Hyperspectral image segmentation with a machine learning model trained using quantum annealer

 $\begin{array}{c} \mbox{Dawid Mazur}^{1[0009-0006-2420-6558]}, \mbox{ Tomasz Rybotycki}^{2,3,4[0000-0003-2493-0459]}, \\ \mbox{ and Piotr Gawron}^{3,4[0000-0001-7476-9160]}, \end{array}$

¹ Faculty of Electrical Engineering, Automatics, Computer Science and Biomedical Engineering, AGH University, al. Mickiewicza 30, 30-059 Cracow, Poland

² Systems Research Institute, Polish Academy of Sciences, ul. Newelska 6, 01-447 Warsaw, Poland

³ Nicolaus Copernicus Astronomical Center, Polish Academy of Sciences, ul. Bartycka 18, 00-716 Warsaw, Poland

⁴ Center of Excellence in Artificial Intelligence, AGH University, al. Mickiewicza 30, 30-059 Cracow, Poland

Abstract. Training of machine learning models consumes large amounts of energy. Since the energy consumption becomes a major problem in the development and implementation of artificial intelligence systems there exists a need to investigate the ways to reduce use of the resources by these systems. In this work we study how application of quantum annealers could lead to reduction of energy cost in training models aiming at pixel-level segmentation of hyperspectral images. Following the results of QBM4EO team, we propose a classical machine learning model, partially trained using quantum annealer, for hyperspectral image segmentation. We show that the model trained using quantum annealer is better or at least comparable with models trained using alternative algorithms, according to the preselected, common metrics. While direct energy use comparison does not make sense at the current stage of quantum computing technology development, we prove that quantum annealing should be considered as a tool for training at least some machine learning models.

Keywords: RBM, QML, Hyperspectral imaging, image segmentation

1 Introduction

The rapid growth of artificial intelligence, especially in the field of generative models [1] and transformer architecture in 2017 [2] has led to a major proliferation of large deep learning models. However, life-cycle of large machine learning (ML) models is known to be problematically energy consuming [3]. In order to mitigate this problem it is important to search for alternative methods of models training. In this work we use a quantum annealer to train a Restricted Boltzmann Machine (RBM). The latter is a generative model that has the ability to learn a probabilistic distribution over its set of inputs. RBM is a widely used

machine learning technique for both unsupervised [4] and supervised tasks [5]. Our motivation for such approach is two-fold. First, such training has already been shown to work [6], [7]. Secondly, quantum annealers have already been considered for reducing energy consumption for specific problems [8]. If the RBM training using quantum annealers proves successful, it would be worth investigating under what circumstances are quantum annealers more energy efficient than classic RBM training approaches.

Among others, RBMs have been used in computer vision, as a part of the image processing systems. For instance, the authors of QBM4EO devised a ML model for multi-label land-use classification [7]. Their model was designed to process hyperspectral data from the Sentinel-2 images dataset, hence we were able to ascertain the model capabilities to process similar data. However, instead of multi-label classification, we looked into a different computer vision task, namely image segmentation.

Segmentation is a core task in computer vision. Its goal is to partition an image into regions representing different objects or materials. Segmentation makes it possible to analyze image's structure accurately, and it has applications in many fields, such as medicine [9] or satellite image analysis [10]. In this work, we focus on the segmentation of multispectral images. Such images differ from ordinary ones in that they capture information in broader electromagnetic wavelength spectrum , not just in the visible range. Such data allow for a deeper analysis of the imaged object properties [11].

We investigated how accurately a model similar to the one proposed by [7] can handle pixel-level multispectral image segmentation task. Since the [7] model was trained using quantum annealers (QA), the problem is essentially a quantum machine learning (QML) one. QML is an intersection of quantum computing (QC) and machine learning (ML). It leverages the principles of quantum mechanics to solve complex ML problems. In our case, the problem is RBM training. Since RBMs were trained far before first QAs became available, another research question arise. Are quantum training techniques better than their classical counterparts, at least in this limited context? If quantum training techniques are at least comparable in terms of results quality, one can hope that they could find wider applications due to their ability to use less energy [8], [12].

This paper is organized as follows. We begin with the problem formulation, where we formally introduce the idea of multi- and hyperspectral images and image segmentation (section 2). Moreover, we discuss previously proposed solutions to the problem therein. In the next section (3), we overview quantum machine learning, focusing especially on the building blocks of the model we propose. Here, we also describe how quantum annealers can be used to train specific ML models. Then, in section 4, we start with describing the experiment. We review the dataset we use and introduce our model. We then discuss our experiments in great detail and conclude this section with results analysis. We finish this paper with conclusions, insights and directions for future research.

