

# Rethinking Data-Free Quantization as a Zero-Sum Game

Biao Qian, Yang Wang\*, Richang Hong, Meng Wang

Key Laboratory of Knowledge Engineering with Big Data, Ministry of Education,  
School of Computer Science and Information Engineering,  
Hefei University of Technology, China  
yangwang@hfut.edu.cn, {hfutqian,hongrc.hfut,eric.mengwang}@gmail.com

## Abstract

Data-free quantization (DFQ) recovers the performance of quantized network (Q) without accessing the real data, but generates the fake sample via a generator (G) by learning from full-precision network (P) instead. However, such sample generation process is totally independent of Q, specialized as failing to consider the adaptability of the generated samples, i.e., beneficial or adversarial, over the learning process of Q, resulting into non-ignorable performance loss. Building on this, several crucial questions — how to measure and exploit the sample adaptability to Q under varied bit-width scenarios? how to generate the samples with desirable adaptability to benefit the quantized network? — impel us to revisit DFQ. In this paper, we answer the above questions from a game-theory perspective to specialize DFQ as a zero-sum game between two players — a generator and a quantized network, and further propose an **Adaptability-aware Sample Generation** (AdaSG) method. Technically, AdaSG reformulates DFQ as a dynamic maximization-vs-minimization game process anchored on the sample adaptability. The maximization process aims to generate the sample with desirable adaptability, such sample adaptability is further reduced by the minimization process after calibrating Q for performance recovery. The Balance Gap is defined to guide the stationarity of the game process to maximally benefit Q. The theoretical analysis and empirical studies verify the superiority of AdaSG over the state-of-the-arts. *Our code is available at <https://github.com/hfutqian/AdaSG>.*

## Introduction

Deep Neural Networks (DNNs) have encountered great challenges when involving the applications (Krizhevsky, Sutskever, and Hinton 2017; Wang 2021) on resource-constrained devices, owing to the increasing demands for computing and storage resources. Network quantization (Lin, Talathi, and Annapureddy 2016; Jacob et al. 2018), which reduces the model size and energy consumption by mapping the floating-point weights and activations to low-bit ones, is a promising approach to improve the efficiency of DNNs for model compression (Han, Mao, and Dally 2015; Qian et al. 2021). Quantization methods generally dedicate themselves to recovering the *performance drop* originating

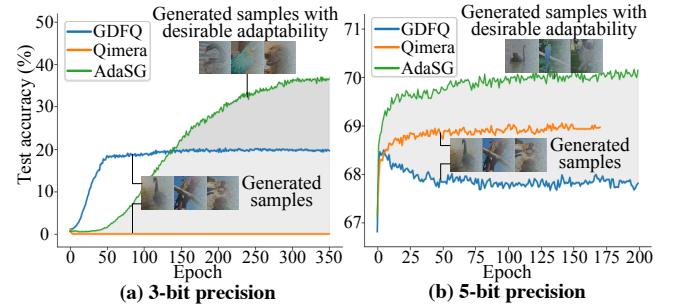


Figure 1: Existing work, e.g., GDFQ (Xu et al. 2020) (the blue) and Qimera (Choi et al. 2021) (the orange), suffers from a non-ignorable accuracy loss, as they fail to consider the sample adaptability. Our AdaSG (the green) generates the sample with desirable adaptability to maximally recover Q with varied bit widths, such as (a) 3-bit and (b) 5-bit precision. The observations are from ResNet-18 on ImageNet.

with the quantization errors, which involves fine-tuning or calibration operations with the original training data.

However, in many real-world scenes, such as medical and military fields, the original data may not be accessible due to privacy and security issues. Fortunately, recently proposed data-free quantization (DFQ), a potential method to quantize models without accessing the original data, aims to synthesize meaningful fake samples instead, which improves quantized network (Q) by knowledge distillation (Hinton, Vinyals, and Dean 2015; Qian et al. 2022) against the pre-trained full-precision model (P). Among the prior researches (Cai et al. 2020; Zhang et al. 2021), the generative fashions (Xu et al. 2020; Choi et al. 2021; Zhu et al. 2021) have recently attracted increasing attention, owing to their superior performance. The generative model is introduced as a generator (G) to capture the distribution of the original data from P for better fake samples, where P is regarded as the discriminator to guide the generation process. For example, Qimera (Choi et al. 2021) generated boundary supporting samples to reduce the gap between the synthetic and real data. Nevertheless, *there still remains a non-ignorable performance loss when encountering various bit-width settings*. The reasons may lie in several aspects:

- (1) Due to the limited capacity of the generator, the generated sample with incomplete distribution is impossible to

\* Yang Wang is the corresponding author.  
Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

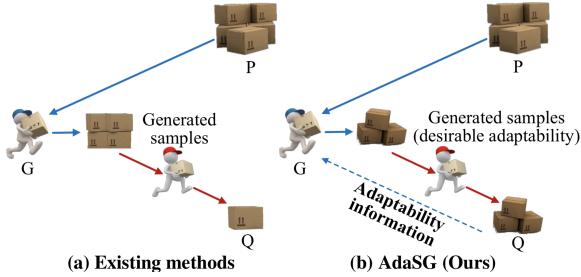


Figure 2: Unlike the existing arts (a), our AdaSG (b) aims to generate the sample with desirable adaptability from G, *i.e.*, the dependence of generated sample on Q, to maximally benefit Q with varied bit widths.

fully recover the original dataset, so that it is a crucial criterion: is the sample *beneficial* or *adversarial* over the learning process of Q? However, in the existing arts, the generated sample customized for P, can't always benefit Q in varied bit-width settings (*e.g.*, 3-bit or 5-bit precision, refer to (2)(3)), where only limited information from P can be utilized to recover Q.

- (2) In 3-bit precision, Q often suffers from a sharp accuracy drop compared to P due to large quantization error, leading to its poor learning ability. Under such case, the generated sample by G may bring an unexpected large *disagreement* between the predictions of P and Q, which makes the optimization loss too large to converge, resulting in no improvement; see the orange flat curve in Fig.1(a).
- (3) In 5-bit precision, Q still possesses comparable recognition ability with P due to a small accuracy drop. Under such case, most of the generated samples by G, for which Q and P give similar predictions (*i.e.*, reach an *agreement*), may not benefit Q. However, under the constraint of the optimization loss, Q receives no improvement, even impairment; see the blue curve (Epoch 0-50) in Fig.1(b).

