

HKUST SPD - INSTITUTIONAL REPOSITORY

Title	Randomized Quantization: A Generic Augmentation for Data Agnostic Self-supervised Learning
Authors	Wu, Huimin; Lei, Chenyang; Sun, Xiao; Wang, Peng-Shuai; Chen, Qifeng; Cheng, Kwang Ting; Lin, Stephen; Wu, Zhirong
Source	Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), 2023, p. 16305-16316
Conference	IEEE/CVF International Conference on Computer Vision (ICCV 2023), Paris, France, 2-6 October 2023
Version	Accepted Version
DOI	
Publisher	IEEE
Copyright	© 2023 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

This version is available at HKUST SPD - Institutional Repository (<https://repository.hkust.edu.hk/ir>)

If it is the author's pre-published version, changes introduced as a result of publishing processes such as copy-editing and formatting may not be reflected in this document. For a definitive version of this work, please refer to the published version.

Randomized Quantization: A Generic Augmentation for Data Agnostic Self-supervised Learning

Huimin Wu^{1*} Chenyang Lei^{2*} Xiao Sun⁴ Peng-Shuai Wang³
Qifeng Chen¹ Kwang-Ting Cheng¹ Stephen Lin⁵ Zhirong Wu⁵

¹HKUST ²CAIR, HKISL-CAS ³Peking University ⁴Shanghai AI Lab ⁵Microsoft Research Asia

Abstract

Self-supervised representation learning follows a paradigm of withholding some part of the data and tasking the network to predict it from the remaining part. Among many techniques, data augmentation lies at the core for creating the information gap. Towards this end, masking has emerged as a generic and powerful tool where content is withheld along the sequential dimension, e.g., spatial in images, temporal in audio, and syntactic in language. In this paper, we explore the orthogonal channel dimension for generic data augmentation by exploiting precision redundancy. The data for each channel is quantized through a non-uniform quantizer, with the quantized value sampled randomly within randomly sampled quantization bins. From another perspective, quantization is analogous to channel-wise masking, as it removes the information within each bin, but preserves the information across bins. Our approach significantly surpasses existing generic data augmentation methods, while showing on par performance against modality-specific augmentations. We comprehensively evaluate our approach on vision, audio, 3D point clouds, as well as the DABS benchmark which is comprised of various data modalities. The code is available at https://github.com/microsoft/random_quantize.

1. Introduction

We are witnessing a convergence of multi-modal AI [23, 6] where the architecture and the learning algorithm are unified for various data modalities. This exciting direction abandons the domain-specific knowledge for an individual data modality, but rather pursues a solution far more generalizable.

For self-supervised representation learning, masked modeling [23] or simply masking input as an augmentation [75] has emerged as an effective approach. The in-

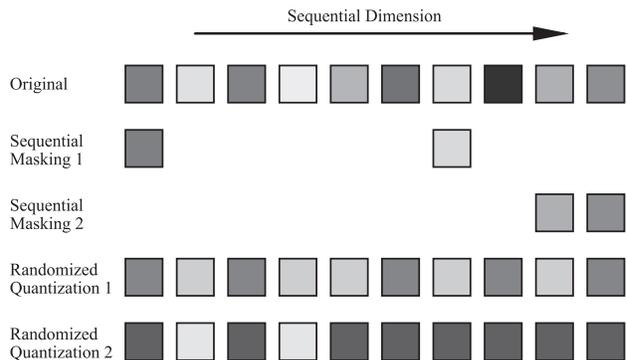


Figure 1. We represent data as a matrix with a sequential dimension and a channel dimension. As a generic data augmentation, masking drops tokens along the sequential dimension. The proposed randomized quantization instead withholds information along the channel dimension. In this figure, we use 1D data of 10 sequential tokens for illustration. Data values are coded in grayscale.

put data is represented by a 2D tensor with a sequential dimension and a channel dimension in a modality-agnostic way [3]. The sequential dimension can be spatial in images, temporal in audio, and syntactic in languages. The masking mechanism withholds information along the sequential dimension, and exploits it for supervision. As a result, models learned from the masking supervision demonstrate strong capability for capturing correlations between sequential tokens [38].

The channel dimension describes the data feature at each sequential location, for example, RGB color at a spatial location or spectrogram frequency at a time step. Despite being generic, masking approaches have neglected to exploit supervision along the channel dimension. While the number of channels for images is as small as three, the channels for audio and tabular data can be as many as hundreds. Formulating the self-supervision from the channel dimension holds much potential for representation learning.

In this paper, we draw a connection between the sequen-

*Equal contribution. Work done during an internship at MSRA.

tial masking operation and quantization by exploring quantization as a novel form of masking along the channel dimension. The data in each channel is dynamically quantized through a non-uniform quantizer, with the quantization value randomly sampled from randomly sampled quantization bins. In this way, information within each quantization bin is masked out, yet information across bins is retained. The information removed by quantization is controlled by the number of bins and the size of the bins, which has been rigorously studied in theory [62]. The larger the distortion rate, the stronger the quantization when it is used as an augmentation for representation learning. The extreme case of using only a single bin is equivalent to dropping the entire channel. We systematically study various quantization configurations for their effects as a data augmentation, for example, with respect to the number bins, uniform or non-uniform bins, and methods to select quantization values.

We apply the randomized quantizer as the only augmentation, or in conjunction with augmentations along the sequential dimension on state-of-the-art Siamese representation learning methods MoCo-v3 [16] and BYOL [34]. In comparisons with domain-agnostic augmentations based on MixUp [85], our approach achieves state-of-the-art results by a large margin on vision, audio, 3d point cloud, and the DABS benchmark. Compared with domain-specific augmentations, our approach achieves competitive performance against handcrafted augmentations on vision, and state-of-the-art performance on audio and 3d point clouds.

Our contributions can be summarized as follows:

- We propose a simple yet effective data-agnostic augmentation for contrastive learning, based on quantization along the channel dimension.
- We design a randomization technique by varying the quantization value and the quantization bins, to enhance data augmentation.
- We demonstrate the generality and strong performance of our channel-wise randomized quantization for vision, audio, 3D point clouds and the DABS benchmark in a data-agnostic way.

