

# MULTI-SCALE GRIDDED GABOR ATTENTION FOR CIRRUS SEGMENTATION

F. Richards, X. Xie

Computer Science dpt,  
Swansea University  
Swansea, UK

A. Paiement

Université de Toulon,  
Aix Marseille Univ,  
CNRS, LIS,  
Marseille, France

E. Sola, P.-A. Duc

Université de Strasbourg,  
CNRS, Observatoire  
astronomique de Strasbourg,  
Strasbourg, France

## ABSTRACT


In this paper, we address the challenge of segmenting global contaminants in large images. The precise delineation of such structures requires ample global context alongside understanding of textural patterns. CNNs specialise in the latter, though their ability to generate global features is limited. Attention measures long range dependencies in images, capturing global context, though at a large computational cost. We propose a gridded attention mechanism to address this limitation, greatly increasing efficiency by processing multi-scale features into smaller tiles. We also enhance the attention mechanism for increased sensitivity to texture orientation, by measuring correlations across features dependent on different orientations, in addition to channel and positional attention. We present results on a new dataset of astronomical images, where the task is segmenting large contaminating dust clouds.

**Index Terms**— Attention, multi-scale, orientation, segmentation, astronomy

## 1. INTRODUCTION

Global context is vital in vision: scenes are understood through key descriptive regions, such as grass or sky, as well as through objects. This is especially relevant when processing contaminants covering large regions, such as clouds in remote sensing [7] and solar imaging, [5], or dust clouds in deep sky imaging [16]. Multi-scale (MS) CNNs were proposed to increase global context, e.g. [25], though context in convolutions remains limited to the final convolutional layers, and achieving larger receptive fields requires downscaling.

Attention proposes to model long range dependencies between feature positions and channels [22, 6]. While attention has been effective in capturing global context, it has a huge computational footprint. Positional attention has squared complexity in relation to image size, which is barely manageable on popular datasets with modern GPUs. This cost of attention can be reduced by downscaling features. Sacrificing texture for gained context is not a worthwhile compromise for some vision tasks, as severe downscaling significantly erodes local textures often to the detriment of model performance.



**Fig. 1:** Cirrus dust of various strengths (top), with uncertain annotations (middle) and predictions (bottom).

Orientational information is also a valuable discriminator in identifying classes of objects. Textures are intrinsically composed of orientation dependent patterns, and thus understanding of orientations has been shown to increase performance on a variety of segmentation tasks e.g. [16].

In this paper, we investigate segmentation of cirrus contamination: dust clouds in the foreground of astronomical images. Cirrus pollution can be difficult to spot, ranging from a slight change in background intensity to total occlusion of galaxies, as shown in Fig. 1. This makes cirrus localisation challenging, leading to frequent disagreements in expert annotations. Background intensity levels vary even in clean regions of images, thus it is necessary to consider the entire image ( $>5000 \text{ px}^2$ ) to maximise global contextual information. Local textural patterns are also necessary discriminating properties of cirrus, which presents as a wispy texture often with filamentary structures sharing a common orientation.

We propose a gridded MS architecture (Section 2.2) that addresses the attention’s high compute. Furthermore, it introduces an MS aspect with no added cost, while MS tends to be computationally expensive. We divide features of different scales into tiles of smaller constant size. Positional and channel attention is computed on these tiles to assess both local and global context in an efficient manner, before reassembling tiles into a final attention map. A closely related work

is [18] where attention is also applied to each scale of MS features, but using whole feature maps in each attention module, resulting in very high computing costs. Additionally, we present a novel attention operator using orientation (Section 2.3), for improved sensitivity to textures. We utilise Gabor modulated convolutions [15, 16] to generate features dependent on different angles. Attention is then computed across these angles, measuring correlations between orientation dependent features. Finally, we propose a loss function to train on annotations that lack a consensus (Section 2.4), that are a particularity of fuzzy and partially transparent contaminants.

## 2. METHODOLOGY

We present enhancements of attention modules for computationally efficient segmentation of large contaminants. These enhancements are compatible with various attention modules (and backbones), and we demonstrate them on [6].

