

Real-time Tropical Cyclone Intensity Estimation by Handling Temporally Heterogeneous Satellite Data

Boyo Chen, Buo-Fu Chen*, Yun-Nung Chen

National Taiwan University
 boyochen722@gmail.com, bfchen777@gmail.com, y.v.chen@ieee.org

Abstract

Analyzing big geophysical observational data collected by multiple advanced sensors on various satellite platforms promotes our understanding of the geophysical system. For instance, convolutional neural networks (CNN) have achieved great success in estimating tropical cyclone (TC) intensity based on satellite data with fixed temporal frequency (e.g., 3 h). However, to achieve more timely (under 30 min) and accurate TC intensity estimates, a deep learning model is demanded to handle temporally-heterogeneous satellite observations. Specifically, infrared (IR1) and water vapor (WV) images are available within every 15 minute period, while passive microwave rain rate (PMW) is available about every 3 hours. Meanwhile, the visible (VIS) channel is severely affected by noise and sunlight intensity, making it difficult to be utilized. Therefore, we propose a novel framework that combines generative adversarial network (GAN) with CNN. The model utilizes all data during the training phase including VIS and PMW information and eventually uses only the high-frequent IR1 and WV data for providing intensity estimates during the predicting phase. Experimental results demonstrate that the hybrid GAN-CNN framework achieves comparable precision to the state-of-the-art models, while possessing the capability of increasing the maximum estimation frequency from 3 hours to less than 15 minutes. Please visit <https://github.com/BoyoChen/CNN-GAN-TC> for codes and implementation details.

1 Introduction

Tropical cyclone (TC) is a type of low-pressure weather system that forms and develops over the warm tropical ocean. It is characterized by intense rotating winds and severe rainfall associated with eyewall clouds and spiral rainbands. A TC hitting the land poses severe threats to society by producing gusty wind, sea surge, flooding, and landslide.

TC intensity (i.e., the maximum sustained surface wind near the center) is one of the most critical factors in disaster management. The state-of-the-art SATCON (Velden and Herndon 2014), widely used in operational forecasting, estimates TC intensity based on consensus decision-making procedures using infrared images from geostationary satellites and other observations from low-Earth-orbit satellites.

*Corresponding author.

Note that high quality SATCON estimates could be obtained at an approximate three-hour frequency.

Recently, several studies have applied convolutional neural networks (CNN) on satellite images to estimate TC intensity (Pradhan et al. 2018; Chen, Chen, and Lin 2018; Chen et al. 2019; Wimmers, Velden, and Cossuth 2019). The previous work released a benchmark dataset "TCIR" for the TC-image-to-intensity regression task (Chen, Chen, and Lin 2018), which consists of satellite images including four channels (fig. 1 (a)-(d)). The CNN-TC network (Chen et al. 2019)¹ utilized IR1 (Infrared) and PMW (passive microwave rain rates) channels and achieved a state-of-the-art performance also at a three-hour frequency.

Notably, the VIS (visible) channel was not utilized in the past, because it only provides meaningful cloud information during the daytime. Besides, PMW channel images can only be collected at the frequency of ≈ 3 hours while the observations from IR1 and water vapor (WV) channel are available almost anytime.² As real-time intensity estimations are required for pragmatic disaster management, we must be able to handle temporally-heterogeneous satellite observations.

This work proposes a novel deep learning model combining CNN and GAN (generative adversarial network) to deal with the temporally heterogeneous datasets. Our goal is to eliminate the dependency on PMW channel and perform the prediction with only IR1 and WV channels, so that 24/7 intensity estimations could be provided. To keep the performance comparable to the state-of-the-art model, we also make use of the most of the information provided by good quality VIS images. We proposed a novel 5-stage training strategy, with which two separated generators are trained for producing simulated VIS and PMW images. Afterward, generated VIS and PMW images can be used along with IR1 images to conduct the estimates. fig. 2 is a schematic of our inference model.

We summarize related work about GAN in section 2 and provide analysis of various used satellite observations in sec-

¹This article is published in *Weather and Forecasting*, one of the best journals in the relevant research field.

²PMW is collected by low-Earth-orbit microwave satellites. Meanwhile, IR1, WV, and VIS are collected by geostationary satellites. Normally, geostationary satellites collect data at the frequency of 15 min. In a super rapid scan mode, a geostationary satellite provides observations every 2 min.

