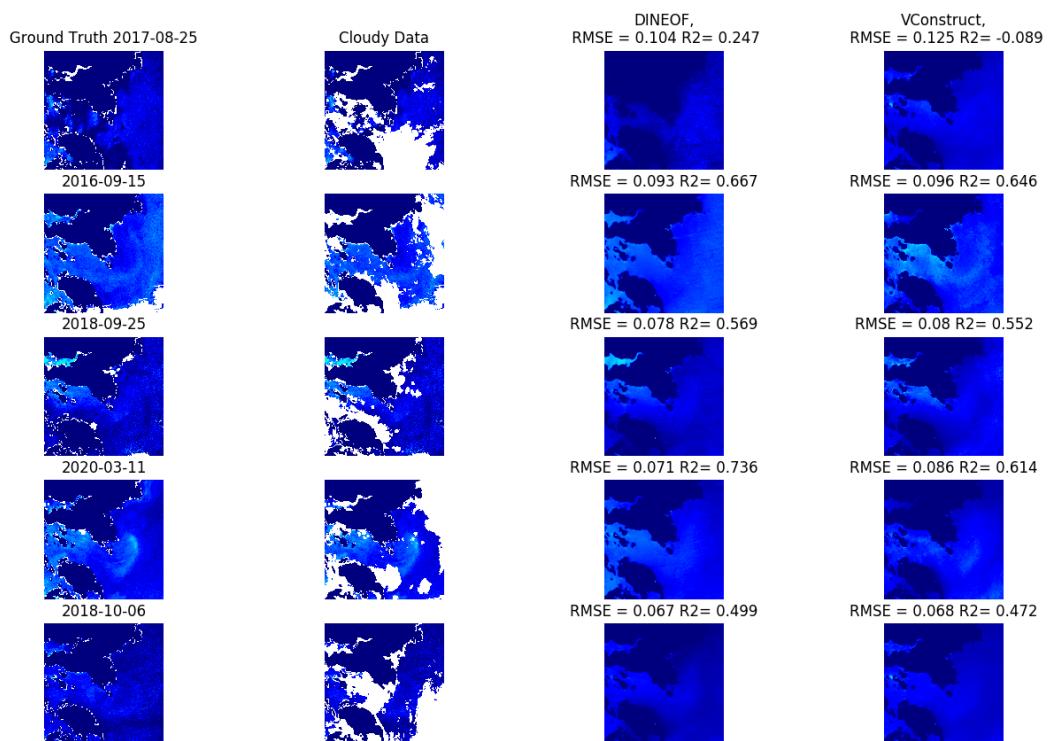


Figure 2: Test Results Victoria Coast

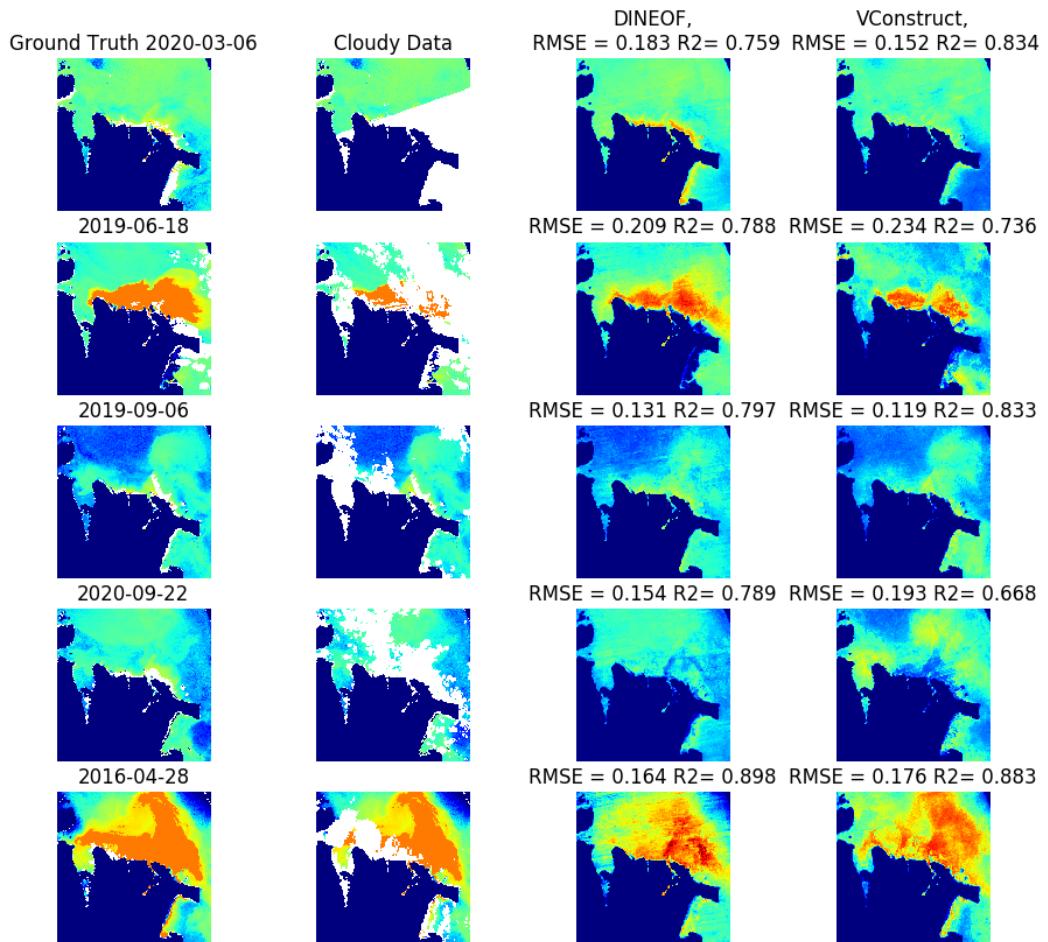


Figure 3: Test Results Fraser River Mouth