2. Related Works

Self-supervised learning extracts labels from the data itself and tasks the network to learn transferable representations. Among the earliest forms of self-supervised models are auto-encoders [42] and generative models [41]. But since the input and the output are identical, a neural network may easily find shortcuts and use memorization to solve the generation task. Advances in recent years show that information needs to be withheld from the input to prevent cheating [25]. Pretext tasks such as colorization [86],

inpainting [58], jigsaw puzzles [55] were proposed in vision, while masked modeling [23], next sentence prediction [46, 45], and replaced word prediction [17] were developed in natural language processing. Speediness [7, 43] and temporal order [73, 52] have been exploited for video representation learning. Due to space limitations, we omit the literature for speech [4], tabular data [1], graph-structured data [65] and many other modalities. The optimal pretext task for each target problem may be different. However, there exists enormous interest in obtaining a single foundation model [9] for all downstream applications.

Instead of withholding data for supervision, contrastive models [78, 56] create new data via data augmentation and compare features extracted using a Siamese network for supervision. Siamese representation learning can be categorized by whether to use explicit negatives [34], ways to define negatives [2], and various loss formulations [13, 83]. However, the main driving signal for learning lies in the augmentations.

Data augmentation enlarges the number of data instances by leveraging prior knowledge of the data and target problem properties. For supervised learning, data augmentation aids in reducing overfitting and regularization [84]. For self-supervised learning, the information gap created by two augmentations provides learning supervision. Typically, the data augmentation function extracts partial information from the data and optionally adds corruptions.

Popular image-specific augmentations include cropping, scaling, color jittering, Gaussian blurring, cut-out [24], cut-mix [82], and auto-augment, which searches for a data augmentation policy [21]. In natural language processing, synonym replacement [74], back translation [11], random word insertion and deletion are most common. For audio and speech, altering the pitch, changing the playback speed, and masking either along the time axis or the frequency axis [57] may improve performance. Additionally, augmenting data through a generative model [10] such as a GAN is a viable approach.

Domain-agnostic augmentation aims to generalize modality-specific and domain-specific augmentations into a unified formulation. Finding such general priors in data is challenging. One line of work follows Mixup [85], which is initially proposed to improve empirical risk minimization of supervised learning by linearly interpolating data and labels. Because of its generality, later works have explored its application on other data modalities [35], a wide range of problems [50, 40], as well as representation learning [66, 48, 68]. Another important line of work generalizes masked modeling [23], which was initially proposed for language modeling, to other data modalities and domains [38, 67, 79]. The masking mechanism samples a subset of the input data, while Mixup introduces additional corruptions which are not observed in the original data

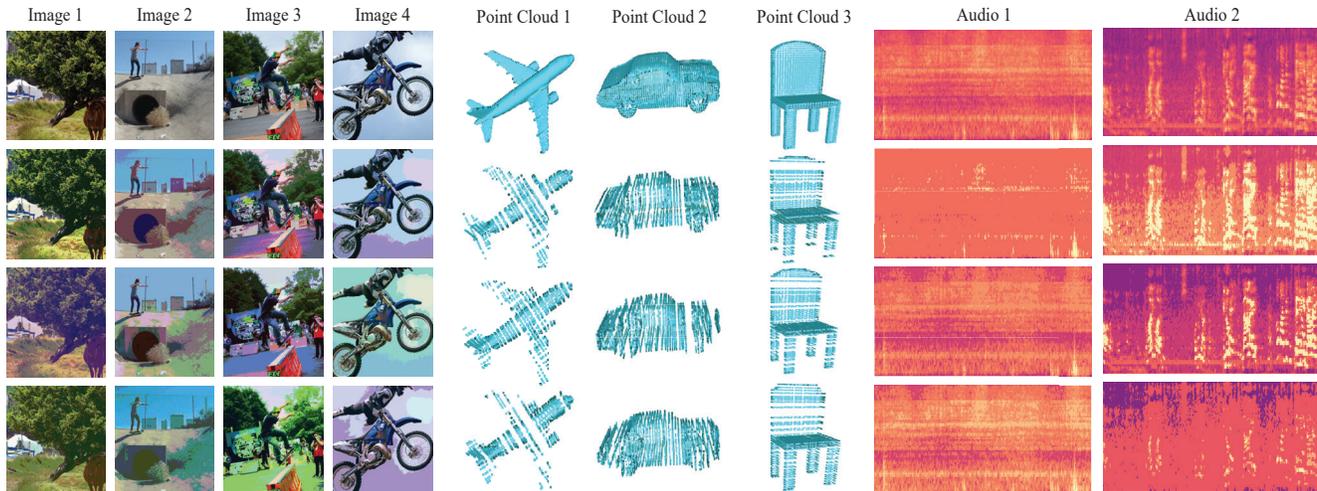


Figure 2. Visualizing randomized quantization augmentation on images, 3d point clouds, and audio. The first row presents the original signal, and the bottom three rows are augmented views. Randomized quantization alters color and enhances edges on images, spatially samples coordinates on point clouds, and enhances frequency channels for audio.

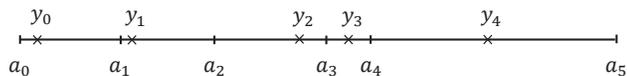


Figure 3. Illustration of a non-uniform quantizer with five bins. a_i denotes the interval of the quantization bin and y_i denotes the reproduction value from each bin.

instance. A third line of work tries to mine inherent regularities and structures in the data. NNCLR [26] and MSF [47] use nearest neighbors derived from the current representations as a novel form of augmentation for contrastive learning. Our randomized quantization differs from these works, and augments data along the channel dimension.

Quantization represents numerical values with a fixed discrete set of numbers so as to reduce communication bandwidth and maintain representation quality. The rounding error was first analyzed a century ago [63], and the theory based on variable-rate quantization [61] and Huffman coding [44] revolutionized the communications industry. We refer readers to a survey [33] that describes this area from a theoretical perspective.

Quantization for efficient neural networks [32] aims to reduce neural network latency while maintaining model accuracy. The advances of half-precision [5, 69] and mixed-precision training [18, 36, 51] has accelerated model training by an order of magnitude. Works have shown that neural networks can be completely binarized [49, 76, 19] with reasonable performance. Stochastic quantization [15, 28, 8] is a technique for learning and compressing model weights in a way that avoids local minima with the low-precision weight representations.

There exist two works [12, 30] that study the use of

quantization in contrastive representation learning. In both, quantization is applied on the intermediate features and model weights. Though effective for model compression, the quantization of features/weights can be viewed as unsuitable for augmentation due to the unpredictability of the corresponding change in the input. Our work additionally differs from these methods [12, 30] in its randomization of a non-uniform quantizer. This technique creates a much more complex augmentation space than uniform quantization, generating valid images of greater diversity.

