SYNTHETIC APERTURE RADAR SHIP DISCRIMINATION, GENERATION AND LATENT VARIABLE EXTRACTION USING INFORMATION MAXIMIZING GENERATIVE ADVERSARIAL NETWORKS

[†‡]*C.P. Schwegmann,*[†‡]*W. Kleynhans,*[⋆]*B.P. Salmon,* [†‡]*L.W. Mdakane,* [†‡]*R.G.V. Meyer*

[†]Department of Electrical, Electronic and Computer Engineering, University of Pretoria, South Africa

[‡]Remote Sensing Research Unit, Meraka Institute, CSIR, Pretoria, South Africa cschwegmann@csir.co.za

[⋆]School of Engineering and ICT, University of Tasmania, Australia

ABSTRACT

A major task in any discrimination scenario requires the collection and validation of as many examples as possible. Depending on the type of data this can be a time consuming process, especially when dealing with large remote sensing data such as Synthetic Aperture Radar imagery. To aid in the creation of improved machine learning-based ship detection and discrimination methods this paper applies a type of neural network known as an Information Maximizing Generative Adversarial Network. Generative Adversarial Networks pit a generating and discriminating network against each other. A generator tries to create samples that are indistinguishable from real data whereas the discriminator tries to identify whether a sample is real or generated. Information Maximizing Generative Adversarial Network extend this idea by extracting untangled latent variables as part of the discrimination process which help to classify the data in terms of categories/classes and properties such as ship rotation. Despite the limited size and class distribution of the dataset, the paper showed that the trained network was able to generate convincing samples from the three given classes as well as create a discriminator that performs similarly to state-of-the-art ship discrimination methods despite using no labels for training.

Index Terms— Synthetic aperture radar, Machine learning, Marine technology

1. INTRODUCTION

The monitoring of large ocean areas is of great concern to the Maritime Domain Awareness effort [1, 2, 3, 4]. Ocean areas typically cannot be monitored in a cost effective manner at-sea so a remote sensing approach such as Synthetic Aperture Radar (SAR) monitoring becomes an effective alternative. Historically SAR has been an expensive means of monitoring areas but has recently become more affordable with the release of free and systematic Sentinel-1 A/B acquisitions.

Medium resolution SAR data typically covers areas of 100s of kilometres and to this end extracting meaningful, accurate and verifiable ship samples from this data can be a time-consuming process that often yields errors. For example, ships in a dataset such as [5] had an average length of 7 pixels versus a single SAR image which contains approximately 22000×17000 pixels. Furthermore, this dataset contained 46 SAR images but only yielded 1596 ship examples. This presents a fundamental limit in terms of the types of machine learning methods that can be applied to this data and so a means of generating more examples would be beneficial. This paper utilizes a new form of neural networks known as Information Maximizing Generative Adversarial Networks (InfoGAN) [6] and applies these to ship SAR data to simultaneously create new samples, an unsupervised classifier and identify latent variables in the data. The paper is organized as follows: In section 2 we discuss the data used in our experiments followed by the detailed discussion of the proposed ship sample generation method in section 3. Sections 4 and 5 provides some results, a concluding discussion about the method presented here and future work.

