Attentive Max Feature Map for Acoustic Scene Classification with Joint Learning considering the Abstraction of Classes

Hye-jin Shim¹, Ju-ho Kim¹, Jee-weon Jung², and Ha-Jin Yu^{1†}

¹School of Computer Science, University of Seoul, Republic of Korea

²Naver Corporation, Republic of Korea

shimhz6.6@gmail.com, wng1187@naver.com, jeejeon.jung@navercorp.com, hjyu@uos.ac.kr

Abstract

The attention mechanism has been widely adopted in acoustic scene classification. However, we find that during the process of attention exclusively emphasizing information, it tends to excessively discard information although improving the performance. We propose a mechanism referred to as the attentive max feature map which combines two effective techniques, attention and max feature map, to further elaborate the attention mechanism and mitigate the abovementioned phenomenon. Furthermore, we explore various joint learning methods that utilize additional labels originally generated for subtask B (3-classes) on top of existing labels for subtask A (10-classes) of the DCASE2020 challenge. We expect that using two kinds of labels simultaneously would be helpful because the labels of the two subtasks differ in their degree of abstraction. Applying two proposed techniques, our proposed system achieves state-of-the-art performance among single systems on subtask A. In addition, because the model has a complexity comparable to subtask B's requirement, it shows the possibility of developing a system that fulfills the requirements of both subtasks; generalization on multiple devices and low-complexity.

Index Terms: acoustic scene classification, attention, joint learning

1. Introduction

Acoustic scene classification (ASC) is the task of recognizing scenes based on environmental sounds. The recent studies on environmental audio analysis are led by the detection and classification of acoustic scenes and events (DCASE) community, which covers a variety of tasks such as ASC, sound event detection, and audio captioning [1–4]. In the DCASE 2020 challenge, ASC is divided into two subtasks: 1) subtask A for generalization across different devices and 2) subtask B, the newly released task, for a low-complexity solution in terms of the model size. Subtask A aims to identify a given audio clip recorded by multiple devices into one of the ten pre-defined acoustic scenes: airport, indoor shopping mall, metro station, pedestrian street, public square, a street with a medium level of traffic, tram, bus, metro, park. Subtask B aims to classify a given audio clip into one of the three higher level classes: outdoor, indoor, and transportation. Hereafter, the task for classifying the class of subtask B is defined as a 3-class classification, whereas that for classifying the class of subtask A is defined as a 10-class classification.

Recent studies on ASC can be largely divided into two strands: data preprocessing and modeling. In data preprocessing, most studies exploit delta and delta-deltas for performance improvements [5–9]. Various data augmentation strategies such as mixup [10], SpecAugment [11], pitch shift, speed change and

adding random noise have also been explored. Many studies have divided the input with several frequency bands and used the separated data as inputs [5, 12, 13]. Feature investigation and temporal division have also been explored [14–16].

In modeling, ResNet and convolutional neural network (CNN) are the most widely used architecture in ASC task. Teacher-student network can be exploited to ease frequently mis-classified classes [17, 18]. The proper size of the receptive field adjusting the kernel size have also been explored [9]. Attention, which is widely used in other domains, has been applied in various ways [19, 20]. Utilizing the representation of other tasks and training a general-purpose network have been investigated considering the characteristics of DCASE dealing with various tasks [21–26]. In the latest DCASE 2020 challenge, top ranking systems in ASC mostly focused on data preprocessing techniques.

In this paper, we propose two modeling techniques independent of data-preprocessing; an attentive max feature map (AMFM) and joint learning of the concrete classes of subtask A and abstract classes of subtask B. First, we find that although attention improves the performance, it excessively loses information owing to its exclusive nature. Our hypothesis is that a fraction of de-emphasized information might further help the discriminative power of a feature if adequately addressed. We find that the max feature map (MFM), a technique that adopts a competitive scheme via element-wise max operation [27], can be used in this perspective. MFM is also effective for ASC as studied in previous works [21, 28, 29]. Specifically, we design a new technique combining attention and MFM and refer to it as the AMFM where attention emphasizes the most informative region and the MFM prevents excessive information loss using a relative max operation. Inspired by MFM operation, AMFM compares the feature maps before and after the attention to mitigate the excessive information loss.