3. Approach

This paper proposes a novel generic data augmentation for representation learning based on quantization. We first provide preliminaries on quantization. We then introduce two factors to inject randomness into the quantization procedure.

3.1. Preliminaries: Quantization

A quantizer is a function which consists of a set of non-overlapping intervals or bins $S = \{S_i = [a_i, a_{i+1})\}$, $i = 0, 1, \dots, n-1$, and a set of reproduction values y_i . n is the number of intervals and reproduction values. Let x be a one-dimensional input signal. The quantizer maps values within an interval to a single scalar, defined as $q(x) = y_i$ for $x \in S_i$. Formally, it can be written as

$$q(x) = \sum_i y_i \cdot 1_{S_i}(x), \quad (1)$$

where the indicator function $1_{S_i}(x) = 1$ if $x \in S_i$ and $1_{S_i}(x) = 0$ otherwise. Figure 3 gives an illustration of a quantizer with five intervals. Quantization represents the

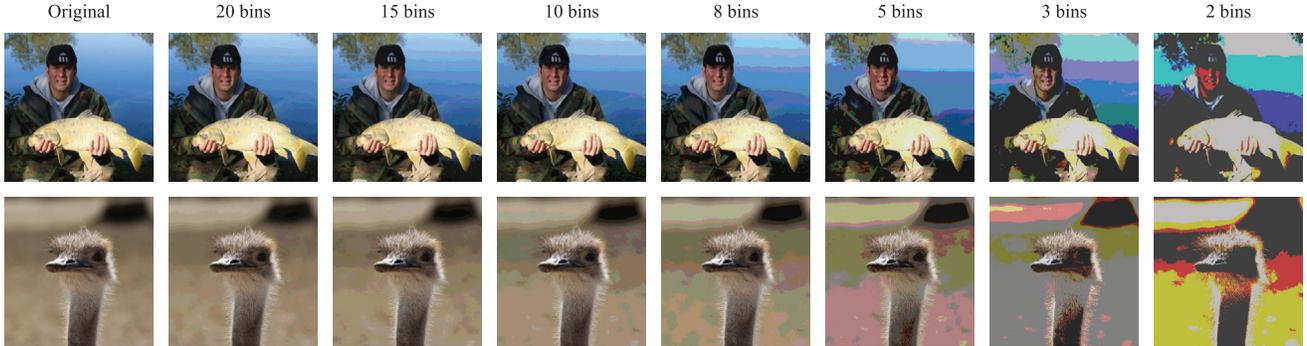


Figure 4. Visualization of quantized images with different numbers of bins. The images are quantized by a uniform quantizer. Fewer than three quantization bins causes severe information reduction, while fifteen or more bins leads to negligible difference from the original image. An intermediate number of bins (e.g., five to ten) is well-suited for image augmentation.

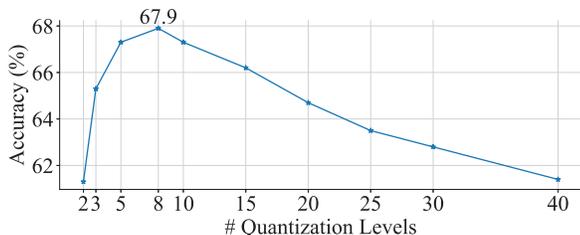


Figure 5. Ablation study on the number of quantization bins. The peak performance is reached at 8 bins. Fewer bins deliver heavier augmentations and larger bins lead to weaker augmentations.

original signal using a finite number of bits and hence introduces error to signal recovery. The central research problem is to find better tradeoffs between communication bandwidths and reproduction errors.

Quantization can be categorized by uniform quantization and non-uniform quantization. A uniform quantizer has intervals and values which are evenly spaced, whereas a non-uniform quantizer allows either intervals or values to be unevenly spaced. Uniform quantizers are amenable to hardware deployment. However, non-uniform quantizers may perform better depending on the probabilistic distribution of x .

3.2. Randomized Quantization as Augmentation

We aim to exploit quantization as a data withholding tool for representation learning. The information within each quantization bin is withheld, while the information across bins are retained. For data augmentation, a key aspect is the complexity of the augmentation space. We design a complex quantization augmentation by randomizing the intervals and the reproduction values. Concretely, given $S_i = [a_i, a_{i+1})$, a_i is generated by

$$a_0, a_1, \dots, a_{n-1} = \text{sort}(a'_0, a'_1, \dots, a'_{n-1}) \quad (2)$$

	random bins	random values	top1 acc
baseline	✗	✗	50.0
+ quantize	✗	✗	54.8
+ quantize	✗	✓	62.6
+ quantize	✓	✗	66.0
+ quantize	✓	✓	67.9

Table 1. Ablation study of the two randomness factors for randomized quantization described in Eq. 3 and Eq. 4. We examine the effect of randomized bins and random reproduction values for each bin. These two factors increase the complexity of the augmentation and significantly improve the performance.

$$a'_i = U(\min(x), \max(x)), i = 0, 1, \dots, n - 1, \quad (3)$$

where U denotes random sampling with a uniform distribution over the interval, and $\min(x)/\max(x)$ refers to the minimum/maximum of the values x per channel. The reproduction value y_i is randomly sampled within the corresponding interval,

$$y_i = U(a_i, a_{i+1}). \quad (4)$$

The resultant randomized quantizer is non-uniform. The number of quantization bins n is the hyperparameter of the augmentation.

3.3. Data-Agnostic Augmentation

The proposed randomized quantization augmentation can be applied to the channel dimension for any arbitrary data modality. The physical interpretation of the augmentation depends on the nature of the data modality. In Figure 2, we visualize the augmentations for images, audio and point clouds. On images, it removes the high frequency details but highlights object boundaries and edges. It also alters color appearance significantly. On audio, we examine the augmented sound acoustically and we find the augmentation tends to enhance specific frequency signals, such as

	100-ep	300-ep	800-ep
MoCo-v3	67.9	71.6	72.1
BYOL	67.2	71.0	71.6

Table 2. Representation learning with randomized quantization augmentation benefits from more training epochs without early saturation.