2. DATA DESCRIPTION

A SAR dataset has been created using 22 Sentinel-1 (Extra Wide Swath) and 3 RADARSAT-2 (Scan-SAR Narrow) acquisitions and a total of 42 dual and 4 single polarized, radiometrically calibrated images [5]. The dataset covers approximately 80% of the South African Exclusive Economic Zone (EEZ) with multiple acquisitions over a number of high ship densities areas (harbors). The dataset contains three classes represented by 21×21 sub-images, containing ships (positives), ship-like areas (false positives) and ocean areas (negatives). The 1596 positive examples were identified using expert analysis and false positives were identified by selecting areas that did not contain ships but were incorrectly highlighted by a low-threshold cell-averaging constant false alarm rate (CA-CFAR) prescreening method [1, 2, 3, 4]. Of all the detected false positives, 3192 sub-images were selected at random for this experiment. This represents a 2 to 1 ra-
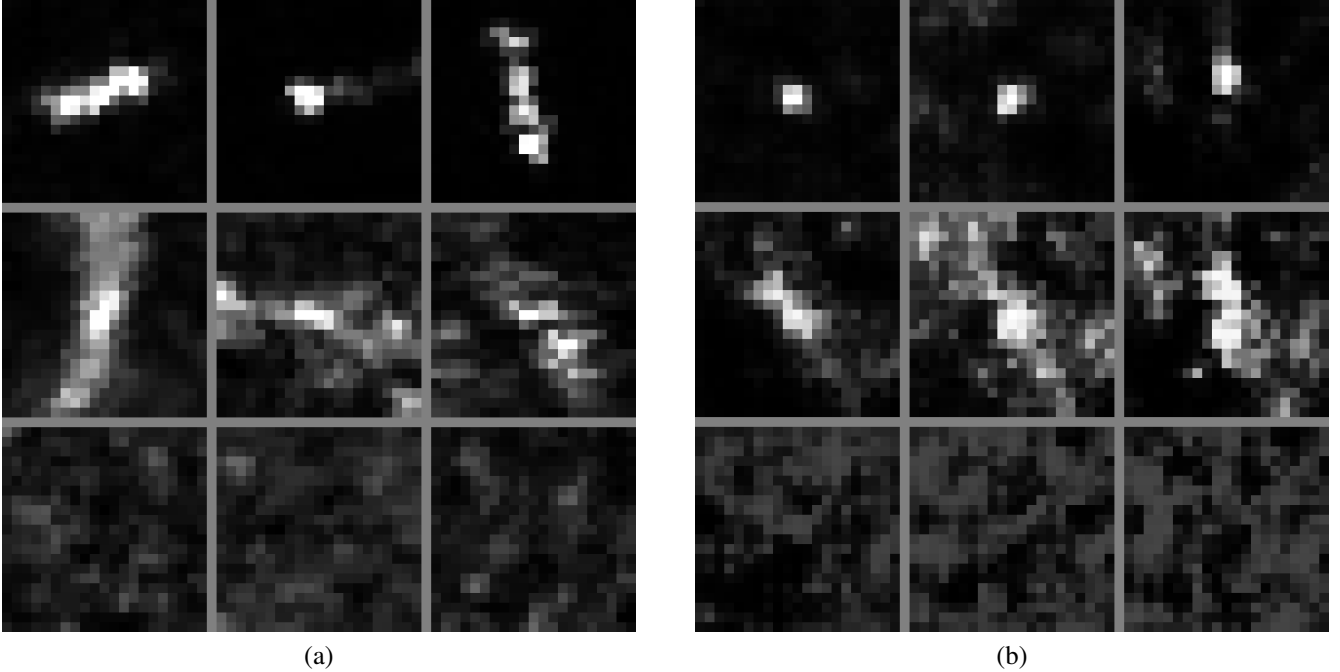
(a)	(b)

Fig. 1: Examples of data from (a) the actual dataset and (b) InfoGAN generator network. The original samples are split into three classes and the unsupervised InfoGAN network was able to identify these distinct classes and generate convincing examples given a uniformly distributed random input z and categorical variable c_1 indicating three classes.

tio of false positives to positives which is a likely ratio that would be encountered with careful prescreening threshold design and selection. Finally, 1596 sub-images that did not contain either a ship or false positive were selected as negative samples (ocean). The negatives samples were added to understand how the proposed method would deal with and generate such samples. No additional feature extraction was performed on the 9588 sub-images resulting in input feature set of 21x21 normalized RCS pixel values per image. Fig. 1 (a) shows 9 examples of the three different classes, across two SAR sensors. This dataset is significantly smaller than datasets that are commonly used to train deep learning systems. For instance, the MNIST dataset has 70000 samples across 10 classes, CIFAR-10 and CIFAR-100 datasets each have 60000 samples, evenly distributed across 10 or 100 classes [5]. The size of the dataset and unbalanced nature presents a unique opportunity to train and test an InfoGAN network within a data-limited SAR environment.