Furthermore, we propose the joint learning of 10-class of subtask A and 3-class of subtask B to improve the performance of subtask A. Many studies related to the DCASE explore the use of auxiliary tasks and general-purpose network [21–26, 30]. Hu et al. [6, 31] utilized the prediction of the classifier of subtask B to improve the performance of subtask A; however, two classifiers were completely trained separately. In this study, since the labels of the two subtasks differ in the degree of abstraction, we assume that training by using the two labels together would be helpful. Therefore, we propose the joint learning of the two subtasks and explore various architectures and balancing between two subtasks. To the best of our knowledge, this is the first approach to simultaneously train an ASC system using both labels of subtask A and B. The proposed system has relatively low-complexity and achieves a performance comparable to that of a state-of-the-art system without complex data preprocessing. The experimental results show that both the gen-

[†] Corresponding author

Figure 1: AMFM block structure (left) and overall pipeline of the joint learning structure of extended MTL considering class abstraction (right). The AMFM block in the left portion consists of both the MFM and the AMFM. The MFM and the AMFM are illustrated as pink and green boxes, respectively. The structure in the right portion shows the overall pipeline of the proposed method, in which a layer for each task is added to perform additional learning in the conventional MTL. Extended MTL structure demonstrates the best result compared to the other joint learning strategies. The conventional MTL structure omits the dotted line blocks, while the extended MTL includes the dotted line blocks.

eralization to multiple devices and the low-complexity can be achieved in a single system.

2. Proposed method

2.1. Attentive max feature map

MFM was originally proposed for image classification where noisy labels are dominant [27]. Wu et al. [27] argued that a rectified linear unit (ReLU) activation separates noisy and informative signals using a threshold (or bias), which might lead to information loss, especially for the first several convolution layers. To overcome this issue, MFM was proposed to replace the non-linear activation function with a competitive scheme via an element-wise max operation. The max operation of MFM selects the optimal feature at each location learned by different filters and helps separate noisy and informative signals. The previous works demonstrate that MFM is also effective for ASC [21, 28, 29].

The implementation of MFM operation can be denoted as follows. Let a be an output feature map of a convolution layer, $a \in \mathbb{R}^{K \times T \times F}$, where K , T , and F refer to the number of output channels, time domain frames, and frequency bins, respectively. a is split into two feature maps, a_1 and a_2 , where $a_1, a_2 \in \mathbb{R}^{\frac{K}{2} \times T \times F}$. The MFM applied feature map is obtained by applying $\text{Max}(a_1, a_2)$, element-wise. In the left portion of Figure 1, the pink box includes the conventional MFM operation.

The attention mechanism has proven its effectiveness on ASC [8, 19, 20]. It can highlight important information and enrich representation. Among various attention mechanisms, the convolutional block attention module (CBAM) considers both channel and spatial attention module and has the advantage of seamless implementation regardless of the architecture. Our previous work confirmed that applying the CBAM to the MFM also improves the performance [28]. However, as a result of visualizing the feature map, we empirically find that excessive

information loss occurs, emphasizing only narrow parts when attention is applied (Figure 2-(b)). We hypothesize that a fraction of de-emphasized information might help the discriminative power of a feature.

Accordingly, we propose the application of the attention mechanism in a competitive manner, referred to as AMFM to mitigate abovementioned excessive deletion. Inspired by MFM operation, AMFM compares two feature maps, before and after attention element-wise. The AMFM can be denoted as $\text{Max}(a_1, \text{CBAM}(a_1))$, element-wisely. We compose an AMFM block, involving MFM and AMFM, and it is illustrated in Figure 1. MFM and AMFM are the pink and green box of left portion of Figure 1, respectively. AMFM is not only applicable to MFM, it can also be used in all architectures to which attention is applied. In this study, the effectiveness is validated only in MFM. We plan to apply AMFM for other architectures in our future work. Figure 2 shows that AMFM could attenuate exaggerated attention and select salient representation by comparing the feature maps at the same time.

