

LABEL-AWARE RANKED LOSS FOR ROBUST PEOPLE COUNTING USING AUTOMOTIVE IN-CABIN RADAR

Lorenzo Servadei^{*◦} Huawei Sun^{*◦} Julius Ott^{*◦} Michael Stephan^{*#} Souvik Hazra^{*}
Thomas Stadelmayer^{*#} Daniela Sánchez Lopera^{*◦} Robert Wille[†] Avik Santra^{*}

^{*}Infineon Technologies AG

[◦]Technical University of Munich

[†]Johannes Kepler University Linz

[#]Friedrich-Alexander-University Erlangen-Nuremberg

ABSTRACT

In this paper, we introduce the Label-Aware Ranked loss, a novel metric loss function. Compared to the state-of-the-art Deep Metric Learning losses, this function takes advantage of the ranked ordering of the labels in regression problems. To this end, we first show that the loss minimises when datapoints of different labels are ranked and laid at uniform angles between each other in the embedding space. Then, to measure its performance, we apply the proposed loss on a regression task of people counting with a short-range radar in a challenging scenario, namely a vehicle cabin. The introduced approach improves the accuracy as well as the neighboring labels accuracy up to 83.0% and 99.9%: An increase of 6.7% and 2.1% on state-of-the-art methods, respectively.

Index Terms— Deep Metric Learning, People Counting, Radar Signal Processing

1. INTRODUCTION

The task of people counting is defined as predicting the number of people present in a given scene. Its direct application in real-life scenarios made it particularly suitable for occupancy estimation [1], surveillance [2], traffic management [3], and several other fields [4, 5]. To this end, several solutions have been developed, particularly in the area of Computer Vision (CV), as in [6–8]. These approaches have shown positive results even on large crowds. Nonetheless, privacy preservation, as well as weather conditions independence are still major concerns for their implementation in real-life scenarios. Accordingly, radar sensors offer a valid alternative that overcome the limitations mentioned above.

However, even though radar-based solutions are resistant to such limitations, their drawbacks are caused by issues such as low-resolution data which decrease the target-identification rate, missed detection as a result of occlusion, and unstable radar signal strength due to the superposition of reflections coming from various parts of the body. These challenges make the use of short-range and low-cost radars ineffective when used in conventional approaches for people counting in dense scenarios.

To overcome this, seminal works have been utilizing Deep Learning (DL) approaches for advancing limitations given by traditional signal processing methods, as [9, 10]. Although the work done shows advantages in terms of people counting

performance, it still presents important challenges. In fact, the methods utilised there do not take advantage of the labels ranking implicit in the people counting task.

Additionally, the experimental scenarios presented are less prone to reflections and false targets w.r.t. small, close environments as automotive cabins, train carriages, etc. Nevertheless, these use-cases are very common in areas such as transportation and have been of particular importance for Heating, Ventilation, and Air Conditioning (HVAC) automated regulation.

In this paper, we present a novel loss function, namely the Label-Aware Ranked (LAR) loss, which leads to an improved DL solution for people counting, performed on a challenging scenario inside a vehicle. The presented loss function takes advantage of recent advancements in supervised Deep Metric Learning (DML), a set of Machine Learning (ML) methods whose goal is to learn such an embedding space in which similar sample pairs stay close while dissimilar ones are far apart. Work done in this area, so far, has been focusing mostly on multi-class classification tasks, enhancing the decision margin among embedded vectors of unrelated classes [11–14]. Instead, this contribution approaches the shaping of the embedding space in a regression problem. To this end, distance information among labels is exploited to have an increasingly ranked embedding space. Additionally, to improve the robustness of the proposed solution, we enhance the generalization phase performance by adding an exponential moving average filter. Taking into account previous predictions enhance the stability of the estimated people counting.

As a result, the approach presented shows an increased accuracy up to 83.0% and a neighboring labels accuracy (i.e. *accuracy +/-1*) up to 99.9%, thus respectively +6.7% and +2.1% w.r.t. the state-of-the-art methods on a real-world, complex people counting dataset.

2. BACKGROUND AND RELATED WORK

In this section, we first review the radar signal processing methods utilized in related problem settings, and afterwards we analyse recent advancements in DML which increase the performance of ML models.