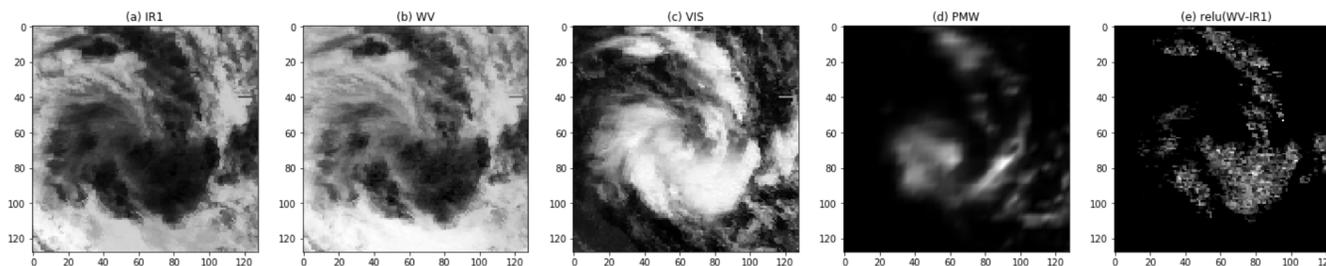


Figure 1: Channels of the TCIR dataset. (a)-(d) The four basic channels: IR1, WV, VIS, and PMW. (e) $\text{relu}(WV - IR1)$. By comparing (d) and (e), we can find the implicit correlations between the PMW and other channels.



Figure 2: A schematic showing how the proposed hybrid GAN-CNN model is used for estimating TC intensity. Note that only the temporally homogeneous IR1 and WV data go into the model as the input while the GAN component generates the PMW and VIS features for the CNN regressor.

tion 3. section 4 discusses the hybrid GAN-CNN architecture and the training strategy. Then section 5 demonstrates the capability of the proposed model dealing with temporally heterogeneous datasets to provide 24/7 intensity estimates in good quality. section 6 provides a quick recap.

2 Background Knowledge

In this section, we summarize several GAN frameworks that inspire our work. In a typical GAN (Goodfellow et al. 2014), there are two opposing players: a *generator* and a *discriminator*. The generator is responsible for generating fake data and trying to confuse the discriminator. A discriminator acts like an umpire, responsible for distinguishing between real and fake data. The competition between both players prompt the generator to generate fake data that is difficult to distinguish from the real data.

CGAN (conditional GAN) (Mirza and Osindero 2014) attaches conditions to the input of the generator. These conditions should be related to several side information provided by the image, such as the class of the object in the picture. The generator is restricted only to generate images that meet the conditions. Additionally, the specified conditions will be disclosed to the discriminator along with the generated data. This formulation allows generators to generate images according to our needs.

AC-GAN (auxiliary classifier GAN) (Odena, Olah, and Shlens 2017) follows the steps of CGAN by setting conditions to the generator. Differently, specified conditions are not exposed to the discriminator. Instead, the discriminator are required to reconstruct the side information of the images on its own. Compared to CGAN, AC-GAN further improves the stability of training and the quality of the generated data.

In most of GAN frameworks, the discriminator only outputs a single probability to determine whether the entire image is generated or not. In contrast, **PatchGAN** (Li and Wand 2016) modifies its discriminator to cut the whole image into multiple small patches with overlap and discriminates them piece by piece. This technique has been proven to be mathematically equivalent to doing blending with data after cutting into patches. This technique is useful for data with even distribution and no distinct boundary, such as satellite images.

U-Net (Ronneberger, Fischer, and Brox 2015) is a type of generators which has a similar structure to an auto-encoder. The input and output of U-Net are both images. It retains the local details from the input image by skip connections, then reconstructs them in the corresponding position. Therefore, it is widely used in situations where the input and output are pictures of the same size.

Proposed in 2017, **Pix2Pix** (Isola et al. 2017) uses pictures as the condition of the generator and completes the style conversion task brilliantly. Because the inputs of the generator are images, Pix2Pix reasonably uses U-Net as its generator. Meanwhile, the PatchGAN discriminator is used. L1 distance between the input picture and the output picture are added to the generator loss. This term in the loss function directly guides the generator to produce the desired image and is useful for fighting against *mode collapse*. However, this framework demands data to be paired before and after conversion.

CycleGAN (Zhu et al. 2017) focuses on training two generators (style A- ζ style B, style B- ζ style A) at the same time instead of training a one-way generator. CycleGAN requires the input image (assuming style A) to be converted back as similar as possible after passing through both generators. We can thus ensure that the generator retains the critical information in the original image when converting the style. The concept in CycleGAN is similar to an auto-encoder. On the other hand, CycleGAN also resolves the limitation of Pix2Pix that requires the paired data. However, due to the characteristics of TC images, we apply Pix2Pix in our framework, but not that of the CycleGAN. The detailed reasons are described in section 3.1.

