

SIGNAL CLUSTERING WITH CLASS-INDEPENDENT SEGMENTATION

Stefano Gasperini^{1,2} Magdalini Paschali¹ Carsten Hopke² David Wittmann² Nassir Navab^{1,3}

¹ Computer Aided Medical Procedures, Technical University of Munich, Germany

² Airbus Defence and Space GmbH, Manching, Germany

³ Computer Aided Medical Procedures, Johns Hopkins University, Baltimore, USA

ABSTRACT

Radar signals have been dramatically increasing in complexity, limiting the source separation ability of traditional approaches. In this paper we propose a Deep Learning-based clustering method, which encodes concurrent signals into images, and, for the first time, tackles clustering with image segmentation. Novel loss functions are introduced to optimize a Neural Network to separate the input pulses into pure and non-fragmented clusters. Outperforming a variety of baselines, the proposed approach is capable of clustering inputs directly with a Neural Network, in an end-to-end fashion.

Index Terms— Clustering, Deep Learning, Radar Signals, Image Segmentation, Class-independence

1. INTRODUCTION

Radars are used for situation awareness in a variety of applications, from weather forecasting to adaptive cruise control. Aircrafts utilize them also as a sensor, inspecting the received signals to get insights about their surroundings. Towards this end, signals are initially isolated and then compared against prior knowledge for identification. In this paper, we focus on separating simultaneous and aligned signals by source.

Radar signals are composed of pulses, described by time of arrival (TOA), radio frequency (RF), pulse width (PW), amplitude (AM) etc. [1]. This task is challenging, especially because the separation is done at pulse level, based on a few parameters, that can be shared across multiple sources.

Furthermore, due to the rapid progress of antenna technologies and electronics [1], traditional methods cannot cope anymore with the increased complexity of incoming signals. Based on statistics and pattern matching, they struggle even with a few concurrent inputs [2]. To overcome this issue, the task has been addressed with data analysis clustering methods, among which DBSCAN reached superior performance and wide applicability [3, 4].

Over the past years, Deep Learning (DL) has achieved state-of-the-art results in a plethora of tasks, including clustering [5]. In the imaging domain, Neural Networks (NNs)




Fig. 1. Domain change from signal to image processing.

are trained to map the inputs to a clustering-friendly representation [5–8]. Afterwards, these newly extracted features are clustered with a commonly used data analysis method [5, 9], such as K-means [10] or Hierarchical agglomerative clustering [11]. The same idea has been applied to the speaker separation task, using K-means [12]. Pairing a NN with traditional clustering methods leads to improved performance, but inherits shortcomings from both approaches, such as requiring a predefined number of clusters and a sufficiently large dataset for training a NN. Additionally, these methods would cluster pulses based on the similarity of their parameters, missing the aim of our task, which is grouping them by source. Furthermore, mapping separable features requires large input dimensions, while our pulses have only a few parameters.

In radar signal processing, NNs have been mostly applied to address classification and identification problems [13–15]. Recently, Recurrent NNs (RNNs) have been deployed to cluster pulses [16]. Each RNN was trained to identify a specific signal and checked at every possible starting pulse of a sequence. Tested on 5 simple signals [16], this method would not be feasible in real world applications with thousands of emitters, each capable of producing multiple signals. It would require long processing times and could not separate those signals for which an RNN has not been trained.

The contribution of this paper can be summarized as follows: 1) We propose an original NN-based clustering method and we evaluate it on the challenging task of aligned radar pulse deinterleaving; 2) We introduce novel loss functions, aimed at delivering pure and non-fragmented clusters, while performing class-independent segmentation; 3) To the best of our knowledge, for the first time, we address clustering as an image segmentation problem.

2. METHOD

The idea is changing the problem domain towards imaging, to benefit from the extensive DL research done in the field. Image segmentation is the task of separating pixels into regions according to some given criteria, alike clustering, which groups elements based on similarity. Towards this end, we aim at exploiting image segmentation techniques to perform clustering and tackle the radar pulses deinterleaving problem.

Our proposed method enables a NN to cluster directly input elements through image segmentation. In this paper, we apply it on radar signals deinterleaving. Towards this end, we first encode the signals into spectrogram-inspired segmentable images (Section 2.1). These are then forwarded to a U-Net [17] trained for image segmentation to group the inputs. The NN is optimized with novel loss functions derived from a newly adjusted confusion matrix, that we named soft confusion matrix (Section 2.3). The objective functions aim at improving clustering performance indicators such as purity and fragmentation [18] (Section 2.2).