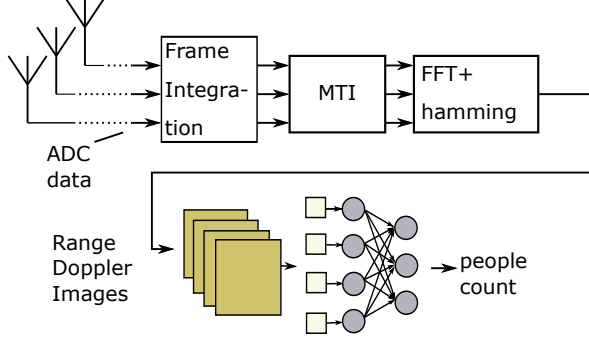


Fig. 1. Block Diagram with Preprocessing

2.1. Radar Signal Processing

Frequency Modulated Continuous Wave (FMCW) radars allow to estimate range, velocity, and angles of targets. For better interpretability, a preprocessing chain is typically used to extract some of these parameters from the sampled intermediate frequency signal before feeding it into the neural network. Different such preprocessing chains are given in [15], in [16] they use Range-Doppler Images (RDIs) as network inputs for object detection, while in [17] they work directly on the time domain data, but implicitly generate RDIs in the first network layer. Fig. 1 shows one such preprocessing chain. First, a $N_S \times N_C$ 2D-dataframe is formed by acquiring the N_S real samples from the intermediate frequency signal for each of the N_C chirps, and stacking them column-wise. To increase the velocity resolution, a slow-time-frame with a larger observation period and therefore a higher velocity resolution is built, by integrating over the chirps of each frame, and by stacking N_C of the resulting vectors to another 2D-dataframe of the same size. Afterwards, the mean values are subtracted along with the chirps as a moving target indication / high-pass-filter, to get rid of the Tx/Rx Leakage and to remove any completely static targets. As a last preprocessing step, both 2D-dataframes are transferred to frequency domain via 2D-FFTs. To reduce the sidelobe levels of each reflecting target, before doing the respective FFTs, the dataframes are multiplied with hamming windows of the respective sizes along the sample and chirp dimension. The 6-channel input to the neural network then consists of the real and the imaginary parts for the Range-Doppler Images of each antenna.

Although the pre-processing method showed good outcomes for different radar data use-cases, the final prediction outcome still suffers from instability. In fact, in tasks like people counting, where the count usually only changes by one or stays the same on a frame-by-frame basis, a temporal smoothing filter will help the network to provide a more stable and reliable prediction. An example of this is the Exponential Smoothing (ES), where the time data is smoothed with an exponential window function as seen in Eq. 1, where $x[k]$ is the current frame, $x_s[k-1]$ the smoothed frame from the last time-step k , and $x_s[k]$ the new smoothed frame.

$$x_s[k] = \alpha x[k] + (1 - \alpha)x_s[k-1] \quad (1)$$

While those methods have been important for many tasks related to radar data as people counting and tracking, the prediction mechanism has been relying more and more on DL.

2.2. Deep Metric Learning

Embedding vectors are lower dimension vectors that represent high dimensional input data, such as images. DML is an area of DL that aims to learn data embedding vectors reducing the distance between samples of the same class, whereas the distance between samples of dissimilar classes is increased. In this paper, we consider the state-of-the-art DML losses for benchmarking our proposed LAR loss. Triplet loss [18] uses the Euclidean distance to measure the similarity between two embedding vectors. However, when the number of samples and classes becomes large, this loss becomes computationally expensive. To overcome this drawback, the Multiclass N-Pair (Mc-N-Pair) loss [19] replaces the Euclidean distance with a dot product as a measurement to improve the computation efficiency. The Constellation loss [20] preserves the triplet structure, but it considers more negative classes than the Triplet loss during the update, for better convergence. In the following, we provide an overview of these DML losses and their properties.

Triplet Loss Triplet loss [18] shown in Eq. 2, considers positive and negative pairs together. In every update step, there is an anchor x_i^a , a positive sample x_i^p , which has the same label as the anchor, as well as a negative sample x_i^n , which has a different label as the anchor. Then, the input triplet $\{x_i^a, x_i^p, x_i^n\}$ is transformed to an embedding vector $\{f_i^a, f_i^p, f_i^n\}$. In this case, the Euclidean norm between the transformed samples w.r.t the anchor sample is given by $E_p = \|f_i^a - f_i^p\|_2$ and $E_n = \|f_i^a - f_i^n\|_2$ for positive and negative samples respectively.

$$L_{tri} = \frac{1}{N} \sum_{i=1}^N \max(0, E_p^2 - E_n^2 + m), \quad (2)$$

where m is the distance margin, and N is the batch size.