### 2.1. Background on Multi-scale Attention

Accurate identification of contaminants requires comparison to surrounding regions. Attention [22] has been proposed to analyse contextual relations in images through computing feature correlations. Correlations are scaled based on the strength of given features, prioritising correlations that have a larger effect on model classification. This process can be performed along any internal organisation (axis) of a feature map, allowing different contextual dependencies to be measured. Thus, attention can be computed with respect to position (correlations between regions) and/or with respect to channels (correlations between learned features) [6]. We utilise multiple attention operators at different spatial scales to enhance local and global contextual understanding.

An approach for generating MS features is to pool initial features into different sizes and pass them through parallel convolutional layers e.g. [25]. Feature pyramid approaches [12] offer compute savings of constant factor by combining intermediate features at different scales generated by various convolutional layers. They have been used in attention works [1, 18], with attention applied to each scale apart. Tao et al. [21] generate features for each image scale separately, and observed a performance gain which outweighed computational cost, therefore we use a similar approach although our proposed attentions are compatible with MS features of [1, 18].

### 2.2. Multi-scale Gridded Attention

We propose a cost effective method for computing attention over multiple scales. Attention is costly, generally with squared complexity with respect to image size. While this effect can be managed through downscaling and cropping, the former compromises key local texture and the latter compromises key global context. It is therefore desirable to minimise




Fig. 2: Proposed MS gridded attention.

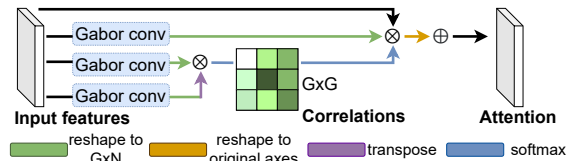


Fig. 3: Proposed Gabor attention operator.  $G$  is number of modulating Gabor filters,  $N$  is the product of other axes.  $\otimes$  denotes matrix multiplication, and  $\oplus$  element-wise addition.

the use of both while introducing a multi-scale analysis.

We illustrate gridded attention in Fig. 2. Our gridded attention module receives input MS feature maps, which we generate as in [21]. Each feature map is concatenated with a fused version of feature maps at all scales, similarly to [18], with the difference that the fused map is produced with same size as the largest feature map (before appropriate rescaling). We divide the feature maps into tiles of consistent smaller size but with multiple underlying spatial scales. Our architecture consists of multiple branches, each handling a spatial scale and comprising of a separate attention module, similar to the multi-branch architecture of [18]. In the example of Fig. 2, attention is computed on 21 tiles with size  $\frac{N}{4^2}$ , decreasing memory usage by  $\frac{21}{256}$ . Due to the massive saving on computational resources, the model can inspect both local texture and global context even on large images (e.g.  $>5000\text{px}^2$  in our case).

We use spatial scales with a common factor so that rescaling operations can be composed by chaining smaller rescalings. After computing attention on smaller scales, tiles are upscaled to realign resolutions. During upscaling, one upscaling convolutional block is used per scale transition, tying weights across scale branches, increasing parallelisation and inducing overlap between different scales. This is a key difference with other gridded attention methods [10] where lack of tile overlap can cause boundary artefacts.

### 2.3. Gabor Attention

We propose a novel attention operator for studying orientational context. Gabor modulated convolutions have been shown to increase performance on oriented textures and rotated samples e.g. [15, 16]. Such layers multiply convolutional weights by Gabor filters with different rotation parameters to generate orientation dependent features.

As illustrated in Fig. 3, we compute Gabor attention by placing independent Gabor modulations before each reshaping operation and calculate correlations across the axis representing orientations. This Gabor attention operator may be used in combination with channel and positional operators, creating an attention module with three separate attentions each measuring different dependencies among features. The three attention maps are combined with element-wise summation. We refer to this new attention module, combining positional, channel and Gabor attention, as tri-attention. By measuring correlations across the additional Gabor axis, we are able to study relationships among orientation rich features.

#### 2.4. Super-Majority Loss function

We propose a super-majority loss (SML) function to train on probabilistic annotations. In this study, targets are generated by four (expert and non-expert) annotators [19]. We create probabilistic targets using a weighted average, where expert annotators are weighted higher. We choose not to use approaches that propose probabilistic networks, such as [9, 23], due to our comparatively small number of expert annotators.