3 Data Analysis and Preprocessing

We conducted our experiment on the benchmark TCIR dataset, which includes 4 channels: Infrared (IR1), Water Vapor (WV), Visible (VIS), and Passive Microwave (PMW). An example is shown in fig. 1 (a)-(d). For more details about the TCIR, please refer to the previous work of Chen, Chen, and Lin (2018).

3.1 Analysis of the PMW Data

According to domain knowledge, PMW is positively correlated with $relu(WV-IR1)$ (Olander and Velden 2009). Therefore, the fig. 1(d) and fig. 1(e) are somewhat similar to each other. This evidence suggests that a properly trained Pix2Pix model can hopefully restore the PMW channel using IR1 and WV information. However, since there is still a certain gap between PMW and $relu(WV-IR1)$, we cannot directly use the latter one to replace the former one.

Besides, we also discovered that the conversion is unidirectional. While we can use IR1 and WV to derive PMW, it is challenging to derive IR1 and WV from PMW. Thus, cycleGAN is not recommend to serve as the PMW generator.

3.2 Analysis of the VIS Data

The VIS channel is the noisiest compared with other three channels. Types of effects includes:

1. VIS images are meaningless at night due to the significant decrease in light intensity.
2. Even under the daylight, about 1/5 of the VIS are noisy or completely black.
3. The intensity of sunlight varies at different hours throughout the day. Cloud is more obvious when the time is closer to noon.
4. Similar to other channels, the signal may be disturbed by noises. Sometimes even half of a VIS image is black. Besides, there could be strip noise occasionally.

By adopting the VIS generator, we want to (1) generate VIS images with IR1 and WV at all times, even during the night, (2) calibrate sunlight level among all VIS images, and (3) remove block and strip noise. For clearer example, please refer to fig. 9 in the experiment section.

3.3 VIS Quality Control

A subset of VIS data with good quality is needed to serve as the labeling data and facilitate the training of the VIS generator. This selection was conducted based on the calculation of the mean and standard deviation of the entire image values.

By filtering unusual values of mean, we can exclude all-black and all-white images. On the other hand, a clear VIS image's standard deviation is likely to fall within a certain interval. After consulting human experts, we subjectively determine the range of reasonable mean and standard deviation in which a high-quality VIS should be (1) $0.1 \geq mean \geq 0.7$ and (2) $0.1 \geq std \geq 0.31$.

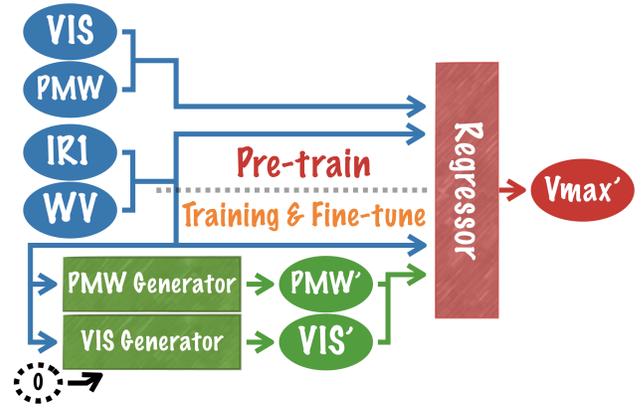


Figure 3: Framework of hybrid GAN-CNN.

Moreover, to further reduce false positives, we limit high-quality visible images to be between 07:00 and 17:00 because it is impossible to have high-quality VIS data during the nighttime, namely our third condition: (3) $7 \geq local.time.hour \geq 17$.

As a result, a total of 12480 TC images (training data) were marked as good quality VIS, accounting for 20.9% of the total 59837 samples in the training data.

4 Proposed Method

To eliminate the dependence on the usage of PMW and solve problems caused by severe noise in the VIS, we demand GANs that can stably generate VIS and PMW channels only by the two temporally homogeneous channels: IR1 and WV (fig. 3). In the dataset, every set of TC images have PMW observation, while only 20% have good-quality VIS. To use available information optimally, generators for PMW and VIS are trained separately (fig. 4).

We use an adjusted U-Net for the **generator**. Thanks to the skip connections in U-Net, local details in the source image can be preserved. Regarding the **discriminator**, we use the idea of **PatchGAN**. Satellite observations are continuous, without concept of objects, foreground, background, and boundaries. Radically speaking, any segmentation of a real TC image can be determined normal because there is no object to be dissected. A PatchGAN divides images into multiple small areas before determining whether it is real, which is suitable for our data and makes training more stable.

