

ORIGINAL RESEARCH

Few-shot learning for satellite characterisation from synthetic inverse synthetic aperture radar images

Friso G. Heslinga  | Faruk Uysal  | Sabina B. van Rooij  | Sven Berberich  | Miguel Caro Cuenca 

TNO - Defence, Safety & Security, the Hague, the Netherlands

Correspondence

Miguel Caro Cuenca.

Email: miguel.carocuenca@tno.nl

Abstract

Space situational awareness systems primarily focus on detecting and tracking space objects, providing crucial positional data. However, understanding the complex space domain requires characterising satellites, often involving estimation of bus and solar panel sizes. While inverse synthetic aperture radar allows satellite visualisation, developing deep learning models for substructure segmentation in inverse synthetic aperture radar images is challenging due to the high costs and hardware requirements. The authors present a framework addressing the scarcity of inverse synthetic aperture radar data through synthetic training data. The authors approach utilises a few-shot domain adaptation technique, leveraging thousands of rapidly simulated low-fidelity inverse synthetic aperture radar images and a small set of inverse synthetic aperture radar images from the target domain. The authors validate their framework by simulating a real-case scenario, fine-tuning a deep learning-based segmentation model using four inverse synthetic aperture radar images generated through the backprojection algorithm from simulated raw radar data (simulated at the analogue-to-digital converter level) as the target domain. The authors results demonstrate the effectiveness of the proposed framework, significantly improving inverse synthetic aperture radar image segmentation across diverse domains. This enhancement enables accurate characterisation of satellite bus and solar panel sizes as well as their orientation, even when the images are sourced from different domains.

KEYWORDS

artificial intelligence, image segmentation, inverse synthetic aperture radar (ISAR), satellite tracking

1 | INTRODUCTION

Space situational awareness (SSA) is usually understood as the ability to monitor, understand, and predict the location, movement, and behaviour of objects in space [1]. Although this definition includes observing celestial bodies, our focus is on man-made objects like satellites and debris.

Radar systems play a very important role in SSA since they are able to provide a wealth of information on remote objects. In particular, imaging radars are able to characterise space objects using inverse synthetic aperture radar (ISAR) techniques, see for example, refs. [2, 3] or [4].

However, understanding the capabilities of an unknown satellite from an ISAR image is not trivial. Satellite capabilities can be estimated simply from its size, shape, and solar panel size. For example, CubeSats whose sizes are just in the order of tens of centimetres usually are passive devices in contrast to earth observations systems which are usually a couple of metres long. Parameters such as the satellite power budget can be inferred from their solar panel size. To achieve this characterisation, one typical strategy is to perform semantic segmentation of the ISAR image of the satellite.

Semantic segmentation is a computer vision task in which the goal is to categorise each pixel in an image into a class. For

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial-NoDerivs](https://creativecommons.org/licenses/by-nc-nd/4.0/) License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

© 2024 TNO. *IET Radar, Sonar & Navigation* published by John Wiley & Sons Ltd on behalf of The Institution of Engineering and Technology.

the development of a computer vision model, we make use of training data that consists of images and corresponding masks with ground truth annotations for the class of each pixel. The training data is used to learn the model parameters such that the model is able to correctly predict the class of each pixel.

Several techniques can be utilised for the segmentation task, with deep learning [5] based methods being the most effective approach. Deep learning-based segmentation methods have shown to yield superior results in terms of accuracy and efficiency compared to other methods for a wide variety of applications [6, 7]. Examples include scene understanding for autonomous driving [8], robot perception [9], semantic segmentation for medical image analysis [10, 11] and automated satellite-to-ground image annotation [12].

The use of deep learning-based segmentation in SSA is still relatively limited. In 2019, Xue et al. [13] proposed StarNet, a deep learning architecture similar to the popular U-Net [14] for dim target detection in star images. In 2023, Mastrofini et al. [15] developed a system for attitude determination of stars and onboard space surveillance and tracking of resident space objects. On the topic of automated segmentation of satellites in ISAR images we only identified a single study: Du et al. [16], that used generative adversarial networks to segment synthetic ISAR images. Next to the segmentation strategies, Wang et al. [17] focus directly on classification and they showed that good classification results can be obtained by using high-resolution ISAR images.

