Glacial cirque identification based on Convolutional Neural Networks

Dongxue Mao[a], Yingkui Li[b*], Qiang Liu[a], Iestyn D. Barr[c], Ian S. Evans[d]

[a] School of Geographic Sciences, Hebei Normal University, Shijiazhuang 050024, China

[b] Department of Geography & Sustainability, University of Tennessee, Knoxville, USA

[c] Department of Natural Sciences, Manchester Metropolitan University, Manchester, UK

[d] Department of Geography, University of Durham, Durham, UK, DH1 3LE, UK

*Corresponding author: yli32@utk.edu (Y. Li).

Abstract

Cirques provide important information about the palaeoclimate conditions that produced past glaciers. However, mapping cirques is challenging, time-consuming, and subjective due to their fuzzy boundaries. A recent study tested the potential of using a deep learning algorithm, Convolutional Neural Networks (CNN), to predict boundary boxes containing cirques. Based on a similar CNN method, RetinaNet, we use a dataset of > 8,000 cirques worldwide and various combinations of digital elevation models and their derivatives to detect these features. We also incorporate the Convolutional Block Attention Module (CBAM) into RetinaNet for training and prediction. The precision of cirque detection with or without the addition of the CBAM is evaluated for various input data combinations, and training sample sizes, based on comparison with mapped cirques in two test areas on the Kamchatka Peninsula and the Gangdise Mountains. The results show that the addition of CBAM increases the average precision by 4-5% ($p < 0.01$), and the trained model can detect the cirque boundary boxes with high precision (84.7% and 87.0%), recall (94.7% and 86.6%), and F_1 score (0.89 and 0.87), for the two test areas, respectively. The inclusion

of CBAM also significantly reduces the number of undetected cirques. The model performance is affected by the quantity and quality of the training samples: the performance generally increases with increasing training samples and a training dataset of 6000 cirques produces the best results. The trained model can effectively detect boundary boxes that contain cirques to help facilitate subsequent cirque outline extraction and morphological analysis.

Keywords: Cirques, Object detection, RetinaNet, CBAM, Deep learning

1. Introduction

Cirques are typical landforms formed by mountain glaciers through erosion at their base (Zhang et al., 2008; Barr and Spagnolo, 2015). They usually have an armchair shape with a gentle, flat or overdeepened floor surrounded by steep headwall and sidewalls, and a convex break of slope that demarcates their lower boundary and creates a separation from the valley below (Evans and Cox, 1974). Cirques are generally believed to form through the rotational flow and basal sliding of relatively small and isolated glaciers during the onset and receding phases of glaciations when glaciers are confined just to the highest areas (Gardner, 1987; Sanders et al., 2012; Evans, 2021). In addition, cirque headwalls are eroded and weathered by paraglacial and periglacial processes both during glacier occupation and following deglaciation (Gardner, 1987; Crest et al., 2017; Jarman and Harrison, 2019).

The presence of cirques has been used as an indicator that a region was formerly occupied by glaciers. The shape and distribution of cirques have also been used to estimate palaeo climate and environmental conditions (Nelson and Jackson, 2003; Principato and Lee, 2014; Barr and Spagnolo, 2015; Li et al., 2023; Pellitero et al., 2024). For example, the morphology and distribution of cirques are related to the intensity and duration of glacial erosion (Barr and Spagnolo, 2013; Bathrellos et al., 2014; Evans, 2006). Cirque floor altitudes have been used as an indicator of palaeo equilibrium line altitudes (ELAs) (Hassinen, 1998; Mitchell and Humphries, 2015). Cirque aspect and hypsometry have been used to assess the interaction between glacial erosion, topography, tectonics, and climate (Anders et al., 2010; Mîndrescu and Evans, 2014).

A reasonable assessment of palaeo climate and environmental information usually requires the analysis of a large population of cirques (Barr and Spagnolo, 2015; Zhang et al., 2020; Barr et al., 2023; Li et al., 2023; Pellitero et al., 2024). Thus, mapping cirque outlines and deriving cirque metrics are of critical importance. Spagnolo et al. (2017) developed an ArcGIS toolbox, ACME, to derive 16 cirque metrics, length, width, circularity, planar and three-dimensional (3D) area, elevation, slope, aspect, plan closure, and hypsometry based on three inputs: cirque outlines, a digital elevation model (DEM), and cirque threshold midpoints. Li et al. (2024) updated this tool to ACME2, which provides new functions to automatically derive cirque foci points and expands the list of metrics to 49. These tools permit analyses of large numbers of cirques where cirque outlines are available. In contrast, cirque outlines have been delineated mainly based on manual digitization from topographic maps, aerial photographs, satellite images, and DEMs (Cui, 1981; Seif and Ebrahimi, 2014; Barr et al., 2017, 2019), which is time-consuming, labor-intensive, and subjective (Li and Zhao, 2022).

In recent years, some automated and semi-automated approaches have been developed to delineate cirque outlines. For example, a classification model was developed by Eisank et al. (2010) to segment cirques based on curvature and spatial context. Li and Zhao (2022) developed an ArcGIS toolbox, AutoCirque, to automatically delineate cirque outlines from DEMs. However, the automatically delineated set of outlines may also include some depressions formed by non-glacial processes, such as landslides and karst depressions. In addition, AutoCirque is computationally expensive because its algorithm searches the source area of each stream to determine the location of probable cirques.

With the development of artificial intelligence (AI), its techniques have started to be applied to detecting geomorphic features. For example, Nagle-McNaughton et al. (2020) used neural networks to detect the surface features of Mars, and Gupta et al. (2020) detected pavement potholes based on deep neural networks. However, the use of AI techniques in glacial geomorphology is still in its initial stage. A pilot study was conducted by Scuderi and Nagle-McNaughton (2022) to identify the boundary boxes of cirques using a convolutional neural network (CNN), RetinaNet, based on a dataset of 1,951 cirques. Williams et al. (2023) applied the same model to compare cirque morphological characteristics between Earth and Mars. These studies demonstrated the potential of using RetinaNet to identify cirques. However, the impacts of input data combination, training sample size, and the effect of neural network structure on cirque identification are still unclear.

In this paper, we assess the performance of cirque identification from different input data combinations, training sample sizes, and with or without the Convolutional Block Attention Module (CBAM) to RetinaNet. We use > 8,000 cirques from High Mountain Asia (Zhang et al., 2020; Li et al., 2023), Kamchatka Peninsula (Barr and Spagnolo, 2013), and Britain and Ireland (Clark et al., 2018) to train the model. The performance is assessed based on two test areas on the Kamchatka Peninsula and the Gangdise Mountains, respectively.

2. Cirque datasets, test Areas, and DEMs

We compile a dataset of 8,207 manually digitized cirque outlines for model training, including 2,831 in High Mountain Asia (Zhang et al., 2020; Li et al., 2023), 3,168 on the Kamchatka Peninsula (Barr and Spagnolo, 2013), and 2,208 from Britain and Ireland (Clark et al., 2018). Most digitized outlines are simple cirques except 98 composite cirques in Britain and Ireland. Table 1 briefly describes the morphometric characteristics of the cirques in the dataset. The cirques in High Mountain Asia are relatively large with an average area of 0.88 km^2, a mean length of 1027 m, and a mean width of 944 m. The cirque length is slightly larger than the width (L/W > 1). The cirques in the Kamchatka Peninsula are slightly smaller with an average area of 0.74 km^2, a mean length of 876 m, and a mean width of 997 m. The cirque width is slightly larger than the length (L/W < 1). The cirques in Britain and Ireland are the smallest with an average area of 0.61 km^2 and similar mean length and width of 775 m and 787 m, respectively (L/W = 1). The circularities of the cirques are similar in the three areas, ranging from 1.05 to 1.08.

Table 1. Morphological parameters of the cirque dataset used for this study

	Parameters	Max	Min	Mean	Median	SD*
High Mountain Asia	L(m)	3341	283	1027	960	385
	W(m)	2528	259	944	896	334
	L/W	2.91	0.39	1.12	1.08	0.31
	Circularity	1.33	1.01	1.08	1.07	0.04
	Area (km^2)	7.03	0.06	0.88	0.71	0.63
Kamchatka Peninsula	L(m)	2110	125	876	850	251
	W(m)	2601	250	997	958	313
	L/W	2.04	0.50	0.91	0.88	0.21
	Circularity	1.29	1.00	1.05	1.04	0.04
	Area (km^2)	4.02	0.03	0.74	0.64	0.43
Britain and Ireland	L(m)	4629	163	775	667	433
	W(m)	4213	156	787	705	370
	L/W	2.50	0.34	1.00	0.97	0.28
	Circularity	1.50	1.01	1.08	1.07	0.05
	Area (km^2)	13.57	0.03	0.61	0.39	0.77

*SD = Standard deviation

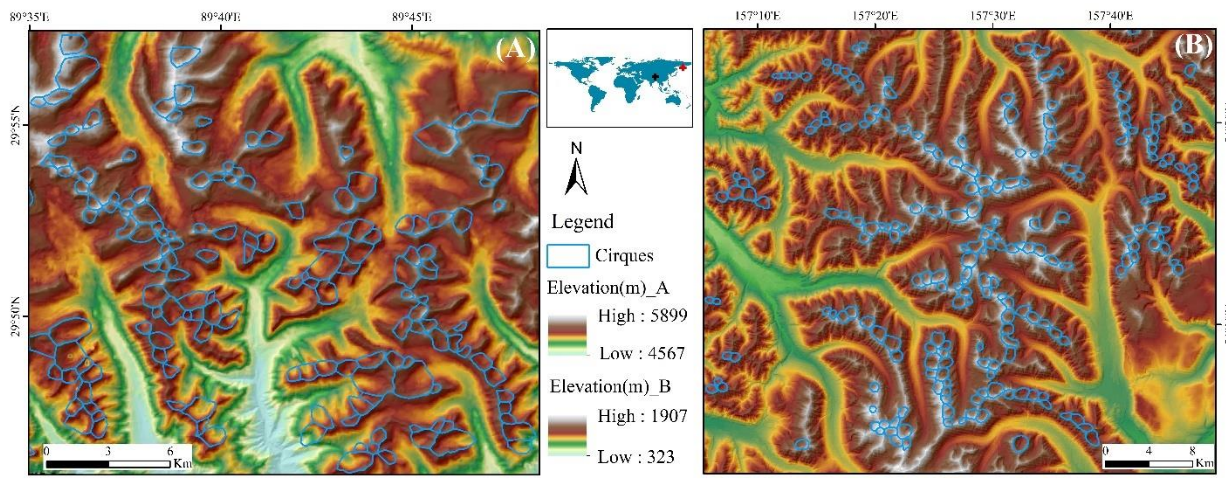

Fig. 1. Maps of the Gangdise Mountains (A) and Kamchatka Peninsula (B) test areas.

We select two sub-regions in the Gangdise Mountains (Tibetan Plateau) and Kamchatka Peninsula to evaluate the performance of the cirque detection (Fig. 1). These two test areas are from different latitude-longitude zones with different glacier types and histories. The cirques were also mapped by different scholars. These two test areas allow for the assessment of the model performance for different glacier types and geographic settings.