Augmentations	MoCo-v3	BYOL
CR	10.1	9.9
CR + DACL [68]	32.7	33.2
CR + i-Mix [48]	30.3	28.7
CR + SSQL [12]	11.1	6.4
CR + Ours	42.9	43.0
RRC	50.0	49.3
RRC + DACL [68]	57.2	57.6
RRC + i-Mix [48]	55.4	49.9
RRC + SSQL [12]	42.9	49.3
RRC + Ours	67.9	67.2

Table 3. Comparisons with alternative domain-agnostic augmentation techniques under the linear classification protocol on ImageNet. CR is short for center crop, and RRC is short for random resized crop. Our randomized quantization approach achieves the best results against prior arts.

low-frequency sounds or high-frequency sounds. On point clouds where the channel dimension represents xyz coordinates, it tends to downsample local structures but highlight the global shape. Some discrete modalities such as language are not directly amenable to quantization, but can be mapped to a continuous representation via a frozen data embedding layer. In Section 5.4, we examine randomized quantization on such input embeddings.

3.4. Siamese Representation Learning

Siamese representation learning or contrastive learning relies heavily on the quality of the augmentations [34, 87]. We apply the proposed augmentation on Siamese representation learning. At each training iteration, we sample two views from a data instance using randomized quantization by itself or in conjunction with other augmentations. Loss terms such as InfoNCE [56] and L2 are applied on the two views. The MoCo-v3 [16] and BYOL [34] frameworks are followed in this paper, and we refer readers to the original papers for details.

4. Ablation Study

We choose visual representation learning for an ablation study. Random resized cropping is taken as the baseline augmentation, and we apply our randomized quantization after it. Following the MoCo-v3 framework [16], we use ResNet-50 [39] as the backbone network. The optimizer

Method	MoCo-v3	BYOL
Ours	42.9	43.0
RRC	50.0	49.4
RRC + CJ	60.1	61.1
RRC + Ours	67.9	67.2
Full	68.9	68.9

Table 4. Comparisons with image-specific augmentations under the linear classification protocol on ImageNet. CJ stands for color jittering, and Full includes random resized crop, color jittering, grayscaling, Gaussian blurring and solarization. Randomized quantization is stronger than color jittering by a large margin. It falls behind the full handcrafted augmentations by just 1%.

is consistent with MoCo-v3, and the network is optimized for 100 epochs. Representation learning is conducted on the ImageNet-1K dataset [22], and linear classification accuracy is reported on the validation set.

We ablate three design factors of the proposed quantization-based augmentation which affect its ability to mask channel-wise information.

Randomizing Bins. The performance of representation learning depends on the complexity of the pretext tasks created from random augmentations. In Table 1, the baseline approach using the random resized crop augmentation obtains 50.0% top-1 accuracy. Using a fixed uniform quantizer improves the performance mildly to 54.8%. Randomizing the locations and sizes of bins allows for uneven masking and creates more effective pretext tasks. It improves the performance significantly to 66.0%.

Randomizing reproduction values. Quantization is also affected by how each bin is represented. Commonly, the values within a bin’s range are represented by the midpoint. As an alternative, we also consider taking a random value in the range. Intuitively, random reproduction values lead to bias in the quantization error, making them no longer zero-mean and bringing a stronger augmentation effect. It is found to benefit representation learning, yielding an increase of 1.9 points upon randomizing the bins in Table 1.

Number of quantization bins. Figure 5 illustrates the effect of various numbers of quantization bins. Intuitively, fewer bins leads to a stronger masking effect and higher cross-view variation. We vary the number of bins and find strong performance with 5-10 bins, peaking at 8 bins. This observation is similar to spatial masking in MAE [38] where an optimal masking ratio is found. In Figure 4, we visualize quantized images with different numbers of bins. To make the visualization consistent and comparable, we use the uniform quantizer in this case. It can be observed that too much information is withheld when using too few bins, and too many bins withholds too little information.

Training epochs. We further study training the augmentations with more epochs. In Table 2, the performance im-



Figure 6. Visual comparisons between color jittering and randomized quantization. Randomized quantization exhibits greater change in visual appearance and stronger edge enhancement.

proves from 67.9% with 100 epochs to 71.6% with 300 epochs and 72.1% with 800 epochs. With this complex augmentation, the network benefits from longer training.

5. Experiments on Various Modalities

We examine pre-training with the proposed augmentation across a variety of modalities including 1) vision (Section 5.1); 2) 3D point clouds (Section 5.2); 3) audio (Section 5.3); and 4) the DABS benchmark [66] (Section 5.4) comprised of data from multiple domains. The hyperparameter n indicating the number of bins is tuned for each modality. We leave the description of corresponding datasets, settings and evaluation metrics to each section.

5.1. Images

We compare the proposed randomized quantization augmentation against domain-agnostic augmentation baselines, as well as domain-specific augmentations designed for images. The number of quantization bins is chosen as $n = 8$. The experimental protocol follows the ablation study.

Comparisons with domain-agnostic augmentations. Recent works on domain-agnostic augmentation are predominantly adapted from Mixup [85]. For example, i-Mix [48] linearly interpolates input data, and their corresponding virtual labels are generated from the current batch. Similarly, DACL [68] interpolates input but uses it as a way of adding noise to the original data. We also consider a baseline SSQ [12] which jointly performs quantization and contrastive learning for lower bit-width deployment.

In Table 3, we compare our approach to these methods on two spatial operations: center crop (CR) and random resize crop (RRC). Center crop amounts to no augmentation, and random resized crop is frequently used in vision applications. Our evaluation is based on two Siamese representation learning frameworks MoCo-v3 and BYOL, since BYOL is said to have different behavior on augmentations.

Randomized quantization performs the best against Mixup-based augmentations. As a standalone augmentation, randomized quantization obtains an accuracy of 42.9% with MoCo-v3, which outperforms DACL and i-Mix by a

Method	Backbone	Epochs	Top1 accuracy
MSF w/w	R-50	200	66.3
Ours	R-50	100	<u>67.9</u>
Ours	R-50	300	71.6

Table 5. Comparison with weakly-augmented SSL baseline MSF.

	1%	2%	5%	10%	20%	100%
FoldingNet (lin)	56.4	66.9	75.6	81.2	83.6	88.4
MID-FC (lin)	61.5	73.1	80.2	84.2	86.9	90.3
Ours (lin)	66.7	74.3	80.0	84.5	87.2	90.5
Scratch	58.5	71.2	80.1	85.4	88.7	92.9
MID-FC (ft)	67.3	76.5	83.6	88.4	90.2	93.0
Ours (ft)	71.3	78.5	84.9	88.6	90.6	93.0

Table 6. Linear probing (lin) and finetuning (ft) results for the shape classification task on the ModelNet40 dataset. Pre-training is conducted on the ShapeNet dataset. Our augmentation improves the classification accuracy substantially on various ratios of data, especially on very limited (1%) data.