Deep learning methods typically require thousands of images for training and model optimisation. However, recording that large amount of ISAR images is costly and requires dedicated SSA systems. A typical strategy to overcome the lack of sufficient real training data is the use of simulated data for deep learning model development. For example, in the medical domain, Al Khalil et al. [18] showed segmentation failures in cardiac magnetic resonance imaging can be reduced by training with realistically synthesised pathology images. Similarly, Chen et al. [19] presented a method to develop a segmentation model for autonomous driving using synthetic images.

The use of simulated data is also promising for SSA. In the previous work [20], we developed a deep learning-based segmentation model using synthetic images. The model was based on the U-Net [14] architecture and trained with images that were simulated by a fast simulator at the image level using a realistic radar response with limited details (low-fidelity). The evaluation was done for other low-fidelity images (same domain), resulting in intersection over union (IoU) scores of 0.64–0.81.

However, when a deep learning model trained on a single (source) domain (e.g. low-fidelity simulated images) is applied to another (target) domain (e.g. measured ISAR imagery), there is typically a strong decrease in segmentation performance, due to the gap in appearance between the domains. One strategy to overcome this is the use of few-shot learning, where only a few examples of the target domain are available. This technique has been studied in the context of image-classification [21–23] and to a limited extent also for semantic segmentation [24–26].

Recently, a novel framework for few-shot learning for semantic segmentation was presented by Tavera et al. called pixel-by-pixel cross-domain alignment (PixDA) [27]. The framework includes a pixel-by-pixel domain adversarial loss that was designed to align the images of the source to the target domain on a pixel-level. The method also includes a strategy to prevent the negative transfer of knowledge learnt on the source domain and a regularisation technique to prevent overfitting. PixDA was shown to achieve state-of-the-art results on synthetic-to-real benchmarks where semantic segmentation is applied to street scenes in German cities.

To the best of our knowledge, few-shot learning for domain adaptation has not been applied in the context of SSA. We propose a framework based on few-shot learning to perform semantic segmentation with only a few ISAR images. We demonstrate the training of a deep learning-based segmentation model using rapidly simulated low-fidelity ISAR images. This model is successfully applied to ISAR images in different domains, such as high-fidelity simulated ISAR images. These high-fidelity simulations closely resemble measured ISAR images, incorporating realistic noise, requiring a significantly longer time for synthesis. We finetune on only a few images from target domain and are able to get useful segmentation results. These, in turn, provide a good basis for the assessment of the satellite's bus and panel sizes. Our experiments demonstrate the effectiveness of domain adaptation in improving results, indicating its potential application with measured data.

The rest of this article is arranged as follows. Section 2 describes the simulated data sets, the models and experiments. Section 3 contains both quantitative and visual results of the segmentation method, which are discussed in Section 4. The article ends with a brief conclusion in Section 5.

2 | MATERIALS AND METHODS

2.1 | ISAR image simulations

Our main focus in this study is to demonstrate the effectiveness of domain adaptation in improving results, indicating its potential application to measured data when accessible. Although measured ISAR images would have been ideal for validation (as well as training), scarcity of those data sets necessitated the use of simulated ISAR images for both training and validation of our segmentation models. To address this challenge, we employed two different sets of simulations: low-fidelity and high-fidelity simulations (or ISAR images). Low-fidelity simulations were utilised for training purposes, providing a foundation for our models. In contrast, high-fidelity simulations (some example images are shown in Figure 1) were reserved for verification due to the absence of measured data sets, ensuring an evaluation of the model's performance.

While ISAR processing uses radar data to generate two-dimensional high-resolution images of a target, where the relative rotation of the target is exploited to create a synthetic

aperture, in our low-fidelity simulations, the generation of ISAR images involved a structured process ensuring that similar ISAR images are quickly simulated for training.