2.2. Joint learning

A few studies have considered the relationship between the two subtasks of ASC. Although Hu et al. [6] utilized joint prediction, classifiers were completely trained separately. In the joint prediction, the classifiers for the two subtasks are trained separately, and the final prediction is calculated using the score fusion of the two classifiers. However, many studies proved that training the related tasks simultaneously helps improve the performance [32, 33]. Since the labels of the two subtasks differ in the degree of abstraction, we assume that training by using the two labels together would be helpful. Therefore, we propose the joint training of two subtasks.

In the experiments, we compare four different methods and vary the weights between the two tasks for the joint learning. First, we adopt the pre-training method. The system is first trained for the 3-class classification and is then fine-tuned with

Figure 2: Each feature map indicates before (a) and after (b) attention, and the output of AMFM (c). These are extracted from the first AMFM block at positions (a) to (c) in the left portion of Figure 1. It demonstrates that attention can cause excessive information loss. After applying the AMFM with a comparison of the two feature maps of (a) and (b), the features are found to enrich the representation preventing the information loss.

the 10-class classification. Second, the conventional multi-task learning (MTL) [34] is applied for the joint learning of the two subtasks. In this case, the network is designed to learn the shared representation and to identify each task only in the last layer. Third, the extended MTL with the additional fine-tuning of each task is exploited. The extended MTL refers to allocating additional layer(s) for each training task after the last hidden layer that the two tasks share. Lastly, we investigate the order of the training of the two tasks, according to the hierarchical relationship of the two tasks. In the conventional and extended MTL methods, the classification layer of the two subtasks is separated only at the end of the system. On the contrary, in the sequential MTL, a layer classifying one subtask is connected to another subtask classification layer. We also compared the order of 3-class and 10-class classifications and find that it is more effective to train the 3-class classification first, followed by the 10-class classification. This is in line with the deep learning structure that trains abstract tasks in the front layer and specific tasks in the back layer [35]. The sequential MTL is premised on the hierarchical relationship between the two tasks, among other MTL structures.

For further improvement, we apply the joint prediction [6]. In terms of adjusting the weight between the two tasks, we explore both intuitive and methodological approaches: a grid search and GradNorm [36].

3. Experimental settings and results

3.1. Experimental settings

For all experiments, DCASE 2020 task 1 subtask A (1-A) dataset is used. A total of 13,965 and 2,970 audio clips are used for the training and validation, respectively. We do not use the DCASE 2020 task 1 subtask B (1-B) audio clips but use only corresponding labels for the 3-class classification. DCASE 2020 task 1-A dataset consists of various audio clips collected from three real devices (A, B, and C) and six simulated devices (s1-s6). Only devices A-C and s1-s3 are used in the training set and s4-s6 are not available in the training phase. Each audio clip has a duration of 10 s with a 44.1kHz sampling rate and 24-bit resolution.

For each data, 256-dimensional Mel-spectrograms are extracted. A short-time Fourier transform with 2048 FFT points is applied, using 40 ms window size and 20 ms hop length.

Table 1: Performance comparison result according to the application of attention and the structure of CNN, MFM, and AMFM.

System	Attention	Acc (%)
CNN w/ LeakyReLU	X	69.6
	O	68.2
MFM	X	69.4
	O	70.4
AMFM	O	70.8

Mixup [10] and SpecAugment [11] are exploited for data augmentation. Other data preprocessing techniques such as applying logarithm, deltas, double-deltas, and sub-band frequency separation have not been used in this work. All models are implemented with Pytorch, a deep learning library in Python. The initial learning rate is set to 0.001 and scheduled with a warm restart of the stochastic gradient descent. SGD optimizer with a momentum of 0.9 is used. The batch size and the number of epochs are set to 24 and 800, respectively. The architecture details are similar to those in our previous work [28]. In the blocks added for the MTL, the parameters of the AMFM block are as same as the other blocks. The last hidden layer for each task has 100 nodes followed by the output layer for each subtask, respectively.

3.2. Result analysis

3.2.1. Attentive max feature map

Table 1 describes the results of comparing the effects of applying attention and DNN structures: CNN, MFM, and AMFM. When attention is applied to the CNN with Leaky ReLU, the performance decreased, but when attention is applied to MFM, the performance significantly improved. When the proposed AMFM is applied, an additional performance improvement is achieved compared to the MFM with an accuracy of 70.8%. Figure 2-(c) illustrates the AMFM result, which can emphasize important information while preventing information loss. The joint learning experiments are conducted using the AMFM structure which shows the best result.