2.1. From Signals to Segmentable Images

To cross from the radar to the imaging domain, we need to encode the signals into images. Spectrograms are graphical representations of signals, showing RF over time. For radars, keeping the PW resolution within reasonable images of 512x512 pixels, would visualize too few pulses per signal, if any, complicating the clustering task.

Signals encoding. Towards this end, our representation is inspired by spectrograms, as we encode concurrent signals into an image, using it as a RF-TOA coordinate system. We keep a time resolution of $5 \mu s$, which allows for the encoding of a reasonable amount of pulses within each 512x512 image. We indicate the pulses by marking in the RF-TOA grid those (rf, toa) cells corresponding to incoming pulses at time toa , parametrized with an RF value of rf . Pixels have gray-scale values representing PW and AM in two separate channels. The encoding is schematized for PW in Fig. 1, the same is done for AM. We maximize the resolution by scaling each parameter individually to cover all its available range: $[0, 512]$ for RF, and $(0, 1]$ for PW and AM. This further diversifies the values, easing the separation task. Moreover, we improve readability, by extending the values across the neighboring 3x3 pixels. Overall, with this transformation no major loss of information occurs, and discretization can help against noise.

Our problem differs substantially from traditional image segmentation, since instead of identifying a preset amount of classes, we deal with an unknown and variable number of clusters. Nevertheless, we can retrieve the cluster assignments from a segmented image, by remembering the pulses location from the encoding step, filtering the background and applying majority voting within each 3x3 region.




Fig. 2. Confusion matrix and properties for evaluation.

2.2. Confusion Matrix-based Metrics

We choose a U-Net architecture [17], due to its state-of-the-art performance in a variety of image segmentation tasks. While typically U-Nets are trained with pixel-wise cross-entropy or Dice loss, these functions are designed for a fixed amount of classes and are not suitable for clustering.

Instead, to address clustering problems, we derive new loss functions, that maximize performance metrics. Such indicators can be formalized with the help of a confusion matrix [18], like the one showed in Fig. 2. The matrix has the predictions along the horizontal axis, with the cells showing the elements count. With C_i^j we refer to the (i, j) cell of the confusion matrix C , sized $G \times P$; with C_i and C^j to its i -th row and j -th column respectively.

Cluster Purity. The first step is identifying the column maxima $M(C) := \{\max(C^j) \forall j\}$, shown green in Fig. 2. This matches each predicted cluster with a ground truth. In the example, what was predicted as cluster "4" corresponds to "B". Cluster purity is defined as $cp = \sum M(C) / \sum C$: it assesses how many elements define the identity of each predicted cluster. A system would score perfect purity predicting a different cluster for each input element, while performing rather poor.

Fragmentation Ratio. To address these cases, we additionally consider fragmentation ratio. It compares the number of fragments, marked with a circle in Fig. 2, with the amount of predicted clusters P . It is defined as $cfr = \sum_i \max(0, |C_i \cap M(C)| - 1) / P$. Fragments occur when a row contains more than one column maxima.

2.3. Overcoming Non-differentiability

The two metrics are non-differentiable: they are extracted from a confusion matrix, which is based on an argmax function over the assignment probabilities. To circumvent this issue, we take a step back and compute the losses directly from the softmax probabilities along the predicted clusters.

Soft Confusion Matrix. Towards this end, we create a new matrix, that we name soft confusion matrix (SCM). Its construction process is shown in Fig. 3. Instead of reporting the discrete argmax assignments, we build the SCM with the softmax probabilities for each input element. We report the averages within its 3x3 pixels along its corresponding SCM row. After repeating this for all input elements, we complete the SCM by summing the probabilities accumulated within each

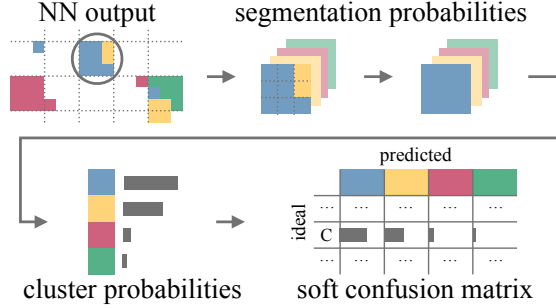


Fig. 3. Construction of the soft confusion matrix.

cell. We construct the SCM using only differentiable operations, so we can compute our losses with its values. Following the notation of Section 2.2, S is the SCM, sized $G \times N$, with N being the amount of output channels of the NN.

