ESANN 2026 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium) and online event, 22-24 April 2026, i6doc.com publ., ISBN 9782875870964. Available from http://www.i6doc.com/en/.

Diminishing Returns - Data Integer Quantization and its Effects on Training Dynamics of Distance Based Classifiers

Thomas Davies[1], Alexander Engelsberger[1], Magdalena Psenickova[1], Thomas Villmann[1,2] *

1- Mittweida University of Applied Sciences
Saxon Institute for Computational Intelligence and Machine Learning
Technikumplatz 17, 09648 Mittweida
2- Technische Universität Freiberg

Abstract. In certain subfields of machine learning, such as those involving homomorphic encryption or quantum computing, it is crucial to estimate the numerical precision of data used for training without compromising model quality. This paper introduces a method for precisely quantifying the loss of accuracy in distance-based classifiers, such as Generalized Learning Vector Quantization, when operating on quantized data represented by bounded integer sets. Our approach employs conditional entropy to measure the information loss induced by quantization, which closely correlates with the model's mean performance.

1 Introduction

With the growing popularity of machine learning, numerous subfields have emerged that address different aspects of the discipline. Particularly relevant to this work are areas such as privacy-preserving machine learning, where homomorphic encryption is applied to machine learning algorithms [1], machine learning on quantum computer circuits [2, 3], and energy-efficient machine learning in edge computing [4]. In each of these domains, the bit depth of numerical data representations is inherently constrained. In homomorphic encryption, the computational complexity and security of the scheme are strongly influenced by the numerical size of the message space. In quantum computing, the architecture of the quantum computer itself constrains the depth of the computational circuit and the length of the input data in a mutually dependent way. Furthermore, with a reduction in the numerical representation of data, more energy-efficient and computational efficient data types can be used on edge computers to enable learning on them.

This paper analyzes the behavior of data quantization into the integers and its effect on the outcome of training (Generalized) Learning Vector Quantization ((G)LVQ) models. Recent work [1] on LVQ under homomorphic encryption has demonstrated that even very small integer ranges are sufficient to train them. We propose a method Entropy-based Integer Quantization Measure (EIQM) for estimating the learning stability across different levels of integer quantization. More precisely, for $b \in \mathbb{N}$ let $\mathcal{B}_b = [-b, b] \cap \mathbb{Z}$ denote the integer range to which

*T.D. and M.P. are supported by the European Social Fund (ESF, 100670478 & 100715238)

ESANN 2026 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium) and online event, 22-24 April 2026, i6doc.com publ., ISBN 9782875870964. Available from http://www.i6doc.com/en/.

data are quantized to, $A_b(\widehat{f}, \mathcal{X}_b)$ the accuracy of a distance based classifier $\widehat{f}$ trained on the integer data $\mathcal{X}_b$ quantized to $\mathcal{B}_b$ and $H(\mathcal{Y} \,|\, \mathcal{X}_b)$ the conditional entropy of the set of class labels $\mathcal{Y}$ given $\mathcal{X}_b$. We aim to experimentally show our hypothesis $A_b(\widehat{f}, \mathcal{X}_b) \propto -H(\mathcal{Y} \,|\, \mathcal{X}_b)$.