4.1 Training Objective

There are three components in this framework: a generator, a discriminator, and a regressor. The loss functions are formulated as:

$$\begin{aligned}
 \mathcal{L}_{G_{target}} &= l_{gen} + \alpha \times l_{L2} + \beta \times l_{regr} \\
 &\quad + \gamma \times l_{m2n} \times [[target = vis]] \\
 \mathcal{L}_{D_{target}} &= l_{disc} + \gamma \times l_{m2n} \times [[target = vis]] \quad (1) \\
 target &\in \{vis, pmw\} \\
 \mathcal{L}_R &= l_{regr}
 \end{aligned}$$

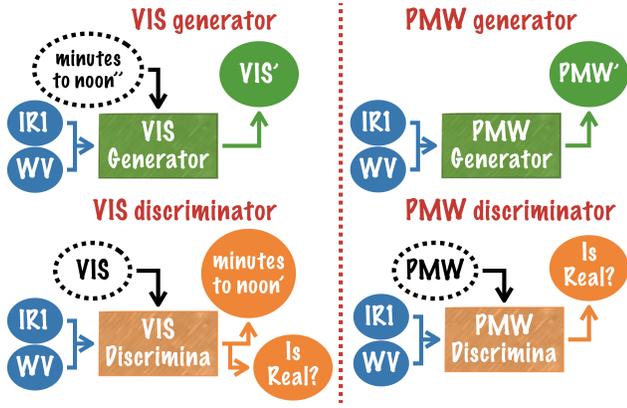


Figure 4: Illustration of two sets of GANs.

α, β, γ used in the experiments are listed in table 1.

Discriminator loss (\mathcal{L}_D) The goal of the discriminator is to correctly identify the target image as real observation or a fake generated by the generator:

$$l_{disc}(G, D) = \mathbb{E}_{ir1, wv, target} [\log D_{disc}(ir1, wv, target) + 1 - \log D_{disc}(ir1, wv, target')],$$

$$target' = G_{target}(ir1, wv),$$

$$target \in \{vis, pmw\}$$
(2)

This concept is the same as that in the Pix2Pix.

Generator loss (\mathcal{L}_G): The loss contains the following requirements associated with our loss function. First, the goal of the generator is to confuse the discriminator:

$$l_{gen}(G, D) = \mathbb{E}_{ir1, wv, target} \log D_{disc}(ir1, wv, target'),$$

$$target' = G_{target}(ir1, wv),$$

$$target \in \{vis, pmw\}$$
(3)

This is equivalent to the loss in the patchGAN.

Second, the generated VIS/PMW must be similar to the input VIS/PMW:

$$l_{L2}(G) = \mathbb{E}_{ir1, wv, target} ||target - target'||_2,$$

$$target' = G_{target}(ir1, wv),$$

$$target \in \{vis, pmw\}$$
(4)

This concept is borrowed from the Pix2Pix. Note, however, that the L2 distance (MSE) is used instead of the L1 distance (MAE) in our model because L2 distance encourages more blurring. Usually, we want to generate images with clear lines and apparent boundaries. However, satellite images have no concept of boundaries, where they are smoother than ordinary pictures, such as dogs, cats, and cars. Therefore, we modify the Pix2Pix architecture based on our need. A comparison of using L1 distance and L2 distance will be shown later in section 5.2.

Besides, we have added the following two innovative designs to our GANs, specializing the GANs to complete tasks

appropriately. The details are described in the following two sub-sections.

Auxiliary time loss (l_{m2n}): As mentioned in section 3.2, we expect the VIS generator to adjust all generated VIS images to the sunlight level at noon. To achieve this goal, we need to first calculate $m2n$ (minutes to noon) for each VIS image: $m2n = |60 \times hour + minute - 60 \times 12|$. We take $m2n$ as an additional condition and apply the concept of the AC-GAN to our VIS generator and discriminator. Then the loss function therefore has an extra term.

Discriminator:

$$l_{m2n}(D) = \mathbb{E}_{vis, m2n} ||m2n - D_{m2n}(vis)||_2$$
(5)

Generator:

$$l_{m2n}(G, D) = \mathbb{E}_{ir1, wv} ||m2n' - D_{m2n}(vis')||_2,$$

$$vis' = G_{vis}(ir1, wv, m2n'),$$

$$0 \leq m2n' \leq 300$$
(6)

VIS images with good quality were obtained during daytime, specifically from 7:00 to 16:59 (see section 3.3), so the largest value of minutes to noon is 300. In training, randomly generated floating-point numbers between $[0, 300]$ are provided to the generator as condition.