Soft Cluster Impurity Loss. From the cluster purity metric (Section 2.2), we derive its differentiable complement from the SCM. This is the cluster impurity loss, which maximizes the column maxima $M(S)$. It is defined as sci_l in Eq. 1.

$$sci_l = \frac{\sum Q(S)}{\sum S} \quad (1)$$

$$Q(S) := \{S_i^j\}: S_i^j \notin M^j(S) \quad \forall i, j$$

Soft Cluster Fragmentation Loss. Analogously, soft fragmentation loss scf_l corresponds to the fragmentation metric (Section 2.2), which is computed with non-differentiable counting functions. We make it suitable for training, by applying the trick in Eq. 2, using only differentiable operations.

$$scf_l = \frac{\sum F(S) \oslash F(S)}{N} \quad (2)$$

$$F(S) := \{S_i^j\}: S_i^j \in R_i(S) \wedge S_i^j \neq \max(R_i(S)) \quad \forall i, j$$

$$R_i(S) := S_i \cap M(S)$$

With reference to Eq. 2: $F(S)$ is the set containing the SCM cells responsible for fragments, which are penalized by this loss; $R_i(S)$ is the set of SCM column maxima within the i -th row; \oslash is the Hadamard element-wise division. Therefore $F(S) \oslash F(S)$ is a unit vector with a length corresponding to the amount of fragments, and the scf_l numerator is the differentiable equivalent of counting them.

Additionally, at each training iteration we randomly swap the cluster target values. With this strategy, we ensure the NN focuses on grouping, rather than learning specific associations. This way we also achieve class-independence.

The key of our method is the combination of the two aforementioned novel loss functions, which enables predicting the clusters directly, and addressing the clustering problem as an image segmentation task.

3. EXPERIMENTAL SETUP

Dataset. Since there was no suitable public dataset available, we created one with the help of a domain expert. 140 real-

sci_l	scf_l	ARI	NMI
0	1	0.177 ± 0.021	0.292 ± 0.043
1	0	0.530 ± 0.021	0.635 ± 0.014
0.5	0.5	0.523 ± 0.018	0.634 ± 0.013
0.7	0.3	0.552 ± 0.014	0.656 ± 0.010

Fig. 4. Ablative testing. Different configurations of the loss functions, indicated on the left columns, are evaluated with ARI and NMI. The last row shows the proposed method.

istic signals were defined and generated with an RF software simulator. Training and test sets were designed separately, increasing the complexity: 75 signals are utilized for training, 65 for testing. To create the dataset, we randomly made concurrent signals from different sources. Since a good clustering method should deal with a variable amount of clusters, random 2.56 ms portions of a signal appear alone, as well as with 1 to 10 other signals a variable amount of times. Concurrent signals are at most 11, all aligned, coming from the same direction, to be distinguished through TOA, RF, PW and AM. Overall, 1200 combinations are generated from the training signals and 300 from the testing, with the amount of pulses varying significantly throughout the dataset (min: 6, max: 1032, average: 467).

Architecture. A U-Net [17] was used for all experiments. To increase its robustness against overfitting and improve its suitability for embedded devices, we reduced its size: feature map channels start at 8 and reach 64 in the bottleneck layer. Since segmentation requires fixed amounts of classes, we set the concurrent predictable clusters to $N = 15$, following [19].

Model Training. Across all experiments, hyperparameters remained constant, to ensure comparability. The U-Net was trained for 300 epochs with Adam optimizer, a learning rate of 10^{-5} and unit batch size. We optimized for the novel losses described in Section 2.3. We implemented our method in PyTorch and trained our models on an NVIDIA Titan Xp GPU.

Evaluation Metrics. The evaluation is based on the criteria presented in Section 2.2: cluster purity (cp) and (non-) fragmentation ratio ($cnfr$). On top of pure and non-fragmented clusters, we want to ensure that each ground truth has a corresponding predicted cluster. Specifically, a cluster remains uncovered (marked red in Fig. 2) when its confusion matrix row contains no column maxima $M(C)$. This is measured by the detection ratio, which is defined as $cdr = 1 - |U|/G$, with $U := \{i\}: C_i \cap M(C) = \emptyset$. These three metrics should be considered combined, since they provide a good assessment of the clustering performance from different important aspects. Additionally, we deployed standard clustering evaluation metrics, namely Adjusted Rand Index (ARI) and Normalized Mutual Information (NMI). We repeated all experiments five times, reporting mean and standard deviation.

	purity	detection	non-fragm.	ARI	NMI
OPTICS	0.958	0.927	0.280	0.254	0.561
AffProp	0.775	0.891	0.383	0.287	0.574
Mean Shift	0.532	0.476	0.976	0.350	0.558
DBSCAN	0.586	0.567	0.958	0.413	0.622
BIRCH	0.632	0.653	0.902	0.463	0.641
Proposed	0.752 ± 0.021	0.840 ± 0.030	0.734 ± 0.034	0.552 ± 0.014	0.656 ± 0.010

Fig. 5. Comparison with baseline methods across five performance metrics.