Triplet loss aims at minimising the distance between the anchor and the positive sample E_p , and maximizing the distance between the anchor and the negative sample E_n at the same time.

Multiclass-N-pair Loss Triplet loss still suffers from slow convergence because in each update step only an anchor, a positive, and a negative sample are considered. Mc-N-Pair loss [19] extends the Triplet loss, using all the labels for each update, as shown in Eq. 3.

$$L_{Mc-N-pr} = \frac{1}{N} \sum_{i=1}^N \log \left(1 + \sum_{j \neq i} \exp(f_i^T f_j^n - f_i^T f_j^p) \right). \quad (3)$$

This property improves the convergence, but if the number of labels is too high, it leads to computational inefficiency. This can further affect the training of the network.

Constellation Loss The Constellation loss [20], shown in Eq. 4, combines the advantages of Mc-N-Pair loss and Triplet loss. Compared to Triplet loss, it takes more negative labels into account, while it follows the formulation of the Mc-N-Pair loss.

$$L_{cons} = \frac{1}{N} \sum_{i=1}^N \log \left(1 + \sum_{j=1}^K \exp(f_i^{aT} f_j^n - f_i^{aT} f_j^p) \right), \quad (4)$$

where K is an optional hyperparameter that indicates how many negative labels are considered for each update.

Although the presented DML losses have been used in multiple tasks, to the best of our knowledge no loss takes distances among labels into account for ordering the embedding space and improving the regression prediction.

3. PROPOSED APPROACH

To address these issues, in this section we present the proposed loss function, namely the LAR loss, and successively demonstrate its theoretical foundations and properties. To this end, we show that it minimises when the rank among labels (i.e. number of people in a scene) is preserved, and the angles among embeddings of different labels are uniform.

LAR takes inspiration from the Constellation loss but additionally takes advantage of the labels' information to reproduce their ranking in the embedding space, thus enhancing the prediction capabilities of the models. The LAR loss is presented in Eq. 5.

$$L_{LAR} = \frac{1}{N} \sum_{i=1}^N \log \left(1 + \sum_{j \neq i} \exp(\log(\Delta_l) f_i^{aT} f_j^n - f_i^{aT} f_j^p) \right), \quad (5)$$

where

$$\Delta_l = \min(|l_a - l_n|, |L - |l_a - l_n||). \quad (6)$$

The loss uses the multiplier $\log(\Delta_l)$ to regulate the ranking of the labels. Here, l_a is the label of the anchor, l_n is the label of the current negative sample and L is the number of different labels. Here, f_i identifies the embedding of the input sample. The multiplier assigns smaller values to neighbouring labels and establishes a distance metric among labels. The logarithm function is applied to it, as it is monotonically increasing and adds numerical stability.

In LAR, we use normalised feature vectors, i.e. $\langle f_i, f_j \rangle = \cos(\theta)$, thus our loss operates on the angles between the feature vectors. Similar to Mc-N-Pair and Constellation loss, in the LAR loss, the embedding vectors of the same label are pushed to an angle of $\theta = 0$, which minimises the loss. In the following, we show the properties of our LAR loss. Regarding different labels, we show that having ranked labels and uniform angles minimises the LAR loss.

By Jensen's inequality for convex functions, i.e.

$$\frac{\sum_{i=1}^I e^{\cos(\theta_i)}}{I} \geq e^{\frac{\sum_{i=1}^I \cos(\theta_i)}{I}}, \quad (7)$$

without loss of generality, we can examine only the angles among the labels that minimise the $\sum \cos(\theta_i)$. Furthermore, the cosine is minimal at $\theta = \pi$. Thus, in the case of an even L , we expect the angle between labels with the highest multiplier to be π . This can be only achieved by a uniform angle $2\pi/L$ among all the labels. In the case of an odd L , we start with the assumption that our points lie on the unit hyper-sphere in uniform angles. Then, we show that the sum of cosines with uniform angles between classes is always smaller than the same sum where one of the points is shifted by an ϵ on the circle which fulfils $0 < \epsilon < 2\pi/L$. Shifting only one angle by ϵ is the smallest perturbation possible and by showing that this will increase the loss, we cover the cases of multiple perturbations as well. Since not all cosine terms are affected by the shift of one point, Eq. 8 states only the difference.