Based on the above, the existing approaches fail to consider the *sample adaptability*, *i.e.*, beneficial or adversarial over the calibration process of Q, to Q with varied bit widths during the sample generation process from G, where Q is independent of the generation process; see Fig.2(a). For example, for (2), the sample with large adaptability may be one with small disagreement between P and Q; while for (3), it may be one with large disagreement. The above naturally elicits the following basic questions: *how to measure and exploit the sample adaptability to quantized network under varied bit-width scenarios? how to generate the sample with desirable adaptability to benefit the quantized network?*

To answer the above questions, we consider to generate the sample with large adaptability to Q by taking Q into account during the generation process; see Fig.2(b). Specifically, G aims to generate the sample with *large* adaptability, which essentially *maximizes* the reward for G, by enlarging the disagreement between P and Q, to benefit Q; while Q is calibrated to improve itself by exploiting the sample with large adaptability. It is apparent that such sample adaptability will not be large to Q after Q is refined; in other words,

such process of benefiting Q leads to *decreasing* the sample adaptability, which essentially *minimizes* the loss for Q, and is adversarial to maximizing the reward goal for G. Based on the above, we rethink date-free quantization process and formulate it as a zero-sum game (Li et al. 2022; Zhang et al. 2022) between two players (an adversarial game process where one player's reward is the other's loss while their sum is zero) — a generator and a quantized network; and we further propose an **Adaptability-aware Sample Generation** (AdaSG) method. Technically, we specify it via a dynamic maximization-vs-minimization game process anchored on sample adaptability, and defined as:

$$\min_{\theta_q \in \Theta_q} \max_{\theta_g \in \Theta_g} \mathcal{R}(\theta_g, \theta_q), \quad (1)$$

where G and Q are parametrized by  $\theta_g \in \Theta_g$  and  $\theta_q \in \Theta_q$ , respectively. The optimization of Eq.(1) consists of fixing  $\theta_q$ , and updating  $\theta_g$  to the optimal  $\theta'_g$  for *maximizing*  $\mathcal{R}(\theta_g, \theta_q)$ ; while alternatively fixing  $\theta'_g$ , and updating  $\theta_q$  to the optimal  $\theta'_q$  for *minimizing*  $\mathcal{R}(\theta'_g, \theta_q)$ . Specifically, the maximization process exploits the sample adaptability to Q upon P to generate two types of samples: disagreement (*i.e.*, P can predict correctly but Q not) and agreement samples (*i.e.*, P and Q have the same prediction), such sample adaptability is further reduced by the minimization process after calibrating Q for performance recovery. To achieve the stationarity to maximally benefit Q, the Balance Gap (BG) for the adaptability of the generated samples within the game process is defined to set up such a stationary objective to balance the maximization versus minimization process. We remark that AdaSG is essentially an *adversarial* game governed by the sample adaptability, which is fundamentally orthogonal to the existing arts that improve Q by transferring knowledge from P to Q. One recent study (Liu, Zhang, and Wang 2021) generates adversarial samples via G, by maximizing the gap between P and Q, and minimizing their gap to benefit Q for calibration. However, it fails to consider the sample adaptability to Q, hence suffers from the non-ideal generated samples. Besides, they focus primarily on the adversarial sample generation rather than adversarial game process perspective for AdaSG to generate the samples with desirable adaptability. The theoretical analysis and empirical studies validate the superiority of AdaSG to the state-of-the-arts.

## Adaptability-aware Sample Generation

Recent generative data-free quantization (DFQ) methods (Xu et al. 2020; Choi et al. 2021; Zhu et al. 2021) aim to reconstruct the original training dataset with a generator (G) by exploiting the distribution from a pre-trained full-precision network (P), which is further exploited to recover the performance of quantized network (Q) via the calibration operation. However, we observe that Q is independent of the generation process by the existing arts and whether the generated sample is beneficial or adversarial, namely *sample adaptability*, over the calibration process of Q, is crucial to the DFQ process. Motivated by the above, we focus primarily on the sample adaptability to Q. Building on this, the DFQ process is formulated as a dynamic maximization-vs-minimization game process governed by the sample adapt-

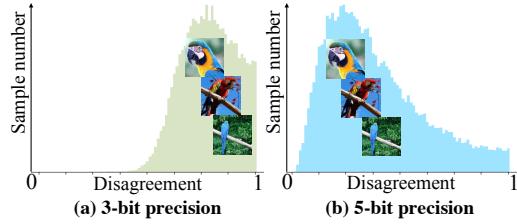


Figure 3: Given the 6000 real images from ImageNet as the input to both P and Q upon ResNet-18, the disagreement between P and Q varies greatly for varied Q with (a) 3-bit and (b) 5-bit precision.

ability, as illustrated in Fig.4. Naturally, one crucial question is how to measure the sample adaptability to Q, which will be discussed in the next section.

### How to Measure the Sample Adaptability to Q?

To measure the sample adaptability to Q, we primarily focus on two crucial issues: 1) the *dependence* of generated sample on Q, and 2) the *disagreement* between the predictions of P and Q, in that for Q with different bit widths (e.g., 3 or 5 bit), the disagreement varies greatly when delivering the same sample (take real data as an example) to both P and Q; see Fig.3. We define the generated sample that depends on Q, namely disagreement sample, below:

**Definition 1** (Disagreement Sample). *Given a random noise vector  $z \sim N(0, 1)$  and an arbitrary one-hot label  $y$ , a generator G generates a sample  $x = G(z|y)$ . Then the logit outputs of pre-trained full-precision model P and quantized model Q are given as  $z_p = P(x)$  and  $z_q = Q(x)$ , respectively. Suppose that the generated sample x can be correctly predicted by P, i.e.,  $\text{argmax}(z_p) = \text{argmax}(y)$ . We say x is the disagreement sample if  $\text{argmax}(z_p) \neq \text{argmax}(z_q)$ . Thus the probability vector encoding the disagreement between P and Q, is formulated as*

$$p_{ds} = \text{softmax}(z_p - z_q) \in \mathbb{R}^C, \quad (2)$$

where  $p_{ds}(c) = \frac{\exp(z_p(c) - z_q(c))}{\sum_{j=1}^C \exp(z_p(j) - z_q(j))}$  ( $c \in \{1, 2, \dots, C\}$ ) represents the  $c$ -th entry of the vector  $p_{ds}$  as the probability that  $x$  is labeled as the disagreement sample of the  $c$ -th class;  $C$  denotes the number of class. Together with the sample dependence to Q, the disagreement sample can viewed as the one with large adaptability to Q.

Another key problem is how to model the disagreement between P and Q, which can be exploited to measure the sample adaptability. Eq.(2) shows that the disagreement reaches the maximum provided  $p_{ds}(c)$  approaches to 1 while with other elements to be 0, which corresponds to the minimum entropy of  $p_{ds}$ ; while the disagreement reaches the minimum if each element in  $p_{ds}$  is equal, indicating the maximum entropy of  $p_{ds}$ . Hence, the disagreement can be computed via the information entropy function  $\mathcal{H}_{info}(\cdot)$ , and formulated as

$$\mathcal{H}_{info}(p_{ds}) = \sum_{c=1}^C p_{ds}(c) \log \frac{1}{p_{ds}(c)}. \quad (3)$$

For different datasets,  $C$  varies greatly, we further normalize

$\mathcal{H}_{info}(p_{ds})$  as

$$\mathcal{H} = 1 - \mathcal{H}'_{info}(p_{ds}) \in [0, 1], \quad (4)$$

where  $\mathcal{H}'_{info}(p_{ds}) = \frac{\mathcal{H}_{info}(p_{ds}) - \min(\mathcal{H}_{info}(p_{ds}))}{\max(\mathcal{H}_{info}(p_{ds})) - \min(\mathcal{H}_{info}(p_{ds}))}$ . The constant  $\max(\mathcal{H}_{info}(p_{ds})) = -\sum_{c=1}^C \frac{1}{C} \log \frac{1}{C}$  represents the maximum value of  $\mathcal{H}_{info}(p_{ds})$ , in the case that each element in  $p_{ds}$  has the same class probability  $\frac{1}{C}$ , where Q perfectly aligns with P (i.e.,  $z_p = z_q$ ); while  $\min(\cdot)$  is utilized to obtain the minimum of  $\mathcal{H}_{info}(p_{ds})$  within a batch. Thus, the sample adaptability is closely related to  $\mathcal{H}$  — the more  $\mathcal{H}$ , the larger the disagreement over the sample to both P and Q, to yield the samples with large adaptability.