The Gangdise Mountains on the southern Tibetan Plateau form an important drainage divide between rivers of inland drainage and those draining into the ocean (Liu et al., 2016; Zhang et al., 2018). They are primarily composed of intensely folded Cretaceous and Jurassic formations, along with extensive intrusions of intermediate and acidic igneous rocks and mixed rocks. The glaciers in the Gangdise Mountains are affected by tectonic uplift and the shadow effect of the Himalayas. The test area (29°46' - 29°57' N, 89°35' - 89°48' E) is in the eastern part of the Gangdise Mountains, ranging from 4567 to 5899 m above sea level (asl) (Fig. 1A). A total of 148 cirques were manually delineated in this area (Zhang et al., 2020).

The second test area (54°3' - 54°24' N, 157°5' - 157°49' E) is in the southwest part of the Kamchatka Peninsula (Fig. 1B) and ranges from 323 to 1907 m asl. The higher part of the Kamchatka Peninsula is covered by glaciers, affected by active volcanic activity and tectonic movements. The glaciers in the Kamchatka Peninsula are close to oceans and do not show a strong precipitation shadow effect. Pleistocene glaciation was the main factor shaping the landscape (Jones and Solomina, 2015; Solomina et al., 2007). A total of 214 cirques were manually delineated in this area (Barr and Spagnolo, 2013).

We use the 30-m Copernicus DEM (COPDEM30) for cirque detection. COPDEM30 was released by the European Space Agency (ESA) and Airbus in 2021. It is a Digital Surface Model (DSM) that represents the Earth's surface, including buildings, infrastructure, and vegetation, with a resolution of 1 arc-second (approximately 30 meters) (Hawker et al., 2022). COPDEM30, with its underlying data from TanDEM-X, is the latest and most accurate global DEM available (Hawker et al., 2022; Li et al., 2022) and has been considered as the gold standard for a global DEM (Ernest et al., 2023; Peter et al., 2023; Trevisani et al., 2023). COPDEM30 is freely available from the OpenTopography website (https://portal.opentopography.org/datasets).

3. Methods

3.1 RetinaNet

We use the RetinaNet model for cirque boundary box detection. RetinaNet is a one-stage object detection model developed by Lin et al. (2017), which uses a focal loss function to address class imbalance during training. As one of the most popular object detection models, RetinaNet is implemented by a single unified network, which is composed of a backbone network and two sub-networks, for object detection. Built upon a ResNet architecture, the backbone network computes the convolutional feature maps for the entire input image. Specifically, a Feature Pyramid Net (FPN) is constructed as a multi-scale feature pyramid across the backbone network, facilitating multi-scale object detection. RetinaNet also uses two sub-networks dedicated to specific tasks: the first is a class sub-network to conduct convolutional object classification on the output of the backbone network; and the second is a location regression sub-network (box subnet) to perform convolutional bounding box regression and prediction (Lin et al., 2017; Fig. 2). Furthermore, RetinaNet introduces a focal loss function (Gupta et al., 2020; Lin et al., 2020) to reduce the loss of easy-to-categorize samples, making the network focus more on difficult and misclassified samples (Gupta et al., 2020; Huang et al., 2020; Lin et al., 2020).

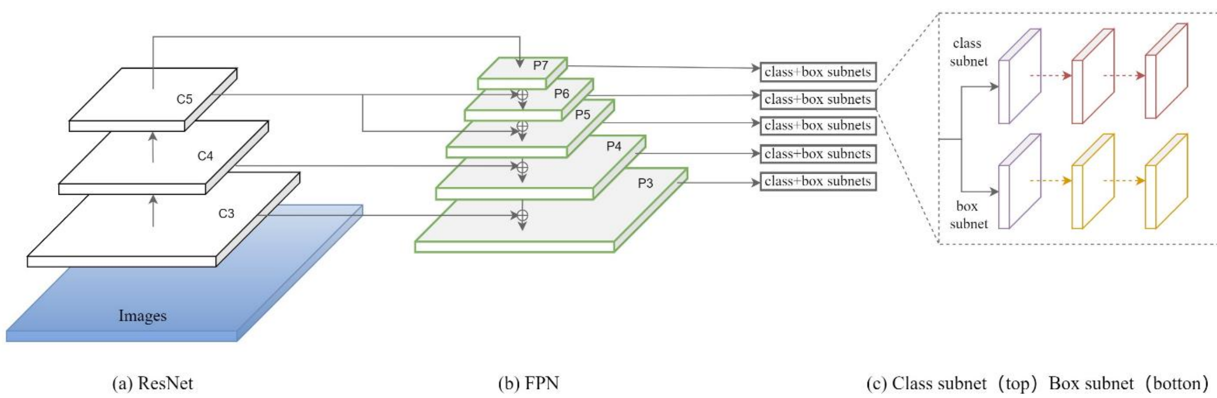

Fig. 2. The algorithm structure of RetinaNet (Lin et al., 2017)

3.2 Attention mechanism

Convolutional Block Attention Module (CBAM) lets the network devote more computations to the important parts and suppress the unimportant information (Woo et al., 2018; Cheng and Yu, 2021). CBAM is a combination of spatial and channel attention modules. The spatial attention module focuses on spatial positions to find the most important regions in the network for processing. The channel attention module focuses on which channels are most important, locating the most meaningful feature layers and deriving the attention weights (Cheng and Yu, 2021; Lu and Hu, 2022). As illustrated in Fig. 3, we add the CBAM in RetinaNet at the feature layer (C3 and C4) of the network and before the class and location regression sub-networks (P6 and P7), to enhance feature extraction (Cheng and Yu, 2021).

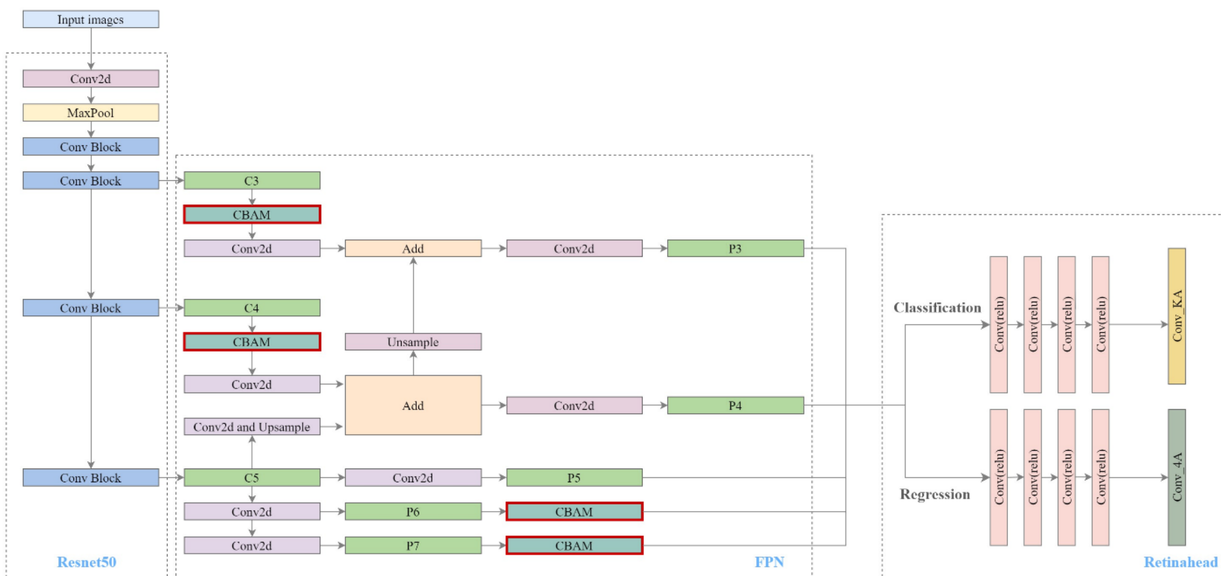

Fig. 3. CBAM-FPN structure diagram with the addition of CBAM. The structure diagram is modified from (https://github.com/bubbliiiing/retinanet-keras/).

3.3 Experimental design

RetinaNet requires input data with a three-band composite format for object detection. For photos and images, these three bands are commonly the RGB combination. For DEMs, the three bands can be elevations or combinations of elevation and its derived layers. We use the elevation data from COPDEM30 and its derived aspect, slope, aspect-slope, and hillshade (viewed from 315° azimuth at a 45° elevation angle based on the default values in ArcGIS Pro) layers to composite the three band combinations for RetinaNet. All derived layers are from the original DEM and can be generated in ArcGIS Pro. These derivatives may be captured within the processes of the convolutional neural networks. Because RetinaNet requires a three-band format as input and the DEM is only treated as one band, our purpose is to check whether adding some of these derivatives as band combinations affects the model performance. Note that hillshade is a commonly used derivative for the grayscale 3D representation of the terrain surface, but it is directional and dependent on the default azimuth and elevation angle values, which may introduce some bias in object detection. In comparison, an aspect-slope (AS) layer simultaneously illustrates the aspect and slope of a terrain surface, which is useful for identifying ridges and valleys (https://www.esri.com/about/newsroom/arcwatch/create-an-aspect-slope-map-quickly-and-easily/). In this study, we test five combinations of bands (Fig. 4): elevation-elevation-elevation (E-E-E), slope-aspect-elevation (S-A-E), AS-aspect-elevation (AS-E-A), AS-hillshade-elevation (AS-H-E), and hillshade-hillshade-hillshade (H-H-H). Each combination produces an image of three 8-bit depth bands for the RetinaNet model.

The 8,207 previously mapped cirque outlines are processed to generate training labels using the 'Export Training Data For Deep Learning' tool in ArcGIS Pro. A 100-m buffer is applied to the cirque outlines to include the immediate surrounds of cirques. The 'Export Training Data For Deep Learning' tool generates a set of 256 x 256 JPEG files and XML label annotations in the Pascal Visual Object Class format. In total, 8,336 JPEG and label annotation files were generated for each of the five combinations.

The H-H-H, AS-H-E, S-A-E, E-E-E, and AS-E-A datasets were trained in RetinaNet (https://github.com/bubbliiiing/retinanet-keras/) separately using the backbone network ResNet50. The images are enhanced by flipping, rotation, and color-gamut transformation to improve training efficiency and help model generalization. The dataset is randomly divided into training (60%), validation (20%), and testing sets (20%). During the training process, the model adjusts its parameters based on the training dataset. After each training epoch, the model is tested on the validation dataset to fine-tune the model and improve its generalization. The model is pretrained with ResNet50 weights (retinanet_resnet50.pth), which are obtained using a variety of datasets to keep the weights universal for different datasets. The use of these weights is necessary for most cases because not using weights for the backbone would result in excessive randomness, leading to unclear feature extraction and poor training results for the network. We chose a "freeze" backbone to limit the number of adjustable parameters and preserve some of the original features of the initial model. The H-H-H, AS-H-E, S-A-E, E-E-E, and AS-E-A datasets were also trained in RetinaNet with the addition of CBAM (CBAM_RetinaNet). The rest of the training processes are the same as the above training using RetinaNet.