 $\mathbf{2}$

2 Problem formulation

Image segmentation is an important area of research and application in the field of computer vision. It is a process of dividing an image into homogeneous regions [13]. Formally, image segmentation can be defined as follows. If F is the set of all pixels and $P: F \to \{\text{true}, \text{false}\}$ is a homogeneity predicate defined on groups of connected pixels, then segmentation is a partitioning of the set F into a set of connected subsets or regions (S_1, S_2, \ldots, S_n) such that:

$$\bigcup_{i=1}^{n} S_{i} = F \quad \text{with} \quad S_{i} \cap S_{j} = \emptyset, i \neq j.$$
(1)

The homogeneity predicate $P(S_i) = \text{true for all regions } S_i \text{ and } P(S_i \cap S_j) = \text{false}$, when S_i is adjacent to S_j . Set F is a digital image defined as

$$F_{W \times H \times B} = [f(x, y, \lambda)]_{W \times H \times B}$$
⁽²⁾

where W and H are image dimensions, B is the band dimension and $f(x, y) \in G_L = \{0, 1, \dots, L-1\}$ is a set of discrete levels of the feature value and (x, y) denotes the spatial coordinate.

In the context of this work, B is of utmost importance, because this dimension is the one used to determine if an image is multi- or hyperspectral. Typically, images with 3–15 spectral bands are considered multispectral, whereas hyperspectral images can have hundreds of spectral bands [14]. One can clearly see how that makes analysis of such images more demanding in comparison to the standard ones.

Since image segmentation task is an old and well established problem there are variety of techniques to deal with it [13]. In the context of this work, the most interesting are the ones using machine learning techniques (ML), especially unsupervised learning methods. Among those cluster analysis techniques are popular [15]. These methods, however, are also known to have limitations [16]. It's therefore reasonable to explore different options.

The decision of which algorithm to select can be further guided by the approach one wishes to implement. Most of the state-of-the-art approaches rely on deep learning methods, which are costly to train [3]. The image can be analyzed using either pixel-level or patch approach. The latter usually involves using deep convolutional neural networks [17]. However, research on hyperspectral imaging at the pixel-level has grown in recent years, leading to an increasing number of scientific publications in this field [11], [17], [18]. We therefore decided to pursue the latter direction.

3 Machine Learning and Quantum Machine Learning

Artificial intelligence is a technology at the intersection of mathematics and computer science. It includes a wide spectrum of methods and algorithms that enable machines to learn, resulting in a wide range of applications [19], [20].

A subset of artificial intelligence that explores machines' capability to learn patterns from the data is called machine learning (ML).

The focus of this work is on unsupervised learning. It is a fundamental approach in machine learning that allows models to learn patterns in data without explicit labels [21]. The goal of learning in this context is to uncover the underlying structure of the dataset. Classic problems in which unsupervised learning is applied are clustering, dimensionality reduction and representation learning. Clustering can be performed using a variety of algorithms [22], including partitioning methods, density-based and hierarchical methods. Dimensionality reduction has traditionally been achieved using principal component analysis (PCA) [20]. Lately, autoencoders have also proven to be highly effective for this task [23]. Representation learning, also called feature learning, is a process during which algorithms extract meaningful patterns from data to create representations usually in a lower dimension than the original data [20].

In the context of this work, unsupervised learning will be used for all the tasks we mentioned. The model we propose, detailed in section 4, employs both dimensionality reduction and clustering. We also use standard clustering algorithms as a baseline for our model quality assessment.

Autoencoders are artificial neural networks used for unsupervised representation learning. Their architecture consists of two separate networks: encoder and decoder. Encoder transforms the data into its latent space representation whereas decoder tries to restore the input data from it. An autoencoder whose latent space dimension is lower than the input data one is called undercomplete [20]. Learning an undercomplete representation forces the autoencoder to capture the most significant features of the training data, thus it essentially implements a representation learning.

Latent Bernoulli Autoencoder (LBAE) — is a Variational Autoencoder [10], that in the context of our work, has a vital property such that that its latent space is binarized and therefore its output can be used as an input to a Restricted Boltzmann Machine.