The low-fidelity simulation starts by calculating precise orbit parameters of the targets derived from user-defined system specifications. The detailed mesh model of the target was transformed into a point target cloud. An in-house electromagnetic (EM) solver was then utilised to analyse the response of each point scatterer for the radar carrier frequency. This analysis accounted for different factors like shadowing effects. The computed point scatterers $r(x, y, z, \sigma)$ were skillfully projected onto the ISAR image grid as $\hat{r}(x, y, \sigma)$, preserving their individual radar cross section (RCS) values σ , where x and y represent the range and cross-range respectively. To simulate ISAR images accurately, the projected scatterers were convolved with the point spread function (PSF) of the system $h(x, y)$ and a noise model is added which takes into account the radar noise level. Based on this procedure, which is also illustrated in Figure 2, an ISAR image can be simulated as follows [20]:

$$I(x, y) = \hat{r}(x, y, \sigma) * h(x, y) + n(x, y) \quad (1)$$

where $*$ is the convolution operator and $n(x, y)$ is the noise signal in the image domain. This convolution process integrated the unique characteristics of the imaging system, translating the point cloud into an ISAR image.

Using the obtained projections, ground truths were also generated. Due to the nature of the projection process, some scatterers might overlap after projection. This intricacy was carefully considered, ensuring the creation of a comprehensive ground truth dataset. For each image, a mask is generated containing the label of each pixel. Pixels can belong to either class 0 (background), class 1 (bus), class 2 (solar panel), or class 3 (both bus and solar panel).

Both simulated images and masks are cropped to a size of 256×256 pixels. Although this way of simulating is fast, some typical noise factors are neglected, such as side lobes due to for example, unknown target motion and detailed noise distribution. A total of 2662 images was simulated, which was split into




FIGURE 1 Examples of simulated ISAR images. Left: Low-fidelity simulated image with corresponding ground truth labels. The red outline represents the solar panel and the yellow outline represents the bus. A small region of overlap is indicated with blue. Right: High-fidelity simulated ISAR with corresponding ground truth labels. Only the outline of the labels are shown.




FIGURE 2 Low-fidelity ISAR image simulation flow diagram.

three sets: 1863 for training, 532 for validation and model selection, and 267 for evaluation.

The second set of synthetic ISAR aims at mimicking a real scenario of data acquisition. This set consists of much more realistic simulations where the noise and side-lobe behaviour is expected to be very close to reality. These images are computationally expensive to generate. We employ a back-projection approach [28] for ISAR imaging using the signal that is simulated at the raw radar level (as output by the analogue-to-digital converter [ADC] of the radar receiver). The complete simulation process includes propagation losses, thermal noise, 3D-mesh models of satellites for RCS modelling, real orbits, and acquisition errors. We refer to these images as high-fidelity ISAR images, which also have a corresponding mask with pixel-wise class labels. A total of five high-fidelity images were simulated and cropped to 256×256 pixels.

The simulator of low and high-fidelity images produces complex values. We convert the image power into a decibel (dB) scale before training the segmentation models. The power level in dB is given by $P_{dB} = 20\log_{10}(|\cdot|)$.

Computational Complexity: For a target object image (of the size of Envisat) utilising a high-quality mesh model, our in-house EM simulations require less than 1 s. The process of generating projections along with ISAR imaging grid, including orbit determinations and shadowing, takes approximately 3 s. This totals to 4 s per low-fidelity image when utilising graphics processing unit parallelised multicore processing in a specialised cluster. In contrast, generating a single high-fidelity ISAR image takes several hours, resulting in significant time requirements for processing a complete training set.

This observation leads to the conclusion that while simulating high-fidelity image sets for training deep learning models is technically possible, it is impractical. It is worth noting that the example presented in this article focuses on a single satellite; however, our ultimate goal is to generalise the method, requiring input from various satellite models to segment any unknown satellites. To address this challenge, we introduced a domain adaptation method for radar image segmentation problems (refer to Section 2.3), which allows us to use low-fidelity simulation for segmentation model training.