The E-E-E of elevation data was randomly divided into seven subsets of various numbers of cirques (2000, 3000, 4000, 5000, 6000, 7000, and 8000). These subsets were trained in CBAM_RetinaNet, respectively, to examine the impact of training sample size on model performance.

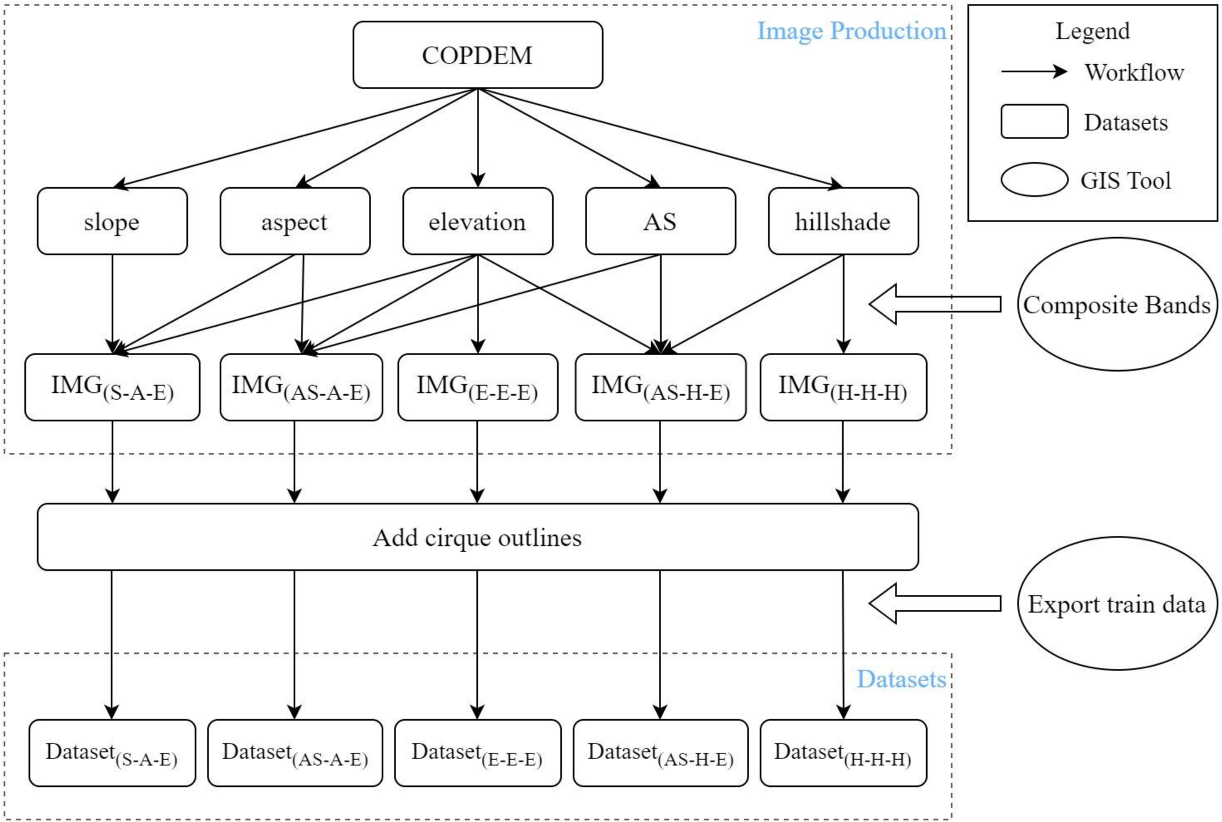

Fig. 4. The flowchart for preparing the datasets for RetinaNet and CBAM_RetinaNet.

The training parameters are set to "freeze" for the first 50 epochs before training, the learning rate is set to 0.001, the batch size is 4, and the model is "unfrozen" after the initial convergence, and trained until the training and validation losses converge to ensure that the model is equally well suited to the training data and unknown (validation) data. We then adjust the learning rate to 0.0001 and conduct further model training with 100 epochs. The training is based on the Windows platform with Inter Core i7-12700H CPU and NVIDIA GTX 3060 graphics card, compiled in torch-1.13.0, torchvision-0.14.0, cuda11.0, cudnn11.0 and Python3.8.

A set of weight files are generated after training. We choose the ones with lower training and validation losses for prediction. The prediction is the same three-band image format as used in the training.

3.4 Performance assessment

The Precision (*P*), Recall (*R*), Average Precision (*AP*), and F_1 score (Eq. 1–4) are used to evaluate the model performance of the cirque detection based on the confusion matrix derived by the comparison between the predicted boundary boxes and the cirque outlines in the test sets (randomly selected 20% of the dataset). All these metrics are derived based on true positives (*TP*, where mapped cirques are correctly detected as cirques), false negatives (*FN*, where mapped cirques are undetected), and false positives (*FP*, where non-cirques are detected as cirques) (Sun et al., 2020).

$$P = \frac{TP}{TP+FP} \tag{1}$$

$$R = \frac{TP}{TP+FN} \tag{2}$$

$$AP = \int_0^1 P\ (R)\ dR \tag{3}$$

$$F_1 = \frac{2PR}{P+R} \tag{4}$$

Precision, *P*, describes the proportion of predicted positive that are actually positive samples (Huang et al., 2020; Tong et al., 2020). It is a measure of the reliability of the network (e.g., 80 out of 100 detections are correct, $P = 0.80$). Recall, *R*, describes the proportion of true positive samples in the sample dataset that are correctly predicted (e.g., 13 out of 15 cirques are detected, R=0.87). In the best case, both *P* and *R* values are 1.00, indicating that the model detects all targets (Recall) and that all detections are correct (Precision). The F_1 score is a measure of the harmonic mean of recall and precision. It gives equal weight to precision and recall, so it is the best single metric of model performance where false positives and false negatives are equally important, for a specific confidence threshold. The overall performance for a range of confidence thresholds can be evaluated graphically by a precision-recall curve, which provides the model performance across many thresholds, rather than a specific value. The area under the precision-recall curve, commonly named as the Average Precision (*AP*), describes the overall stability of the recall and precision for a range of confidence thresholds(Huang et al., 2020; Sun et al., 2020; Tong et al., 2020). *AP* is high when both precision and recall are high, and low when either of them is low for a range of confidence thresholds.

The model produces three outputs for each input data combination: a label (the prediction class), a score (confidence), and the boundary box predicted to contain a target (cirque). A minimum confidence threshold and an intersection/union value (IOU), calculated as a ratio between the intersection and union of the predicted boundary box and its corresponding boundary box of the ground truth (mapped cirque) (Fig. 5), are needed to evaluate the model performance(Huang et al., 2020). A perfect detection should have an IOU of 1.0 (the predicted boundary box is the same as the boundary box of the mapped cirque), but 0.5 leaves room for variation. We set IOU to 0.5: IOU < 0.5 is considered as an incorrect detection, while IOU > 0.5 as a correct detection (Fig.5). For example, a confidence threshold of 0.4 and an IOU of 0.5 consider a detection where the network is 40% confidence and the predicted boundary box overlaps 50% with the boundary box of a target (cirque) as a correct detection.

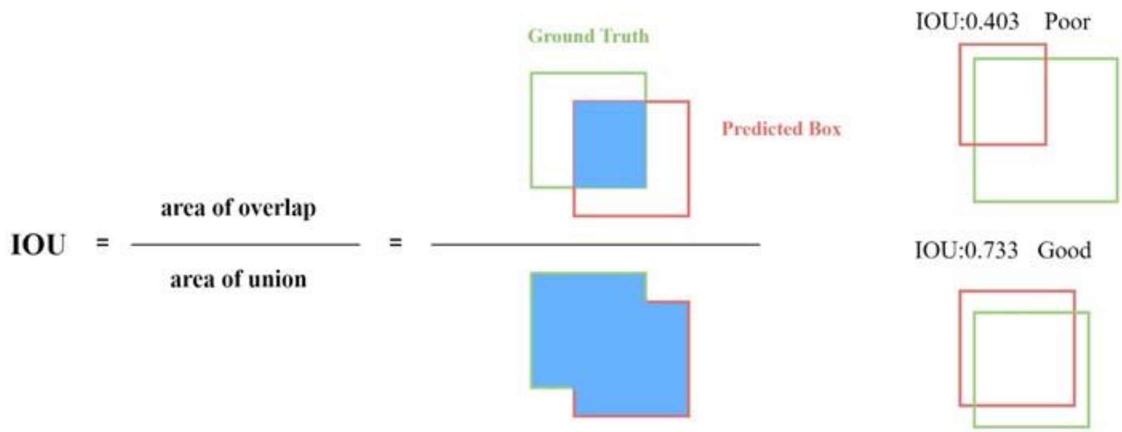

Fig. 5. A sketch to show the definition of IOU based on the intersection and union of the predicted boundary box and its corresponding boundary box of a target cirque.

4. Results

4.1 Overall model performance

Table 2 compares the overall performance of the RetinaNet and the CBAM_RetinaNet models for the AS-E-A, E-E-E, S-A-E, AS-H-E, and H-H-H combinations using the 8,207 cirques. For all five combinations, the addition of CBAM improves performance on every metric. Overall, the addition of CBAM improves the average precision by 4-5%. A paired t-test shows statistically significant differences ($p < 0.01$) between the metrics of RetinaNet and CBAM-RetinaNet (Table 3). F_1 score shows improvement, especially for the AS-E-A, H-H-H, and AS-H-E combinations. The precision-recall curve demonstrates an inverse relationship between precision and recall: as precision increases, recall decreases. When the confidence threshold is set higher, the predicted boundary boxes are more accurate. However, due to the higher threshold, fewer predictions meet the requirement, resulting in a lower recall. Conversely, when the threshold is lower, precision decreases while recall increases. The comparison of the precision-recall curves (Fig. 6) also indicates that CBAM_RetinaNet has better performance, with a larger under-curve area. From Table 2, AS-E-A is the best of the five combinations in CBAM_RetinaNet, with the highest AP, recall, precision, and F_1 score.

Table 2. The performance metrics of RetinaNet and CBAM-RetinaNet for the whole dataset of 8,207 cirques.

Input combination	Model	AP (%) (IOU50)	Recall* (%)	Precision* (%)	F_1 score*
AS-E-A	RetinaNet	65.81	48.82	85.50	0.63
	CBAM-RetinaNet	70.18	55.77	88.85	0.68
E-E-E	RetinaNet	62.08	50.80	80.91	0.62
	CBAM-RetinaNet	66.54	52.94	87.77	0.66
H-H-H	RetinaNet	63.84	45.15	83.54	0.59
	CBAM-RetinaNet	68.21	55.52	87.28	0.68
S-A-E	RetinaNet	62.74	50.46	82.01	0.62
	CBAM-RetinaNet	66.19	50.81	86.30	0.64
AS-H-E	RetinaNet	57.52	38.46	78.19	0.51
	CBAM-RetinaNet	60.86	41.98	82.15	0.55

*Precision, Recall, and *F1* score are derived for the confidence thresholds of > 0.2.

Table 3. Paired t-test between the metrics of RetinaNet and CBAM-RetinaNet.