2 Generalized Learning Vector Quantization

In this section we briefly introduce the prototype- and distance-based model GLVQ [5]. Let $\mathcal{X} \subseteq \mathbb{R}^n$ denote the dataset with dimensionality $n > 0$ and cardinality $|\mathcal{X}| = N$ where each $\boldsymbol{x} \in \mathcal{X}$ is assigned a class label $c(\boldsymbol{x}) \in \mathcal{Y}$. Furthermore, denote $\mathcal{P} \subset \mathbb{R}^n$ a set of prototypes such that there exists at least one prototype $\boldsymbol{p} \in \mathcal{P}$ for each class $c(\boldsymbol{p}) \in \mathcal{Y}$ of the dataset. The GLVQ cost function is given as $E = \sum_{i=1}^{N} g(\mu(\boldsymbol{x}_i))$ where g is a monotonically increasing activation function and $\mu(\boldsymbol{x}) = \frac{d(\boldsymbol{x},\boldsymbol{p}^+) - d(\boldsymbol{x},\boldsymbol{p}^-)}{d(\boldsymbol{x},\boldsymbol{p}^+) + d(\boldsymbol{x},\boldsymbol{p}^-)}$ is the relative distance difference function. Here, $d(\boldsymbol{x}, \boldsymbol{p}) = ||\boldsymbol{x} - \boldsymbol{p}||_2^2$ denotes the squared Euclidean distance and $\boldsymbol{p}^{\pm} = \min_{\boldsymbol{p} \in \mathcal{P}} d(\boldsymbol{x}, \boldsymbol{p})$, where $\boldsymbol{p}^+$ with $c(\boldsymbol{x}) = c(\boldsymbol{p})$ and $\boldsymbol{p}^-$ with $c(\boldsymbol{x}) \neq c(\boldsymbol{p})$ are the closest prototypes of the same and different class, respectively. During training, for each randomly selected $\boldsymbol{x}$ the closest prototypes $\boldsymbol{p}^{\pm}$ are determined and the gradients are computed as $\nabla_{\boldsymbol{p}^{\pm}} = \frac{\partial g}{\partial \mu} \frac{4d(\boldsymbol{p}^{\mp}, \boldsymbol{x})}{(d(\boldsymbol{x},\boldsymbol{p}^+) + d(\boldsymbol{x},\boldsymbol{p}^-))^2} (\boldsymbol{x} - \boldsymbol{p}^{\pm})$. For some learning rate $0 < \varepsilon \ll 1$ the update of prototypes at time $t > 0$ is given as $\boldsymbol{p}^{\pm}(t+1) = \boldsymbol{p}^{\pm}(t) \pm \varepsilon \nabla_{\boldsymbol{p}^{\pm}(t)}$.

3 Entropy-based Integer Quantization Measure

First, we introduce the method of quantizing data $\boldsymbol{x} \in \mathcal{X}$ into integers. To this end we apply feature scaling. Let $X_i = (\boldsymbol{x}_{1,i}, \boldsymbol{x}_{2,i}, \ldots, \boldsymbol{x}_{N,i})$ denote the vector of all values of the i-th feature in the dataset. Feature scaling to integer range $\mathcal{B}_b$ is computed as

$$\overline{\boldsymbol{x}}_{k,i} = \frac{2b(\boldsymbol{x}_{k,i} - \min X_i)}{\max X_i - \min X_i} - b.$$

The quantized data point is given by $\mathring{\boldsymbol{x}} = \lceil \overline{\boldsymbol{x}} \rfloor$ where $\lceil \cdot \rfloor$ denotes the rounding function to the nearest integer. Note that as b decreases, the more significant the change in distribution of the data is. The probability of data collapsing, i.e. $\mathring{\boldsymbol{x}}_i = \mathring{\boldsymbol{x}}_j$ for distinct $\boldsymbol{x}_i \neq \boldsymbol{x}_j$, increases as $b \searrow 1$, since the number of possible integer realizations is $|\mathcal{B}_b^n| = (2b+1)^n$. We interpret the process of data collapsing as a reduction of randomness within the data space with respect to entropy. A natural way to quantify the degree of randomness or structure in a system is through entropy [6]. To capture the distortion in learning dynamics we utilize conditional entropy:

$$H(\mathcal{Y} \,|\, \mathcal{X}_b) = -\sum_{\mathring{\boldsymbol{x}} \in \mathcal{X}_b, y \in \mathcal{Y}} p(\mathring{\boldsymbol{x}}, y) \log\left(p(y \,|\, \mathring{\boldsymbol{x}})\right) = -\sum_{\mathring{\boldsymbol{x}} \in \mathcal{X}_b, y \in \mathcal{Y}} p(\mathring{\boldsymbol{x}}, y) \log\left(\frac{p(\mathring{\boldsymbol{x}}, y)}{p(\mathring{\boldsymbol{x}})}\right)$$