2.2 | Baseline model

There is a wide variety of deep learning architectures available for the task of semantic segmentation. Here, we build upon previous work by van Rooij et al., [20] using the U-Net architecture [14], originally introduced for segmentation of biomedical images. U-Net has been successfully used in a variety of applications [29, 30]. Van Rooij et al. showed that it is possible to effectively distinguish between the solar panel and the main bus in images that are simulated under optimal observing conditions. The U-Net model is implemented with the *segmentation_models_pytorch* library based on PyTorch and we use a ResNet-18 [31] backbone. The model is trained with patches of 64×64 pixels (Figure 3) to output a 64×64 probability map for each of the classes background, bus, and solar panel. Ten patches with corresponding masks are extracted from each of the low-fidelity images of the training set (18,650 patches in total). The cropping is done such that most patches contain a part of the satellite. All image patch intensities are scaled to have zero mean and unit standard deviation.

The model is iteratively optimised as follows. The per-pixel loss for each class is computed as follows:

$$l_{i,c} = \begin{cases} z_{i,c} \log(\hat{z}_{i,c}) & \text{if } c < 3 \\ z_{i,c} \log(\max(\hat{z}_{i,1}, \hat{z}_{i,2})) & \text{if } c = 3 \end{cases} \quad (2)$$

where $z_{i,c}$ refers to the ground truth label per pixel for class c , either 0 or 1. On the other hand, $\hat{z}_{i,c}$ is the predicted probability that pixel i belongs to class c which is only obtained for classes 0, 1, and 2. The total loss $\mathcal{L}_{tot}$ is computed by summing over the per-pixel losses per class:

$$\mathcal{L}_{tot} = \frac{\sum_i \sum_c l_{i,c}}{N} \quad (3)$$

where N is the total number of pixels. The baseline model is trained for 20 epochs with batch size of 32. The Adam optimiser [32] was used and the learning rate was set 0.0005.




FIGURE 3 Patch extraction for the training of the baseline segmentation model. On the left, a low-fidelity simulated image is shown with a corresponding ground truth segmentation map. The red, light blue, and light green boxes represent locations where (64×64 pixels) patches are extracted. On the right these patches are shown in more detail. Yellow pixels = bus, red pixels = solar panel, blue pixels = both.

Evaluation of the baseline model is done for the 267 images of the evaluation set. The metric for evaluation is IoU, which we calculate for the bus class ($i = 1$) and solar panel class ($i = 2$):

$$IoU = \frac{1}{2} \sum_{i=1}^2 \frac{TP_i}{TP_i + FP_i + FN_i} \quad (4)$$

where TP = true positives, FP = false positives, and FN = false negatives. Without overlapping classes the TP, FP, and FN can be calculated by using a soft-max operator on predicted class probabilities; however, we have to take into account that our ground truth labels contain class 3 (both bus and solar panel), while our model only outputs probabilities for classes 0, 1 and 2. One solution is to assign both classes 1 and 2 to a prediction when a certain threshold is met for a pixel. We experimented with various ratios (using the validation set) and found that we can simply assign both classes 1 and 2 to the prediction for a pixel if both of the predicted output for class 1 and 2 are positive. For the ground truth annotations we assign pixels with class 3 to both class 1 and 2. The baseline model is also directly applied to the high-fidelity simulated ISAR images to show how the segmentation model performs without domain alignment.

2.3 | Proposed framework for ISAR image segmentation

For segmenting ISAR images, we propose a framework (Figure 4) based on PixDA by Taverna et al. [27]. PixDA is a domain alignment technique that uses a domain adversarial loss. This means that, apart from predicting the correct class for each pixel, the model is incentivised to learn a feature representation that is similar to the source and target domain [33, 34]. Domain adversarial networks have shown to be effective for learning from different domains for various applications [35, 36], including semantic segmentation [37]. Novel for PixDA is that the domain adversarial loss is implemented on a pixel-level to promote good performance for each class, including those that are underrepresented. To achieve this, a standard adversarial loss is modulated by a term that represents the model's ability to predict the class correctly and weight based on the frequency of the class in the target data set.




FIGURE 4 Framework for few-shot learning with PixDA.