Mean difference	Standard deviation	Difference at the 95% confidence level	t	p
-3.27	2.76	-4.56 ~ -1.98	-5.31	0.00

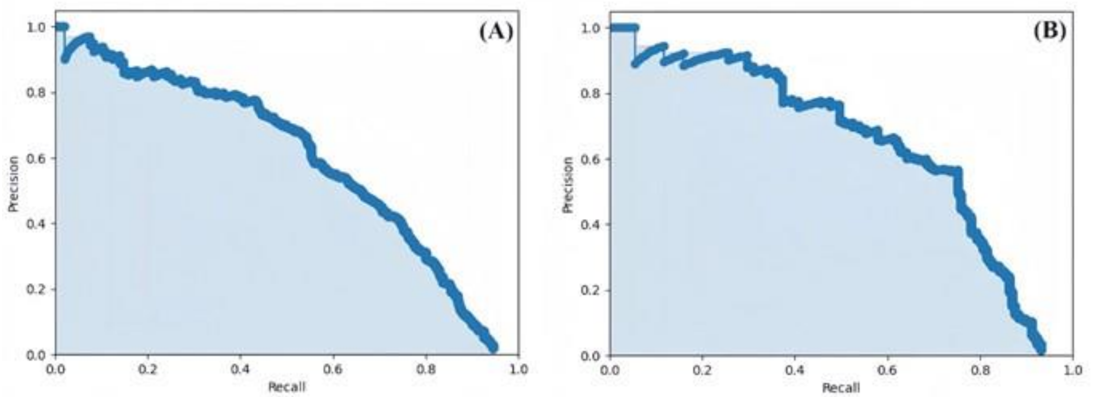

Fig. 6. The precision-recall curves (AS-E-A) of RetinaNet (A) and CBAM_RetinaNet (B), plotted by the derived precision and recall pairs for a set of confidence thresholds of > 0.2.

The model identifies a set of boundary boxes with different confidence levels. The performance metrics can be derived for each confidence level. Table 4 shows the performance metrics for a set of confidence levels using the AS-E-A combination and CBAM_RetinaNet based on the 8,207 cirques. As an example, CBAM_RetinaNet achieves a precision of 89.6%, a recall of 35.9%, and an F_1 score of 0.51 if using a confidence threshold of 0.6 and an IOU of 0.5. The high precision indicates that the model produces few false positives (most features detected are true cirques), while the low recall suggests large numbers of cirques are not detected (false negatives). Generally, as the confidence level increases, precision improves while recall decreases. As the F_1 score integrates precision and recall, a higher F_1 score indicates a better model performance. For all confidence levels, the model produces an *AP* of 70.2% for IOU_{50} and 56.9% for IOU_{75}, indicating that a relatively higher accuracy can be achieved if using a lower IOU ratio (IOU_{50}), while the accuracy decreases for a higher IOU ratio.

Table 4. The performance metrics (AS-E-A) of CBAM-RetinaNet for the whole dataset of 8,207 cirques.

Confidence Threshold	AP (%) (IOU_{50})	AP (%) (IOU_{75})	Recall (%) (IOU_{50})	Precision (%) (IOU_{50})	F_1 score (IOU_{50})
0.2			48.89	74.47	0.59
0.3			48.89	74.47	0.59
0.4			48.89	74.47	0.59
0.5	70.18	56.88	38.88	84.49	0.57
0.6			35.87	89.56	0.51
0.7			26.91	93.25	0.42
0.8			19.83	97.05	0.33

In addition to the overall performance metrics derived from the whole dataset, the comparison of the model-predicted boundary boxes with mapped cirques in the test areas also suggests the better performance of CBAM-RetinaNet over RetinaNet. Specifically, Fig. 7 shows the results of the identified cirques in the test area of the Gangdise Mountains, which contains 148 cirques manually mapped by Zhang et al. (2020). We also manually checked and identified 24 additional cirques, increasing the number of mapped cirques in this area to 172. Note that some detected boundary boxes overlap with each other, making the number of correctly detected cirque boxes higher than that of mapped cirques. RetinaNet detects a total of 163 cirque boundary boxes (P = 89.5%, R = 63.4%, F_1 = 0.74): 146 boxes are correctly detected (IOU > 0.5), representing 88 identified cirques (51%) because one cirque may correspond to multiple detected boundary boxes, and 17 incorrectly detected (10%). However, 84 cirques are not detected (49%), probably because they are with gently sloping surfaces between their heads and sidewalls (see detailed discussion in Section 4.4). In comparison, CBAM-RetinaNet detects 244 cirque boundary boxes (P = 87.0%, R

= 86.6%, F_1 = 0.87): 195 boxes correctly detect cirques, representing 142 identified cirques (83%). Among the 244 detected boundary boxes, 49 are incorrect detections (20%). Among the 172 manually mapped cirques, only 30 cirques are not detected (17%). The above results show the considerable improvement of CBAM-RetinaNet over RetinaNet in detecting cirques, especially the significant reduction from 49% to 17% in the number of undetected cirques. The confidence level of the detected cirques also improved with CBAM-RetinaNet (the average confidence level increases from 0.36 to 0.40, Fig. 8).

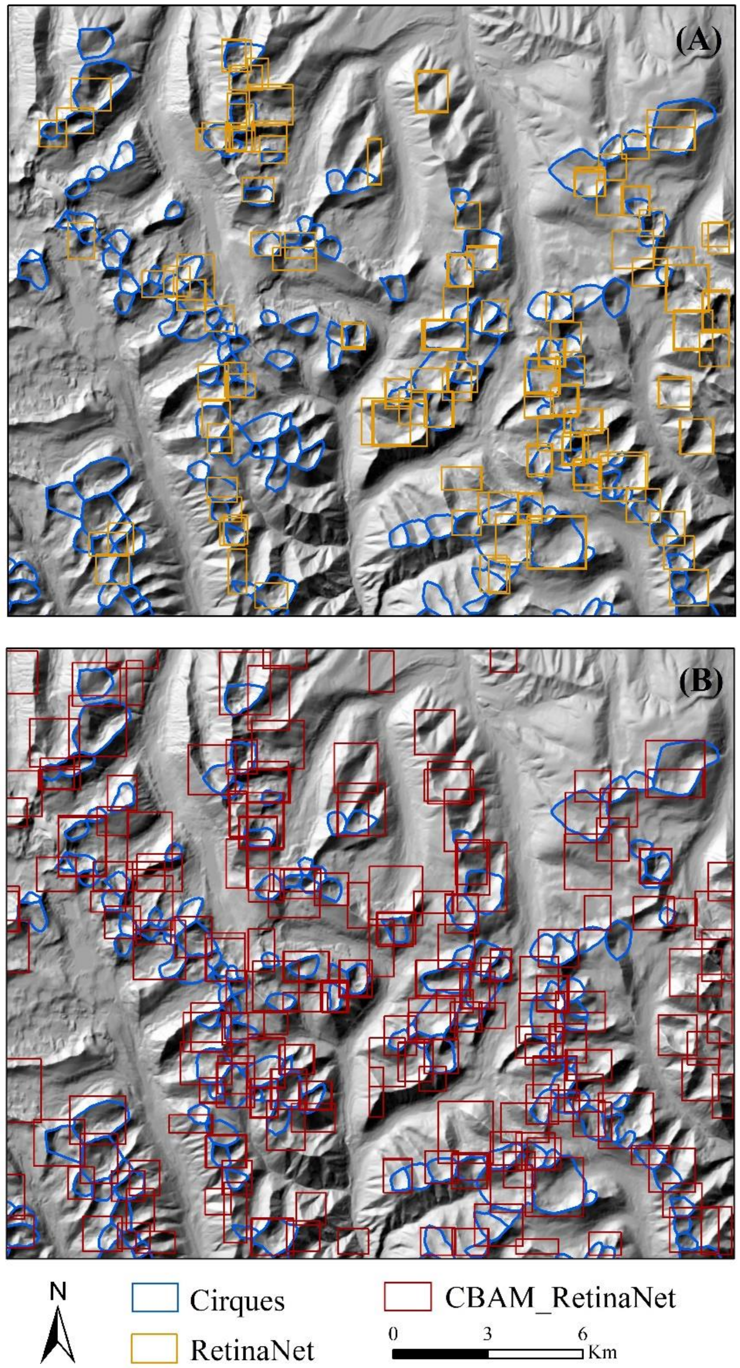

Fig. 7. Cirque detection results in the test area of the Gangdise Mountains. The blue polygons are manually drawn cirques. The yellow-orange boxes (in A) are the RetinaNet-detected boundary boxes and the red boxes (in B) are the CBAM_RetinaNet-detected boundary boxes.

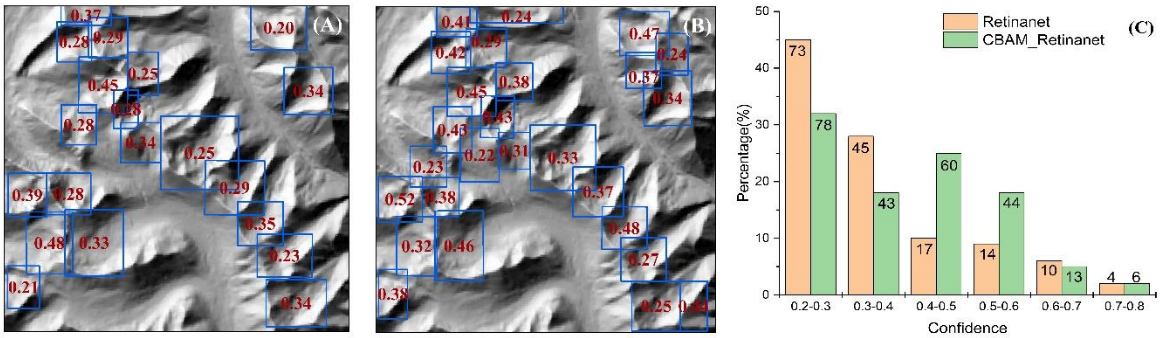

Fig. 8. Comparison of the confidence scores of detected cirque boundary boxes in part of the test area of the Gangdise Mountains. (A) RetinaNet with 20 detected boundary boxes. (B) CBAM-RetinaNet with 26 detected boundary boxes. The background is a hillshade raster. (C) The distribution of predicted boundary boxes with different confidence levels.

4.2 Input combinations

Tables 5 and 6 list the performance metrics of the five input/band combinations for CBAM-RetinaNet in the test areas of the Kamchatka Peninsula and the Gangdise Mountains, respectively. Among the 242 manually mapped cirques in the test area of the Kamchatka Peninsula, the AS-E-A combination can correctly detect 225 of them (93%), followed by the E-E-E combination (220 cirques) and the S-A-E combination (218 cirques). The E-E-E combination incorrectly predicted 46 boxes (11%), then the AS-E-A combination (55 boxes) and the S-A-E combination (57 boxes). The F_1 score is high in E-E-E (0.91), followed by the AS-E-A and the S-A-E (0.89). Among the 172 manually mapped cirques in the test area of the Gangdise Mountains, the E-E-E combination can correctly detect 146 of them (85%), followed by the AS-E-A combination (142 cirques). The AS-E-A combination has the lowest rate of incorrectly predicted boundary boxes (20%), while the other combinations predict 24-28% of incorrect boxes. The F_1 score reaches the highest in the AS-E-A combination (0.87), followed by the E-E-E (0.81) and the AS-H-E (0.76). These results indicate that the AS-E-A and E-E-E combinations have relatively higher performance than other combinations in the two test areas.