As per Eq.(2)(4),  $p_{ds}$  distributes in a  $C$ -dimensional vector space, while  $\mathcal{H}$  is a real value. We hence map that to a  $C$ -dimensional space  $\mathbb{R}^C$  while characterizing  $\mathcal{H}$ , which is achieved via the unit vector  $\frac{p_{ds}}{\|p_{ds}\|}$  ( $\|\cdot\|$  denotes  $\ell_2$  norm), then  $\mathcal{H}$  is reformulated as

$$\mathcal{H}_C = \frac{p_{ds}}{\|p_{ds}\|} \mathcal{H} = \frac{p_{ds}}{\|p_{ds}\|} (1 - \mathcal{H}'_{info}(p_{ds})) \in \mathbb{R}^C, \quad (5)$$

where the category information (i.e.,  $\frac{p_{ds}}{\|p_{ds}\|}$ ) of the generated sample apart from the disagreement (i.e.,  $\mathcal{H}$ ) is exploited to well measure the sample adaptability.

The measurement of sample adaptability inspires us to revisit the DFQ process based on  $\mathcal{H}$ : the goal of G is to generate the sample with *large* adaptability to Q, i.e., *increasing* the sample adaptability, by *maximizing*  $\mathcal{H}$  (the gain of  $\mathcal{H}$  is *positive*, serving as the reward for G), to benefit Q; while for the calibration process, Q is optimized to recover itself with the generated sample. It infers that such sample adaptability will not be large to Q, while the sample can no longer benefit Q after Q is refined; in other words, such process of benefiting Q by learning from P results in *decreasing* the sample adaptability, which is achieved by *minimizing*  $\mathcal{H}$  (the gain of  $\mathcal{H}$  is *negative*, serving as the loss for Q), *adversarial* to maximizing  $\mathcal{H}$ , which encourages the loss for Q to cancel out the reward for G, such that the sum of reward and loss tends to be zero. Such fact makes the overall changing (summation) of the disagreement ( $\mathcal{H}$ ) *close* to 0 (see Fig.4 for the intuition), which is in line with the principle of zero-sum game (Shoham and Leyton-Brown 2008; v. Neumann 1928), as discussed in the next section.

### Zero-Sum Game: Adversarial Game for Data-Free Quantization

We formally revisit the DFQ process from a game-theory perspective (Li et al. 2022; Zhang et al. 2022), and formulate the DFQ as a zero-sum game between two players — a generator and a quantized network, which is an expansion of Eq.(1), as follows:

$$\min_{\theta_g \in \Theta_g} \max_{\theta_q \in \Theta_q} \mathcal{R}(\theta_g, \theta_q) = \min_{\theta_g \in \Theta_g} \max_{\theta_q \in \Theta_q} \mathbb{E}_{z,y} [1 - \mathcal{H}'_{info}(p_{ds})], \quad (6)$$

where G and Q are parameterized by  $\theta_g \in \Theta_g$  and  $\theta_q \in \Theta_q$ , respectively. In particular, Eq.(6) is iteratively optimized via gradient descent during the training process, where each iteration consists of two steps: on one hand,  $\theta_q$  is fixed,  $\mathcal{R}(\theta_g, \theta_q)$  in Eq.(6) is *maximized* to update  $\theta_g$ , which is equivalent to generating the sample with *large* adaptability,

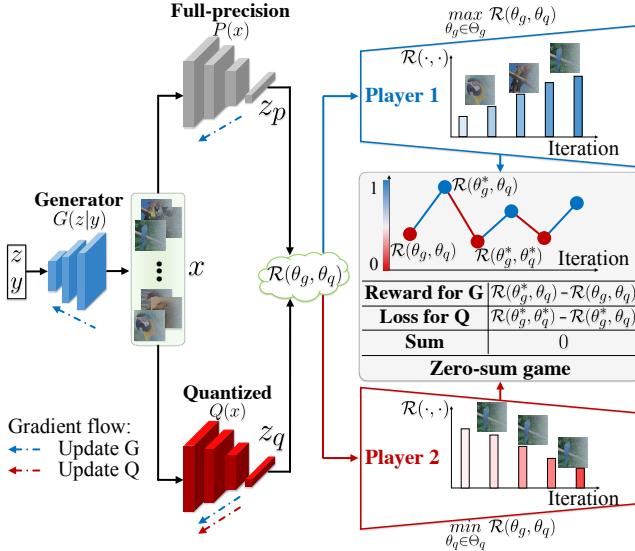


Figure 4: Illustration of data-free quantization process as a zero-sum game between G (player 1) and Q (player 2). G generates the sample with large adaptability by maximizing  $\mathcal{R}(\cdot, \cdot)$  during the maximization process, such sample adaptability is further reduced by minimizing  $\mathcal{R}(\cdot, \cdot)$  during the minimization process when calibrating Q.

i.e., increasing the sample adaptability; on the other hand,  $\theta_g$  is fixed,  $\mathcal{R}(\theta_g, \theta_q)$  in Eq.(6) is minimized to update  $\theta_q$ , which is equivalent to calibrating Q with the generated sample, benefiting from decreasing the sample adaptability. During the zero-sum game, the overall changing (summation) of the sample adaptability to Q incurred by maximum and minimum optimization is close to zero. One critical question is how to converge Eq.(6) to achieve the equilibrium (i.e., the sample adaptability no longer changes) for the zero-sum game. Thanks to the classical Nash equilibrium (Cardoso et al. 2019) defined as a state, where no player can improve its individual gain during the zero-sum game, the optimization process will reach an equilibrium  $(\theta_g^*, \theta_q^*)$  when (1) for the fixed  $\theta_q^*$ , G fails to maximize  $\mathcal{R}(\theta_g, \theta_q^*)$  to obtain a  $\theta_g \in \Theta_g$ , yielding  $\mathcal{R}(\theta_g, \theta_q^*) > \mathcal{R}(\theta_g^*, \theta_q^*)$ ; (2) for the fixed  $\theta_g^*$ , Q can no longer improve itself by minimizing  $\mathcal{R}(\theta_g^*, \theta_q)$  to obtain a  $\theta_q \in \Theta_q$ , yielding  $\mathcal{R}(\theta_g^*, \theta_q) < \mathcal{R}(\theta_g^*, \theta_q^*)$ . Based on that, for all  $\theta_g \in \Theta_g$  and  $\theta_q \in \Theta_q$ ,  $(\theta_g^*, \theta_q^*)$  satisfies the following inequality:

$$\mathcal{R}(\theta_g, \theta_q^*) \leq \mathcal{R}(\theta_g^*, \theta_q^*) \leq \mathcal{R}(\theta_g^*, \theta_q), \quad (7)$$

where  $\theta_g^*$  and  $\theta_q^*$  are the parameters of G and Q under an equilibrium state, respectively.

We remark that maximizing  $\mathcal{R}(\cdot, \cdot)$  in Eq.(6) is equivalent to generating the sample with the largest adaptability, which may incur too large disagreement (i.e., too small  $\mathcal{H}'_{info}(p_{ds})$ ) between P and Q, while Q (especially for Q with low bit width) has no sufficient ability to learn informative knowledge from P, which, in turn, encourages G to generate the sample with the lowest adaptability by alternatively maximizing  $\mathcal{R}(\cdot, \cdot)$  in Eq.(6). However, encouraging the sample with the lowest adaptability may incur too

large agreement (i.e., too large  $\mathcal{H}'_{info}(p_{ds})$ ) between P and Q, where the generated sample may not be informative to calibrate Q (especially for Q with high bit width). The above facts indicate that the sample with either largest or lowest adaptability generated by maximizing  $\mathcal{R}(\cdot, \cdot)$  in Eq.(6) is not necessarily the best. To address the issues, we further refine the maximization objective of Eq.(6) during the zero-sum game by the cooperation between the disagreement and agreement sample along with the bound constraints on sample adaptability, which will be elaborated in the next.