The PixDA method also includes a sample selection procedure to find samples from the source domain that are more similar to the target. In addition, knowledge distillation is used as a technique to prevent overfitting. With knowledge distillation, knowledge from a large model is transferred to a model with less parameters [38, 39], a strategy that has been shown to be useful in image classification [40], biomarker approximation [41], and image segmentation [42].

We use the implementation of PixDA provided by the authors [43]. The starting point for the PixDA framework is our U-Net baseline model. Our source data set consists of the 1863 training images from the low-fidelity simulations. For the fine-tuning with PixDA, four out of five high-fidelity images are used. We use all default settings, except for the batch size set to 4 and early stopping after 5000 iterations. Evaluation of the fine-tuned model is done using leave-one-out cross-validation: five models are trained, each with one of the high-fidelity simulated images left out. Evaluation is done for the left-out image by calculating the IoU as described previously.

3 | RESULTS

When the baseline model is applied to the low-fidelity simulated ISAR evaluation set images an IoU of 0.72 (± 0.14) is obtained, where $\pm$ represent one standard deviation. Figure 5 displays two examples of the results. For the top row example, the predicted segmentation is very close to the ground truth mask (IoU = 0.86). We see some minor mistakes on the borders of the segmentation, but the satellites' shape, size, and orientation can be correctly identified. The bottom row shows an example where the segmentation is below average (IoU = 0.60). Although it is still possible to determine the orientation of the satellite, the shape and the size of the bus are more difficult to determine.

On the right side of Figure 5, we include uncertainty maps for the predictions. To measure uncertainty we use the amplitude of the predicted confidence for the class that is predicted for each pixel. We invert the confidences and normalise with respect to the pixel that has the highest predicted value. We end up with a map that shows the most certain regions in black and the most uncertain regions in white/yellow. For the top row in Figure 5 the uncertainty maps show that the most uncertain regions overlay with the border of the bus and solar panel. For the example in the bottom row, it is not only the border regions which are uncertain, but most of solar panel area is highlighted. Examining the solar panel area in the simulated image clearly highlights the challenging nature of assessing this region, even for human experts.

Figure 6 shows the results for two examples of the high-fidelity simulated ISAR images. When the baseline model is applied to the high-fidelity simulated ISAR, the resulting segmentations are very poor (IoU = 0.13 ± 0.04), as can be seen in the third column of Figure 6. These segmentations do not provide useful information about the satellite.




FIGURE 5 Segmentation results for the baseline model on the low fidelity simulated ISAR images. IoU, intersection over union. Yellow = bus, red = solar panel, blue = both.




FIGURE 6 Segmentation results for on the high-fidelity simulated ISAR images. The third column shows the segmentations by the baseline model. The fourth columns shows the segmentations with the PixDA model. Yellow = bus, red = solar panel, blue = both.

Segmentation results on the high-fidelity simulated ISAR images with the PixDA model are much better with an IoU of $0.55 (\pm 0.06)$. In Figure 6, the fourth column shows the segmentation with PixDA. Although the segmentations contain some errors, the overall structure is much closer to the ground truth mask. An overview of the numerical results is provided in Table 1.

4 | DISCUSSION

In this study, we showed that a typical deep learning-based segmentation model can be trained to segment satellite sub-structures in synthetic ISAR images from a single source. The direct application of such a model on ISAR images from a different domain (high-fidelity simulations in our case) results

TABLE 1 Overview of results. Scores represent intersection over union (IoU).

	Low-fidelity simulations (Source)	High-fidelity simulations (Target)
Baseline model	0.72 ($\pm$ 0.14)	0.13 ($\pm$ 0.04)
PixDA (4-shot)	–	0.55 ($\pm$ 0.06)

in poor segmentations. To overcome this knowledge gap, the deep learning model can be finetuned with only a few images of the target domain to get useful segmentations on the target domain. We tested with simulated ISAR images with large side-lobes to include cases with acquisition and processing errors (not fully compensated trajectory). Although the segmentations are not perfect, they do allow for detailed analysis, including the size and shape of the solar panel and bus, and the satellite's orientation, which can otherwise be difficult for a user to visually interpret. These results are in line with our objectives for SSA, where we are interested in more than just the detection and orbit estimation of satellites.