3. METHODOLOGY

This section will describe the basic premise of GAN [7] and explain how the extended InfoGAN [6] help to identify classes and pertinent features such as rotation or brightness within the SAR data.

3.1. Generative Adversarial Networks

GAN networks are a relatively new development in the machine learning community which tries to pit two competing deep neural networks against one another using a minimax

approach [7, 6], see Fig. 2 for an example of such a network. The goal is to create a generator distribution $P_G(x)$ which closely matches that of real world data $P_{\text{data}}(x)$. GAN creates a generator $G(z)$ that uses a noisy variable $z \sim P_{\text{noise}}(z)$ to generate a sample $x \sim p_z(x|z)$. The training and evaluation of this generator proceeds against a competing discriminator network $D(x)$ which tries to learn whether any given sample is real or generated. For a given generator, the discriminator $D(x) = \frac{P_{\text{data}}(x)}{P_{\text{data}}(x) + P_G(x)}$ is the optimal discriminator [7, 6]. The minimax procedure of this can be defined as

$$\min_G \max_D V_{\text{GAN}}(D, G) = \mathbb{E}_{x \sim p_{data}(x)}[\log D(x)] \qquad (1)$$
$$+ \mathbb{E}_{z \sim p_z(z)}[\log(1 - D(G(z)))].$$

Using a stochastic gradient descent approach the networks are trained using steps with the only knowledge available to the $D(x)$ being whether it correctly guessed the type of sample and for $G(z)$ whether it correctly fooled the $D(x)$. Errors in the $D(x)$ are propagated back to $G(x)$ in order to improve the generated examples and in this way improvements to the discriminator create better samples. The system tries to find an equilibrium where mutual information is maximized between the two (i.e. the discriminator cannot tell the difference between real or generated data).

3.2. Information Maximizing GAN

An extension to the GANs was introduce in 2016 which add a simple step to the minimax procedure where a discriminator paired network $Q(c|x)$ tries to identify latent variables c
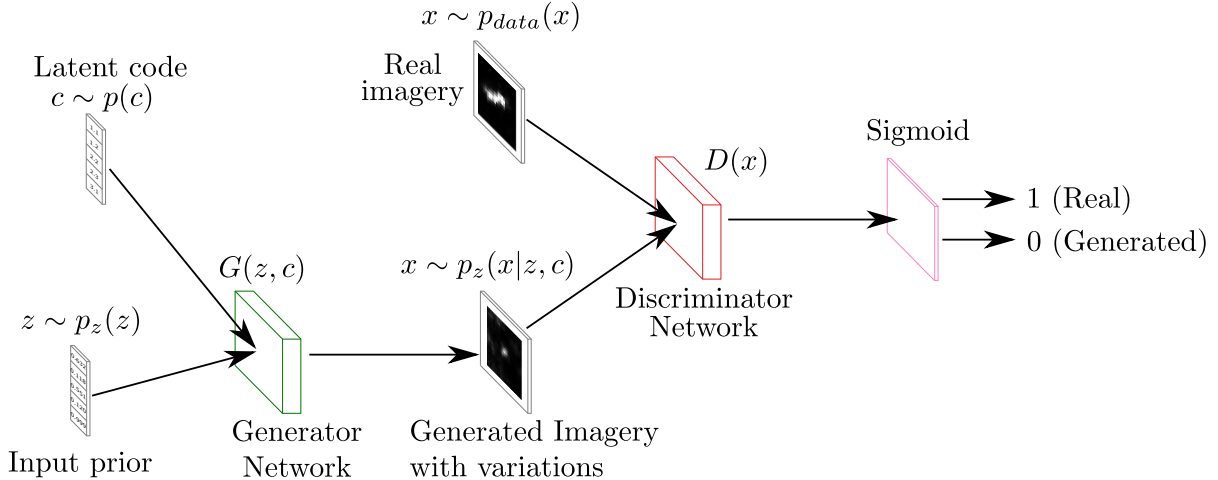
Fig. 2: Information Maximizing Generative Adversarial Network as described in [7, 6]. The generator network $G(z, c)$ generates examples given uniform random input z and latent codes c. The discriminator network $D(x)$ then tries to determine if any given sample is real or generated. Errors gradients for the discriminator are propagated back to the generator network in the hopes of teaching it to generate more convincing samples. The addition of the latent variable c allows the system to vary intrinsic properties of the samples generated by the GAN.