Fig. 2. Constellation Loss Embeddings on MNIST

$$\begin{aligned} & \sum_{l=3}^L 2 \sum_{j=1}^{\frac{l-1}{2}} 2 \log(j) \cos\left(j \frac{2\pi}{l}\right) \\ & < \sum_{l=3}^L \sum_{j=1}^{\frac{l-1}{2}} 2 \log(j) \cos\left(j \frac{2\pi}{l} - \epsilon\right) + 2 \log(j) \cos\left(j \frac{2\pi}{l} + \epsilon\right), \quad (8) \end{aligned}$$

Now we can apply the properties of the cosine and obtain

$$\sum_{l=3}^L 2 \sum_{j=1}^{\frac{l-1}{2}} 2 \log(j) \cos\left(j \frac{2\pi}{l}\right) < \sum_{l=3}^L 2 \cos(\epsilon) \sum_{j=1}^{\frac{l-1}{2}} 2 \log(j) \cos\left(j \frac{2\pi}{l}\right). \quad (9)$$

By the symmetry of the cosine, the uniform angles and the ordering of our multiplier, $\sum_{j=1}^{\frac{l-1}{2}} 2 \log(j) \cos\left(j \frac{2\pi}{l}\right) < 0$ holds for any odd $l \geq 3$. Hence, we can divide by the inner sum which yields

$$\sum_{l=3}^L 1 > \sum_{l=3}^L \cos(\epsilon). \quad (10)$$

As known, $\cos(0) = \cos(2\pi) = 1$ are the maxima of the cosine, and these upper bounds are never reached by any $\cos(\epsilon)$ with $0 < \epsilon < 2\pi/L$ and $L \geq 3$. For an even L , the inequality changes to

$$\begin{aligned} & \sum_{l=4}^L 2 \sum_{j=1}^{\frac{l}{2}-1} 2 \log(j) \cos\left(j \frac{2\pi}{l}\right) + \frac{l}{2} \cos(\pi) \\ & < \sum_{l=4}^L 2 \cos(\epsilon) \sum_{j=1}^{\frac{l}{2}-1} 2 \log(j) \cos\left(j \frac{2\pi}{l}\right) + \frac{l}{2} \cos(\pi - \epsilon). \quad (11) \end{aligned}$$

Since, $\cos(\pi) < \cos(\pi - \epsilon)$ is always true for any $0 < \epsilon < 2\pi/L$ we can exclude it from the inequality which yields

$$\sum_{l=6}^L 2 \sum_{j=1}^{\frac{l}{2}-1} 2 \log(j) \cos\left(j \frac{2\pi}{l}\right) < \sum_{l=6}^L 2 \cos(\epsilon) \sum_{j=1}^{\frac{l}{2}-1} 2 \log(j) \cos\left(j \frac{2\pi}{l}\right). \quad (12)$$

Now the same reasoning from the odd case applies here. This shows that our loss is minimal for uniform angles and when the ranking is preserved.

To show the effect of the LAR loss onto a simple dataset, we applied it to the MNIST dataset [21]. The embedded space resulting from the Constellation loss is presented in Fig. 2, while the effect of the LAR loss is shown in Fig. 3. There we observe that the LAR loss ensures uniform angles and order-




Fig. 3. LAR Loss Embeddings on MNIST

ing between the class mean and shows higher discriminative power between the classes.

To encourage stability on the people counting task, we add ES to stabilize the inference output of the network.

In this repository¹, we show the effect of LAR on uniform angles and labels ranking, on a randomly generated dataset and MNIST.

In the next section, we apply the proposed LAR loss on a real-world dataset and show how the implicit ordering of the embedding would help in a regression problem. In fact, by constructing an ordered representation on the embedding space, we implicitly express the ranking of the labels and support the model in the final prediction.

4. EXPERIMENTAL

In this section, we first review the implementation settings, successively we benchmark our proposed loss function with and without temporal smoothing in an ablation study.