### Maximization Process: Generating Samples with Desirable Adaptability

According to Definition 1, the category (label) information is crucial to establish the dependence of generated sample on Q. Hence, to generate the disagreement sample with large adaptability and further maximize  $\mathcal{R}(\cdot, \cdot)$  in Eq.(6), we exploit the category (label) information to generate the sample. To achieve it, given the label  $y$ , the generated sample should be classified as disagreement sample with the same label  $y$ . Thereby, we present the following Cross-Entropy loss  $\mathcal{H}_{CE}(\cdot, \cdot)$  to match  $p_{ds}$  and  $y$ , formulated as

$$\mathcal{L}_{ds} = \mathbb{E}_{z,y}[\mathcal{H}_{CE}(p_{ds}, y)]. \quad (8)$$

We aim to minimize Eq.(8) to encourage G to generate the disagreement sample that P can predict correctly but Q fails, which, however, may incur too large disagreement (i.e., too small  $\mathcal{H}'_{info}(p_{ds})$ ) between P and Q, thus fail to yield the desirable sample adaptability. To remedy such issue, complementary to the disagreement sample, we further define the agreement sample below:

**Definition 2** (Agreement Sample). *Based on Definition 1, we say the generated sample  $x$  is the agreement sample if  $\text{argmax}(z_p) = \text{argmax}(z_q)$ .*

Thus, similar in spirit to  $p_{ds}$ , the probability vector that describes the agreement between P and Q, is formulated as

$$p_{as} = \text{softmax}(z_p + z_q) \in \mathbb{R}^C, \quad (9)$$

where  $p_{as}(c) = \frac{\exp(z_p(c) + z_q(c))}{\sum_{j=1}^C \exp(z_p(j) + z_q(j))}$  is the  $c$ -th entry of the vector  $p_{as}$  as the probability that  $x$  is the agreement sample of the  $c$ -th class. Following Eq.(8), the loss function for generating agreement sample is given as

$$\mathcal{L}_{as} = \mathbb{E}_{z,y}[\mathcal{H}_{CE}(p_{as}, y)], \quad (10)$$

which is minimized to encourage to generate the agreement sample that both P and Q can correctly predict to possess larger  $\mathcal{H}'_{info}(p_{ds})$ .

**Cooperation for desirable sample adaptability** Upon the above, let  $\mathcal{L}_{ds}$  and  $\mathcal{L}_{as}$  cooperate with each other to maximize  $\mathcal{R}(\cdot, \cdot)$  in Eq.(6), the sample with desirable adaptability can be generated. Hence, the generation loss is given as

$$\mathcal{L}_s = \mathcal{L}_{ds} + \mathcal{L}_{as}. \quad (11)$$

With  $\mathcal{L}_s$ , when a large disagreement incurs,  $\mathcal{L}_{as}$  dominates to generate agreement sample, to reduce the gap between P and Q, so as to enlarge  $\mathcal{H}'_{info}(p_{ds})$  for disagreement sample; to be analogous,  $\mathcal{L}_{ds}$  dominates to reduce  $\mathcal{H}'_{info}(p_{ds})$  for agreement sample; see Fig.5. Nevertheless, in some cases, there still exist the samples with either too small or large  $\mathcal{H}'_{info}(p_{ds})$  beyond the cooperation ability.

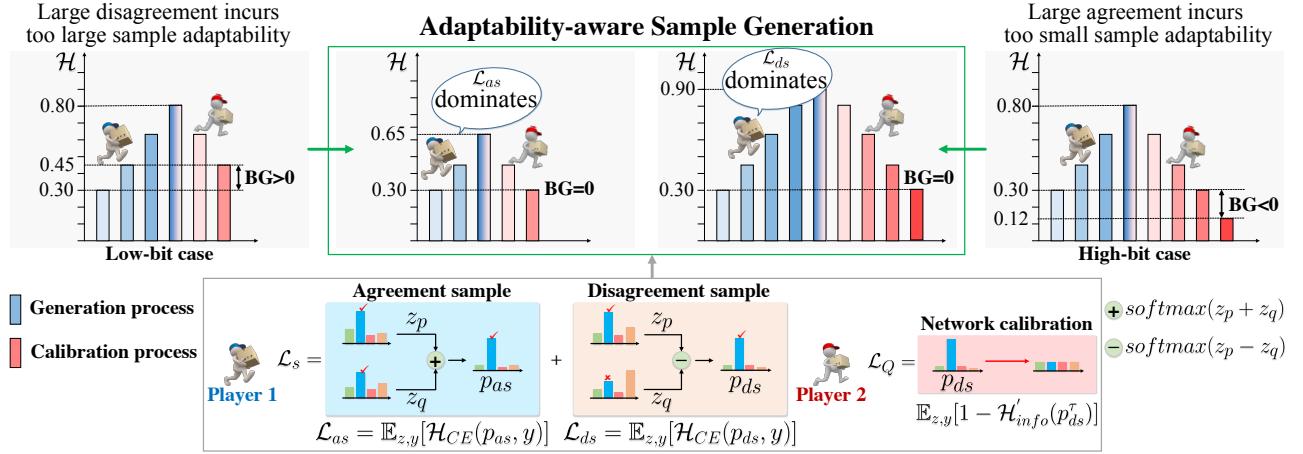


Figure 5: Illustration of AdaSG on achieving the stationarity of the zero-sum game via the balance gap (BG). When  $BG > 0$  ( $BG < 0$ ), the generated sample leads to a large disagreement (agreement) between P and Q, owing to the weak (strong) learning ability of Q, e.g., Q with low-bit (high-bit) precision, resulting into too large (small) sample adaptability. When  $BG = 0$ , the DFQ game process is stationary, revealing that the sample adaptability increased by the maximization process can be fully exploited by the minimization process to maximally benefit Q with varied bit widths.

**Bound constraint on sample adaptability** To address the above problem, we further impose the bound constraints for  $\mathcal{H}'_{info}(p_{ds})$  via the hinge loss (Lim and Ye 2017) below:

$$\mathcal{L}_b = \mathbb{E}_{z,y}[\max(\lambda_l - \mathcal{H}'_{info}(p_{ds}), 0)] + \mathbb{E}_{z,y}[\max(\mathcal{H}'_{info}(p_{ds}) - \lambda_u, 0)], \quad (12)$$

where  $\lambda_l$  and  $\lambda_u$  denote the lower and upper bound of  $\mathcal{H}'_{info}(p_{ds})$ , such that  $0 \leq \lambda_l < \lambda_u \leq 1$ . Specifically,  $\lambda_l$  serves to prevent G from generating the sample with too large adaptability (i.e., too small  $\mathcal{H}'_{info}(p_{ds})$ ) via maximizing Eq.(6), while  $\lambda_u$  aims to avoid the sample with too small adaptability (i.e., too large  $\mathcal{H}'_{info}(p_{ds})$ ), which, in turn, offer the guarantee for the cooperation between  $\mathcal{L}_{ds}$  and  $\mathcal{L}_{as}$ .