We assume an empirical distribution $p(\mathring{\boldsymbol{x}}) = \frac{N_{\mathring{\boldsymbol{x}}}^{\circ}}{N}$, where $N_{\mathring{\boldsymbol{x}}}^{\circ} = |\{\mathring{\boldsymbol{x}}_j \mid \boldsymbol{x}_j \in \mathcal{X} \text{ with } \mathring{\boldsymbol{x}} = \mathring{\boldsymbol{x}}_j\}|$ denotes the count of samples quantized to the same grid point. With abuse of notation, $\mathcal{Y}$ and $\mathcal{X}_b$ denote random variables with empirical joint

ESANN 2026 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium) and online event, 22-24 April 2026, i6doc.com publ., ISBN 9782875870964. Available from http://www.i6doc.com/en/.

distribution. The joint probability $p(\overset{\circ}{\boldsymbol{x}}, y) = \frac{N_{(\overset{\circ}{\boldsymbol{x}}, y)}}{N}$ is analogously defined for labeled data $(\overset{\circ}{\boldsymbol{x}}, y)$. The conditional probability $p(y|\overset{\circ}{\boldsymbol{x}})$ quantifies the relative distribution of class labels within each grid point. If only a single class accumulates at a grid point, then $\log(p(y|\overset{\circ}{\boldsymbol{x}})) = 0$ and the corresponding summand vanishes. The mixing of classes on a grid point is reflected by $\log(p(y|\overset{\circ}{\boldsymbol{x}})) < 0$ which can be considered the class label mixing error. The more data points of distinct classes collapse to the same grid point the larger the joint probability $p(\overset{\circ}{\boldsymbol{x}}, y)$ becomes and therefore increasing the class label mixing error. Generally, the larger the conditional entropy is, the higher the chance of class-confusion in the update of the prototype is, thus distorting the hypothesis margin of GLVQ. The EIQM is given as the sequence $(H(\mathcal{Y}|\mathcal{X}_b))_{b=1}^{u}$ for some upper bound $u \in \mathbb{N}$, where lower values reflect a low loss of optimal accuracy while higher values reflect a high loss of optimal accuracy.

4 Methodology

An experiment are conducted on datasets $\mathcal{X}_b$ quantized to a set of integer ranges $(\mathcal{B}_b)_{b=1}^{u}$, with $u > 0$ denoting an upper bound. We define a quantization sweep as the training of GLVQ models with a fixed random seed over the quantized datasets $\mathcal{X}_b$ for each of the integer ranges. The fixed random seed $s \in S \subseteq \mathbb{N}$ ensures consistent training and test splits for a 5-fold cross validation $\widehat{f}_s$ within each sweep. Furthermore, for a fixed seed, the random sampling of training data $\overset{\circ}{\boldsymbol{x}}$ during each epoch is also consistent. Those sweeps are repeated 100 times with different seeds, where for each bound the classifier's accuracy $A_b^s(\widehat{f}_s, \mathcal{X}_b)$ is measured. The mean accuracy over all accuracies is computed as $\overline{A}_b = \frac{1}{|S|}\sum_{s\in S} A_b^s(\widehat{f}_s, \mathcal{X}_b)$. For each bound $\mathcal{B}_b$ we additionally compute the conditional entropy $H(\mathcal{Y}|\mathcal{X}_b)$.

Experiments are conducted over various datasets, where for each dataset we perform experiments on both the original and an outlier-augmented dataset. Outliers introduce stronger distortions under equal-binning integer quantization, as the resulting distributions become skewed due to larger gaps in the integer grid, thus not fully utilizing the integer lattice [7]. The outliers are generated by a set of minimal data points $\mathcal{O}_{min} = \bigcup_{i=1}^{n} \arg\min_{\boldsymbol{x}\in\mathcal{X}} x_i$ and analogously the maximal data points $\mathcal{O}_{max}$ across all dimensions $1 \leq i \leq n$. The augmented dataset is then defined as $\mathcal{X}^k = \{\frac{\boldsymbol{x}}{k} | \boldsymbol{x} \in \mathcal{X}\} \cup \mathcal{O}_{min} \cup \mathcal{O}_{max}$ which is subsequently quantized into $\mathcal{X}_b^k$ over the range $\mathcal{B}_b$. To evaluate the quality of the hypothesis we compute the correlation $\mathrm{Corr}((\overline{A}_b)_{b=1}^{u}, (H(\mathcal{Y}|\mathcal{X}_b^k))_{b=1}^{u})$. We regularized GLVQ to have prototypes $\boldsymbol{p}$ stay within the hypercube $\mathcal{B}_b^n$, i.e. $p_i = \max(p_i, -b)$ and $p_i = \min(p_i, b)$.