Restricted Boltzmann Machines are undirected graphical, energy-based models that contain a layer of observable variables and a layer of latent variables. Typically they're referred to as visible and hidden layer, respectively [20]. A key feature of RBM is that there are no connections between neurons in the same layer, hence the "restricted" prefix. RBMs are used for tasks such as classification [24], feature extraction [25] and multispectral image processing [7]. In addition, RBMs are commonly used as building blocks in other architectures such as Deep Belief Networks [20], [25]. RBM training is based on maximizing the log-likelihood of the training data, and is typically done using gradient-based techniques. However, parts of the log-likelihood gradient function, such as socalled negative-phase [5] are hard to compute. Fortunately, their approximate values can be obtained using the Monte Carlo Markov Chain methods (MCMC) [20]. One family of such algorithms is the contrastive divergence (CD). While widely used, CD relies on limited Gibbs sampling steps and can get stuck in local minima.

Alternatively, the negative-phase can be calculated using samples drawn from an annealer [5]. The use of annealing can lead to more efficient sampling than MCMC algorithms. That's because the latter can struggle to correctly approximate the negative phase of the gradient [20]. Annealing also explores more global configurations, improving sampling accuracy and learning dynamics in challenging scenarios.

Simulated annealing (SA) is an optimization algorithm inspired by annealing process employed in metallurgy. SA treats the optimization problem as a physical system, with energy representing the objective function value and temperature controlling the probability of accepting inferior solutions. The algorithm starts at a high temperature, where energy-increasing movements are allowed, and then the temperature slowly decreases to find the global minimum of the objective function [26].

Quantum annealing (QA) is an extension of the idea of simulated annealing that uses quantum effects to search the solution space. Unlike SA, which relies on classical temperature perturbations, QA uses quantum fluctuations to overcome energy barriers. Optimization problems, solvable by QA, are most often represented in terms of an Ising model [27]. Although Quantum Annealing, in principle, follows similar scheme as Adiabatic Quantum Computing [28], the difference is that system evolution in QA is not necessarily adiabatic [29].

The use of quantum annealing for RBM training represents a novel approach [5], [30]. It starts with determining the probability distribution of the hidden RBM layer for a given set of input data. The RBM coefficients are then transformed into a Quadratic Unconstrained Binary Optimization (QUBO) problem. The QUBO is defined by a quadratic function, where the linear elements correspond to the biases, and the quadratic elements correspond to the weights between the neurons. The generated QUBO is then sampled by an annealer, which finds the variables' values corresponding to the objective function's minimization. In this context, the annealer acts as a probabilistic sampler, providing samples of the values of hidden and visible RBM units based on energy minimization. The values of these samples are then used to update the RBM weights and biases. The algorithm gradually adjusts the model by comparing the distributions sampled by the annealer with those inferred from the data, allowing the hidden structures to be correctly represented in the data. Using the annealer in this context can help more efficient sampling in cases where gradient methods become insufficient or when the problem has a complex energy space, making sampling more difficult. At the beginning of the annealing process, the system exhibits a significantly quantum behavior. As time passes, and the system cools down, we arrive at the systems final state, which corresponds to the low-energy solution of the assigned problem.

Cluster analysis is a category of unsupervised learning algorithms that seek to divide a given set of objects into homogeneous clusters [15]. In the context of image segmentation, cluster analysis divides an image into regions with similar properties. 5

The Agglomerative Hierarchical Clustering (AHC) algorithm relies on a bottomup approach. It starts the clustering process by treating each data point as an individual cluster and then iteratively merging the most similar clusters based on a selected distance metric. The process continues until all data points are combined into a single cluster or the desired number of clusters is reached.

The AHC algorithm requires specifying inter- and intra-cluster distance measures. The latter are commonly called linkage methods, whereas the first are usually standard distance measures used in analyzed objects processing. In this work we used Euclidean and Spectral Angle distances [31]. We also used complete and average linking criteria [15]. In the context of this work, an interesting property of the AHC algorithm is that it can be used to further cluster the preclustered data. In other words it can be used to continue initial clustering. This initial clustering could, in particular, use different algorithm.

The k-means algorithm is a clustering algorithm that partition a dataset into k clusters [22], aiming to minimize the differences within the clusters by assigning data points to the nearest clusters. A centroid is the center of a cluster, a point that represents the average value of all points assigned to a given cluster. The algorithm begins by initializing centroids, which represent the centers of the clusters. Next, the distance from each data point to each centroid is computed, and each point is assigned to the cluster with the nearest centroid. Once all points have been assigned to their respective clusters, the centroids are updated by computing the arithmetic mean of all points assigned to cluster. This update moves the centroids. The algorithm works iteratively, and the cycle of assigning points to clusters and updating centroids repeats until the centroids stop changing, which means that the algorithm converged.