Table 5. The performance metrics for different input combinations by CBAM_RetinaNet in the test area of the Kamchatka Peninsula.

Input combinations	Precision (%)	Recall (%)	F_1 score	Correctly predicted cirque number	Incorrectly predicted boxes
AS-E-A	84.7	94.7	0.89	225 (92.9%)	55 (15.2%)
E-E-E	88.6	94.7	0.91	220 (90.9%)	46 (11.4%)
S-A-E	85.6	93.4	0.89	218 (90.1%)	57 (14.3%)
H-H-H	83.4	92.4	0.88	216 (89.3%)	63 (17 %)
AS-H-E	81	91.2	0.86	213 (88.0%)	71 (19.0%)

Table 6. The performance metrics for different input combinations by CBAM_RetinaNet in the test area of the Gangdise Mountains.

Input combinations	Precision (%)	Recall (%)	F_1 score	Correctly predicted cirque number	Incorrectly predicted boxes
AS-E-A	87.05	86.6	0.87	142（82.55%）	49（20.08%）
E-E-E	75.45	88.79	0.81	146（84.88%）	67（24.45%）
S-A-E	74.31	76.41	0.75	122（70.93%）	56（25.68%）
H-H-H	73.36	76.58	0.75	124（72.09%）	57（26.63%）
AS-H-E	72.22	81.25	0.76	133（77.32%）	65（27.77%）

4.3 Training Sample Size

Fig. 9 shows the variations of performance metrics using different numbers of training cirques in the CBAM_RetinaNet model. When the confidence threshold is > 0.2, as the training dataset increases, *AP* and F_1 scores fluctuate, *P* gradually increases, and *R* fluctuates. It seems that the random selection of the sub-dataset for the training led to fluctuations: both data quantity and quality affect the model learning and performance. In general, more cirque boundary boxes can be detected with increased training samples, improving the precision and F_1 score. Our results suggest that *P* and F_1 score improve until the training samples reach 6,000 cirques. However, when the training dataset increases to 7000 cirques, the model detects a large number of small and contiguous targets, which are not cirques, leading to the reduction of the performance metrics. This reduction may be caused by the inclusion of poor-quality cirque outlines in the training. When the training dataset increases to 8000, the above-mentioned issue is minimized, rebounding the model performance. More work is necessary to examine the impact of the quantity and quality of the training dataset on model performance.

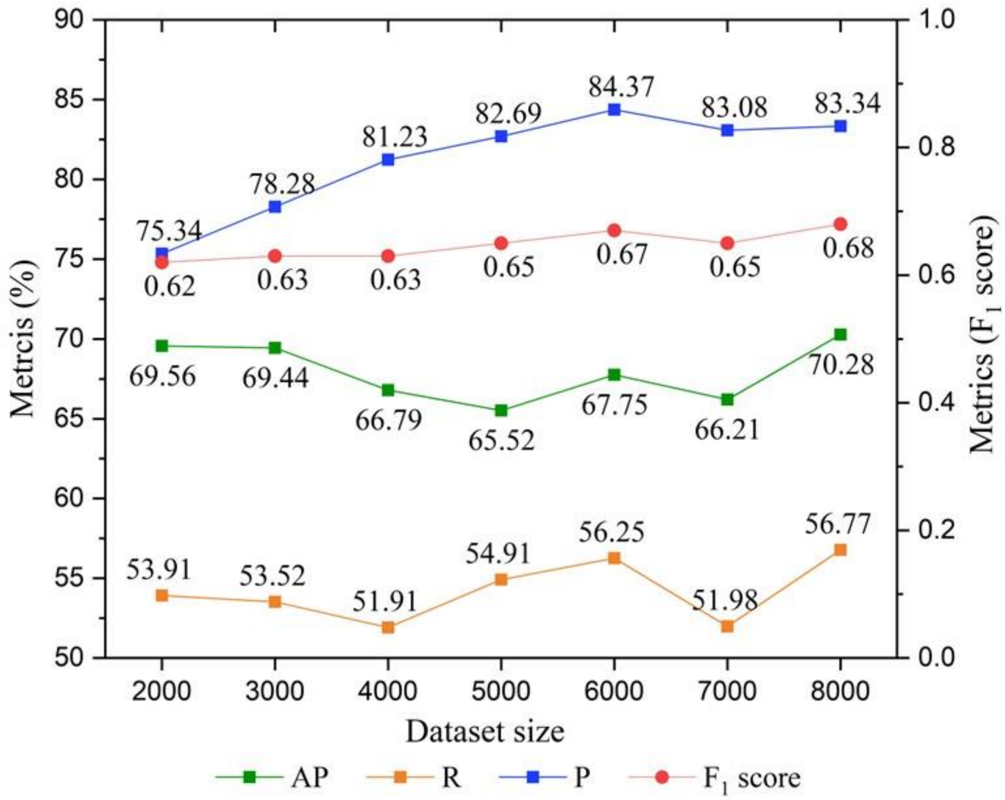

Fig. 9. Variations in the performance metrics for different dataset sizes.

Table 7 and Fig. 10 indicate that the number of detection boxes increases with increasing training dataset size. The number of correctly detected cirques increases. However, as the number of detection boxes increases, the number of incorrect detection boxes also slightly increases. These results indicate that increasing training samples can improve both the model's generalization and the accuracy of cirque detection. The improvement becomes minor when the dataset exceeds 6,000, indicating that training the model with 6,000 samples would yield a good performance. Fig. 11 shows the prediction results of using 8,000 samples in CBAM_RetinaNet: some boundary boxes are completely not related to any cirques, and some just include parts of cirque edges. Most of the undetected cirque boundaries are blurred and not easy to distinguish.

Table 7. Prediction results for different dataset samples by CBAM_RetinaNet in the test area of the Gangdise Mountains.

Dataset samples	**Manually drawn number**	**Total number of detect boxes**	**Correctly predicted cirque number**	**Incorrect detection boxes**	**Undetected cirques**
Dataset2000	172	61	52 (30.2%)	2 (3%)	120 (69.8%)
Dataset4000	172	195	111 (64.5%)	29 (14%)	61 (35.5%)
Dataset6000	172	211	148 (86.0%)	49 (23%)	24 (14.0%)
Dataset8000	172	244	142 (82.6%)	49 (20%)	30 (17.4%)

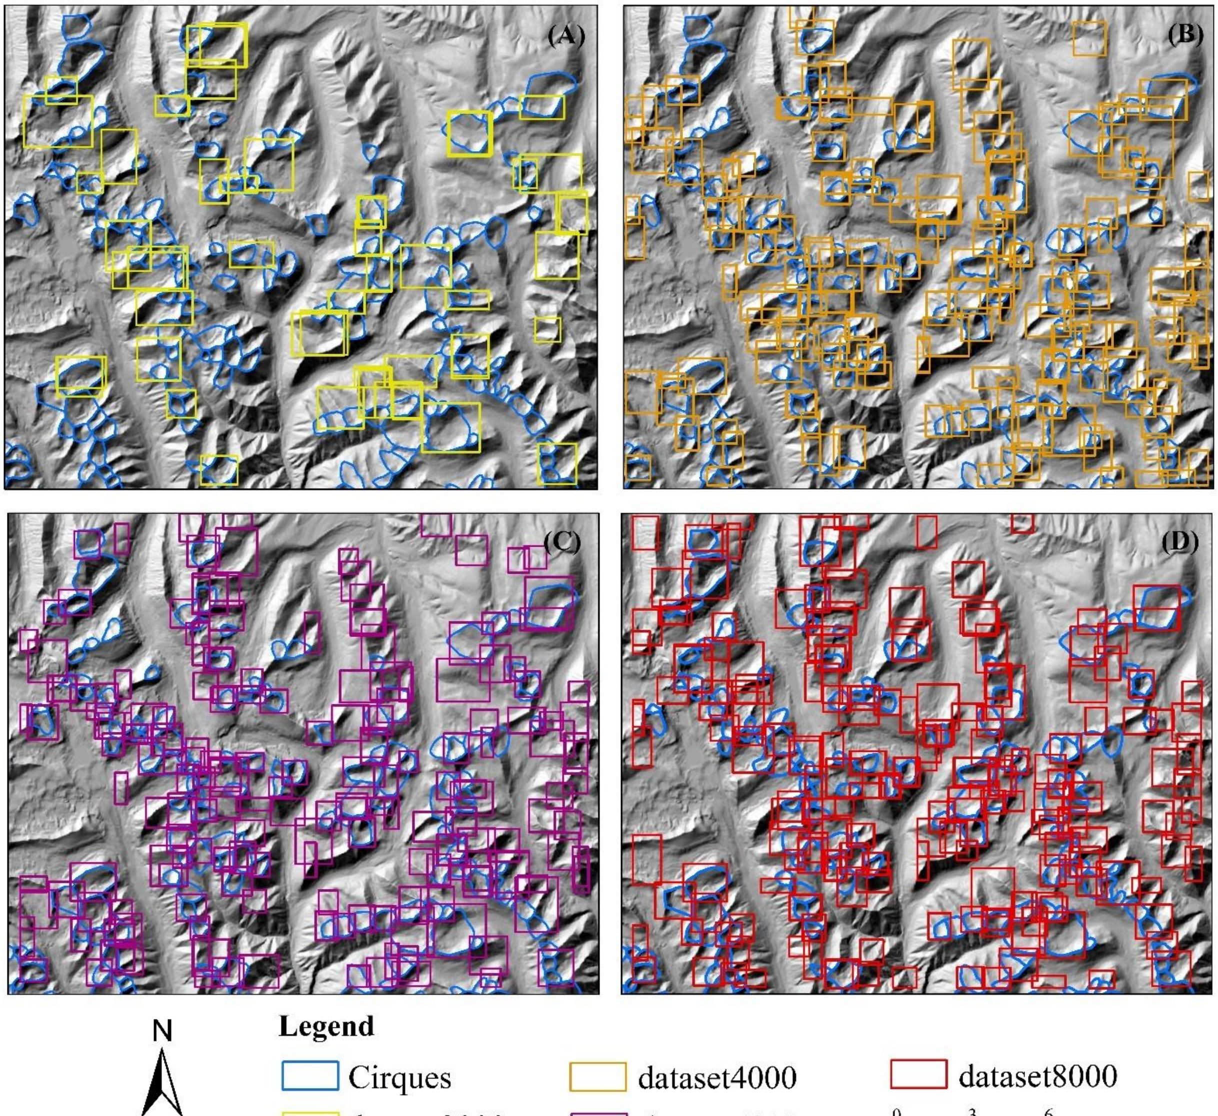

Fig. 10. Cirque detection results in the test area of the Gangdise Mountains. Panels A-D show the results using 2000, 4000, 6000, and 8000 datasets in CBAM_RetinaNet.