**BNS information** The above fact discusses the adaptability of single sample to Q, while calibrating Q requires the generated samples within a batch, motivating us to exploit the distribution information of the training data. We consider the batch normalization statistics (BNS) information (Xu et al. 2020; Choi et al. 2021) about the training data contained in P, which is learned by

$$\mathcal{L}_{BNS} = \sum_{m=1}^M (\|\mu_m^g - \mu_m\|_2^2 + \|\sigma_m^g - \sigma_m\|_2^2), \quad (13)$$

where  $\mu_m^g/\sigma_m^g$  are the mean/variance of the generated sample's distribution at the  $m$ -th BN layer of the total  $M$  layers; and  $\mu_m/\sigma_m$  are the corresponding mean/variance parameters stored in the  $m$ -th BN layer of P.

**Overall loss** To this end, we finalize the loss function for the maximization (sample generation) process as

$$\mathcal{L}_G = \alpha(\mathcal{L}_{ds} + \mathcal{L}_{as}) + \beta\mathcal{L}_b + \gamma\mathcal{L}_{BNS}, \quad (14)$$

where  $\alpha, \beta$  and  $\gamma$  are the balance hyperparameters. Throughout minimizing  $\mathcal{L}_G$ , which is equivalent to maximizing  $\mathcal{R}(\cdot, \cdot)$  in Eq.(6), the sample with desirable adaptability can be generated, such sample is exploited by calibrating Q during the minimization process according to Eq.(6).

### Minimization Process: Calibrating Q by Reducing Sample Adaptability

With the above generated sample, the goal of the minimization process is to calibrate Q for the performance recovery by minimizing  $\mathcal{R}(\cdot, \cdot)$  in Eq.(6), which indicates that the disagreement between P and Q is decreased to reach the agreement. In particular, we further introduce a temperature parameter  $\tau$  to soften the output, thus the loss function for the minimization (calibration) process is formulated as

$$\mathcal{L}_Q = \mathbb{E}_{z,y}[1 - \mathcal{H}'_{info}(p_{ds}^\tau)], \quad (15)$$

where  $p_{ds}^\tau = \text{softmax}((z_p - z_q)/\tau)$ . By alternatively optimizing Eq.(14) and Eq.(15) during a zero-sum game, the samples with desirable adaptability can be generated by G to maximally recover Q until reaching a Nash equilibrium.

### Theoretical Analysis: a Balance Gap for Sample Adaptability

One may wonder whether the above maximization-vs-minimization process can keep stationary during the optimization process over Eq.(6), which is critical for sample generation with desirable adaptability and Q's recovery. We confirm that by our theoretical analysis based on the balance gap (BG) over sample adaptability, which is defined as:

**Definition 3** (Balance Gap). *Considering the objective  $\mathcal{R}(\cdot, \cdot)$  for the DFQ game process. Assume that, after one iteration, the parameters  $(\theta_g^1, \theta_q^1)$  of the game are updated to  $(\theta_g^2, \theta_q^2)$  via the gradient descent. Then, the balance gap (BG) is defined as*

$$BG = \mathcal{R}(\theta_g^2, \theta_q^2) - \mathcal{R}(\theta_g^1, \theta_q^1), \quad (16)$$

which encodes the deviation for the value of  $\mathcal{R}(\cdot, \cdot)$  before and after each iteration. When  $BG > 0$  or  $BG < 0$ , the generated sample leads to a large disagreement or agreement between P and Q, owing to the weak or strong learning ability of Q, where the sample with too large or small adaptability

fails to benefit Q during the calibration process. We remark that BG is well bounded to avoid the large deviation, which facilitates generating the samples with desirable adaptability. Such fact is validated via the following proposition:

**Proposition 1.** *Considering the stationarity of DFQ game process. Then, the balance gap (BG) is bounded below:*

$$|BG| = |\mathcal{R}(\theta_g^2, \theta_q^2) - \mathcal{R}(\theta_g^1, \theta_q^1)| \leq L \|\theta_g^2 - \theta_g^1\| \quad (17)$$

where  $L$  is the Lipschitz constant.

**Proof.** Considering the Lipschitz continuity of the objective function  $\mathcal{R}(\cdot, \cdot)$ , we divide the proof into two parts:

**Part 1. Lipschitz Continuity.** We say a function  $f(x, y)$  is  $L$ -Lipschitz continuous with respect to a norm  $\|\cdot\|$  if

$$|f(x_2, y_2) - f(x_1, y_1)| \leq L \|[x_2; y_2] - [x_1; y_1]\| \quad (18)$$

for any  $x_1, x_2 \in X$  and any  $y_1, y_2 \in Y$ . The previous inequality holds if and only if

$$\|[\nabla_x f(x, y); \nabla_y f(x, y)]\| \leq L \quad (19)$$

for any  $x \in X$  and any  $y \in Y$ .  $L$  is the Lipschitz constant.

**Part 2.** According to the **first-order Taylor expansion** and **Cauchy inequality**, we have:

$$\begin{aligned} |BG| &= |\mathcal{R}(\theta_g^2, \theta_q^2) - \mathcal{R}(\theta_g^1, \theta_q^1)| \\ &= \|[\nabla_{\theta_g} \mathcal{R}(\theta_g, \theta_q); \nabla_{\theta_q} \mathcal{R}(\theta_g, \theta_q)]^T [\theta_g^2 - \theta_g^1; \theta_q^2 - \theta_q^1]\| \\ &\leq \|[\nabla_{\theta_g} \mathcal{R}(\theta_g, \theta_q); \nabla_{\theta_q} \mathcal{R}(\theta_g, \theta_q)]^T\| \|\theta_g^2 - \theta_g^1; \theta_q^2 - \theta_q^1\|, \end{aligned}$$

then Eq.(17) holds if and only if  $\nabla_{\theta_g} \mathcal{R}(\theta_g, \theta_q)$  and  $\nabla_{\theta_q} \mathcal{R}(\theta_g, \theta_q)$  are upper bounded. First, we note that

$$\mathcal{R}(\theta_g, \theta_q) = \mathbb{E}_{z, y} [1 - \mathcal{H}'_{\text{info}}(p_{ds})], \quad (20)$$

where  $\mathcal{H}'_{\text{info}}(p_{ds})$  is obtained by normalizing  $\mathcal{H}_{\text{info}}(p_{ds}) = \sum_{c=1}^C p_{ds}(c) \log \frac{1}{p_{ds}(c)}$ , and  $p_{ds} = \text{softmax}(z_p - z_q)$ ; and the corresponding gradients are computed by

$$\begin{aligned} \nabla_{\theta_g} \mathcal{R}(\theta_g, \theta_q) &= \frac{\partial \mathcal{R}(\theta_g, \theta_q)}{\partial p_{ds}} \frac{\partial p_{ds}}{\partial x} \frac{\partial x}{\partial \theta_g}, \\ \nabla_{\theta_q} \mathcal{R}(\theta_g, \theta_q) &= \frac{\partial \mathcal{R}(\theta_g, \theta_q)}{\partial p_{ds}} \frac{\partial p_{ds}}{\partial z_q} \frac{\partial z_q}{\partial \theta_q}, \end{aligned} \quad (21)$$

where  $\frac{\partial x}{\partial \theta_g}$  and  $\frac{\partial z_q}{\partial \theta_q}$  are the gradients produced by  $G$  and  $Q$ , respectively. It is apparent that the gradients of the information entropy function  $\mathcal{H}_{\text{info}}(\cdot)$ , the softmax function and the activation function (e.g., *Sigmoid*, *Relu*) in deep neural network are **all bounded**. Therefore,  $\forall L \in \mathbb{R}$ , make  $\|[\nabla_{\theta_g} \mathcal{R}(\theta_g, \theta_q); \nabla_{\theta_q} \mathcal{R}(\theta_g, \theta_q)]^T\|$  be upper bounded. To sum up, the balance gap (BG) is bounded by Eq.(17).  $\square$