4 Experiments and results

We used HyperBlood dataset [32] in our the experiments. It consists of fourteen hyperspectral images showing a mock-up crime scene with bloodstains and other visually similar substances. The images, collected over three weeks, vary in background composition and lighting conditions. Each image consists of 120 spectral bands, and is annotated with class labels indicating the presence of respective substances.

The dataset documentation [32] describes some of the hyperspectral channels as "noisy". We decided that the indicated spectral bands should be removed from further processing. Additionally, we excluded pixels whose position (x, y)corresponded to the label "0" (which is the background) in the ground truth image. Our goal was to improve homogeneity of the training data and thus quality of the model.

Further data transformations were to perform normalization, shuffle the data, and divide the data into training, validation, and test sets. Scaling to the interval $\langle 0, 1 \rangle$ was chosen as a standard normalization method in machine learning [19] but was also used for processing hyperspectral images [33]. Data shuffling was

6

used to reduce the influence of sequential pixel order. Data partitioning was done according to the Pareto principle: first, the data were divided into a training set and a test set in a ratio of 0.8 : 0.2, and then the training set was further split into a training and validation set, also in a ratio of 0.8 : 0.2 [34].

4.1 Proposed Model

The model we propose consists of two parts. The initial one is encoder of a LBAE. It processes the input and returns its binarized latent space representation. Binarization is vital for our model to work, because RBM, which constitutes the second layer of our model, accepts only binarized inputs. By setting respective neurons of the RMB visible layer to 1, we update probabilities of neuron activation in the hidden layer or the RBM. We compute and binarize these probabilities, thus obtaining a label for the inputted hyperspectral pixel. We present our model in the figure 1. Notice that such pipeline is basically a repurposed QBM4EO pipeline [7], which was shown to work well for similar problem.



Fig. 1: Proposed model pipeline. The pipeline processes spectral pixel data through the LBAE encoder, which consists of one-dimensional convolutional layers followed by a *tanh* activation and binarization. The resulting binary representation is forwarded into the visible layer of the RBM. Then we compute each hidden layer neuron's activation probability and binarize it. The resulting binary vector is considered the input pixel's label.

LBAE is an autoencoder which was introduced in [10] and implemented in [7]. It contains convolutional layers in both encoding and decoding parts. Our approach relies on the work done in the [7]. However, our approach explores the possibility of pixel-level analysis in the context of hyperspectral images, so we had to adjust [7] implementation of the LBAE. This adjustment consisted of changing the convolution layers from two- to one-dimensional. The rest of the network and its layers will stay the same as in [7] as this project obtained satisfactory performance. LBAE was trained using standard back-propagation algorithm. Our LBAE transforms the input pixel data into its binary representation. Due to the undercomplete architecture of this autoencoder, the dimension of input vector shrinks accordingly. We show the LBAE convolutional layers parameters in Table 1. Each latent space vector has 28 binary elements. After the encoding, the

Layer	Padding	Dilation	Kernel Size	Stride
Conv1	1	1	3	1
Conv2	1	1	4	2
Conv3	1	1	4	2
Conv4	1	1	3	1

Table 1: Parameters of the convolutional layers in the LBAE encoder.

data is processed by an RBM. The number of neurons in the visible layer of our RBM is determined by input data dimension and LBAE encoder convolutional layers. In the case of our experiments, that would be 28 neurons.

We trained the RBM with a fixed number of neurons in its visible layer and a changing number of neurons in its hidden layer $h_i \in \{3, 4, \ldots, 28\}$ using CD-1 algorithm. Our approach is to take RBM's hidden layer neurons activation probabilities and binarize them. The obtained binary vector is a label assigned to the input pixel. The binarization threshold was selected as follows: for each $th_i \in \{0.1, 0.2, \ldots, 0.9\}$, we computed the adjusted Rand score (ARS) between true labels and predicted labels. Then we checked for which threshold we obtained the highest ARS, and for this threshold, we compute other metrics. Using the V-measure, we then compare the models. We selected the $\beta \in [0, 1]$ parameter of the V-measure such that the metric promotes homogeneity — a metric we deemed more important in the image segmentation task. For each *i*-th model, we manually tuned $\beta_i \in \{0, 0.01, \ldots, 1\}$, to find such β that yield the best V-measure scores was selected as the best one. Thus concluding the model architecture design.