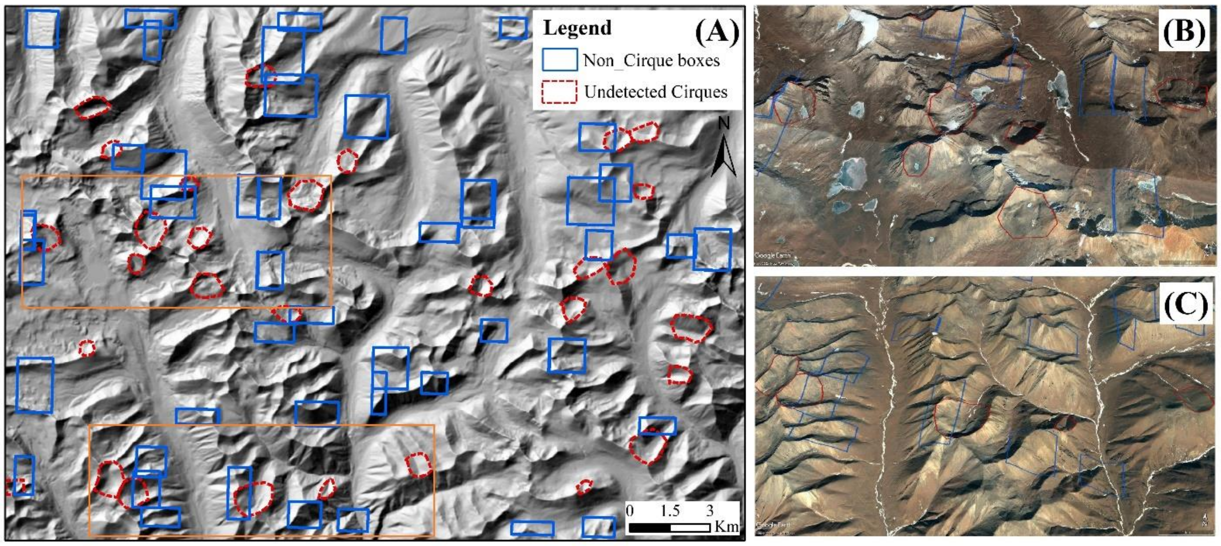

Fig. 11. Prediction results for incorrectly detected boxes and undetected cirques using 8000 datasets to train CAMB_RetinaNet in the test area of the Gangdise Mountains. (B) and (C) are two enlarged parts, showing detailed views of some undetected cirques in Google Earth. The extents of these two enlarged parts are marked as the orange boxes in (A).

4.4 Category confusion

In the test area of the Kamchatka Peninsula, CBAM_RetinaNet generated 361 detection boxes (Fig. 12), including 306 that correctly detected cirques and 55 incorrect detections (P = 84.7%, R = 94.7%, F_1 = 0.89). After visual inspection, 225 of the 242 manually delineated cirques are correctly detected (92.9%), while 17 cirques (7.1%) are undetected. Among the 55 incorrectly detected boxes, 14 are overlapping detections. As illustrated in Fig.13, 30 incorrectly detected boxes include the edges of some cirques, and 11 are completely unrelated to cirques. These boundary boxes also have relatively low confidence levels and less discernible features.

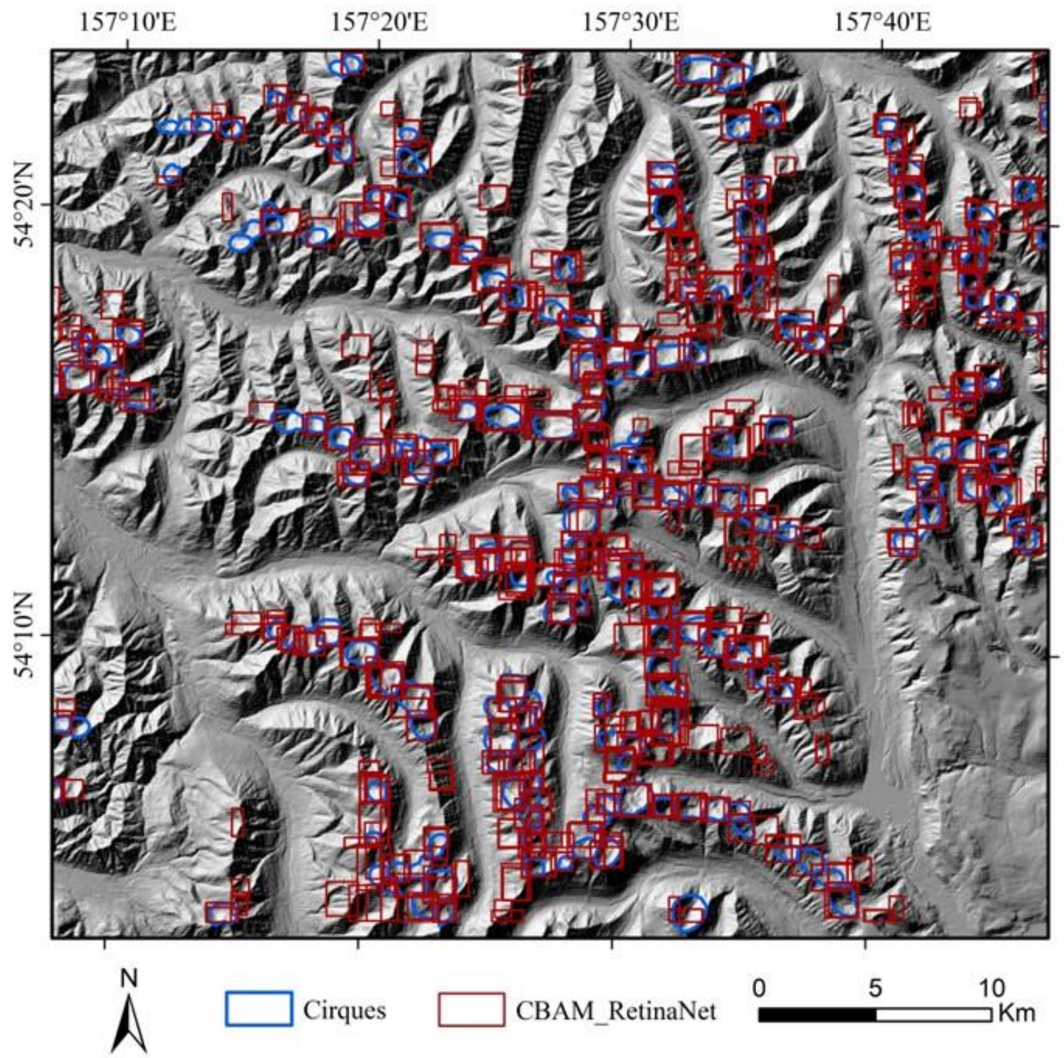

Fig. 12. Predicted results of the test area of Kamchatka Peninsula in CBAM_RetinaNet.

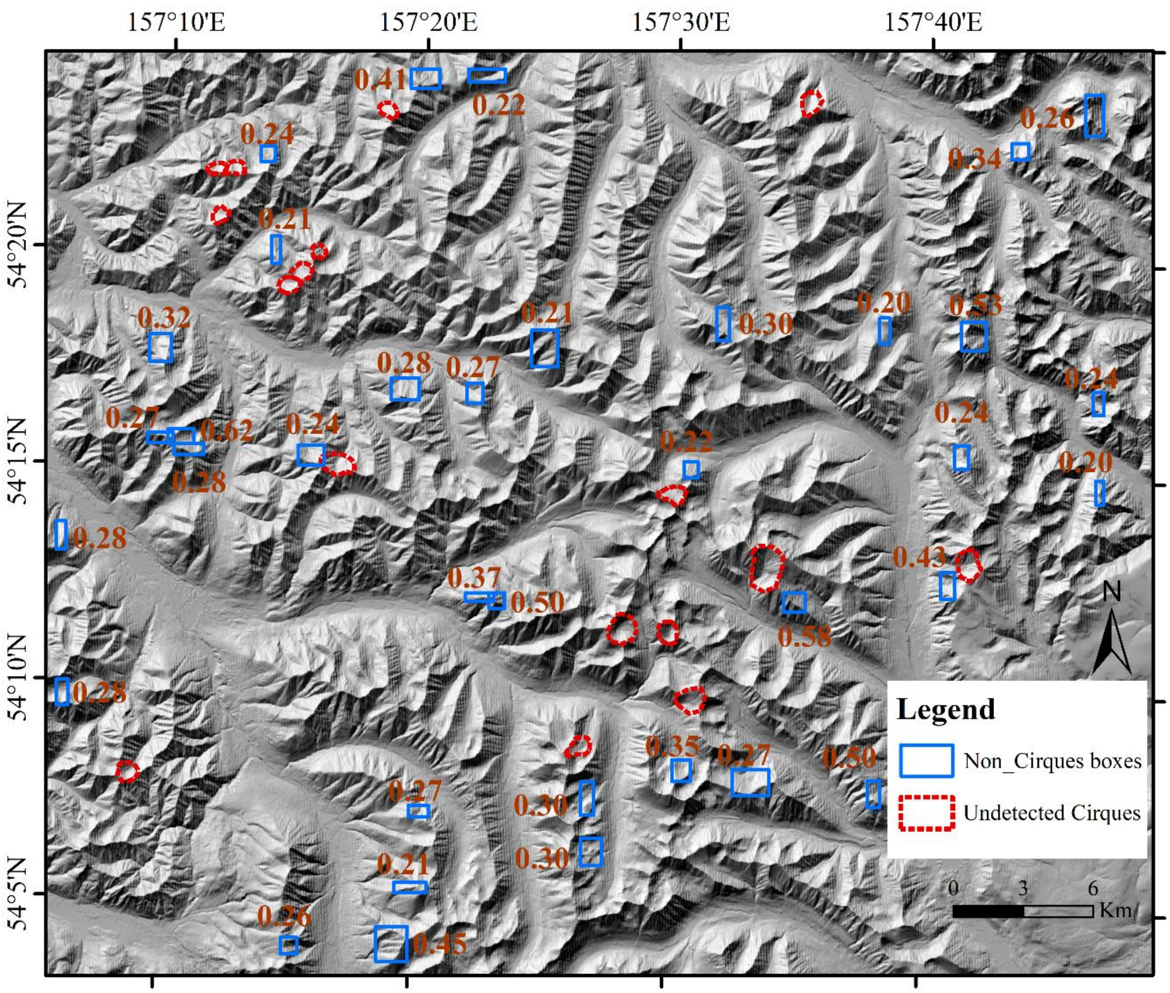

Fig. 13. Non-cirque boxes (blue boxes with confidence scores) and undetected cirques (red polygons) after visual inspection in the test area of the Kamchatka Peninsula.