Ablative Testing. In order to showcase the effectiveness of the main components of our method, we performed ablative testing. We evaluated the impact on the clustering performance of different loss function configurations.

Baseline Comparison. Furthermore, we compared our method against various clustering approaches. Despite being popular and effective in other domains, several methods, such as K-means, are not applicable to the task at hand, since they require a predefined number of clusters, which in our case is unknown and varies over time. We deployed five suitable baseline methods, namely DBSCAN, OPTICS, BIRCH, Affinity propagation and Mean shift [20, 21].

4. RESULTS AND DISCUSSION

Ablative Testing. Fig. 4 showcases that combining the losses can improve the results regarding both ARI and NMI. The purity loss achieves satisfying results when deployed by itself too, since its optimization acts on every cell of the confusion matrix (Section 2.2), trying to improve all predictions. On the other hand, the fragmentation loss focuses only on a few cells, rendering it unable to deliver good overall performance by itself. However, optimizing them jointly in a weighted fashion of 0.7 for purity and 0.3 for fragmentation delivers higher ARI and NMI. The fragmentation loss is important: it contributes by penalizing the prediction of unnecessary clusters.

Comparative Methods. Fig. 5 highlights the superior performance of our method compared to the other clustering approaches. Most baseline methods suffer from either over predicting the amount of clusters (high cp , low $cn.fr$), or underestimating it (low cp , high $cn.fr$). BIRCH stands out among the other baseline approaches. Nevertheless, our method outperforms BIRCH as well with an improvement of 1.5% for NMI and 8.9% for ARI. This is accounted to the learning capacity of our method, which is trained to improve both cp and $cn.fr$ jointly. Repeating the experiments with the baseline approaches leads to the same clusters. Our method, despite relying on a random initialization of its model weights, is able to deliver consistent results with a small standard deviation.

Metrics Trade-off. As can be seen in Fig. 5, there is often a

trade-off between purity and non-fragmentation ratio scores. The two metrics are evaluating the performance from different perspectives and it is challenging to achieve high scores for both, especially in a task such as radar pulse deinterleaving. Instead, getting a perfect score for only one of the two is trivial, although decreasing the other metric substantially. Such unbalanced performances can be achieved by predicting a single cluster for maximum $cn.fr$, or one per input for cp . The strength of the proposed method lies in its ability to optimize both metrics simultaneously, improving two complementary aspects.

Cross Domain Applicability. The results achieved show good potential for a previously unexplored way of solving clustering problems. Despite being applied on the pulses deinterleaving task, the method can be extended to other domains: the novel soft confusion matrix-based loss functions could pave the way towards new NN-based clustering approaches that could operate in an end-to-end fashion, without requiring conventional data analysis methods. This would alleviate shortcomings of traditional clustering approaches, such as dealing with a predefined number of clusters, or dependencies to cluster shapes and point densities.

5. CONCLUSION AND FUTURE WORK

In this paper we proposed a new DL-based clustering method, which we applied on the challenging task of aligned radar pulse deinterleaving. For the first time, clustering was targeted as an image segmentation problem. We changed domain by encoding the concurrent signals into segmentable images. The NN was trained to predict the clusters directly in an end-to-end fashion, aiming at pure and non-fragmented clusters, thanks to new loss functions derived from a novel probability-based confusion matrix.

Furthermore, we plan on scheduling different loss weightings over training and exploring other input representations, suitable for a even larger amount of simultaneous and aligned signals. Moreover, we will adapt the proposed method and losses to be suitable for different domains, such as image clustering, comparing ours against other DL-based approaches.