Our approach based on binarized values of neuron activation probability in RBM's hidden layer leads to a maximum number of unique returned labels equal to 2^N , where $N \in \{3, 4, \ldots, 28\}$. Since there are seven classes in the HyperBlood dataset after our preprocessing, we see that the RBM will return more labels than in the dataset. Accordingly, to [4], RBM may be used to build relational trees and then use these hierarchies to divide the data into groups and subgroups. We will use the structure clustered by RBM to create a distance matrix, which will serve as an input to the Agglomerative Hierarchical Clustering (AHC) algorithm. By specifying the target number of clusters for the AHC, we aim to obtain clusterization that will be useful in our segmentation task. This final phase of the segmentation takes place only for the best model; after training algorithms are compared.

4.2 Classical Model Training and Evaluation

The selection of appropriate evaluation metrics is crucial for an objective analysis of the quality of the machine learning models. We decided to use the Homogeneity Score, the Completeness Score, the Rand Score, and the Adjusted Rand Score, which are commonly used in evaluating clustering tasks [35], [36]. Additionally, Euclidean distance and Spectral Angle Distance were employed to evaluate the capability of the autoencoder to reconstruct the data. We trained the model using three different approaches: the traditional contrastive divergence (CD-1) algorithm, an approach based on simulated annealing (SA) and quantum annealing (QA).

At first, we begin with LBAE training. Since, compared to [7], we only changed convolutional layers from two to one-dimensional, we kept the hyperparameters values. We investigated the influence of changing batch size and learning rate.

The following values of batch size $b_i \in \{4, 8, 16\}$ and learning rate $\eta_i \in \{10^{-2}, 10^{-3}, 10^{-4}\}$ were chosen as a standard values [19]. For each combination of these parameters, LBAE training was conducted, ultimately yielding nine trained models. Each model's performance was evaluated on a test dataset using Euclidean distance and spectral angle distance (SAD). We used LBAE model that obtained the lowest values of both metrics, that is the model trained with hyperparameters $b_i = 4$ and $\eta_i = 10^{-3}$.

Having trained the LBAE model, we determined a baseline for Restricted Boltzmann Machines (RBM) clusterization that we aim to surpass. For this purpose, we used the k-means algorithm to perform clusterization on the test dataset, on both raw data and its latent space representation. To avoid the impact of how the k-means centroids are initialized, we conducted the clustering ten times using different random seeds. Results of the baseline clustering metrics on the test dataset are included in Table 2.

Table 2: Comparison of clustering evaluation metrics for k-means and LBAE+k-means. Values represent the mean \pm standard deviation.

Metric	k-means	LBAE+k-means
Homogeneity	0.509 ± 0.020	0.596 ± 0.056
Completeness	0.443 ± 0.027	0.520 ± 0.046
ARS	0.362 ± 0.065	0.395 ± 0.090
Rand Score	0.779 ± 0.021	0.789 ± 0.031

We then proceeded with the RBM training using CD-1. Similar to the case of LBAE, we keep hyperparameters as they were in [7], except for one of them — a number of neurons in the hidden layer. RBM training experiments were conducted for the following number of neurons in hidden layer $h_i \in \{3, 4, \ldots, 28\}$, and those experiments were repeated ten times for different weights initialization. This experiment concluded that the RBM model with 23 neurons in the hidden layer was the most promising for the segmentation task. Figure 2a shows the learning curve for that model. We noticed that our model is returning more unique labels than the target number of labels on ground-truth images. Following the idea that the RBM returned structure is hierarchical [4], we could pass this structure to another algorithm, known as Agglomerative Hierarchical Clustering (AHC), and specify the target number of clusters.

We want to compare the final segmentation with a reliable baseline. We, again, used the standard k-means algorithm to obtain it. Figure 3a illustrates the pixels clustering using k-means, and Figure 3b illustrates the ground truth image. Table 3 shows metrics comparison for created segmentation images. Then, we created a segmentation images using our model. The results are presented in Figure 4a and Figure 4d.