The amount of ISAR data collected when looking into space is growing rapidly and manual analysis of ISAR images will be unfeasible and too slow in most cases. At the same time, annotating ISAR data is costly and time-consuming. The few-shot learning approach that we applied indicates that only a limited amount of data is needed to finetune a model for another domain. Although we demonstrate the benefits of domain adaptation by using high-fidelity simulations, the concept and improvement remains almost the same when we use measured images.

We chose PixDA as our few-shot domain adaptation technique since this framework was shown to do well for another application. However, there are other methods available which we will consider in future work. For example, Zhang et al. [44] also developed a method for synthetic-to-real domain alignment for street scene segmentation. Their method includes a two-stage adversarial network which contains a scene parser and two discriminators. Keaton et al. [45] used few-shot domain adaptation for the case of cellular instance segmentation in microscopy images by introducing a specialised contrastive loss. Kalluri and Chandraker [46] proposed a clustering-based adaptation approach for outdoor scene to indoor scene segmentation.

In future work, we will first test with measured satellite images to understand further the limitations of our proposed framework. We will then extend this work with a variety of satellites. In that way, the size and estimations can be evaluated for a wide range of appearances. One additional interesting question would be to find the relation between the number of images used in the few-shot domain-adaptation and the quality of segmentation and subsequent classification.

5 | CONCLUSION

Obtaining ISAR data from space objects, for training deep learning segmentation models or other usage, is time-consuming, costly and requires very dedicated hardware.

However, fast simulations enable the generation of a large number of synthetic ISAR images containing similar features, which can be used to train a baseline segmentation model. With only a few images from a target domain, a segmentation model can be fine-tuned to perform well on the specific domain. Our experiments demonstrated the effectiveness of domain adaptation in improving IoU scores for segmenting the satellite's bus and solar panel, indicating its potential application with measured data. Based on accurate segmentations, satellite characteristics can be determined, allowing for the classification of satellite types and providing us with better SSA.

AUTHOR CONTRIBUTIONS

Miguel Caro Cuenca, Faruk Uysal, Sabina B. van Rooij, and Friso G. Heslinga conceptualised and designed the study. Faruk Uysal and Miguel Caro Cuenca created the simulations. Friso G. Heslinga, Sabina B. van Rooij, and Sven Berberich implemented the models and conducted the experiments. Friso G. Heslinga analysed the results and wrote the original draft. Miguel Caro Cuenca, Faruk Uysal, Sabina B. van Rooij, and Friso G. Heslinga reviewed the analysis. All authors reviewed the manuscript.

CONFLICT OF INTEREST STATEMENT

The authors declare no conflicts of interest.

DATA AVAILABILITY STATEMENT

Research data are not shared.

ORCID

Friso G. Heslinga  <https://orcid.org/0000-0002-8383-307X>

Faruk Uysal  <https://orcid.org/0000-0002-4518-5649>

Sabina B. van Rooij  <https://orcid.org/0009-0001-0948-8175>

Sven Berberich  <https://orcid.org/0009-0008-5058-5103>

Miguel Caro Cuenca  <https://orcid.org/0009-0006-9327-0043>

REFERENCES

1. Console, A.: Command and control of a multinational space surveillance and tracking network. Joint Air Power Competence Centre (2019)
2. Mehrholz, D.: Radar techniques for the characterization of meter-sized objects in space. *Adv. Space Res.* 28(9), 1259–1268 (2001). [https://doi.org/10.1016/s0273-1177\(01\)00395-7](https://doi.org/10.1016/s0273-1177(01)00395-7)
3. Ender, J., et al.: Radar techniques for space situational awareness. In: 12th International Radar Symposium, pp. 21–26 (2011)
4. Czerwinski, M.G., Usoff, J.M.: Development of the Haystack ultra-wideband satellite imaging radar. *Linc. Lab. J.* 21(1), 28–44 (2014)
5. Lecun, Y., Bengio, Y., Hinton, G.: Deep learning. *Nature.* 521(7553), 436–444 (2015). <https://doi.org/10.1038/nature14539>
6. Dohare, N., Nigam, M.J.: Comparative analysis of image segmentation techniques. In: 2019 6th International Conference on Signal Processing and Integrated Networks (SPIN), pp. 214–218. IEEE (2019)
7. Senthilnath, J., Sivaramakrishnan, R., Omkar, S.: Image segmentation using deep learning: a survey. *J. Imag.* 4(12), 130 (2018)
8. Yurtsever, E., et al.: A survey of autonomous driving: common practices and emerging technologies. *IEEE Access* 8, 58443–58469 (2020). <https://doi.org/10.1109/access.2020.2983149>
9. Xie, C., et al.: Unseen object instance segmentation for robotic environments. *arXiv* (2021), abs/2007.08073