4.1. Implementation Settings

In the implementation, we used PyTorch v1.8.0TM - GPU v2.4.0 with CUDA[®] Toolkit v11.1.0 and cuDNN v8.0.5. As a processing unit, we used the Nvidia[®] Tesla[®] P40 GPU, Intel[®] Core i7-8700K CPU, and DIMM 16GB DDR4-3000 module of RAM. In order to count people, one Infineon’s XENSIVTM 60 GHz sensor has been utilized, on the internal-upper side of the front window of the vehicle cabin. The radar input data has been pre-processed as explained in Section 2, and results in 95000 frames of scenes of people counting inside a vehicle, performed on zero to five people, and divided into recordings with an average length of 350 frames. The frames are recorded with a frame rate of 10 Hz. The dataset has been split into training and test set, dividing it into 76000 (training) and 19000 (testing) frames. In order to benchmark state-of-the-art losses, we create a smart batch structure, as mentioned in [19]. This means, every batch contains two samples per label, thus accounting for a batch size of 12. Additionally, we take advantage of an equal number of samples for each label. This corresponds to around 16k samples for each class. For the people counting task, we utilize a network with the same architecture as proposed in the encoder of [9],

¹LAR: <https://github.com/2Geeks2/LabelAwareRanked-Loss.git>

composed of three convolutional layers with ReLU activation and a pooling layer after each convolutional layer. Each of the convolutional layers has 32 feature maps and a kernel size of 3×3 . As the last layer we use a fully connected ReLU layer where we round the output for the prediction.

4.2. Ablation Study

In order to show the outcome of our approach, we implement the proposed methods in an ablation study. To this end, we analyse different losses and benchmark them on the aforementioned people counting dataset, as shown in Table 1. Here, both *accuracy* and *accuracy +/-1* are shown. While *accuracy* points out the overall accuracy on the test set, the *accuracy +/-1* score takes into account the two neighboring labels (for the label 0 and 5, only one is considered). To obtain the accuracy measure, we round the regression predictions obtained using the Mean Squared Error (MSE) loss. In our benchmark, we first compare our LAR loss to the state-of-the-art DML losses. Additionally, we compare the same losses by enhancing and stabilizing their outcome using ES.

Table 1. Benchmark of DML Losses

	Accuracy	Accuracy +/-1
MSE	68.6%	94.5%
MSE + Triplet	70.6%	95.3%
MSE + Mc-N-Pair	68.9%	95.7%
MSE + Constellation	74.6%	97.2%
MSE + LAR (Ours)	80.8%	98.5%
MSE + ES	71.9%	96.4%
MSE + Triplet + ES	72.7%	97.1%
MSE + Mc-N-Pair + ES	72.8%	97.3%
MSE + Constellation + ES	76.3%	97.8%
MSE + LAR + ES (Ours)	83.0%	99.9%

Table 1 shows the results for our proposed approach against the state-of-the-art. Our proposed method shows an accuracy of 80.8% and a +/-1 accuracy of 98.5%, respectively +6.2% and +1.3% towards the second best performing loss (i.e. MSE + Constellation loss). Adding the ES, our method reaches an accuracy of 83.0% and a +/-1 accuracy of 99.9%, respectively +6.7% and +2.1% towards the second best performing loss (i.e. MSE + Constellation loss + ES). By ranking the latent space and ordering the embedding, we show the advantages in prediction performance, even when the neighboring labels are considered.

5. CONCLUSION

In this work, we formulated a novel loss function, namely LAR loss. First demonstrated that it minimises when samples are ranked in label orders in the embedding space, and they are separated by uniform angles between labels. Afterward, we show how these properties can help in the people counting task on a difficult scenario of a vehicle cabin. To this end, we first train a CNN with LAR and MSE losses. Afterward, we increase the robustness of the prediction using ES. As a result, our approach presents an increment in accuracy and accuracy +/- 1 of respectively +6.7% and +2.1% towards the second-best performing loss.