within the data. The minimax procedure of eq. 2 becomes

$$\min_{G,Q} \max_{D} V_{\text{InfoGAN}}(D, G, Q) = V(D, G) - \lambda L_I(G, Q) \quad (2)$$

for the InfoGAN network where λ is a hyperparamter normalised during training, and $L_I(G, Q)$ is the variational lower bound of the mutual information between the input variables z and latent variables c. These latent variables can be in two forms $c \sim$ Categorical $\|$ Uniform. Categorical variables are discovered during training as any part of the data that forms discontinuities, such as different classes. Uniformly distributed latent variables describe continuous properties of the generated samples such as rotation, size and brightness [6]. For a full mathematical treatment of GANs and InfoGANs readers are referred to [6].

4. RESULTS

One of the major difficulties when training an InfoGAN is the metric by which its overall performance is measured. Currently results are determined in two ways - visual inspection and the usage of the network in downstream tasks such as classification. For the purpose of this study both will be done but future work should include a metric by which generated samples can be objectively measured. The system was trained 5 times using cross validation which partitioned the data into 5 equal segments with one set (20%) used as testing data while the rest (80%) was used as training data. No labels were used to train the system, only three types of input were given to the system: the random uniform input prior z_L where $L = 21 \times 21 = 441$, a categorical variable $c_1 \sim Cat(K = 3, p = 0.\overline{333})$ to model the discontinuous class variations and $c_N \sim$ Uniform$(-1, 1)$ where $N = 2$ for this study but can be increased to other numbers given larger datasets/more train-

ing. The results of using these three inputs are shown in Fig. 1 (b), Fig. 3 and Fig. 4.

Fig. 1 (b) shows that given the three class categories, the system is able to separate data into 3 classes and generate samples that are visually quite similar (despite being drawn from randomly generated values). Two flaws are immediately apparent, one for the positive class and one for the negative class. The generator seems to fail at consistently generating larger ships but this can be explained by the fact that the majority ($> 70\%$) of ships presented to the discriminator are small with an average ship length of 7 pixels across all examples. The second error lies in the undersaturation of the "ocean clutter" generated in the negative samples. The reasons for this error are not clear and could stem from scaling by the network as output. For each of the 5-fold cross validation steps the discriminator network was tested to determine its performance given c_1 and the results for a single fold are shown in Fig. 3. These results compare well with the previously proposed Deep Highway Networks for discrimination (76.8% versus 84% positive class accuracy) with the added benefit of not needing labelled data. However, it should be noted that the networks will fail to converge to a meaningful result in approximately 10% of all cases. This can be caused by a number of factors including generator "collapse", small sample size or incorrect gradient descent learning rates for the discriminator versus the generator [8]. Finally, Fig. 4 shows the results of varying two separate latent variables c_2 and c_3 to determine which intrinsic parameter is being modelled. The Fig. 4 (a) shows how the network can capture the positive classes variance in terms of rotation and how this parameter can be smoothly varied to give different output images with different rotations. In a similar manner, Fig. 4 (b) shows that c_3 seems to control the relative brightness of the generated ship sample to the ocean clutter.