Fig. 14 shows the detected boundary boxes using CBAM_RetinaNet in three enlarged parts of the two test areas. The cirques with fuzzy boundaries are more prone to misdetection, usually have relatively low confidence scores, and may even be undetected entirely. Panels (A1) and (B1) show that cirques with clear shapes are correctly detected. Panels (A2) and (B2) show some cirques with gently sloping surfaces between their heads and sidewalls are missing or incorrectly detected, and some cirques are not fully enclosed within the boundary boxes. Panels (A3) and (B3) show that some large cirques are detected as multiple cirques. During the dataset creation process, some cirques were truncated by the edge of the training area. Consequently, including these truncated edges as training boxes may lead to the detection of some edges as part of the cirques, resulting in boxes of edges with low confidence.

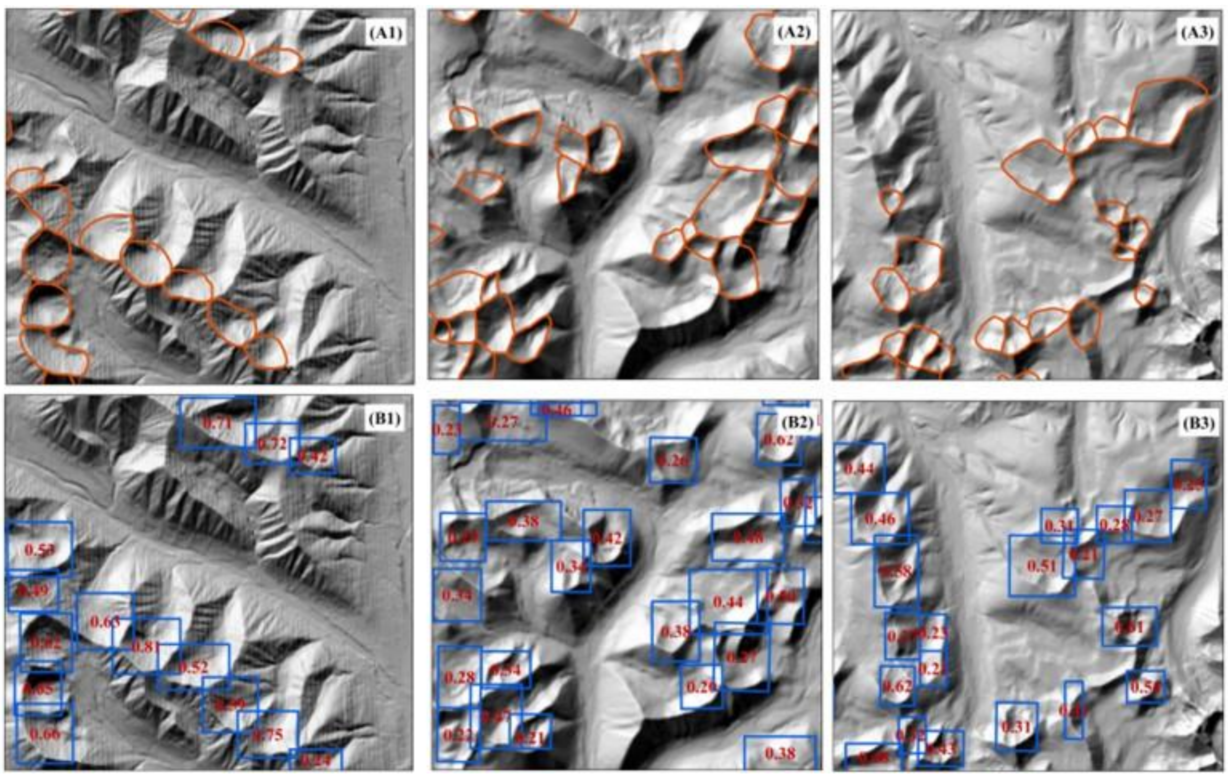

Fig. 14. Three enlarged areas show the performance of cirque detection. (A1) and (B1) are the mapped cirques by Barr and Spagnolo (2013) and detected cirque boundary boxes with their confidence levels from the Kamchatka Peninsula. (A2), (B2), (A3), and (B3) are the mapped cirques by Zhang et al. (2020) and detected cirque boundary boxes with their confidence levels from the Gangdise Mountains.

5. Discussion

Williams et al. (2023) used a RetinaNet model trained and validated using the cirque dataset on Earth to detect potential cirques on Mars. This work reported that the majority (99%) of detected boundary boxes are false positives, reducing the efficiency of using RetinaNet to detect cirque boundary boxes. We trained and validated RetinaNet with the addition of the CBAM module for cirque boundary box detection. The model performance was evaluated in two test areas in the Kamchatka Peninsula and the Gangdise Mountains. The results show that the addition of CBAM not only improves the model performance but also significantly reduces the percentage of undetected cirques. Specifically, CBAM_RetinaNet can correctly detect 83% of the cirques in the test area of the Gangdise Mountains and 93% in the test area of the Kamchatka Peninsula. In the test area of the Gangdise Mountains, about 20% of detected boundary boxes are not related to cirques, and about 17% of manually mapped cirques are not detected. In the test area of the Kamchatka Peninsula, about 15% of detected boundary boxes are not associated with cirques, and about 7% of manually mapped cirques are not detected. Cirques with gentle slopes between their heads and sidewalls were prone to misclassification. Some erroneous detection boxes may be generated at the edge of the test area because small parts of cirques may be incorrectly identified as the entire cirque when preparing the training datasets. In addition, the model sometimes detected similar landforms as cirques, such as local depressions and rock basins.

The dataset used for cirque detection included > 8,000 cirques. The majority are simple cirques. Therefore, the model exhibited strong generalization for simple cirques but weaker generalization for compound cirques. Williams et al., (2023) also found that RetinaNet failed to detect larger composite cirques possibly because of disruption of morphological signature by smaller interior cirques. A separated training dataset for compound cirques is required to generate the model to detect their boundary boxes.

The model-detected results are similar for different input data combinations, with somewhat better results for the AS-E-A and E-E-E combinations. This similarity can be attributed to the fact that all layers were derived from the same DEM. In the future, it may be beneficial to incorporate other types of DEMs and data, such as satellite imagery, for testing and application. In addition, we visually assessed the model performance only in the two test areas on the Kamchatka Peninsula and the Gangdise Mountains. Further evaluation of model performance is recommended in other areas.

Manual methods may not always map all cirques in an area. CBAM_RetinaNet has the potential to compensate for these omissions, although it also detects some incorrect bounding boxes. In summary, both manual and automated identification methods have limitations. The automated approach helps highlight the potential "missed" cirques, enabling the mapping of all potential cirques in an area. We acknowledge that CBAM_RetinaNet only detects the boundary boxes, not the outlines of the cirques. However, CBAM_RetinaNet can be used to quickly detect the boundary boxes of potential cirques, to help further delineation of cirque outlines. In particular, applying automated cirque delineation tools, such as AutoCirque (Li and Zhao, 2022), only to these detected boundary boxes would significantly improve the efficiency and accuracy of the cirque outline delineation. Future work is therefore needed to incorporate the model-detected boundary boxes with AutoCirque or other automated cirque delineation methods to extract cirque outlines.

6. Conclusions

In a pilot study, Scuderi and Nagle-McNaughton (2022) demonstrated the potential of using Convolutional Neural Networks, RetinaNet, to detect the boundary boxes of cirques based on a dataset of 1,951 cirques. However, the impacts of sample size, input data combination, and model algorithm were not fully explored. We compile a dataset of > 8,000 cirques to examine the impacts of training sample size and input data combination on the performance of cirque detection. We also incorporate CBAM into RetinaNet to improve the model performance. The model performance is evaluated based on the comparison with manually mapped cirques in two test areas: one on the Kamchatka Peninsula and the other in the Gangdise Mountains. The major findings are summarized below:

(1) The addition of CBAM in RetinaNet improves the average precision of cirque detection by approximately 4-5%, leading to more accurate boundary box detection and higher confidence levels. Specifically, the model performance reaches 84.7% and 87.0% in precision, 94.7% and 86.6% in recall, and 0.89 and 0.87 in F_1 scores, respectively, in the two test areas on the Kamchatka Peninsula and the Gangdise Mountains. More importantly, the addition of CBAM significantly reduces the percentage of undetected cirques, from 49% to 17% in the test area of Gangdise

Mountains. The proportion of undetected cirques is even lower (7%) in the test area of the Kamchatka Peninsula.

(2) The model performance is affected by training sample sizes and quality, improving with increasing dataset size from 2,000 to 6,000, but decreases at 7000, likely due to the inclusion of poor-quality cirques in the training dataset.

(3) The model performance is affected by the input data combinations. The test of five combinations of DEM and its derived layers shows that the average precision is higher for AS-E-A and H-H-H combinations, while the precision is higher for AS-E-A, E-E-E, and H-H-H combinations for the whole dataset of >8000 cirques. The validation in the two test areas suggests that the AS-E-A and E-E-E combinations have relatively higher performances than other combinations.

(4) CBAM_RetinaNet can detect 83% and 93% of the cirques in the two test areas, respectively. Some cirques on the edges of the DEM are only partly visible, resulting in incorrectly detected boxes.

CBAM_RetinaNet has the potential to be used to detect other landforms. With the fast development of AI technology, further studies are recommended to explore other AI methods for detecting cirques and other landforms with fuzzy boundaries.

Acknowledgments: This study was supported by the National Natural Science Foundation of China (No. 41971075). The authors thank Xin Chen, Zimeng Zhu, and Zixiong Wang from Hebei Normal University and Jinhua Liu from East China University of Technology for their help in the work. We also thank the editor and three anonymous reviewers for their insightful comments and suggestions to greatly improve the manuscript.

References

Anders, A.M., Mitchell, S.G., Tomkin, J.H., 2010. Cirques, peaks, and precipitation patterns in the Swiss Alps: Connections among climate, glacial erosion, and topography. Geology 38, 239–242.

Barr, I.D., Ely, J.C., Spagnolo, M., Clark, C.D., Evans, I.S., Pellitero, X.M., Rea, B.R., 2017. Climate patterns during former periods of mountain glaciation in Britain and Ireland: inferences from the cirque record. Palaeogeogr. Palaeoclimatol. Palaeoecol 485, 466–475.

Barr, I.D., Ely, J.C., Spagnolo, M., Evans, I.S., Tomkins, M.D., 2019. The dynamics of mountain erosion: cirque growth slows as landscapes age. Earth Surface Processes and Landforms 44, 2628–2637.

Barr, I.D., Spagnolo, M., 2015. Glacial cirques as palaeoenvironmental indicators: Their potential and limitations. Earth-Science Reviews 151, 48–78. https://doi.org/10.1016/j.earscirev.2015.10.004

Barr, I.D., Spagnolo, M., 2013. Palaeoglacial and palaeoclimatic conditions in the NW Pacific, as revealed by a morphometric analysis of cirques upon the Kamchatka Peninsula. Geomorphology 192, 15–29. https://doi.org/10.1016/j.geomorph.2013.03.011

Barr, I.D., Spagnolo, M., Tomkins, M.D., 2023. Cirques in the Transantarctic Mountains reveal controls on glacier formation and landscape evolution. Geomorphology 445, 108970. https://doi.org/10.1016/j.geomorph.2023.108970

Bathrellos, G.D., Skilodimou, H.D., Maroukian, H., 2014. The Spatial Distribution of Middle and Late Pleistocene Cirques in Greece. Geografiska Annaler:Series A,Physical Geography 96, 323–338.