10. Heslinga, F.G., et al.: Quantifying graft detachment after Descemet's membrane endothelial keratoplasty with deep convolutional neural networks. *Transl. Vis. Sci. Technol.* 9(2), 48 (2020). <https://doi.org/10.1167/tvst.9.2.48>
11. Mehrtash, A., et al.: Automatic needle segmentation and localization in MRI with 3-D convolutional neural networks: application to MRI-targeted prostate biopsy. *IEEE Trans. Med. Imag.* 38(4), 1026–1036 (2019). <https://doi.org/10.1109/tmi.2018.2876796>
12. Mohanty, S.P., et al.: Deep learning for understanding satellite imagery: an experimental survey. *Front. Artif. Intell.* 3, 1–21 (2020). <https://doi.org/10.3389/frai.2020.534696>
13. Xue, D., et al.: Dim small target detection based on convolutional neural network in star image. *Multimed. Tool. Appl.* 79(7–8), 4681–4698 (2019). <https://doi.org/10.1007/s11042-019-7412-z>
14. Ronneberger, O., Fischer, P., Brox, T.: U-net: convolutional networks for biomedical image segmentation. In: *International Conference on Medical Image Computing and Computer-assisted Intervention*, pp. 234–241. Springer (2015)
15. Mastrofini, M., Agostinelli, I., Curti, F.: Design and validation of a U-net-based algorithm for star sensor image segmentation. *Appl. Sci.* 13(3), 1947 (2023). <https://doi.org/10.3390/app13031947>
16. Du, L., Lyu, G., Shi, Y.: ISAR image semantic segmentation based on GAN. *Radar Sci. Technol.* 19(5), 479–484 (2021)
17. Wang, F., et al.: Satellite target recognition method based on deep learning. In: *2022 IEEE 5th Advanced Information Management, Communicates, Electronic and Automation Control Conference (IMCEC)*, vol. 5, pp. 112–116 (2022)
18. Al Khalil, Y., et al.: Reducing segmentation failures in cardiac MRI via late feature fusion and GAN-based augmentation. *Comput. Biol. Med.* 161, 106973 (2023). <https://doi.org/10.1016/j.combiomed.2023.106973>
19. Chen, Y., et al.: Learning semantic segmentation from synthetic data: a geometrically guided input-output adaptation approach. In: *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1841–1850 (2019)
20. van Rooij, S.B., et al.: Determination of satellite solar panel and bus size from radar images with deep learning. In: *3rd IAA Conference on Space Situational Awareness (ICSSA)*, International Academy of Astronautics (2022)
21. Hariharan, B., Girshick, R.B.: Low-shot visual recognition by shrinking and hallucinating features. In: *IEEE International Conference on Computer Vision (ICCV)*, pp. 3037–3046 (2016)
22. Tian, Y., et al.: Rethinking few-shot image classification: a good embedding is all you need? In: *ECCV 2020: Computer Vision*, 266–282. Springer-Verlag (2020)
23. Sun, X., et al.: Few-shot learning for domain-specific fine-grained image classification. *IEEE Trans. Ind. Electron.* 68(4), 3588–3598 (2021). <https://doi.org/10.1109/tie.2020.2977553>
24. Dong, N., Xing, E.P.: Few-shot semantic segmentation with prototype learning. In: *British Machine Vision Conference* (2018)
25. Zhang, C., et al.: CANet: class-agnostic segmentation networks with iterative refinement and attentive few-shot learning. In: *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5212–5221 (2019)
26. Rakelly, K., et al.: Few-shot segmentation propagation with guided networks. *arXiv* (2018), abs/1806.07373
27. Tavera, A., et al.: Pixel-by-pixel cross-domain alignment for few-shot semantic segmentation. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 1626–1635 (2022)