Inspired by works on edge detection with CNNs [14, 24], we separate consensus values into quartiles and conditionally adjust the loss function per quartile. By coarsely dividing probabilities into ranges of confidence, we mitigate against any uncertainty associated with our probabilistic consensus. We also use focal loss [13]  $L_f$  to encourage the model to focus on difficult examples rather than easy negatives or ‘clean’ pixels which dominate the class balance. Target pixels  $y$  with uncertain labels ( $0.25 < y < 0.5$ ) are ignored, whereas super-majority consensus ( $y \geq 0.75$ ) are prioritised by multiplying with a boosting coefficient  $\beta = 1.25$ . For output  $x$ :

$$L_{\text{SML}} = \begin{cases} \beta \cdot L_f(x, y) & \text{if } y \geq 0.75. \\ L_f(x, y) & \text{if } 0.5 \leq y < 0.75. \\ 0 & \text{if } 0.25 < y < 0.5. \\ L_f(x, y) & \text{otherwise.} \end{cases} \quad (1)$$

#### 2.5. Final Segmentation Model

We use the proposed MS gridded and Gabor attention mechanisms in combination, by using the Gabor tri-attention module in the MS gridded framework. We implement an MS attention with three spatial scales. The original feature maps and resulting attention maps are then passed through a fully connected layer to generate segmentation predictions.

As in [18], we generate six segmentation predictions. During training, we consider each segmentation separately and compute six loss terms for all predictions before summing them. The backbone is then explicitly forced to preserve spatial locations of features, relieving the attention module of any realigning effort. For inference, we take the average of the three segmentation computed from the attention maps.

### 3. RESULTS

#### 3.1. Data

We produced a dataset of 186 images from the MATLAS survey [4] (avg. image size 7000 px<sup>2</sup>). 25% of images have cirrus contamination, and in contaminated images, 60% of pixels contain cirrus. We use 70% of samples for training, 15% for validation and testing. We also run experiments on the subset of contaminated images, and on an additional dataset of 300 synthesised cirrus images (size 512 px<sup>2</sup>) generated similarly to [16] with some added simulated imaging artefacts for greater realism. In these two additional datasets, class imbalance is eased, allowing the model to focus further on discriminative cirrus features and thus more clearly revealing findings during ablation studies. We augment images with geometric transformations: flips, rotations and translations, and apply element-wise Gaussian noise with zero mean and variance of 0.1. We also pretrain on synthesised cirrus images.

#### 3.2. Comparative Analyses on Proposed Techniques

We verify the effectiveness of the proposed methods individually. For control model, we use non-gridded MS dual attention similar to [18] in our MS architecture, though with a simple backbone with four convolutional layers. All networks are trained for 200 epochs with the Adam optimiser [11] with learning rate  $10^{-3}$  and L2 weight regularisation  $10^{-7}$ . Learning rate is decreased every epoch by a factor of 0.98. We score models using the intersection over union (IoU) averaged over five splits, and also report standard error across splits. We fix batch size and overall parameter size across modifications to be roughly equal and downscale images to the maximum size our GPUs could accommodate for each network.

We perform an ablation study on proposed attention techniques, shown in Table 1. The effectiveness of the proposed loss is first compared against a focal loss with rounded probabilistic labels 1. On our dataset of real images containing probabilistic annotations, ignoring pixels with uncertain annotation and prioritising very certain pixels is helpful, with the proposed super-majority loss increasing performance. On the synthetic dataset, the probabilities of simulated labels do not correlate well with actual prediction difficulty, and the

	All images	Cirrus only	Synth.
Control	0.469(0.005)	0.814(0.013)	0.831(0.001)
SML	0.483(0.006)	0.822(0.013)	0.830(0.001)
SML+Gabor	0.497(0.007)	0.861(0.015)	0.867(0.002)
SML+Grid	0.542(0.003)	0.886(0.016)	0.844(0.003)
All	<b>0.548</b> (0.003)	<b>0.892</b> (0.013)	0.871(0.002)
U-Net [17]	0.381(0.128)	0.685(0.096)	0.794(0.001)
LGCN [16]	0.414(0.049)	0.741(0.033)	<b>0.891</b> (0.001)

**Table 1:** Segmentation IoU in format mean(std).

	Single	Ensemble
IoU	0.745	<b>0.790</b>
Dice	0.766	<b>0.814</b>

**Table 2:** Proposed cirrus segmentation results

	Ours	Dev et al. [2]	Dev et al. [3]	Song et al. [20]
IoU	<b>0.90</b>	0.69	0.80	0.86
Dice	<b>0.95</b>	0.82	0.89	0.92

**Table 3:** Segmentation scores on the SWIMSEG sky/cloud segmentation dataset.

proposed loss was not able to provide an improvement. We see that the control MS attention model offers a sizeable benefit on real data, indicating that dual attention is well suited to large contaminant segmentation. The proposed MS gridded attention significantly compounds this positive effect, showing that the ability to study long range dependencies across multiple scales is beneficial. The proposed tri-attention increases accuracy on all cirrus segmentation scenarios, indicating that sensitivity to texture orientation is helpful in localising the fuzzy and textured contaminants.