The proposition provides an upper and lower bound for BG, to well avoid the sample generation encoding too large or small adaptability, in line with the intuition of  $\mathcal{L}_b$ , which offers a guarantee for the stationarity of DFQ game, we further achieve the stationarity via the following proposition:

**Proposition 2.** *During each iteration, the DFQ game is stationary, that is, the sample adaptability increased by the maximization process can be fully exploited by the minimization process to maximally benefit Q, when BG = 0.*

**Proof.** As aforementioned, each iteration contains the maximization and minimization process, then we have:

$$\begin{aligned} BG &= \mathcal{R}(\theta_g^2, \theta_q^2) - \mathcal{R}(\theta_g^1, \theta_q^1) \\ &= (\mathcal{R}(\theta_g^2, \theta_q^1) - \mathcal{R}(\theta_g^1, \theta_q^1)) - (\mathcal{R}(\theta_g^2, \theta_q^1) - \mathcal{R}(\theta_g^2, \theta_q^2)) \\ &= \Delta_g - \Delta_q, \end{aligned}$$

where  $\Delta_g$  and  $\Delta_q$  denote the **deviation** for the value of  $\mathcal{R}(\cdot, \cdot)$  during the **maximization** and **minimization** process, respectively. Thus, if  $\Delta_g > \Delta_q$  (i.e.,  $BG > 0$ ), the generated sample leads to a large disagreement between  $P$  and  $Q$ , owing to the weak learning ability of  $Q$ ; if  $\Delta_g < \Delta_q$  (i.e.,  $BG < 0$ ), the generated sample leads to a large agreement between  $P$  and  $Q$ , owing to the strong learning ability of  $Q$ . Under such cases, the sample with too large or small adaptability fails to benefit  $Q$  during the calibration process. To sum up, if and only if  $\Delta_g = \Delta_q$  (i.e.,  $BG = 0$ ), the DFQ game is stationary, that is, the sample adaptability increased by the maximization process can be fully exploited by the minimization process to maximally benefit  $Q$  by minimizing  $\mathcal{L}_Q$ , which is equivalent to minimizing  $\mathcal{R}(\cdot, \cdot)$  in Eq.(6).  $\square$

The proposition discloses the effectiveness of the coordination between  $\mathcal{L}_{ds}$  and  $\mathcal{L}_{as}$  during the maximization process. Specifically, when  $BG > 0 (< 0)$ , the sample adaptability is too large or small; therefore, for next iteration,  $\mathcal{L}_{as}$  ( $\mathcal{L}_{ds}$ ) will dominate to generate the sample with desirable adaptability by minimizing  $\mathcal{L}_G$  (equivalent to maximizing  $\mathcal{R}(\cdot, \cdot)$  in Eq.(6)), to achieve the stationarity of the DFQ game, i.e.,  $BG = 0$ ; see Fig.5.

## Experiment

### Experimental Settings and Implementation Details

We validate AdaSG over three typical image classification datasets, including **CIFAR-10**, **CIFAR-100** (Krizhevsky 2009) and **ImageNet (ILSVRC2012)** (Russakovsky et al. 2015). CIFAR-10 and CIFAR-100 contain 10 and 100 classes of images, respectively. Both of them are split into 50K training images and 10K testing images. ImageNet consists of 1.2M samples for training and 50k samples for validation with 1000 categories. For data-free setting, only validation sets are adopted to evaluate the performance of the quantized models (Q). In the experiments, we quantize pre-trained full-precision networks (P) including ResNet-20 for CIFAR, and ResNet-18, ResNet-50, and MobileNetV2 for ImageNet, via the following quantizer to yield Q:

**Quantizer.** Following (Xu et al. 2020; Choi et al. 2021), we quantize both full-precision (float32) weights and activations into  $n$ -bit precision via a symmetric linear quantization method based on (Jacob et al. 2018) below:

$$\theta_q = \text{round}\left((2^n - 1) * \frac{\theta - \theta_{\min}}{\theta_{\max} - \theta_{\min}} - 2^{n-1}\right), \quad (22)$$

where  $\theta$  and  $\theta_q$  are the full-precision and quantized value.  $\text{round}(\cdot)$  returns the nearest integer value to the input.  $\theta_{\min}$  and  $\theta_{\max}$  are the minimum and maximum of  $\theta$ . Regarding the bit width  $n$ , we select  $n = \{3, 4, 5\}$ , which are representative for low-bit and high-bit cases, particularly: 3-bit quantization actually leads to a huge accuracy loss, which is

a major challenge for the existing DFQ methods; while 5-bit or higher-bit quantization produces a small performance loss, which is selected to validate the generalization ability.

For the *maximization* process, we construct the architecture of the generator (G) following ACGAN (Odena, Olah, and Shlens 2017), while P and Q play the role of discriminator, where G is trained with the loss function Eq.(14) using Adam (Kingma and Ba 2014) as an optimizer with a momentum of 0.9 and a learning rate of 1e-3. For the *minimization* process, Q is optimized with the loss function Eq.(15), where SGD with Nesterov (Nesterov 1983) is adopted as an optimizer with a momentum of 0.9 and weight decay of 1e-4. For CIFAR, the learning rate is initialized to 1e-4 and decayed by 0.1 for every 100 epochs, while it is 1e-5 (1e-4 for ResNet-50) and divided by 10 at epoch 350 (at epoch 200 and 300 for ResNet-50) on ImageNet. The generator and quantized model are totally trained for 400 epochs. The batch size is set to 16. For the hyperparameters,  $\alpha$ ,  $\beta$  and  $\gamma$  in Eq.(14);  $\lambda_l$  and  $\lambda_u$  in Eq.(12) are empirically set to be 0.1, 1, 1, 0.3 and 0.8. All experiments are implemented with pytorch (Paszke et al. 2019) based on the code of GDFQ (Xu et al. 2020) and run on an NVIDIA GeForce GTX 1080 Ti GPU and an Intel(R) Core(TM) i7-6950X CPU @ 3.00GHz.

To validate how AdaSG generates the sample with desirable adaptability to maximally benefit Q, we empirically validate why AdaSG works, including the comparisons with the state-of-the-arts, ablation study as well as visual analysis.

## Why does AdaSG Work?

We experimentally verify the core idea of AdaSG — generating the sample with desirable adaptability to recover the performance of Q with varied bit widths. We perform the data-free quantization experiments with ResNet-18 (3-bit and 5-bit precision) serving as both P and Q on ImageNet. Fig.6(a1)(b1) illustrates that the disagreement ( $\mathcal{H}_{info}(p_{ds})$ ) in Eq.(3)) between P and Q for AdaSG fluctuates stably within a small range compared to GDFQ (Xu et al. 2020) and Qimera (Choi et al. 2021), confirming that the generated sample with desirable adaptability by AdaSG is fully exploited to benefit Q, and the bound constraint loss (Eq.(12)) can avoid generating the sample with too large or small adaptability, which leads to ridiculously large disagreement or agreement. Fig.6(a2)(b2) reveals that the balance gap (Definition 3) keeps on par with zero (either  $BG > 0$  or  $BG < 0$ ) during the iteration process unlike GDFQ and Qimera, which confirm the principle that AdaSG achieves the stationarity of the DFQ game and the coordination between  $\mathcal{L}_{ds}$  (Eq.(8)) and  $\mathcal{L}_{as}$  (Eq.(10)) can generate the sample with desirable adaptability to maximally benefit Q.