	C.mIoU			I.mIoU		
	1%	5%	100%	1%	5%	100%
Multi-Task (lin)	-	73.9	-	68.2	80.7	-
MID-FC (lin)	66.2	76.5	82.8	72.4	80.9	84.1
Ours (lin)	70.6	76.9	82.9	77.4	81.9	84.3
MID-FC (ft)	67.6	77.8	84.3	76.2	82.1	85.5
Ours (ft)	69.5	78.4	84.4	77.8	82.3	85.5

Table 7. Linear probing (lin) and finetuning (ft) results for the shape segmentation task on the ShapeNet Part dataset. Pre-training is conducted on ShapeNet. Our augmentation improves the performance substantially on various ratios.

large margin. In conjunction with random resized crop, a 10% margin is maintained. SSQ heavily relies on domain-specific augmentations, and in many cases, it fails to improve the baseline when handcrafted augmentations are removed. The results using MoCo-v3 and BYOL training objectives are similar. Overall, randomized quantization achieves state-of-the-art results against domain-agnostic baselines in the vision domain.

Comparisons with domain-specific augmentations. We further compare with image-specific augmentations for visual representation learning in Table 4. We find that randomized quantization is much stronger than color jittering, which is heavily designed with prior knowledge such as brightness, contrast, and saturation for pixels. In Figure 6, we visualize color jittering and our augmentation. It can be observed that our augmentation leads to stronger and more diverse visual appearances than color jittering. Our augmentation is 1% weaker than the full augmentation, which includes random resized crop, color jittering, grayscaling, Gaussian blurring and solarization successively.

Name	Task	#Classes	Data size	Avg duration (s)
NSynth (NS) [27]	Musical instrument classification	11	305,979	4.0
UrbanSound8K (US8K) [60]	Urban sound classification	10	8,732	4.0
VoxCeleb1 (VC1) [53]	Speaker identification	1,211	153,514	8.2
VoxForge (VF) [from Voxforge.org]	Language identification	6	176,438	5.8
Speech Commands V2 (SPCV2) [72]	Command classification	35	105,829	1.0
Speech Commands V2 (SPCV2/12) [72]	Command classification	12	105,829	1.0

Table 8. Downstream dataset details for audio representation learning.

Method	NS	US8K	VC1	VF	SPCV2/12	SPCV2	Average
TRILL [64]	-	-	17.9	88.1	74.9	-	-
COLA [59]	63.4	-	29.9	71.3	71.7	62.4	-
OpenL3 [20]	-	78.2	-	-	-	-	-
COALA [29]	73.1	72.7	-	-	-	-	-
COLA [59]	70.2	78.5	30.4	79.5	76.7	76.8	68.7
BYOL-A [54]	74.1	79.1	40.1	90.2	91.0	92.2	77.8
Ours	74.2	78.0	45.7	92.6	95.1	92.1	79.6

Table 9. Linear probing results for audio representation learning on six downstream datasets. Pre-training is conducted on the AudioSet dataset. Our model outperforms BYOL-A on four of the six datasets, with an average improvement of 1.8%.

Comparison with weakly-augmented SSL baselines. We compare the performance with MSF [47] in a weak augmentation setting. MSF implicitly utilizes close nearest neighbors as a form of augmentation. In Table 5, our method is better, and it benefits from longer training (67.9 vs. 71.9). Additionally, we demonstrate our approach on diverse modalities, while MSF only focuses on vision.

5.2. 3D Point Clouds

We explore self-supervised representation learning on point clouds, represented by a disordered set of xyz coordinates. The pretraining is conducted on the ShapeNet [14] dataset consisting of 57,449 3D shapes. Octree-based Sparse CNN [70] is used as the backbone network, which takes 3D point clouds as input and extracts point features as well as shape features. We follow the MID-Net [71] model as the baseline, which is trained by a point-wise and instance-wise contrastive loss. The model is trained by a SGD optimizer with a batch size of 32 and a weight decay of $5e-4$. The initial learning rate is 0.03 and decreases by a factor of 10 after 200 and 300 epochs, and the training process terminates after 400 epochs. We apply the randomized quantization augmentation after the base augmentations used in MID-Net. Unlike images and audio which are snapped to grids, strong quantization of point cloud coordinates drastically degrades 3d point data. We thus choose to use a larger number of bins, $n=30$, in order to maintain more information. In practice, since the 3d points are sparsified by quantization, we observe a substantial training speedup as a side benefit. For evaluation, we experiment on two downstream tasks: shape classification and segmentation.

Shape classification is conducted on ModelNet40 [77] which is composed of 13,834 3D models from 40 categories. For each shape, we extract a global feature with the pre-trained backbone then train a linear classifier, or finetune the network, and report the average classification accuracy in Table 6. We do comparison with FoldingNet [80] and MID-FC [71]. With our augmentation, we improve the classification accuracy over the baseline MID-FC [71] substantially, especially when the training data is limited as shown in Table 6. For example, with 1% of the training data, we improve the classification accuracy by 5.2 and 4.0 points on linear probing and finetuning, respectively.

Shape segmentation is conducted on ShapeNet Part [81] with 16,881 3D point clouds from 16 categories. For each shape, we extract point-wise features with the pre-trained backbone then train a segmentation head composed of two fully connected layers, or finetune the network, and report the mean IoU across all categories (C.mIoU) and the mean IoU across all instances (I.mIoU) in Table 7. We do comparison with two unsupervised pretraining methods: MID-FC [71] and Multi-Task [37]. Our results are consistently better than the baselines across different ratios of training data. And similarly to ModelNet40 classification, we observe significant improvements with limited (1% and 5%) training data. For example, with 1% of the training data, our method improves the segmentation performance by 4.4 and 1.9 IoU on linear probing and finetuning, respectively.