Regressor loss (\mathcal{L}_R): Since our ultimate goal is to use the generated VIS / PMW channels as the regressor's inputs, the generator is requested to create useful features that can facilitate estimations.

The generator is first asked to generate PMW and VIS given the condition $m2n = 0$. The obtained results are provided to the regressor, along with IR1 and WV. Finally we can obtain our intensity estimation using generated VIS / PMW (e.g. lower part of fig. 3). The precision of the prediction is also added to one of the terms of the generator loss:

$$l_{regr}(G_{vis}, G_{pmw}, R) = \mathbb{E}_{ir1, wv, v_{max}} ||v_{max} - v_{max}'||_2,$$

$$v_{max}' = R(ir1, wv, vis', pmw'),$$

$$vis' = G_{vis}(ir1, wv, 0),$$

$$pmw' = G_{pmw}(ir1, wv),$$
(7)

The term v_{max} stands for the maximum wind velocity, the definition of TC intensity.

4.2 Strategy of Three-Stage Training

To fairly calculate l_{regr} , we need to pre-train the regressor. Therefore, a novel three-stage training illustrated in fig. 3 and fig. 5 is proposed.

Stage	VIS / PMW used in regressor	Optimize target
Pre-train regressor	original	Regressor
Training generator	generated	Generator
Fine-tune regressor	generated	Regressor

Figure 5: 3-stage training.

	Stage	Data used	VIS used in regressor	PMW used in regressor	Optimize target
1	Pre-train regressor	good quality VIS only	original	original	Regressor
	Training VIS generator	good quality VIS only	generated	original	VIS Generator
2	Pre-train regressor	all data	generated	original	Regressor
	Training PMW generator	all data	generated	generated	PMW Generator
	Finetune regressor	all data	generated	generated	Regressor

Figure 6: The steps of the extended 5-stage training.

- Pre-train regressor:** Pre-train the regressor with original VIS and PMW along with other regressor inputs and thereby let the regressor learn how to extract essential features from real VIS / PMW in advance.
- Train generator:** The generator and discriminator are optimized together. The loss mentioned in eq. (7) ensures that significant feature regressor is using will be generated.
- Fine-tune regressor:** With well-trained generator, we use the generated VIS / PMW features to fine-tune the regressor.

4.3 Strategy of Five-Stage Training

Good results are achieved through the above procedures of three-stage training. To be even better, VIS generator should be trained with only data with good quality VIS images. Meanwhile, the training process becomes more stable when 2 generators are only optimized when another one is fixed.

Therefore, based on the three-stage training described above, a more detailed five-stage training process illustrated in fig. 6 is proposed. In the five-stage training, the first two stages of three-stage training are repeated for two loops:

- Loop 1:** Only data with good quality VIS images are used in this loop. PMW generator is fixed while we focus on optimizing the VIS generator. When calculating loss in eq. (7) we use the original PMW instead of a generated one.
- Loop 2:** All data are used. VIS generator is fixed while we pay attention to PMW generator.
- After training both generators, regressor get fine-tuned using the generated PMW' and VIS'.

By applying five-stage training, we obtain the final operational predicting model (fig. 2).

5 Experiments and Analysis

In the following section, two techniques proposed in the previous work will be explained briefly, including **auxiliary features** and **rotation-blending**. Please refer to Chen et al. (2019) for more details. Next, we will compare our performance with related works which also focuses on estimating TC intensity, including both operational meteorological models and deep learning techniques. Finally, we qualitatively analyze the quality of the proposed model.

In the experiments, we split the dataset into three parts: (1) Training data: TCs during 2000-2014, (2) Validation data: TCs during 2015 and 2016, (3) Testing data: TCs during 2017.

Auxiliary Features: In addition to the output from convolution layers, additional features are passed into the regressor. The auxiliary features are demonstrated to be helpful in improving the precision of estimation (Chen et al. 2019). These features provide clues such as (1) day of year: stands for seasonal information, (2) local time, and the most influential one: (3) One-hot encoded region codes: region codes is in {WPAC, EPAC, CPAC, ATLN, IO, SH}, representing 6 different basins.