5 Experiments[1]

We used a larger collection of datasets found in the UCI database (Iris, Wine, Breast Cancer, Rice, Seeds) [8], a real-world spectral dataset on coffee

[1]Experiment code is available at `https://github.com/lvlanson/data_quantization`

ESANN 2026 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium) and online event, 22-24 April 2026, i6doc.com publ., ISBN 9782875870964. Available from http://www.i6doc.com/en/.

Dataset	Data Points	Features	Classes	Number Prototypes	Outlier Factors k	Upper Bounds u
Iris	150	4	3	1	$(0, 25, 50)$	$(50, 150, 300)$
Wine	178	13	3	1	$(0, 25, 50)$	$(50, 150, 300)$
Breast Cancer	569	30	2	1	$(0, 25, 50)$	$(50, 150, 300)$
Spiral	300	2	2	25	$(0, 25, 50)$	$(50, 150, 300)$
Coffee	210	7	3	1	$(0, 25, 50)$	$(50, 150, 300)$
Rice	3810	7	2	1	$(0, 25, 50)$	$(50, 150, 300)$
Seeds	210	7	3	3	$(0, 25, 50)$	$(50, 150, 300)$

Table 1: Overview of the datasets (left), their characteristics (middle), and experimental configurations (right).

beans[2] (Coffee) and a highly non-linear artificial dataset (Spiral) (Figure 1). Table 1 lists all dataset attributes and their corresponding experimental configurations, where values in the same position align across experiments, for instance an outlier factor $k = 50$ corresponds to an upper bound $u = 300$. We first noticed that even relatively low numerical bounds such as $\mathcal{B}_{10}$ for the Spiral dataset (Figure 1) yield distributions that still reflect the original geometrical structure of the data. This effect becomes even more pronounced for higher dimensional datasets, as discussed in section 3. Most datasets were noticeably stable in their accuracies throughout quantization levels down to $3 \leq b \leq 11$, where higher-dimensional datasets tend to be more stable for lower bounds. In Table 2 we observe that Breast Cancer remained stable across all quantization levels, which can be explained by its high dimensionality. Due to this anomaly, the correlation for Breast Cancer with $(k, u) = (0, 50)$ was the only experiment with low correlation. Added outliers clearly destabilized accuracies in lower bounds. Figure 2 illustrates this dynamic and shows how well EIQM can mirror the average loss in accuracy over various quantization levels. Since breaking-points shifted differently across datasets, we additionally report the minimum entropy as a measure to highlight how class-mixing errors are already noticeable within the fixed bounds in an experiment. Datasets such as Spiral and Rice show that quantization effects are already present at the highest bound. Therefore, breaking-points are to be found outside the given range, which is indicated by the greater-than relation. We found that first significant effects on accuracy appear around $H \approx 0.05$, with drops in accuracy of roughly $0.02 - 0.05$, while for $H \approx 0.1$ drops of roughly 0.1 in accuracy can be observed. It is important to note that the drop is relative to the best expected accuracy, hence, its

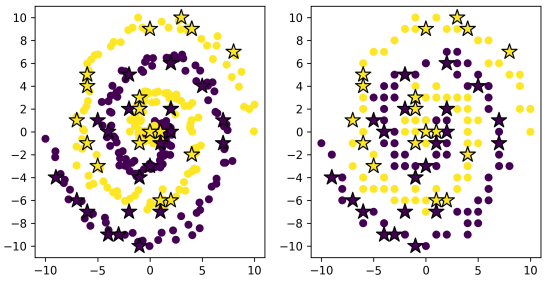

Figure 1: Left the scaled dataset without rounding with the prototypes (stars) as result from quantized learning. On the right the quantized dataset $\mathcal{X}_{10}$ with the respective prototypes (stars).