6. REFERENCES

- [1] J.A. Scheer, M.A. Richards, W.A. Holm, and W.L. Melvin, *Principles of Modern Radar*, SciTech Publishing, 2013.
- [2] Kaveer Manickchand, Johannes J Strydom, and AK Mishra, “Comparative study of TOA based emitter deinterleaving and tracking algorithms,” in *2017 IEEE AFRICON*. IEEE, 2017, pp. 221–226.
- [3] Xiaofeng Wang, Xuzhou Zhang, Runlan Tian, and Xinglong Qi, “A new method of unknown radar signals sorting,” in *Proceedings of 2013 Chinese Intelligent Automation Conference*. Springer, 2013, pp. 727–733.
- [4] Shad Mahmood, “Deinterleaving pulse trains with DB-SCAN and FART,” Master thesis, Uppsala Universitet, Uppsala, Sweden, 2019.
- [5] Elie Aljalbout, Vladimir Golkov, Yawar Siddiqui, Maximilian Strobel, and Daniel Cremers, “Clustering with deep learning: Taxonomy and new methods,” *arXiv preprint arXiv:1801.07648*, 2018.
- [6] Junyuan Xie, Ross Girshick, and Ali Farhadi, “Unsupervised deep embedding for clustering analysis,” in *Proceedings of the 33rd International Conference on Machine Learning (ICML)*, 2016, pp. 478–487.
- [7] Peihao Huang, Yan Huang, Wei Wang, and Liang Wang, “Deep embedding network for clustering,” in *2014 22nd International Conference on Pattern Recognition (ICPR)*. IEEE, 2014, pp. 1532–1537.
- [8] Kamran Ghasedi Dizaji, Amirhossein Herandi, Cheng Deng, Weidong Cai, and Heng Huang, “Deep clustering via joint convolutional autoencoder embedding and relative entropy minimization,” in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 5736–5745.
- [9] Erxue Min, Xifeng Guo, Qiang Liu, Gen Zhang, Jianjing Cui, and Jun Long, “A survey of clustering with deep learning: From the perspective of network architecture,” *IEEE Access*, vol. 6, pp. 39501–39514, 2018.
- [10] Bo Yang, Xiao Fu, Nicholas D Sidiropoulos, and Mingyi Hong, “Towards K-Means-friendly spaces: Simultaneous deep learning and clustering,” in *Proceedings of the 34th International Conference on Machine Learning (ICML)*, 2017, pp. 3861–3870.
- [11] Jianwei Yang, Devi Parikh, and Dhruv Batra, “Joint unsupervised learning of deep representations and image clusters,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 5147–5156.
- [12] John R Hershey, Zhuo Chen, Jonathan Le Roux, and Shinji Watanabe, “Deep clustering: Discriminative embeddings for segmentation and separation,” in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 31–35.
- [13] Lindsay Cain, Jeffrey Clark, Eric Pauls, Ben Ausdenmoore, Richard Clouse, and Ted Josue, “Convolutional neural networks for radar emitter classification,” in *2018 IEEE Annual Computing and Communication Workshop and Conference (CCWC)*. IEEE, 2018, pp. 79–83.
- [14] Jan Matuszewski, “Radar signal identification using a neural network and pattern recognition methods,” in *2018 International Conference on Advanced Trends in Radioelectronics, Telecommunications and Computer Engineering (TCSET)*. IEEE, 2018, pp. 79–83.
- [15] Hongbo Li, Wei Jing, and Yang Bai, “Radar emitter recognition based on deep learning architecture,” in *2016 CIE International Conference on Radar (RADAR)*. IEEE, 2016, pp. 1–5.
- [16] Zhang-Meng Liu and S Yu Philip, “Classification, denoising and deinterleaving of pulse streams with recurrent neural networks,” *IEEE Transactions on Aerospace and Electronic Systems*, 2018.
- [17] Olaf Ronneberger, Philipp Fischer, and Thomas Brox, “U-Net: Convolutional networks for biomedical image segmentation,” in *International Conference on Medical image computing and computer-assisted intervention (MICCAI)*. Springer, 2015, pp. 234–241.
- [18] Michael J Thompson, Samuel Lin, and John C Sciortino Jr, “Measures of effectiveness for analysis of radar pulse train deinterleavers,” in *Signal Processing, Sensor Fusion, and Target Recognition XV*. International Society for Optics and Photonics, 2006, vol. 6235, p. 62351I.
- [19] Yanick Lukic, Carlo Vogt, Oliver Dürr, and Thilo Stadelmann, “Speaker identification and clustering using convolutional neural networks,” in *2016 IEEE 26th international workshop on machine learning for signal processing (MLSP)*. IEEE, 2016, pp. 1–6.
- [20] Rui Xu and Donald C. Wunsch II, “Survey of clustering algorithms,” *IEEE Trans. Neural Networks*, vol. 16, no. 3, pp. 645–678, 2005.
- [21] Dongkuan Xu and Yingjie Tian, “A comprehensive survey of clustering algorithms,” *Annals of Data Science*, vol. 2, no. 2, pp. 165–193, 2015.