Rotation Blending: Considering the nature of TCs as a rotating weather system, TC data is rotation invariant. That is, rotations with respect to the center usually do not affect the estimation of the TC intensity. (Chen, Chen, and Lin 2018) demonstrated that the idea of using rotation for augmentation leads to a significant improvement in performance.

During the training phase, each image will be randomly rotated by any degree before feeding into our model. When it comes to inference, images will be rotated by evenly distributed ten angles ranged from 0 to 360 to collect 10 different estimations. Afterward, these intensity estimations are blended to obtain the final estimate.

5.1 Intensity Estimation Performance

The main task of this work is to accurately estimate the TC intensity, which is the output of our model. The unit of TC intensity is knot (kt) defined as its maximum wind speed (V_{max}). The value of V_{max} is usually ranged in [30, 180], and TCs with a V_{max} larger than 96 kt are considered as intense TCs.

Table 2 compares our performance to other works. ADT (Advanced Dvorak Technique) (Olander and Velden 2007) is a common used method to estimate TC intensity, which extract features from IR1 images before applying linear regression. SATCON (Velden and Herndon 2014) is the state-of-the-art model used by meteorologists in operational forecast. It rely highly on observations from low-Earth-orbit satellites.

The performance of our proposed model is comparable to the state-of-the-art model in both deep learning and meteorology, while our model can provide much more timely estimations.

In fig. 7, we compare the validation MSE score over the first 100 epoch of training. The blue line represents the state-of-the-art model, which provides intensity estimates every 3 hours. The orange and green line shares the same inputs, IR1 and WV, which is available every 15 min. The former is the model that directly uses them for estimations while the latter is our proposed model. In contrast, our adequately trained GAN model helps us further improve the performance of intensity estimation, bringing it closer to the state-of-the-art model. Most importantly, our proposed model can provide intensity estimates every 15 min.

Stage	$\alpha(L2)$	$\beta(\text{regr})$	$\gamma(\text{m2n})$	Max epochs
Pre-training Regressor	-	-	-	70
VIS Generator	1000	0.0001	0.002	500
Pre-training Regressor	-	-	-	100
PMW Generator	10	0.001	-	200
Fine-tune regressor	-	-	-	300

Table 1: Hyper-parameters used in the 5-stage training.

	no smoothing		w/ smoothing ³		input	frequency
	valid	test	valid	test		
ADT (Olander and Velden 2007)			12.65		IR1	30 min ⁴
SATCON (Velden and Herndon 2014)			8.59		IR1, PMW ⁵	\approx 3H
CNN-TC (Chen et al. 2019)	10.38	-	8.74	8.39	IR1, PMW	\approx 3H
<i>CNN-TC</i> ⁶	10.13	10.13	8.62	8.89	IR1, PMW	\approx 3H
Proposed model	10.43	10.19	9.01	9.33	IR1, WV	\leq 15 min

Table 2: The comparison between RMSEs of our proposed models and state-of-the-art models.

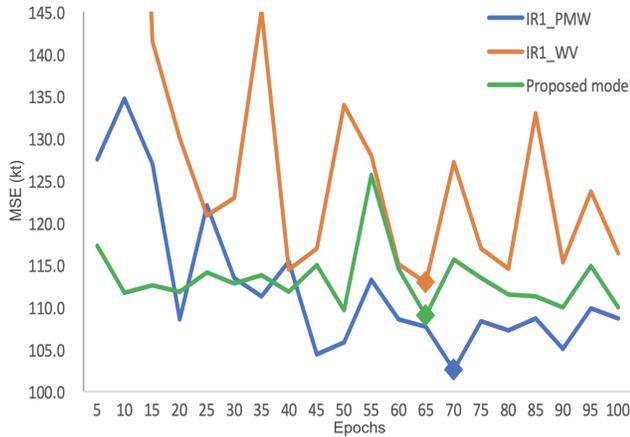


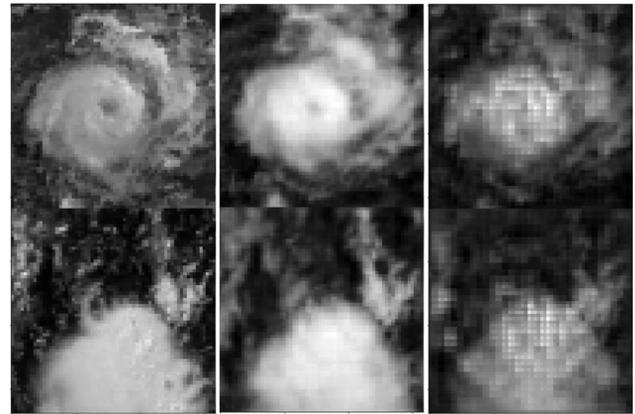
Figure 7: Learning curves in MSEs for models with different channel combinations as the input.