## Comparison with State-of-the-arts

To verify the superiority of AdaSG, we compare it with typical DFQ fashions, including: GDFQ (Xu et al. 2020), ARC (Zhu et al. 2021) and Qimera (Choi et al. 2021): reconstructing the original data from P; ZAQ (Liu, Zhang, and Wang 2021) focuses primarily on the adversarial sample generation rather than adversarial game process for AdaSG; IntraQ (Zhong et al. 2022) optimizes the noise to obtain fake sample without a generator; AIT (Choi et al. 2022) focuses on

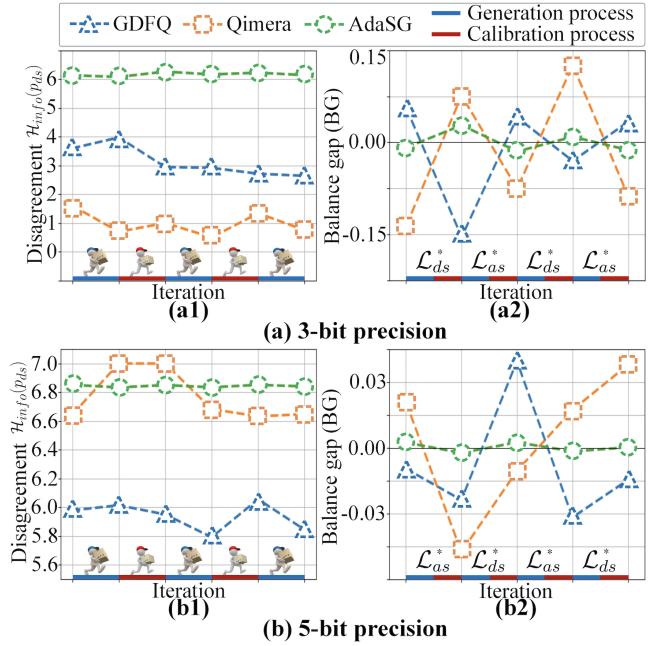


Figure 6: Illustration of AdaSG on generating the samples with desirable adaptability to Q under (a) 3-bit and (b) 5-bit precision. (a1)(b1) Disagreement (i.e.,  $\mathcal{H}_{info}(p_{ds})$ ) between P and Q during the generation and calibration process. (a2)(b2) Balance gap (BG).

improving the loss function and manipulating the gradients for ARC to generate better sample, denoted as ARC+AIT.

Table 1 summarizes our following observations: 1) AdaSG offers a significant and consistent performance gain over the state-of-the-arts, in line with our purpose of generating the sample with desirable adaptability to maximally benefit Q. Impressively, AdaSG achieves at most 9.71%, 6.63% and 35.87% accuracy gains on CIFAR-10, CIFAR-100 and ImageNet, respectively. Especially, compared with GDFQ, ARC and Qimera where Q is independent of the generation process, AdaSG obtains accuracy improvement with a large margin, *e.g.*, at least 0.3% gain (ResNet-20 with 5w5a on CIFAR-10), confirming the necessity of AdaSG focusing on the sample adaptability to Q. Specifically, ZAQ suffers from a large performance gap compared to AdaSG, since many unexpected samples are generated without considering the sample adaptability, which are harmful to calibrating Q. AdaSG shows obvious advantages over AIT despite of the combination with ARC. 2) AdaSG achieves the substantial gains for Q with varied bit-widths, confirming the desirable adaptability of our generated sample to varied Q. Note that, for 3-bit case, most of the existing methods suffer from a poor accuracy, even fail to converge, while AdaSG obtains at most 35.87% (ResNet-18 with 3w3a) and at least 4.21% (ResNet-20 with 3w3a on CIFAR-100) performance gains.

## Ablation Study

**Validating adaptability with disagreement and agreement samples** As aforementioned, the disagreement and agreement samples play a critical role on the sample adapt-

Dataset	Model (FP.)	Bit width (nwna)	ZAQ (CVPR 2021)	IntraQ (CVPR 2022)	ARC+AIT (CVPR 2022)	GDFQ (ECCV 2020)	ARC (IJCAI 2021)	Qimera (NeurIPS 2021)	AdaSG (Ours)
CIFAR-10	ResNet-20 (93.89)	3w3a	-	77.07	-	75.11 <sup>†</sup>	-	74.43 <sup>†</sup>	<b>84.14</b>
		4w4a	92.13	91.49	90.49	90.11	88.55	91.26	<b>92.10</b>
		5w5a	93.36	-	92.98	93.38	92.88	93.46	<b>93.76</b>
CIFAR-100	ResNet-20 (70.33)	3w3a	-	48.25	-	47.61 <sup>†</sup>	-	46.13 <sup>†</sup>	<b>52.76</b>
		4w4a	60.42	64.98	61.05	63.75	62.76	65.10	<b>66.42</b>
		5w5a	68.70	-	68.40	67.52	68.40	69.02	<b>69.42</b>
ImageNet	ResNet-18 (71.47)	3w3a	-	-	-	20.23 <sup>†</sup>	23.37	1.17 <sup>†</sup>	<b>37.04</b>
		4w4a	52.64	66.47	65.73	60.60	61.32	63.84	<b>66.50</b>
		5w5a	64.54	69.94	70.28	68.49	68.88	69.29	<b>70.29</b>
	MobileNetV2 (73.03)	3w3a	-	-	-	1.46 <sup>†</sup>	14.30	-	<b>26.90</b>
		4w4a	0.10	65.10	<b>66.47</b>	59.43	60.13	61.62	65.15
		5w5a	62.35	71.28	<b>71.96</b>	68.11	68.40	70.45	71.61
	ResNet-50 (77.73)	3w3a	-	-	-	0.31 <sup>†</sup>	1.63	-	<b>16.98</b>
		4w4a	53.02	-	68.27	54.16	64.37	66.25	<b>68.58</b>
		5w5a	73.38	-	76.00	71.63	74.13	75.32	<b>76.03</b>

Table 1: Accuracy (%) comparison with the state-of-the-arts on CIFAR-10, CIFAR-100 and ImageNet. <sup>†</sup>: the results implemented by author-provided code. -: no results are reported. nwna indicates the weights and activations are quantized to  $n$  bit. FP.: the accuracy of full-precision model. The best results are reported with boldface.