Table 2: Comparison of various joint learning strategies considering the difference in the level of abstraction for the two different subtasks.

System	Joint prediction	# Params	Acc (%)
Pre-training	X	1.5M	69.2
Conventional MTL	X	1.5M	69.7
Extended MTL	X	0.6M	71.3
	O	0.6M	70.0
Sequential MTL	X	0.7M	71.0
	O	0.7M	69.1
Separated Classifier [6]	O	1.5M	69.4

Table 3: Experimental results adjusting the joint learning weight ratio between the 3-class and 10-class classifications.

System	Ratio	Acc (%)
Proposed method	1 : 1	70.3
	1 : 2	69.6
	1 : 3	70.3
	1 : 4	70.7
	1 : 5	71.3
GradNorm [36]	-	70.1

3.2.2. joint learning

Table 2 shows the comparison of various joint learning strategies. Pre-training has first explored among various methods. The pre-training needs to perform two training steps, training for 3-class task followed by the fine-tuning for 10-class task, and the optimization of adjusting the learning rate and model selection are time-consuming. Therefore, we did no more adjustment of other variables on this method. With the conventional MTL, there is no performance improvement. Using extended MTL, which has additional layers for each training task after the last hidden layer that the two tasks share, the best result without joint prediction was achieved with fewer parameters. The sequential MTL without joint prediction showed a slight improvement in performance compared to the 70.8% AMFM. Although the joint prediction showed its effectiveness in [6], the performance improvement could not be confirmed in our experiments on joint learning strategies applied to the AMFM structure. The experiment results imply that the conventional MTL may not be able to learn each task enough because each task output has only one classification layer, and the sequential MTL has a high dependence between the two tasks because they share many hidden layers. This proves that even if there is a relationship between tasks, performance can decrease when excessive interference occurs as reported in [35]. On the contrary, the extended MTL additionally learns about individual tasks with more hidden layers, thus it can reduce task interference and yields the best performance.

Table 3 indicates the results of adjusting the weight ratio between the two subtasks for the joint learning. We adjust the weight ratio between the two tasks using both manual and automatic approaches. The best result is achieved when the ratio is 1 : 5 for the abstract and specific labels.

Table 4 shows the accuracy of 3-class classification. The MTL is well known to be effective when related tasks are jointly

Table 4: Accuracy of the 3-class classification.

System	MTL	Extended MTL	Sequential MTL	Separated Classifiers
Acc (%)	89.9	92.2	90.3	91.4

Table 5: Comparison with the state-of-the-art systems. It only considers a single system, not an ensemble system. Suh et al. to Kountini et al. represent the first to fifth ranking systems in order on the DCASE2020 task 1-A.

System	Acc (%)	# Params
Proposed Method	71.3	0.6M
DCASE2020 Baseline [4]	54.1	5M
Suh et al. [5]	73.7	13M
Hu et al. [6]	74.6	-
Gao et al. [7]	71.8	4M
Liu et al. [8]	72.1	3M
Koutini et al. [9]	71.8	225M

trained, but it is difficult to show if the two tasks are related. Therefore, when the relevance between tasks is low, MTL often causes a performance degradation compared to training each task individually. To verify the relationship between 3-class and 10-class classifications, we also investigate the performance of the 3-class classification, although our goal is to improve the 10-class classification task. The accuracy of the DCASE2020 baseline of subtask B is 88% and the average of the submitted system is 87.3% [4]. In our experiment, when training with AMFM structure without joint learning, the accuracy is 91.3% for 3-class and 70.8% for 10-class. Compared with the aforementioned result, when joint learning is applied, it is proved that the two subtasks are related in that we can get the best accuracy of 92.2% for 3-class and 71.3% for 10-class.

3.2.3. Comparison with the recent studies

Table 5 presents a comparison with the top five state-of-the-art systems, without the application of ensemble. Our proposed system demonstrates comparable performance with state-of-the-art systems without complex data preprocessing techniques. In addition, our method achieves such performance with a low-complexity architecture. Although it is outside the scope of this study, we expect that applying more data preprocessing methods could lead to further improvements in the future.