5.2 Effectiveness of the L2 Distance

L2 distance is chosen instead of L1 distance in eq. (4), which is different from an ordinary Pix2Pix framework. Compared to commonly seen pictures, satellite observations are continuous and have smoother boundaries. As described in section 4.1, using L2 distance encourages more blurring. In fig. 8, we take VIS channels as examples, compare the generated results from models using L1 distance and L2 distance. As shown, the model using L1 distance generates images less smoothly.

5.3 Qualitative Study

fig. 9 shows the generated images from the proposed model. Compared with the original VIS images, the generated VIS image is slightly blurred, and the eye is not as clear as the original one. Nevertheless, these images can be generated stably anytime, with the removal of most noises, and with adjustment of sunlight intensity.



(a) original VIS (b) VIS generator using L2 distance (c) VIS generator using L1 distance

Figure 8: VIS examples generated by generators using L1 and L2 distance for l_{regr} , described in eq. (4). Validation data is used for generating VIS in this figure.

Interestingly, most of the generated PMW seems to be slightly rotated. This is presumably because observations can only be obtained when the low-Earth-orbit satellites pass across the TCs. Therefore, the time PMW channel being collected could be slightly misaligned with other channels.

6 Conclusion

This paper focuses on improving the utility of deep learning in TC intensity estimation for practical scenarios. A modified Pix2Pix GAN framework is presented to better fit in our unusual TC data that is temporally-heterogeneous, and we combine the GAN with a CNN regressor, which eventually becomes our proposed hybrid GAN-CNN model. By properly dealing with temporally-heterogeneous data, our proposed model not only achieves comparable performance with the state-of-the-art model in estimating precision but

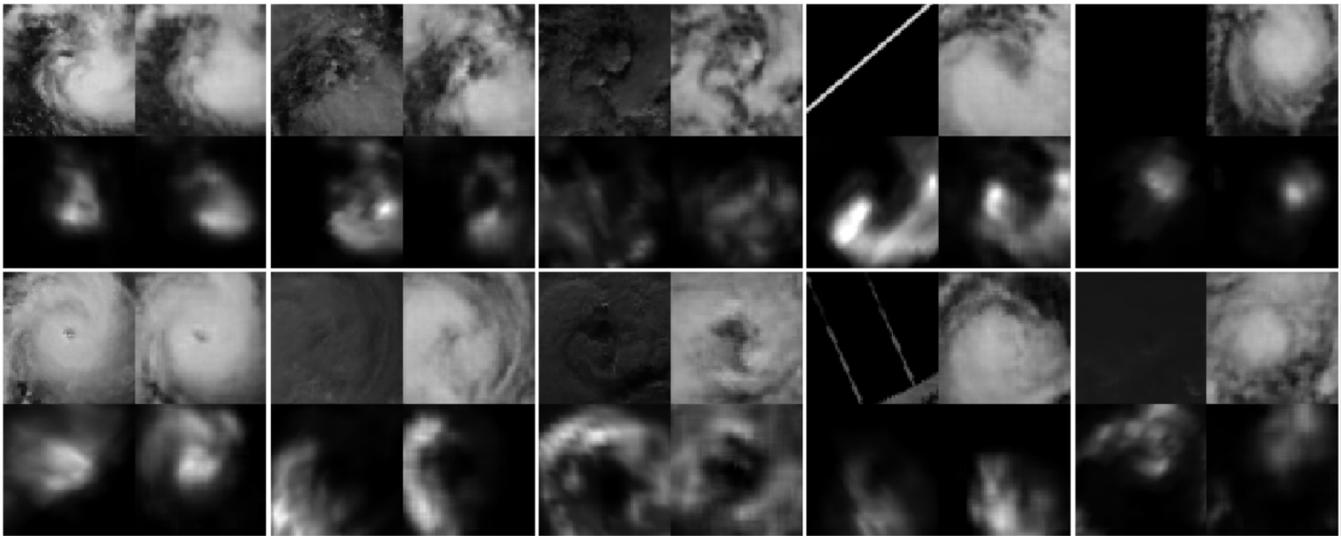


Figure 9: The generated VIS and PMW data. For each block, there are original VIS (upper left), generated VIS (upper right), original PMW (lower left), and generated PMW (lower right). Validation data is used for generating this figure.

also provides the estimates much more timely, which leads to better TC warning operations and could eventually save more lives.