Cheng, X., Yu, J., 2021. RetinaNet With Difference Channel Attention and Adaptively Spatial Feature Fusion for Steel Surface Defect Detection. IEEE Trans. Instrum. Meas. 70, 1–11. https://doi.org/10.1109/TIM.2020.3040485

Clark, C.D., Ely, J.C., Greenwood, S.L., Hughes, A.L.C., Meehan, R., Barr, I.D., Bateman, M.D., Bradwell, T., Doole, J., Evans, D.J.A., Jordan, C.J., Monteys, X., Pellicer, X.M., Sheehy, M., 2018. BRITICE Glacial Map, version 2: a map and GIS database of glacial landforms of the last British–Irish Ice Sheet. Boreas 47, 11-e8.

Crest, Y., Delmas, M., Braucher, R., Gunnell, Y., Calvet, M., 2017. Cirques have growth spurts during deglacial and interglacial periods: Evidence from 10Be and 26Al nuclide inventories in the central and eastern Pyrenees. Geomorphology 278, 60–77. https://doi.org/10.1016/j.geomorph.2016.10.035

Cui, Z., 1981. Glacial erosion landforms and development of trough at the head of Urumqi River, Tian Shan. Journal of Glaciology and Geocryology 3, 1–15.

Eisank, C., Drăguţ, L., Götz, J., Blaschke, T., 2010. Developing a semantic model of glacial landforms for object-based terrain classification—the example of glacial cirques. In: Addink, E.A., Van Coillie, F.M.B. (Eds.), GEOBIA-Geographic Object-Based Image Analysis, pp. 1682–1777 ISPRS Vol. No. XXXVIII-4/C7.

Ernest, F., Virginia, H., Hanne, P., Henning, S., 2023. Evolution of the Copernicus DEM: beyond today's elevation data with WorldDEM Neo. Geomorphometry, Iasi, Romania.

Evans, I.S., 2021. Glaciers, rock avalanches and the 'buzzsaw' in cirque development: Why mountain cirques are of mainly glacial origin. Earth Surf Processes Landf 46, 24–46. https://doi.org/10.1002/esp.4810

Evans, I.S., 2006. Local aspect asymmetry of mountain glaciation: a global survey of consistency of favoured directions for glacier numbers and altitudes. Geomorphology 73 (1), 166–184.

Evans, I.S., Cox, N.J., 1974. Geomorphometry and the operational definition of cirques. Area 6, 150–153.

Gardner, J.S., 1987. Evidence for headwall weathering zones, Boundary Glacier, Canadian Rocky Mountains. Journal of Glaciology 33, 60–67.

Gupta, S., Sharma, P., Sharma, D., Gupta, V., Sambyal, N., 2020. Detection and localization of potholes in thermal images using deep neural networks. Multimed Tools Appl 79, 26265–26284. https://doi.org/10.1007/s11042-020-09293-8

Hassinen, S., 1998. A morpho-statistical study of cirques and cirque glaciers in the Senja–Kilpisjärvi area, northern Scandinavia. Nor. Geogr. Tidsskr. Nor. Journal of Glaciology 52 (1), 27–36.

Hawker, L., Uhe, P., Paulo, L., Sosa, J., Savage, J., Sampson, C., Neal, J., 2022. A 30 m global map of elevation with forests and buildings removed. Environ. Res. Lett. 17, 024016. https://doi.org/10.1088/1748-9326/ac4d4f

Huang, L., Luo, J., Lin, Z., Niu, F., Liu, L., 2020. Using deep learning to map retrogressive thaw slumps in the Beiluhe region (Tibetan Plateau) from CubeSat images. Remote Sensing of Environment 237, 111534. https://doi.org/10.1016/j.rse.2019.111534

Jarman, D., Harrison, S., 2019. Rock slope failure in the British mountains. Geomorphology 340, 202-233. https://doi.org/10.1016/j.geomorph.2019.03.002.

Jones, V., Solomina, O., 2015. The geography of Kamchatka. Global and Planetary Change 134, 3–9. https://doi.org/10.1016/j.gloplacha.2015.06.003

Li, H., Zhao, J., Yan, B., Yue, L., Wang, L., 2022. Global DEMs vary from one to another: an evaluation of newly released Copernicus, NASA and AW3D30 DEM on selected terrains of China using ICESat-2 altimetry data. International Journal of Digital Earth 15, 1149–1168. https://doi.org/10.1080/17538947.2022.2094002

Li, Y., Evans, I.S., Spagnolo, M., Pellitero, R., Barr, I.D., Ely, J.C., 2024. ACME2: An extended toolbox for automated cirque metrics extraction. Geomorphology 445, 108982. https://doi.org/10.1016/j.geomorph.2023.108982

Li, Y., Zhao, Z., 2022. AutoCirque: An automated method to delineate glacial cirque outlines from digital elevation models. Geomorphology 398, 108059. https://doi.org/10.1016/j.geomorph.2021.108059

Li, Y., Zhao, Z., Evans, I.S., 2023. Cirque morphology and palaeo-climate indications along a south-north transect in High Mountain Asia. Geomorphology 431, 108688. https://doi.org/10.1016/j.geomorph.2023.108688

Lin, T.-Y., Goyal, P., Girshick, R., He, K., Dollar, P., 2020. Focal Loss for Dense Object Detection. IEEE Trans. Pattern Anal. Mach. Intell. 42, 318–327.

https://doi.org/10.1109/TPAMI.2018.2858826

Lin, T.-Y., Goyal, P., Girshick, R., He, K., Dollar, P., 2017. Focal loss for dense object detection.Proceedings of the IEEE International Conference on Computer Vision. Piscataway: IEEE, 2980-2988.

Liu, X., Xu, Q., Ding, L., 2016. Differential surface uplift: Cenozoic paleoelevation history of the Tibetan Plateau. Sci. China Earth Sci. 59, 2105–2120. https://doi.org/10.1007/s11430-015-5486-y

Lu, E., Hu, X., 2022. Image super-resolution via channel attention and spatial attention. Appl Intell 52, 2260–2268. https://doi.org/10.1007/s10489-021-02464-6

Mîndrescu, M., Evans, I.S., 2014. Cirque form and development in Romania: Allometry and the buzzsaw hypothesis. Geomorphology 208, 117–136. https://doi.org/ 10.1016/j.geomorph.2013.11.019.

Mitchell, S.G., Humphries, E.E., 2015. Glacial cirques and the relationship between equilibrium line altitudes and mountain range height. Geology 43 (1), 35–38.

Nagle-McNaughton, T., McClanahan, T., Scuderi, L., 2020. PlaNet: A Neural Network for Detecting Transverse Aeolian Ridges on Mars. Remote Sensing 12, 3607. https://doi.org/10.3390/rs12213607

Nelson, F.E.N., Jackson, Jr.L.E., 2003. Cirque forms and alpine glaciation during the Pleistocene, west-central Yukon. In: Emond, D.S., Lewis, L.L. (Eds.), Yukon Exploration and Geology 2, Exploration and Geological Services Division. Yukon Region, Indian and Northern Affairs Canada, pp. 183–198.

Pellitero, R., Barr, I., Spagnolo, M., Tomkins, M., 2024. Morphometric analysis of cirques on the Iberian Peninsula provides insights into climate during past glaciations. Geomorphology 461, 109318. https://doi.org/10.1016/j.geomorph.2024.109318.

Peter, G., Carlos, G., Sebastiano, T., 2023. Geomorphometric and Geospatial Patterns in Differences Between ALOS and COPDEM. Geomorphometry, Iasi, Romania.

Principato, S.M., Lee, J.F., 2014. GIS analysis of cirques on Vestfirðir, northwest Iceland: implications for palaeoclimate. Boreas 43, 807–817. https://doi.org/10.1111/bor.12075.

Sanders, J.W., Cuffey, K.M., Moore, J.R., MacGregor, K.R., Kavanaugh, J.L., 2012. Periglacial weathering and headwall erosion in cirque glacier bergschrunds. Geology 40 (9), 779–782.

Scuderi, L.A., Nagle-McNaughton, T., 2022. Automated neural network identification of cirques. Physical Geography 43, 24–51. https://doi.org/10.1080/02723646.2021.1928871

Seif, A., Ebrahimi, B., 2014. Combined use of GIS and experimental functions for the morphometric study of glacial cirques, Zardkuh Mountain, Iran. Quaternary International 353, 236–249. https://doi.org/10.1016/j.quaint.2014.07.005

Solomina, O., Wiles, G., Shiraiwa, T., D’Arrigo, R., 2007. Multiproxy records of climate variability for Kamchatka for the past 400 years. Climate of the Past 3, 119–128. https://doi.org/10.5194/cp-3-119-2007

Spagnolo, M., Pellitero, R., Barr, I.D., Ely, J.C., Pellicer, X.M., Rea, B.R., 2017. ACME, a GIS tool for Automated Cirque Metric Extraction. Geomorphology 278, 280–286.

https://doi.org/10.1016/j.geomorph.2016.11.018
Sun, P., Chen, G., Shang, Y., 2020. Adaptive Saliency Biased Loss for Object Detection in Aerial Images. IEEE Trans. Geosci. Remote Sensing 58, 7154–7165. https://doi.org/10.1109/TGRS.2020.2980023
Tong, K., Wu, Y., Zhou, F., 2020. Recent advances in small object detection based on deep learning: A review. Image and Vision Computing 97, 103910. https://doi.org/10.1016/j.imavis.2020.103910
Trevisani, S., Skrypitsyna, T.N., Florinsky, I.V., 2023. Global digital elevation models for terrain morphology analysis in mountain environments: insights on Copernicus GLO-30 and ALOS AW3D30 for a large Alpine area. Environ Earth Sci 82, 198. https://doi.org/10.1007/s12665-023-10882-7
Williams, J.M., Scuderi, L.A., McClanahan, T.P., Banks, M.E., Baker, D.M.H., 2023. Comparative planetology – Comparing cirques on Mars and Earth using a CNN. Geomorphology 440, 108881. https://doi.org/10.1016/j.geomorph.2023.108881
Woo, S., Park, J., Lee, J.-Y., Kweon, I.S., 2018. CBAM: Convolutional Block Attention Module.
Zhang, Q., Fu, P., Yi, C., Wang, N., Wang, Y., Capolongo, D., Zech, R., 2020. Palaeoglacial and palaeoenvironmental conditions of the Gangdise Mountains, southern Tibetan Plateau, as revealed by an ice-free cirque morphology analysis. Geomorphology 370, 107391. https://doi.org/10.1016/j.geomorph.2020.107391
Zhang, Q., Yi, C., Fu, P., Wu, Y., Liu, J., Wang, N., 2018. Glacier change in the Gangdise Mountains, southern Tibet, since the Little Ice Age. Geomorphology 306, 51–63. https://doi.org/10.1016/j.geomorph.2018.01.002
Zhang, W., Yan, L., Cui, Z., 2008. Factors Controlling the Development of Mountain Cirques and Climate Change. Journal of glaciology and geocryology 30, 266–273.