	POS	FP	NEG
POS	235	61	10
FP	44	1255	3
NEG	5	22	283

(a)

	POS	FP	NEG
POS	257	48	1
FP	44	1258	0
NEG	11	0	299

(b)

Fig. 3: Confusion matrices for the 2nd K-fold version of the (a) InfoGAN discriminator network and (b) 50 layer Highway Network [5]. The headings are POS, FP and NEG which represent the positive, false positive and negative classes respectively. The $D(x)$ using categorical variable c_1 was able to obtain a 76.8% positive class accuracy without labelled data versus the highway network's 84% with it.

(a)

(b)

Fig. 4: The effects of varying c_2 and c_3 after the network has been trained. The latent variable c_2 seems to capture the rotation of the generated sample whereas c_3 has captured the relative brightness of the "ocean clutter" versus the ship at the center.

5. CONCLUSION

Ship detection and discrimination in SAR imagery is an on-going research area for any sea-bordering country. New SAR satellite sensors are constantly being launched and with the advent of freely accessible data from Sentinel-1 creating effective ship discrimination networks is important to maintain a high level of MDA. One of the major tasks when designing any new ship detection system is the generation of verifiable SAR data to test the system with. Creating datasets, specifically on large spatial data such as SAR imagery, can be time consuming and error prone. One approach in the machine learning community is to use generative networks to create samples that closely match real-world data to supplement data already obtained. New samples are often generated using rules that are set up given observations of the data. Another means of generating new sample data is by applying machine learning generative networks such as InfoGAN. This paper showed how an InfoGAN can be applied, in an unsupervised manner, to medium resolution SAR ship data to: generate convincing sample for all three classes present; create a discriminator that performs similarly to state-of-the-art solutions; and, finally, extract latent variables of the ship such as rotation, brightness and size. Future work includes using these latent variables to better separate false positives and ships as well as training a generator network based on input data such as AIS ship information.

6. REFERENCES

[1] W. Kleynhans W., B. P. Salmon, C. P. Schwegmann, and V. Seotlo, "Ship Detection in South African oceans using a combination of SAR and historic LRIT data," in *IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, July 2013, pp. 1521–1524.

[2] C. Schwegmann, W. Kleynhans, and B. Salmon, "Manifold Adaptation for Constant False Alarm Rate Ship Detection in South African Oceans," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 8, no. 7, pp. 3329–3337, July 2015.

[3] D. J. Crisp, "The State-of-the-Art in Ship Detection in Synthetic Aperture Radar Imagery," Tech. Rep. DSTO-RR-0272, DSTO Information Sciences Laboratory, Edinburgh, South Australia, May 2004.

[4] K. El-Darymli, P. McGuire, D. Power, and C. Moloney, "Target detection in synthetic aperture radar imagery: a state-of-the-art survey," *Journal of Applied Remote Sensing*, vol. 7, no. 1, pp. 071598–071598, Mar. 2013.

[5] C. P. Schwegmann, W. Kleynhans, B. P. Salmon, L. W. Mdakane, and R. G. V. Meyer, "Very Deep Learning for Ship Discrimination in Synthetic Aperture Radar Iimagery," in *IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, July 2016, pp. 104–107.

[6] X. Chen et al., "InfoGAN: Interpretable Representation Learning by Information Maximizing Generative Adversarial Nets," *Journal of Computing Research Repository*, vol. abs/1606.03657, no. 1, pp. 1–14, July 2016.

[7] I. Goodfellow, "Generative adversarial nets," in *Advances in Neural Information Processing Systems 27 (NIPS 2014)*, Sept. 2014, pp. 2672–2680.

[8] L. Metz, B. Poole, D. Pfau, and J. Sohl-Dickstein, "Unrolled Generative Adversarial Networks," *Journal of Computing Research Repository*, vol. abs/1611.02163, no. 1, pp. 1–19, Dec. 2016.