Moreover, being capable of estimating TC intensity by only IR1 and WV channels, the new model enables us to reanalyze the intensity of TCs further back to the time only simple IR1 and WV images were available (1980s) instead of starting from the 2000s, providing the potential of a breakthrough in the research about climate change and global warming from the aspect of TC activities.

Acknowledgements

This work was funded by Grant MOST 109-2111-M-002-009 of the Ministry of Science and Technology, Taiwan.

Broader Impact

Timely warnings are critical to a better disaster warning/management system. From Hurricane Katrina to Hurricane Harvey and even to recent COVID-19, we should learn that early awareness of these disasters could result in saving many lives.

Take Hurricane Harvey as an example. Back on August 25, 2017, Hurricane Harvey was estimated as a category 2 hurricane during 02:00 to 14:00. After it became a category 3 hurricane at 17:00, all in a sudden, it grew into a category

³Simple smoothing techniques are applied here to obtain a boost in estimation precision (Chen, Chen, and Lin 2018; Chen et al. 2019).

⁴They are currently providing estimations every 30 mins. But it could be ≤ 15 min as well.

⁵SATCON depends on low-Earth-orbit satellites observations, which is somehow similar to the PMW.

⁶Our reproduction of CNN-TC. We add additional batch normalization layers in our reproduced CNN-TC, which leads to a minor improvement.

4 hurricane only 3 hours later at 20:00. Eventually, Hurricane Harvey caused record-breaking damage in the U.S. To timely monitor this rapid intensification is challenging for the conventional TC intensity estimation techniques, which provides estimates only once per 3 hours. Imagine if there was a timely estimation system which can provide intensity estimates every 15 minutes. The insane increasing trend could be detected earlier, precious time could be bought, people could be better prepared, and lives could be saved.

On the downside, a significant pitfall of applying machine learning techniques in disaster management is that models are naturally conservative when facing extreme cases. For example, TC intensity ranges from 20 to 180 kt, while about 70% of the value distributed within 35 kt to 64 kt, which makes it almost unavoidable for machine learning models to under-estimate unprecedented extreme values. Therefore, an auxiliary warning system should exist no matter how precise we humans can be with our machine learning techniques. After all, disaster management is a cost-sensitive scenario in which we should always keep in mind that while false positive is annoying, false negative is deadly.

References

- Chen, B.; Chen, B.-F.; and Lin, H.-T. 2018. Rotation-blended CNNs on a new open dataset for tropical cyclone image-to-intensity regression. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 90–99.
- Chen, B.-F.; Chen, B.; Lin, H.-T.; and Elsberry, R. L. 2019. Estimating tropical cyclone intensity by satellite imagery utilizing convolutional neural networks. *Weather and Forecasting* 34(2): 447–465.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y.

2014. Generative adversarial nets. In *Advances in neural information processing systems*, 2672–2680.

Isola, P.; Zhu, J.-Y.; Zhou, T.; and Efros, A. A. 2017. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1125–1134.

Li, C.; and Wand, M. 2016. Precomputed real-time texture synthesis with markovian generative adversarial networks. In *European conference on computer vision*, 702–716. Springer.

Mirza, M.; and Osindero, S. 2014. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784* .

Odena, A.; Olah, C.; and Shlens, J. 2017. Conditional image synthesis with auxiliary classifier gans. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 2642–2651. JMLR. org.

Olander, T. L.; and Velden, C. S. 2007. The advanced Dvorak technique: Continued development of an objective scheme to estimate tropical cyclone intensity using geostationary infrared satellite imagery. *Weather and Forecasting* 22(2): 287–298.

Olander, T. L.; and Velden, C. S. 2009. Tropical cyclone convection and intensity analysis using differenced infrared and water vapor imagery. *Weather and forecasting* 24(6): 1558–1572.

Pradhan, R.; Aygun, R. S.; Maskey, M.; Ramachandran, R.; and Cecil, D. J. 2018. Tropical Cyclone Intensity Estimation Using a Deep Convolutional Neural Network. *IEEE Transactions on Image Processing* 27(2): 692–702.

Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, 234–241. Springer.

Velden, C.; and Herndon, D. 2014. Update on the SATellite CONsensus (SATCON) algorithm for estimating TC intensity. *Poster session I. San Diego* .

Wimmers, A.; Velden, C.; and Cossuth, J. H. 2019. Using deep learning to estimate tropical cyclone intensity from satellite passive microwave imagery. *Monthly Weather Review* 147(6): 2261–2282.

Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, 2223–2232.