4. Conclusions

In this paper, we proposed the AMFM technique and joint learning considering information loss and the abstraction level of classes. First, we proposed the AMFM which selects salient features in a competitive manner, avoiding information loss. Second, we proposed the joint learning of the 10-class classification of subtask A and the 3-class classification of subtask B. Experimental results demonstrated the effectiveness of our proposed methods. Although our works aimed to improve the performance of subtask A, we achieved comparable performance with the state-of-the-art systems even in low-complexity comparable to subtask B. Data preprocessing in various ways could be complementary with our proposed method and we plan to utilize both for further improvement in our future work.

5. References

- [1] A. Mesaros, T. Heittola, E. Benetos, P. Foster, M. Lagrange, T. Virtanen, and M. D. Plumbley, "Detection and classification of acoustic scenes and events: Outcome of the dcase 2016 challenge," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 2, pp. 379–393, 2017.
- [2] A. Mesaros, T. Heittola, and T. Virtanen, "A multi-device dataset for urban acoustic scene classification," *arXiv preprint arXiv:1807.09840*, 2018.
- [3] —, "Acoustic scene classification in dcase 2019 challenge: Closed and open set classification and data mismatch setups," 2019.
- [4] T. Heittola, A. Mesaros, and T. Virtanen, "Acoustic scene classification in dcase 2020 challenge: generalization across devices and low complexity solutions," *arXiv preprint arXiv:2005.14623*, 2020.
- [5] Y. Liu, A. Neophytou, S. Sengupta, and E. Sommerlade, "Cross-modal spectrum transformation network for acoustic scene classification," Tech. Rep., 2021.
- [6] H. Hu, C.-H. H. Yang, X. Xia, X. Bai, X. Tang, Y. Wang, S. Niu, L. Chai, J. Li, H. Zhu *et al.*, "Device-robust acoustic scene classification based on two-stage categorization and data augmentation," *arXiv preprint arXiv:2007.08389*, 2020.
- [7] W. Gao, M. McDonnell, and S. UniSA, "Acoustic scene classification using deep residual networks with focal loss and mild domain adaptation," Technical Report, Tech. Rep., 2020.
- [8] J. Liu, "Acoustic scene classification with residual networks and attention mechanism," Tech. Rep., DCASE2020 Challenge, Tech. Rep., 2020.
- [9] K. Koutini, F. Henkel, H. Eghbal-zadeh, and G. Widmer, "Cp-jku submissions to dcase'20: Low-complexity cross-device acoustic scene classification with rf-regularized cnns," DCASE2020 Challenge, Tech. Rep., Tech. Rep., 2020.
- [10] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, "mixup: Beyond empirical risk minimization," *arXiv preprint arXiv:1710.09412*, 2017.
- [11] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, "SpecAugment: A simple data augmentation method for automatic speech recognition," *Proc. Interspeech 2019*, pp. 2613–2617, 2019.
- [12] S. S. R. Phayre, E. Benetos, and Y. Wang, "Subspectralnet—using sub-spectrogram based convolutional neural networks for acoustic scene classification," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 825–829.
- [13] H. Wang, D. Chong, and Y. Zou, "Acoustic scene classification with multiple decision schemes," Technical Report, DCASE2020 Challenge, Tech. Rep., 2020.
- [14] S. Mun, S. Shon, W. Kim, D. K. Han, and H. Ko, "Deep neural network based learning and transferring mid-level audio features for acoustic scene classification," in *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2017, pp. 796–800.
- [15] L. Pham, H. Phan, T. Nguyen, R. Palaniappan, A. Mertins, and I. McLoughlin, "Robust acoustic scene classification using a multi-spectrogram encoder-decoder framework," *Digital Signal Processing*, vol. 110, p. 102943, 2021.
- [16] Y. Sakashita and M. Aono, "Acoustic scene classification by ensemble of spectrograms based on adaptive temporal divisions," *Detection and Classification of Acoustic Scenes and Events (DCASE) Challenge*, 2018.