6. REFERENCES

- [1] Oliver Shih and Anthony Rowe, "Occupancy estimation using ultrasonic chirps," in *Proceedings of the ACM/IEEE Sixth International Conference on Cyber-Physical Systems*, 2015, pp. 149–158.
- [2] Enwei Zhang and Feng Chen, "A fast and robust people counting method in video surveillance," in *2007 International Conference on Computational Intelligence and Security (CIS 2007)*, 2007.
- [3] Fei Liu, Zhiyuan Zeng, and Rong Jiang, "A video-based real-time adaptive vehicle-counting system for urban roads," *PLOS ONE*, 11 2017.
- [4] Mengxiao Tian, Hao Guo, Hong Chen, Qing Wang, Chengjiang Long, and Yuhao Ma, "Automated pig counting using deep learning," *Computers and Electronics in Agriculture*.
- [5] Prithvi N. Amin, Sayali S. Moghe, Sparsh N. Prabhakar, and Charusheela M. Nehete, "Deep learning based face mask detection and crowd counting," 2021.
- [6] Khalil Khan, Rehan Ullah Khan, Waleed Albattah, Durre Nayab, Ali Mustafa Qamar, Shabana Habib, and Muhammad Islam, "Crowd counting using end-to-end semantic image segmentation," *Electronics*, 2021.
- [7] Youmei Zhang, Chunluan Zhou, Faliang Chang, Alex C. Kot, and Wei Zhang, "Attention to head locations for crowd counting," in *Image and Graphics*, Yao Zhao, Nick Barnes, Baoquan Chen, Rüdiger Westermann, Xi-angwei Kong, and Chunyu Lin, Eds. 2019, Springer International Publishing.
- [8] Junyu Gao, Qi Wang, and Yuan Yuan, "Scar: Spatial/channel-wise attention regression networks for crowd counting," *Neurocomputing*, vol. 363, pp. 1–8, 2019.
- [9] Cem Yusuf Aydogdu, Souvik Hazra, Avik Santra, and Robert Weigel, "Multi-modal cross learning for improved people counting using short-range fmcw radar," in *2020 IEEE International Radar Conference (RADAR)*, 2020.
- [10] Jae-Ho Choi, Ji-Eun Kim, Nam-Hoon Jeong, Kyung-Tae Kim, and Seung-Hyun Jin, "Accurate people counting based on radar: Deep learning approach," in *2020 IEEE Radar Conference (RadarConf20)*, 2020.
- [11] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song, "Sphereface: Deep hypersphere embedding for face recognition," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 6738–6746.
- [12] Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Dihong Gong, Jingchao Zhou, Zhifeng Li, and Wei Liu, "Cosface: Large margin cosine loss for deep face recognition," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018.
- [13] Wanping Zhang, Yongru Chen, Wenming Yang, Guijin Wang, Jing-Hao Xue, and Qingmin Liao, "Class-variant margin normalized softmax loss for deep face recognition," *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–6, 2020.
- [14] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou, "Arcface: Additive angular margin loss for deep face recognition," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4690–4699.
- [15] Avik Santra and Souvik Hazra, *Deep Learning Applications of Short-Range Radars*, Artech House, 2020.
- [16] Guoqiang Zhang, Haopeng Li, and Fabian Wenger, "Object detection and 3d estimation via an FMCW radar using a fully convolutional network," in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. may 2020, IEEE.
- [17] Michael Stephan, Thomas Stadelmayer, Avik Santra, Georg Fischer, Robert Weigel, and Fabian Lurz, "Radar image reconstruction from raw ADC data using parametric variational autoencoder with domain adaptation," in *2020 25th International Conference on Pattern Recognition (ICPR)*. jan 2021, IEEE.
- [18] Kilian Q Weinberger, John Blitzer, and Lawrence Saul, "Distance metric learning for large margin nearest neighbor classification," in *Advances in Neural Information Processing Systems*, Y. Weiss, B. Schölkopf, and J. Platt, Eds. 2006, vol. 18, MIT Press.
- [19] Kihyuk Sohn, "Improved deep metric learning with multi-class n-pair loss objective," in *Advances in Neural Information Processing Systems*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, Eds. 2016, vol. 29, Curran Associates, Inc.
- [20] Alfonso Medela and Artzai Picón, "Constellation loss: Improving the efficiency of deep metric learning loss functions for optimal embedding," *CoRR*, vol. abs/1905.10675, 2019.
- [21] Li Deng, "The MNIST database of handwritten digit images for machine learning research," *IEEE Signal Processing Magazine*, vol. 29, no. 6, pp. 141–142, 2012.