The two attention techniques benefit from being combined, with further improved results. On real images, when compared against a popular U-Net [17] baseline and the state-of-the-art cirrus segmentation method LGCN [16], it outperforms both by a large margin. On synthetic images, LGCN outperforms attention models due to no long range correlations existing within the small images.

### 3.3. Proposed Cirrus Segmentation

We construct a final model from the best performing components of the previous section and analyse the predictions. We significantly increase the parameter size of our segmentation model and swap the simple backbone network out for a ResNet-50 [8] network. To account for this larger network, we increase training epochs to 400. We also increase the size of both the training and testing sets by dividing the validation samples between them. Cirrus dust is faint in comparison to bright stars. To compensate for this we add an initial layer to the network implementing arcsinh scaling, popular in astronomical image processing:  $X_s = \text{arcsinh}(aX + b)$ , where  $a, b$  are learned. We also use ensemble predictions, where 5 models are trained, and predictions from each model are averaged over to give the final segmentation. Inference takes an average time of  $\sim 0.4s$  per sample on a single GTX 1080 Ti, which could be significantly reduced with optimisation efforts such as model compilation, pruning or half precision weights.

In Table 2 we report both IoU and Dice scores, with Dice being more biased to precision. We observe that there is a very large performance increase from models in the compar-

ative analysis, owing to the larger parameter space, additional training samples and longer training period. We see that a single model is outperformed by the model ensemble by a significant margin. Aggregating predictions also appears to handle the class imbalance issue better than the single model, when comparing the predicted distributions of cloud coverage across images against target groundtruth distribution, with Kullback-Leibler divergences being respectively 0.40 and 0.07 for the single and ensemble models. Prediction aggregation appears to also increase the gap between IoU and Dice scores, indicating that precision is increased through averaging over predictions. This may be due to the aggregation smoothing over-eager positive predictions from single models within the ensemble. Fig. 1 shows some examples of predictions by the best performing ensemble model.

### 3.4. Cloud segmentation in natural images

We evaluate the proposed methodology on cloud segmentation in natural images from the Singapore Whole sky IMaging SEGmentation database [2] (SWIMSEG). This task is highly relevant to the proposed methodology: both local and global features are key to good performance on this dataset, as difficult positive regions of light cloud can only be correctly identified based on subtle textural patterns and comparison with surrounding regions. SWIMSEG contains 1013  $600px^2$  images of sky patches and corresponding cloud segmentations. Training, validation and testing sets each contain 861, 101 and 51 samples, respectively. The model and training setup used is identical to those of Section 3.3. Results are detailed in Table 3, where we report the segmentation IoU and Dice score on the testing set as in [2, 20, 3]. The tri-attention model achieves respective IoU and Dice scores of 0.90 and 0.95, representing a significant improvement over state of the art.

## 4. CONCLUSION

We presented a computationally efficient multi-scale attention architecture that is sensitive to texture orientation for segmentation of global contaminants in large images. Efficiency is achieved through a gridded architecture, and orientation sensitivity is provided by a new Gabor attention operator. To address the challenge of uncertain groundtruth labelling, we proposed a simple consensus loss. These contributions are combined into a new state-of-the-art model for the segmentation of cirrus contaminants from astronomical images. Reliable performance was achieved with only a small training dataset of images. Our model can process an entire image in one pass with minimal downscaling, meaning that the proposed method can be easily integrated into data processing pipelines for imaging instruments to obtain contamination masks. In future work it would be interesting, using the methodologies presented in this work, to craft a deep generative model capable of removing global contaminants.