5.3. Audio

We apply randomized quantization to audio representation learning. We use AudioSet [31] as the pretraining

Method	Natural Images	Text	Speech	Sensors	Chest x-rays	Images & Text	Average
Scratch	10.1	42.3	24.9	69.8	68.1	57.5	45.5
e-Mix	27.9	44.1	41.8	79.5	72.4	48.9	52.4
Ours	32.1	44.7	44.5	84.9	73.4	54.5	55.6

Table 10. Evaluation of the representation performance over six modalities in the DABS benchmark. Representations are trained on a single primary dataset for each modality and evaluated on a number of downstream datasets. The performance for each modality is averaged across the downstream datasets and shown in the table.

dataset, with 1,963,807 audio samples of 527 classes. The pretrained representation is evaluated on six downstream audio classification datasets, summarized in Table 8.

We largely follow the experimental settings of BYOL-A [54] and treat it as our baseline.

We convert audio clips into the commonly used log-scaled spectrogram representation. Random resized crop is used to extract a 64×96 frequency-temporal segment for training. We replaced the Mixup augmentation used in BYOL-A with our randomized quantization, with the number of bins set to 5. We follow prior works [54] by using a lightweight 2D convolutional network as the backbone. We train the network using the Adam optimizer with a base learning rate of $3e-4$ and a batch size of 256 for 100 epochs.

Table 9 summarizes the results on the six downstream classification tasks. Compared against BYOL-A with the Mixup augmentation, our randomized quantization outperforms it in four out of the six tasks. Our approach is particularly stronger by a margin of 5.6% on the VoxCeleb1 dataset, which is the hardest classification task with 1211 classes among all six tasks. Our improvements tend to be smaller for tasks with fewer classes. On average, the proposed augmentation surpasses the current state-of-the-art BYOL-A by a margin of 1.8%.

5.4. DABS

We additionally conduct experiments on the public benchmark DABS [66] which is designed to study domain-agnostic self-supervised representation learning. Since some of its domains are discrete in nature (e.g., language), we first embed the data with a frozen layer and then augment the data embeddings for all the modalities consistently. It contains six data modalities¹, covering natural RGB images, multichannel sensor data, English text, audio, chest x-ray images, as well as captioned images. In each domain, pre-training is conducted on a large-scale dataset, and the learned representations are evaluated with linear probing on various in-domain downstream datasets. The average performance for the in-domain downstream datasets is reported. We refer the reader to the benchmark for a full description of the pretraining datasets and in-domain evaluation datasets.

¹The benchmark also provides an additional multi-lingual text modality. However, it is not evaluated in the original paper. We thus omit this.

We follow a leading method e-Mix [66] with a Transformer architecture on this domain-agnostic benchmark. The network is optimized with the Adam optimizer with a learning rate of $1e-4$ and weight decay of $1e-4$. The training protocol follows e-Mix, and all modalities share the same recipe.

We apply randomized quantization on the token embeddings before the Transformer. Since the quantization function has zero gradients everywhere, we randomly initialize the token embedding module without updating it. The straight-through estimator can be potentially useful, but it is not the focus of this work.

Table 10 summarizes the results for this benchmark. Our model outperforms the baseline e-Mix on all modalities. The improvements on natural images, speech, and sensors are larger than 3%, while the improvements on text and chest x-rays are relatively smaller, less than 1%. Both e-Mix and our pretraining seem to hurt the representation quality for captioned images. We hypothesize that the two modalities of images and texts pose significant challenges for a naive contrastive learning approach.

6. Conclusion

We propose randomized quantization as a novel data augmentation tool for self-supervised representation learning. Quantization effectively withholds information within the quantization bins but retains the information across bins. It could be applied on arbitrary data along the channel dimension without domain-specific knowledge. Randomized quantization significantly outperforms existing domain-agnostic augmentations based on Mixup. It compares favorably against domain-specific augmentations on vision, and attains state-of-the-art results on audio and 3D point clouds. We also explored randomized quantization on input data embeddings in a neural network for a wide range of data modalities. Experimental results on the DABS benchmark demonstrates state-of-the-art results for speech, text, images and multiple sensors. Randomized quantization could potentially be applied in a masked modeling framework, where the original images are reconstructed from quantized ones. This direction will be explored in future work.

Broader Impacts

Although the proposed augmentation is generic in its formulation, it is not guaranteed to work beyond the modalities investigated in this paper. Application of the augmentation for other self-supervised learning frameworks such as masked modeling or generalization to other downstream tasks remains under-explored.

Acknowledgements

We sincerely appreciate Reviewer #5 and AC’s unwavering efforts in saving this work from rejection. This research was supported in part by National Natural Science Foundation of China/HKSAR Research Grants Council Joint Research Scheme under Grant N_HKUST627/20, and in part by HKSAR RGC General Research Fund (GRF) #16203319.

References

- [1] Sercan Ö Arik and Tomas Pfister. Tabnet: Attentive interpretable tabular learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 6679–6687, 2021. 2
- [2] Philip Bachman, R Devon Hjelm, and William Buchwalter. Learning representations by maximizing mutual information across views. *Advances in neural information processing systems*, 32, 2019. 2
- [3] Alexei Baevski, Wei-Ning Hsu, Qiantong Xu, Arun Babu, Jiatuo Gu, and Michael Auli. Data2vec: A general framework for self-supervised learning in speech, vision and language. *arXiv preprint arXiv:2202.03555*, 2022. 1
- [4] Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in Neural Information Processing Systems*, 33:12449–12460, 2020. 2
- [5] Ron Banner, Itay Hubara, Elad Hoffer, and Daniel Soudry. Scalable methods for 8-bit training of neural networks. *Advances in neural information processing systems*, 31, 2018. 3
- [6] Hangbo Bao, Li Dong, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021. 1
- [7] Sagie Benaim, Ariel Ephrat, Oran Lang, Inbar Mosseri, William T Freeman, Michael Rubinstein, Michal Irani, and Tali Dekel. Speednet: Learning the speediness in videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9922–9931, 2020. 2
- [8] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013. 3
- [9] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021. 2
- [10] Christopher Bowles, Liang Chen, Ricardo Guerrero, Paul Bentley, Roger Gunn, Alexander Hammers, David Alexander Dickie, Maria Valdés Hernández, Joanna Wardlaw, and Daniel Rueckert. Gan augmentation: Augmenting training data using generative adversarial networks. *arXiv preprint arXiv:1810.10863*, 2018. 2
- [11] Richard W Brislin. Back-translation for cross-cultural research. *Journal of cross-cultural psychology*, 1(3):185–216, 1970. 2
- [12] Yun-Hao Cao, Peiqin Sun, Yechang Huang, Jianxin Wu, and Shuchang Zhou. Synergistic self-supervised and quantization learning. In *ECCV*, 2022. 3, 5, 6
- [13] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9650–9660, 2021. 2
- [14] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. 7
- [15] Jianfei Chen, Yu Gai, Zhewei Yao, Michael W Mahoney, and Joseph E Gonzalez. A statistical framework for low-bitwidth training of deep neural networks. *Advances in Neural Information Processing Systems*, 33:883–894, 2020. 3
- [16] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9640–9649, 2021. 2, 5
- [17] Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. Electra: Pre-training text encoders as discriminators rather than generators. *arXiv preprint arXiv:2003.10555*, 2020. 2
- [18] Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. Training deep neural networks with low precision multiplications. *arXiv preprint arXiv:1412.7024*, 2014. 3
- [19] Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. Binaryconnect: Training deep neural networks with binary weights during propagations. *Advances in neural information processing systems*, 28, 2015. 3
- [20] Jason Cramer, Ho-Hsiang Wu, Justin Salamon, and Juan Pablo Bello. Look, listen, and learn more: Design choices for deep audio embeddings. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3852–3856. IEEE, 2019. 7
- [21] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation policies from data. *arXiv preprint arXiv:1805.09501*, 2018. 2
- [22] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 5
- [23] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional

- transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 1, 2
- [24] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017. 2
- [25] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE international conference on computer vision*, pages 1422–1430, 2015. 2
- [26] Debidatta Dwibedi, Yusuf Aytar, Jonathan Tompson, Pierre Sermanet, and Andrew Zisserman. With a little help from my friends: Nearest-neighbor contrastive learning of visual representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9588–9597, 2021. 3
- [27] Jesse Engel, Cinjon Resnick, Adam Roberts, Sander Dieleman, Mohammad Norouzi, Douglas Eck, and Karen Simonyan. Neural audio synthesis of musical notes with wavenet autoencoders. In *International Conference on Machine Learning*, pages 1068–1077. PMLR, 2017. 7
- [28] Angela Fan, Pierre Stock, Benjamin Graham, Edouard Grave, Rémi Gribonval, Herve Jegou, and Armand Joulin. Training with quantization noise for extreme model compression. *arXiv preprint arXiv:2004.07320*, 2020. 3
- [29] Xavier Favory, Konstantinos Drossos, Tuomas Virtanen, and Xavier Serra. Coala: Co-aligned autoencoders for learning semantically enriched audio representations. *arXiv preprint arXiv:2006.08386*, 2020. 7
- [30] Yonggan Fu, Qixuan Yu, Meng Li, Xu Ouyang, Vikas Chandra, and Yingyan Lin. Contrastive quant: quantization makes stronger contrastive learning. In *Proceedings of the 59th ACM/IEEE Design Automation Conference*, pages 205–210, 2022. 3
- [31] Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing Moore, Manoj Plakal, and Marvin Ritter. Audio set: An ontology and human-labeled dataset for audio events. In *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 776–780. IEEE, 2017. 7
- [32] Amir Gholami, Sehoon Kim, Zhen Dong, Zhewei Yao, Michael W Mahoney, and Kurt Keutzer. A survey of quantization methods for efficient neural network inference. *arXiv preprint arXiv:2103.13630*, 2021. 3
- [33] Robert M. Gray and David L. Neuhoff. Quantization. *IEEE transactions on information theory*, 44(6):2325–2383, 1998. 3
- [34] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020. 2, 5
- [35] Hongyu Guo. Nonlinear mixup: Out-of-manifold data augmentation for text classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 4044–4051, 2020. 2
- [36] Suyog Gupta, Ankur Agrawal, Kailash Gopalakrishnan, and Pritish Narayanan. Deep learning with limited numerical precision. In *International conference on machine learning*, pages 1737–1746. PMLR, 2015. 3
- [37] Kaveh Hassani and Mike Haley. Unsupervised multi-task feature learning on point clouds. In *ICCV*, 2019. 7
- [38] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022. 1, 2, 5
- [39] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 5
- [40] Dan Hendrycks, Norman Mu, Ekin D Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshminarayanan. Augmix: A simple data processing method to improve robustness and uncertainty. *arXiv preprint arXiv:1912.02781*, 2019. 2
- [41] Geoffrey E Hinton. Deep belief networks. *Scholarpedia*, 4(5):5947, 2009. 2
- [42] Geoffrey E Hinton and Richard Zemel. Autoencoders, minimum description length and helmholtz free energy. *Advances in neural information processing systems*, 6, 1993. 2
- [43] Deng Huang, Wenhao Wu, Weiwen Hu, Xu Liu, Dongliang He, Zhihua Wu, Xiangmiao Wu, Minghui Tan, and Errui Ding. Ascnet: Self-supervised video representation learning with appearance-speed consistency. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8096–8105, 2021. 2
- [44] David A Huffman. A method for the construction of minimum-redundancy codes. *Proceedings of the IRE*, 40(9):1098–1101, 1952. 3
- [45] Yacine Jernite, Samuel R Bowman, and David Sontag. Discourse-based objectives for fast unsupervised sentence representation learning. *arXiv preprint arXiv:1705.00557*, 2017. 2
- [46] Ryan Kiros, Yukun Zhu, Russ R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. Skip-thought vectors. *Advances in neural information processing systems*, 28, 2015. 2
- [47] Soroush Abbasi Koohpayegani, Ajinkya Tejanekar, and Hamed Pirsiavash. Mean shift for self-supervised learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10326–10335, 2021. 3, 7
- [48] Kibok Lee, Yian Zhu, Kihyuk Sohn, Chun-Liang Li, Jinwoo Shin, and Honglak Lee. i-mix: A domain-agnostic strategy for contrastive representation learning. *arXiv preprint arXiv:2010.08887*, 2020. 2, 5, 6
- [49] Xiaofan Lin, Cong Zhao, and Wei Pan. Towards accurate binary convolutional neural network. *Advances in neural information processing systems*, 30, 2017. 3
- [50] Thomas Lucas, Corentin Tallec, Yann Ollivier, and Jakob Verbeek. Mixed batches and symmetric discriminators for gan training. In *International Conference on Machine Learning*, pages 2844–2853. PMLR, 2018. 2
- [51] Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory Diamos, Erich Elsen, David Garcia, Boris Ginsburg, Michael Houston, Oleksii Kuchaiev, Ganesh Venkatesh, et al. Mixed