[2]Special thanks to Udo Seiffert and his team from the Fraunhofer Institute IFF Magdeburg for the collection of the data.

ESANN 2026 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium) and online event, 22-24 April 2026, i6doc.com publ., ISBN 9782875870964. Available from http://www.i6doc.com/en/.

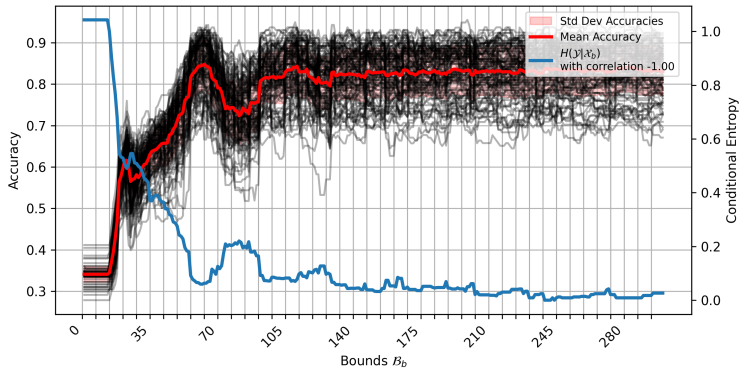

Figure 2: Accuracies for Iris Dataset with outlier factor $k = 50$ and upper bound $b = 300$. Gray lines represent a sweep from $b = 300$ down to $b = 1$.

absolute magnitude can't be determined in general. As mentioned earlier, this observation explains why the breaking-point for Rice and Spiral lies outside of the experimental bounds. Further, datasets such as Seeds show that although some class-mixing error is already present, it is not significant enough to significantly impact the expected accuracy. We also observed for breaking points close to minimal quantization levels the entropy value can break the correlation, while still giving relatively good estimates. A special case here is the result for Breast Cancer with $(k, u) = (0, 50)$ with a correlation of 0.09, where the mean accuracy slightly dips at $b = 2$ but remains stable throughout the sweep, hence the low correlation. Finally, the correlation is only meaningful if a classification problem can be solved meaningful. For Spiral with $(50, 300)$ the dataset already strongly deteriorated where accuracies hovered around $0.5 - 0.6$, which returns an unreliable classifier and therefore a weak correlation.

Dataset	Config. (k, u)	Breaking-Point	Min $H(\mathcal{Y}\|\mathcal{X}_b^k)$	Corr	Dataset	Config. (k, u)	Breaking-Point	Min $H(\mathcal{Y}\|\mathcal{X}_b^k)$	Corr
Iris	$(0, 50)$	2	0	**−0.94**	Seeds	$(0, 50)$	3	0	**−0.99**
	$(25, 150)$	46	0	**−0.99**		$(25, 150)$	56	0.03	**−0.99**
	$(50, 300)$	92	0	**−1.00**		$(50, 300)$	112	0.04	**−0.99**
Wine	$(0, 50)$	1	0	**−0.96**	Rice	$(0, 50)$	4	0	**−0.83**
	$(25, 150)$	18	0	**−0.99**		$(25, 150)$	$150 <$	0.18	**−0.99**
	$(50, 300)$	36	0	**−0.99**		$(50, 300)$	$300 <$	0.18	**−0.99**
Breast Cancer	$(0, 50)$	–	0	**−0.09**	Coffee	$(0, 50)$	2	0	**−1.00**
	$(25, 150)$	21	0	**−0.99**		$(25, 150)$	96	0.03	**−0.98**
	$(50, 300)$	44	0	**−0.99**		$(50, 300)$	167	0.04	**−0.98**
Spiral	$(0, 50)$	23	0	**−0.94**					
	$(25, 150)$	$150 <$	0.08	**−0.92**					
	$(50, 300)$	$300 <$	0.07	**−0.72**					

Table 2: Overview of the results. The Breaking-point refers to the observed bound b at which the accuracy begins to noticeably change. The value Min $H(\mathcal{Y}|\mathcal{X}_b^k)$ denotes the minimum entropy observed over the given bounds.