28. Otten, M.P.G., et al.: Multi-channel processing for digital beam forming SAR. In: *21st European Signal Processing Conference (EUSIPCO 2013)*, pp. 1–5 (2013)
29. Lucassen, R.T., et al.: Deep learning for detection and localization of B-lines in lung ultrasound. *IEEE J. Biomed. Health Inform.* 27(9), 1–10 (2023). <https://doi.org/10.1109/jbhi.2023.3282596>
30. Heslinga, F.G., et al.: Corneal pachymetry by AS-OCT after Descemet's membrane endothelial keratoplasty. *Sci. Rep.* 11(1), 13976 (2021). <https://doi.org/10.1038/s41598-021-93186-9>
31. He, K., et al.: Deep residual learning for image recognition. In: *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778 (2016)
32. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv:1412.6980.2014* (2014)
33. Ganin, Y., et al.: Domain-adversarial training of neural networks. *J. Mach. Learn. Res.* 17(59), 1–35 (2016)
34. Kamnitsas, K., et al.: Unsupervised domain adaptation in brain lesion segmentation with adversarial networks. *Inform. Process. Med. Imag.* 25, 597–609 (2016)
35. Lafarge, M.W., et al.: Domain-adversarial neural networks to address the appearance variability of histopathology images. *Lect. Notes Comput. Sci.* 10553 LNCS, 83–91 (2017)
36. Wang, S., et al.: Progressive adversarial networks for fine-grained domain adaptation. In: *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9210–9219 (2020)
37. Hoffman, J., et al.: FCNs in the wild: pixel-level adversarial and constraint-based adaptation. *arXiv* (2016), abs/1612.02649
38. Buciluă, C., Caruana, R., Niculescu-Mizil, A.: Model compression. In: *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Association for Computing Machinery, pp. 535–541 (2006) <https://doi.org/10.1145/1150402.1150464>
39. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. 1–9 (2015), [arXiv:1503.02531](https://arxiv.org/abs/1503.02531)
40. Zagoruyko, S., Komodakis, N.: Paying more attention to attention: improving the performance of convolutional neural networks via attention transfer. In: *5th International Conference on Learning Representations (ICLR)*, pp. 1–13 (2017)
41. Heslinga, F.G., et al.: Approximation of a pipeline of unsupervised retina image analysis methods with a CNN. In: *Proceedings of SPIE 10949, Medical Imaging 2019: Image Processing*, p. 10949N (2019)
42. Noothout, J.M.H., et al.: Knowledge distillation with ensembles of convolutional neural networks for medical image segmentation. *J. Med. Imag.* 9(5), 052407 (2022). <https://doi.org/10.1117/1.jmi.9.5.052407>
43. Tavera, A.: Taveraantonio/PixDA: Official Code for pixel-by-pixel cross-domain alignment for few-shot semantic segmentation. <https://github.com/taveraantonio/PixDA> (2022). Accessed 1 Dec 2022
44. Zhang, J., et al.: Few-shot structured domain adaptation for virtual-to-real scene parsing. In: *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, pp. 9–17 (2019)
45. Keaton, M.R., Zaveri, R.J., Doretto, G.: CellTranspose: few-shot domain adaptation for cellular instance segmentation. In: *2023 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 455–466 (2023)
46. Kalluri, T., Chandraker, M.: Cluster-to-adapt: few shot domain adaptation for semantic segmentation across disjoint labels. In: *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pp. 4120–4130 (2022)

How to cite this article: Heslinga, F.G., et al.: Few-shot learning for satellite characterisation from synthetic inverse synthetic aperture radar images. *IET Radar Sonar Navig.* 1–8 (2024). <https://doi.org/10.1049/rsn2.12516>