Model (Full precision)	$\mathcal{L}_{ds}$	$\mathcal{L}_{as}$	$\mathcal{L}_b$	$\mathcal{L}_{BNS}$	Bit width	
					3w3a	5w5a
ResNet-18 (71.47)	✓	✓	✓	✓	14.86	69.15
	✓	✓	✓	✓	26.41	69.81
	✓	✓	✓	✓	32.75	70.13
	✓	✓	✓	✓	15.01	70.06
	✓	✓	✓	✓	20.98	67.35
	✓	✓	✓	✓	<b>37.04</b>	<b>70.29</b>

Table 2: Ablation study about varied components of AdaSG. nwna indicates the weights and activations are quantized to  $n$  bit. The best results are reported with boldface.

ability to Q, which serve as a pivotal role between the maximization and minimization process during the zero-sum game. We conduct the ablation study on  $\mathcal{L}_{ds}$  (Eq.(8)) and  $\mathcal{L}_{as}$  (Eq.(10)) over ImageNet. Table 2 suggests the great superiority (37.04% and 70.29%) of AdaSG (including the both) to other cases. It is worth noting that, removing either or both of  $\mathcal{L}_{ds}$  and  $\mathcal{L}_{as}$  obtains a large performance degradation (at most 22.18% and 1.14%), implying the intuition of the cooperation between  $\mathcal{L}_{ds}$  and  $\mathcal{L}_{as}$ . Interestingly, the case without  $\mathcal{L}_b$  (Eq.(12)) receives the minimal accuracy loss (4.29% and 0.16%), confirming the importance of  $\mathcal{L}_b$  on the basis of  $\mathcal{L}_{ds}$  and  $\mathcal{L}_{as}$ .

### Visualization of Generated Samples

To further show the desirable adaptability of the generated samples by AdaSG to Q, we conduct the visual analysis over MobileNetV2 serving as both P and Q on ImageNet by the similarity matrix (each element is obtained by computing the  $\ell_1$  norm between the probability distribution  $p_{ds}$  of every two samples), along with the visualization of generated samples in Fig.7. Fig.7(a) illustrates that the generated sam-

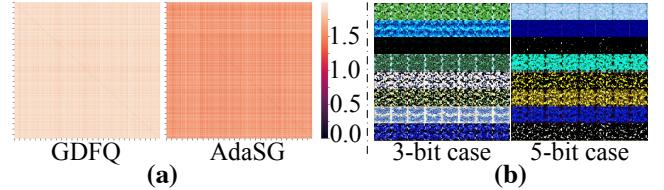


Figure 7: (a) The similarity comparison between the generated samples. (b) Visualization of the generated samples, where each row denotes one of 8 randomly chosen classes from ImageNet.

ples by AdaSG possess a much larger similarity (the darker, the larger) than those by GDFQ, implying that the generated sample by GDFQ varies greatly, where a lots of samples with undesirable adaptability exist against AdaSG. Fig.7(b) shows that the samples generated for 3-bit and 5-bit precision vary greatly, while the samples from varied categories also differ greatly from each other, verifying that the samples possess desirable adaptability to varied Q, upon the fact that the category (label) information is fully exploited.

### Conclusion

In this paper, we rethink date-free quantization process as a zero-sum game between two players — a generator and a quantized network, then further develop an Adaptability-aware Sample Generation (AdaSG) method, which features a dynamic maximization-vs-minimization game process anchored on sample adaptability. The maximization process generates the sample with desirable adaptability, which is further reduced by the minimization process after recovering Q. Balance Gap is defined to achieve the stationarity for the zero-sum game. The theoretical analysis and empirical studies validate the advantages of AdaSG to the existing arts.

## Acknowledgments

This work are supported by National Natural Science Foundation of China under the grant no U21A20470, 62172136, 72188101, U1936217. Key Research and Technology Development Projects of Anhui Province (no.202004a5020043).

## References

Cai, Y.; Yao, Z.; Dong, Z.; Gholami, A.; Mahoney, M. W.; and Keutzer, K. 2020. Zeroq: A novel zero shot quantization framework. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13169–13178.

Cardoso, A. R.; Abernethy, J.; Wang, H.; and Xu, H. 2019. Competing against Nash equilibria in adversarially changing zero-sum games. In *International Conference on Machine Learning*, 921–930. PMLR.

Choi, K.; Hong, D.; Park, N.; Kim, Y.; and Lee, J. 2021. Qimera: Data-free Quantization with Synthetic Boundary Supporting Samples. *Advances in Neural Information Processing Systems*, 34.

Choi, K.; Lee, H. Y.; Hong, D.; Yu, J.; Park, N.; Kim, Y.; and Lee, J. 2022. It’s All In the Teacher: Zero-Shot Quantization Brought Closer to the Teacher. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8311–8321.

Han, S.; Mao, H.; and Dally, W. J. 2015. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149*.

Hinton, G.; Vinyals, O.; and Dean, J. 2015. Distilling the knowledge in a neural network. In *NIPS*.

Jacob, B.; Kligys, S.; Chen, B.; Zhu, M.; Tang, M.; Howard, A.; Adam, H.; and Kalenichenko, D. 2018. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2704–2713.

Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Krizhevsky, A. 2009. Learning Multiple Layers of Features from Tiny Images. *Master’s thesis, University of Tront*.

Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2017. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6): 84–90.

Li, J.; Ren, T.; Yan, D.; Su, H.; and Zhu, J. 2022. Policy learning for robust markov decision process with a mismatched generative model. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 7417–7425.

Lim, J. H.; and Ye, J. C. 2017. Geometric gan. *arXiv preprint arXiv:1705.02894*.

Lin, D.; Talathi, S.; and Annapureddy, S. 2016. Fixed point quantization of deep convolutional networks. In *International conference on machine learning*, 2849–2858. PMLR.

Liu, Y.; Zhang, W.; and Wang, J. 2021. Zero-shot adversarial quantization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1512–1521.

Nesterov, Y. E. 1983. A method for solving the convex programming problem with convergence rate  $O(1/k^2)$ . In *Dokl. akad. nauk Sssr*, volume 269, 543–547.

Odena, A.; Olah, C.; and Shlens, J. 2017. Conditional image synthesis with auxiliary classifier gans. In *International conference on machine learning*, 2642–2651. PMLR.

Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.

Qian, B.; Wang, Y.; Hong, R.; Wang, M.; and Shao, L. 2021. Diversifying inference path selection: Moving-mobile-network for landmark recognition. *IEEE Transactions on Image Processing*, 30: 4894–4904.

Qian, B.; Wang, Y.; Yin, H.; Hong, R.; and Wang, M. 2022. Switchable Online Knowledge Distillation. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XI*, 449–466.

Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; et al. 2015. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3): 211–252.

Shoham, Y.; and Leyton-Brown, K. 2008. *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*. Cambridge University Press.

v. Neumann, J. 1928. Zur theorie der gesellschaftsspiele. *Mathematische annalen*, 100(1): 295–320.

Wang, Y. 2021. Survey on deep multi-modal data analytics: Collaboration, rivalry, and fusion. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 17(1s): 1–25.

Xu, S.; Li, H.; Zhuang, B.; Liu, J.; Cao, J.; Liang, C.; and Tan, M. 2020. Generative low-bitwidth data free quantization. In *European Conference on Computer Vision*, 1–17. Springer.

Zhang, M.; Zhao, P.; Luo, H.; and Zhou, Z.-H. 2022. No-regret learning in time-varying zero-sum games. In *International Conference on Machine Learning*, 26772–26808. PMLR.

Zhang, X.; Qin, H.; Ding, Y.; Gong, R.; Yan, Q.; Tao, R.; Li, Y.; Yu, F.; and Liu, X. 2021. Diversifying sample generation for accurate data-free quantization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15658–15667.

Zhong, Y.; Lin, M.; Nan, G.; Liu, J.; Zhang, B.; Tian, Y.; and Ji, R. 2022. IntraQ: Learning Synthetic Images with Intra-Class Heterogeneity for Zero-Shot Network Quantization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12339–12348.

Zhu, B.; Hofstee, P.; Peltenburg, J.; Lee, J.; and Alars, Z. 2021. AutoReCon: Neural Architecture Search-based Reconstruction for Data-free. In *International Joint Conference on Artificial Intelligence*.