- [17] H.-S. Heo, J.-w. Jung, H.-j. Shim, and H.-J. Yu, "Acoustic scene classification using teacher-student learning with soft-labels," *arXiv preprint arXiv:1904.10135*, 2019.
- [18] J.-W. Jung, H.-S. Heo, H.-J. Shim, and H.-J. Yu, "Knowledge distillation in acoustic scene classification," *IEEE Access*, vol. 8, pp. 166 870–166 879, 2020.
- [19] Z. Ren, Q. Kong, J. Han, M. D. Plumbley, and B. W. Schuller, "Attention-based atrous convolutional neural networks: Visualisation and understanding perspectives of acoustic scenes," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 56–60.
- [20] H. Phan, O. Y. Chén, L. Pham, P. Koch, M. De Vos, I. McLoughlin, and A. Mertins, "Spatio-temporal attention pooling for audio scene classification," *arXiv preprint arXiv:1904.03543*, 2019.
- [21] H.-j. Shim, J.-h. Kim, J.-w. Jung, and H.-j. Yu, "Audio tagging and deep architectures for acoustic scene classification: Uos submission for the dcase 2020 challenge," Tech. Rep., DCASE2020 Challenge, Tech. Rep., 2020.
- [22] J.-w. Jung, H.-j. Shim, J.-h. Kim, S.-b. Kim, and H.-J. Yu, "Acoustic scene classification using audio tagging," *arXiv preprint arXiv:2003.09164*, 2020.
- [23] J.-h. Kim, J.-w. Jung, H.-j. Shim, and H.-j. Yu, "Audio tag representation guided dual attention network for acoustic scene classification," in *Proceedings of the Detection and Classification of Acoustic Scenes and Events 2020 Workshop (DCASE2020)*, 2020.
- [24] E. Fonseca, M. Plakal, F. Font, D. P. Ellis, X. Favory, J. Pons, and X. Serra, "General-purpose tagging of freesound audio with audioset labels: Task description, dataset, and baseline," *arXiv preprint arXiv:1807.09902*, 2018.
- [25] Q. Kong, T. Iqbal, Y. Xu, W. Wang, and M. D. Plumbley, "Dcase 2018 challenge surrey cross-task convolutional neural network baseline," *arXiv preprint arXiv:1808.00773*, 2018.
- [26] J.-w. Jung, H.-j. Shim, J.-h. Kim, and H.-J. Yu, "Dcasenet: A joint pre-trained deep neural network for detecting and classifying acoustic scenes and events," *arXiv preprint arXiv:2009.09642*, 2020.
- [27] X. Wu, R. He, Z. Sun, and T. Tan, "A light cnn for deep face representation with noisy labels," *IEEE Transactions on Information Forensics and Security*, vol. 13, no. 11, pp. 2884–2896, 2018.
- [28] H.-j. Shim, J.-w. Jung, J.-h. Kim, and H.-j. Yu, "Capturing scattered discriminative information using a deep architecture in acoustic scene classification," *arXiv preprint arXiv:2007.04631*, 2020.
- [29] Y. Lee, S. Lim, and I.-Y. Kwak, "Cnn-based acoustic scene classification system," *Electronics*, vol. 10, no. 4, p. 371, 2021.
- [30] S. Deshmukh, B. Raj, and R. Singh, "Multi-task learning for interpretable weakly labelled sound event detection," *arXiv preprint arXiv:2008.07085*, 2020.
- [31] H. Hu, C.-H. H. Yang, X. Xia, X. Bai, X. Tang, Y. Wang, S. Niu, L. Chai, J. Li, H. Zhu *et al.*, "A two-stage approach to device-robust acoustic scene classification," *arXiv preprint arXiv:2011.01447*, 2020.
- [32] A. R. Zamir, A. Sax, W. Shen, L. J. Guibas, J. Malik, and S. Savarese, "Taskonomy: Disentangling task transfer learning," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 3712–3722.
- [33] K. Dwivedi and G. Roig, "Representation similarity analysis for efficient task taxonomy & transfer learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 12 387–12 396.
- [34] R. Caruana, "Multitask learning," *Machine learning*, vol. 28, no. 1, pp. 41–75, 1997.
- [35] K.-K. Maninis, I. Radosavovic, and I. Kokkinos, "Attentive single-tasking of multiple tasks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 1851–1860.
- [36] Z. Chen, V. Badrinarayanan, C.-Y. Lee, and A. Rabinovich, "Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks," in *International Conference on Machine Learning*. PMLR, 2018, pp. 794–803.