- precision training. *arXiv preprint arXiv:1710.03740*, 2017. 3
- [52] Ishan Misra, C Lawrence Zitnick, and Martial Hebert. Shuffle and learn: unsupervised learning using temporal order verification. In *European conference on computer vision*, pages 527–544. Springer, 2016. 2
- [53] Arsha Nagrani, Joon Son Chung, and Andrew Zisserman. Voxceleb: a large-scale speaker identification dataset. *arXiv preprint arXiv:1706.08612*, 2017. 7
- [54] Daisuke Niizumi, Daiki Takeuchi, Yasunori Ohishi, Noboru Harada, and Kunio Kashino. Byol for audio: Self-supervised learning for general-purpose audio representation. In *2021 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2021. 7, 8
- [55] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European conference on computer vision*, pages 69–84. Springer, 2016. 2
- [56] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 2, 5
- [57] Daniel S Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D Cubuk, and Quoc V Le. SpecAugment: A simple data augmentation method for automatic speech recognition. *arXiv preprint arXiv:1904.08779*, 2019. 2
- [58] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2536–2544, 2016. 2
- [59] Aaqib Saeed, David Grangier, and Neil Zeghidour. Contrastive learning of general-purpose audio representations. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3875–3879. IEEE, 2021. 7
- [60] Justin Salamon, Christopher Jacoby, and Juan Pablo Bello. A dataset and taxonomy for urban sound research. In *Proceedings of the 22nd ACM international conference on Multimedia*, pages 1041–1044, 2014. 7
- [61] Claude Elwood Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948. 3
- [62] Claude E Shannon et al. Coding theorems for a discrete source with a fidelity criterion. *IRE Nat. Conv. Rec.*, 4(142-163):1, 1959. 2
- [63] William Fleetwood Sheppard. On the calculation of the most probable values of frequency-constants, for data arranged according to equidistant division of a scale. *Proceedings of the London Mathematical Society*, 1(1):353–380, 1897. 3
- [64] Joel Shor, Aren Jansen, Ronnie Maor, Oran Lang, Omry Tuval, Felix de Chaumont Quitry, Marco Tagliasacchi, Ira Shavitt, Dotan Emanuel, and Yinnon Haviv. Towards learning a universal non-semantic representation of speech. *arXiv preprint arXiv:2002.12764*, 2020. 7
- [65] Fan-Yun Sun, Jordan Hoffmann, Vikas Verma, and Jian Tang. Infograph: Unsupervised and semi-supervised graph-level representation learning via mutual information maximization. *arXiv preprint arXiv:1908.01000*, 2019. 2
- [66] Alex Tamkin, Vincent Liu, Rongfei Lu, Daniel Fein, Colin Schultz, and Noah Goodman. Dabs: A domain-agnostic benchmark for self-supervised learning. *arXiv preprint arXiv:2111.12062*, 2021. 2, 6, 8
- [67] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *arXiv preprint arXiv:2203.12602*, 2022. 2
- [68] Vikas Verma, Thang Luong, Kenji Kawaguchi, Hieu Pham, and Quoc Le. Towards domain-agnostic contrastive learning. In *International Conference on Machine Learning*, pages 10530–10541. PMLR, 2021. 2, 5, 6
- [69] Naigang Wang, Jungwook Choi, Daniel Brand, Chia-Yu Chen, and Kailash Gopalakrishnan. Training deep neural networks with 8-bit floating point numbers. *Advances in neural information processing systems*, 31, 2018. 3
- [70] Peng-Shuai Wang, Yang Liu, Yu-Xiao Guo, Chun-Yu Sun, and Xin Tong. O-cnn: Octree-based convolutional neural networks for 3d shape analysis. *ACM Transactions On Graphics (TOG)*, 36(4):1–11, 2017. 7
- [71] Peng-Shuai Wang, Yu-Qi Yang, Qian-Fang Zou, Zhirong Wu, Yang Liu, and Xin Tong. Unsupervised 3d learning for shape analysis via multiresolution instance discrimination. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 2773–2781, 2021. 7
- [72] Pete Warden. Speech commands: A dataset for limited-vocabulary speech recognition. *arXiv preprint arXiv:1804.03209*, 2018. 7
- [73] Donglai Wei, Joseph J Lim, Andrew Zisserman, and William T Freeman. Learning and using the arrow of time. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8052–8060, 2018. 2
- [74] Jason Wei and Kai Zou. Eda: Easy data augmentation techniques for boosting performance on text classification tasks. *arXiv preprint arXiv:1901.11196*, 2019. 2
- [75] Zhirong Wu, Zihang Lai, Xiao Sun, and Stephen Lin. Extreme masking for learning instance and distributed visual representations. *arXiv preprint arXiv:2206.04667*, 2022. 1
- [76] Zhirong Wu, Dahua Lin, and Xiaoou Tang. Adjustable bounded rectifiers: Towards deep binary representations. *arXiv preprint arXiv:1511.06201*, 2015. 3
- [77] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1912–1920, 2015. 7
- [78] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3733–3742, 2018. 2
- [79] Hu Xu, Juncheng Li, Alexei Baevski, Michael Auli, Wojciech Galuba, Florian Metze, Christoph Feichtenhofer, et al. Masked autoencoders that listen. *arXiv preprint arXiv:2207.06405*, 2022. 2

- [80] Yaoqing Yang, Chen Feng, Yiru Shen, and Dong Tian. FoldingNet: Point cloud auto-encoder via deep grid deformation. In *CVPR*, 2018. [7](#)
- [81] Li Yi, Lin Shao, Manolis Savva, Haibin Huang, Yang Zhou, Qirui Wang, Benjamin Graham, Martin Engelcke, Roman Klokov, Victor Lempitsky, et al. Large-scale 3d shape reconstruction and segmentation from shapenet core55. *arXiv preprint arXiv:1710.06104*, 2017. [7](#)
- [82] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6023–6032, 2019. [2](#)
- [83] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *International Conference on Machine Learning*, pages 12310–12320. PMLR, 2021. [2](#)
- [84] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning (still) requires rethinking generalization. *Communications of the ACM*, 64(3):107–115, 2021. [2](#)
- [85] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. [2](#), [6](#)
- [86] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *European conference on computer vision*, pages 649–666. Springer, 2016. [2](#)
- [87] Nanxuan Zhao, Zhirong Wu, Rynson WH Lau, and Stephen Lin. Distilling localization for self-supervised representation learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 10990–10998, 2021. [5](#)