6 Conclusion

In this paper, we introduced the measure EIQM to determine suitable minimal integer bounds for GLVQ without sacrificing potential model performance. Our

ESANN 2026 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Bruges (Belgium) and online event, 22-24 April 2026, i6doc.com publ., ISBN 9782875870964. Available from http://www.i6doc.com/en/.

experiments reveal a strong link between increasing class-mixing error and decreasing expected accuracy for a bound $\mathcal{B}_b$. Further work should investigate the theoretical relationship between the expected accuracy of a GLVQ classifier and the conditional entropy of a dataset. Earlier work by Sarajalew already demonstrated the high robustness of GLVQ models during inference against adversarial attacks [9], where the author assumes the robustness is linked to GLVQ maximizing the hypothesis margin. We assume that this margin maximization also makes GLVQ robust to quantization-induced noise during training. Consequently, class-mixing error is assumed to be proportional to the deterioration of the hypothesis margin. We also showed the magnitude of the influence of outliers on the effective range of quantization levels. This underlines the importance of proper data preprocessing.

While GLVQ has a predictive behavior due to its simple dynamics, Generalized Matrix LVQ (GMLVQ) offers a more sophisticated learning dynamic. Here, a matrix Λ is learned which rotates and scales the input space [10], such that a better separation of the data with respect to its discriminative features is achieved. Therefore, further work can investigate the effects of the GMLVQ matrix mapping on the class-mixing error during integer quantization. Since the learning dynamics of GMLVQ are not designed for integers an appropriate design has to be presented to maintain effective learning. This investigation could reveal how information-theoretical properties behave in the learning dynamics of GMLVQ and if it reduces the magnitude of class-mixing error.

References

[1] Ronny Schubert, Thomas Davies, Mandy Lange-Geisler, Klaus Dohmen, and Thomas Villmann. Towards learning vector quantization in the setting of homomorphic encryption. In *ESANN 2025 proceedings*, ESANN 2025, pages 407–412. Ciaco - i6doc.com, 2025.

[2] Alexander Engelsberger and Thomas Villmann. Quantum Computing Approaches for Vector Quantization—Current Perspectives and Developments. 25(3):540, 2023.

[3] M. Gross, M. Lange, B. Bujnowski, and H.-M. Rieser. Expressivity vs. generalization in quantum kernel methods. In *Proc. 33rd European Symposium on Artificial Neural Networks (ESANN)*, pages 543–548. i6cod.com, 2025.

[4] Suyog Gupta, Ankur Agrawal, Kailash Gopalakrishnan, and Pritish Narayanan. Deep Learning with Limited Numerical Precision, 2025.

[5] Atsushi Sato and Keiji Yamada. Generalized Learning Vector Quantization. In *Advances in Neural Information Processing Systems*, volume 8. MIT Press, 1995.

[6] C. E. Shannon. A mathematical theory of communication. 27(3):379–423, 1948.

[7] James Dougherty, Ron Kohavi, and Mehran Sahami. Supervised and Unsupervised Discretization of Continuous Features. In Armand Prieditis and Stuart Russell, editors, *Machine Learning Proceedings 1995*, pages 194–202. Morgan Kaufmann, 1995.

[8] Markelle Kelly, Rachel Longjohn, and Kolby Nottingham. The uci machine learning repository.

[9] Sascha Saralajew, Lars Holdijk, Maike Rees, and Thomas Villmann. Robustness of Generalized Learning Vector Quantization Models against Adversarial Attacks. In A. Vellido, K. Gibert, C. Angulo, and J.D.M. Guerrero, editors, *Advances in Self-Organizing Maps, Learning Vector Quantization, Clustering and Data Visualization – Proceedings of the 13th International Workshop on Self-Organizing Maps and Learning Vector Quantization, Clustering and Data Visualization, WSOM+2019, Barcelona*, volume 976 of *Advances in Intelligent Systems and Computing*, pages 189–199. Springer Berlin-Heidelberg, 2020.

[10] P. Schneider, B. Hammer, and M. Biehl. Adaptive relevance matrices in learning vector quantization. *Neural Computation*, 21:3532–3561, 2009.