## References

- [1] Psanet: Point-wise spatial attention network for scene parsing. In *ECCV*, pages 267–283, 2018.
- [2] Soumyabrata Dev, Yee Hui Lee, and Stefan Winkler. Color-Based Segmentation of Sky/Cloud Images From Ground-Based Cameras. *IEEE J-STARS*, 10(1):231–242, 2017.
- [3] Soumyabrata Dev, Atul Nautiyal, Yee Hui Lee, and Stefan Winkler. CloudSegNet: A deep network for nychthemeron cloud image segmentation. *IEEE Geoscience and Remote Sensing Letters*, 16(12):1814–1818, 2019.
- [4] Pierre-Alain Duc. MATLAS: a deep exploration of the surroundings of massive early-type galaxies. *arXiv preprint arXiv:2007.13874*, 2020.
- [5] Cong Feng and Jie Zhang. SolarNet: A sky image-based deep convolutional neural network for intra-hour solar forecasting. *Solar Energy*, 204:71–78, 2020.
- [6] Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. In *CVPR*, pages 3146–3154, 2019.
- [7] Jianhua Guo, Jingyu Yang, Huanjing Yue, Hai Tan, Chunping Hou, and Kun Li. CDnetV2: CNN-based cloud detection for remote sensing imagery with cloud-snow coexistence. *IEEE Transactions on Geoscience and Remote Sensing*, 59(1):700–713, 2020.
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [9] Shi Hu, Daniel Worrall, Stefan Knekt, Bas Veeling, Henkjan Huisman, and Max Welling. Supervised uncertainty quantification for segmentation with multiple annotations. In *MICCAI*, pages 137–145, 2019.
- [10] Yifan Jiang, Shiyu Chang, and Zhangyang Wang. Transgan: Two pure transformers can make one strong gan, and that can scale up. *Advances in Neural Information Processing Systems*, 34, 2021.
- [11] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2014.
- [12] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *ICPR*, pages 2117–2125, 2017.
- [13] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal Loss for Dense Object Detection. In *ICCV*, pages 2999–3007, 2017.
- [14] Yun Liu, Ming-Ming Cheng, Xiaowei Hu, Kai Wang, and Xiang Bai. Richer convolutional features for edge detection. In *CVPR*, pages 3000–3009, 2017.
- [15] Shangzhen Luan, Chen Chen, Baochang Zhang, Jungong Han, and Jianzhuang Liu. Gabor convolutional networks. *IEEE Transactions on Image Processing*, 27(9):4357–4366, 2018.
- [16] Felix Richards, Adeline Paiement, Xianghua Xie, Elisabeth Sola, and Pierre-Alain Duc. Learnable Gabor modulated complex-valued networks for orientation robustness. *arXiv preprint arXiv:2011.11734*, 2020.
- [17] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, pages 234–241, 2015.
- [18] Ashish Sinha and Jose Dolz. Multi-scale self-guided attention for medical image segmentation. *IEEE J. Biomed. Health Inform.*, 25(1):121–130, 2020.
- [19] Elisabeth Sola, Pierre-Alain Duc, Felix Richards, Adeline Paiement, Mathias Urbano, Julie Klehammer, Michal Bílek, Jean-Charles Cuillandre, Stephen Gwyn, and Alan McConnachie. Characterization of low surface brightness structures in annotated deep images. *Astronomy & Astrophysics*, 662:A124, 2022.
- [20] Qianqian Song, Zhihui Cui, and Pu Liu. An Efficient Solution for Semantic Segmentation of Three Ground-based Cloud Datasets. *Earth and Space Science*, 7(4), 2020.
- [21] Andrew Tao, Karan Sapra, and Bryan Catanzaro. Hierarchical multi-scale attention for semantic segmentation. *arXiv preprint arXiv:2005.10821*, 2020.
- [22] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, pages 5998–6008, 2017.
- [23] Mike Walmsley, Lewis Smith, Chris Lintott, Yarin Gal, Steven Bamford, Hugh Dickinson, Lucy Fortson, Sandor Kruk, Karen Masters, Claudia Scarlata, and others. Galaxy Zoo: probabilistic morphology through Bayesian CNNs and active learning. *MNRAS*, 491(2): 1554–1574, 2020.
- [24] Saining Xie and Zhuowen Tu. Holistically-nested edge detection. In *ICCV*, pages 1395–1403, 2015.
- [25] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *CVPR*, pages 2881–2890, 2017.