4.3 Quantum Model Training and Evaluation

The next step of the project was to test the implemented training algorithm using the annealer-based algorithms. First, we used D-Wave's implementation of a simulated annealing sampler [37]. The training was repeated ten times for initialization with different weights, this time only for an RBM model containing 23 neurons in the hidden layer. We decided to save a model after every other hundred training epochs to execute an insightful analysis of model performance in the context of clusterization metrics. The best model, according to the Vmeasure, was the one after 200 training epochs. The learning curve of the RBM trained using the simulated annealer is shown in Figure 2b.

Next, we trained the model using quantum annealing. Again, we used the sampler provided by D-Wave [37]. We also used the automatic embedding of our problem into the target QPU — D-Wave Advantage 5.4 system. The training was repeated ten times for initialization with different weights, again only for an RBM model containing 23 neurons in the hidden layer. We decided to save a model after every other hundred training epochs. The best model, according to V-measure, was obtained after 300 training epochs. The learning curve of the RBM trained using the quantum annealer is shown in Figure 2c.



Fig. 2: Learning curve of the RBM with respective algorithms. The blue line represents the loss value on the training dataset, and the red line represents the loss value on the validation dataset.

4.4 Segmentation Results

We present the comparison our experiments results with our baseline segmentation obtained by using k-means algorithm on the Figure 3a and ground truth image on the Figure 3b. They are followed by the segmentations obtained using our model trained with contrastive divergence (CD-1) finalized with AHC using both linkage methods — complete and average, those will be respectively Figure 4a and Figure 4d. Similarly, we present segmentations for models trained with simulated annealing (SA) and quantum annealing (QA). For simulated annealing obtained images are Figure 4b and Figure 4e, and for quantum annealing obtained images are Figure 4f.



(a) k-means segmentation.

(b) Ground truth.

Fig. 3: Baseline segmentation by k-means and ground truth.

A summary of the computed evaluation metrics such as homogeneity, completeness, adjusted Rand score (ARS) and Rand score (RS) for each segmentation obtained is presented collectively in Table 3.

Analyzing obtained metrics on the segmentation images, we notice that almost all metrics improve their value for RBM trained with QA, and if they are not improved, they are close to the best-obtained value. For a model that does not use AHC, we note the consistent improvement of the metrics while changing the RBM training algorithm from CD-1 through SA to QA. For the metrics homogeneity and ARS, the RBM model trained with QA surpassed the baseline value, which was set by the k-means algorithm. The use of the AHC algorithm in combination with RBM affects clustering characteristics. First, AHC improves the completeness metric, especially when using the RBM version trained with Quantum Annealing. At the same time, adding AHC reduces the clusters' homogeneity.



Fig. 4: Comparison of segmentation results. Subfigure captions contain RBM training and clustering algorithms used to obtain the segmentation.

5 Conclusions

We proposed a hybrid neural network architecture consisting of Latent Bernoulli Autoencoder encoder connected with Restricted Boltzmann Machines for hyperspectral image segmentation. Firstly, the LBAE's encoder part is used for dimensionality reduction and spectral data binary representation encoding, necessary for RBM processing. Secondly, we use RBM to perform clustering on the encoded spectral pixels.

The proposed architecture and experiments confirmed the effectiveness of applying quantum annealing techniques and RBM models training. One can clearly see that since we achieved better segmentation results, in terms of observed metrics, than the baseline segmentation obtained using standard algorithm — k-means. The results suggest the potential for further research on using quantum annealing in the image processing field. It's also vital to analyze cost and power feasibility of quantum annealers as RBM training tools.

However, other quantum data processing techniques should also be explored in this field. Particularly interesting are quantum kernel methods, performance of which is yet to be tested on real gate-based quantum devices.

Acknowledgments

We gratefully acknowledge Poland's high-performance Infrastructure PLGrid ACK Cyfronet AGH for providing computer facilities and support within computational grant no. PLG/2024/017155 through and no. PLG/2024/017550.

Metric	Training	k-means	RBM	RBM+AHC-C	RBM+AHC-A
Homogeneity	_	0.368	_	_	_
	CD-1	_	0.492	0.232	0.243
	SA	_	0.510	0.231	0.171
	QA	_	0.505	0.254	0.260
Completeness –		0.354	_	_	_
	CD-1	_	0.239	0.257	0.416
	SA	_	0.244	0.266	0.381
	QA	_	0.282	0.319	0.445
ARS	_	0.244	_	_	_
	CD-1	_	0.163	0.129	0.242
	SA	_	0.180	0.199	0.181
	QA	_	0.308	0.284	0.246
Rand Score	_	0.769	_	_	_
	CD-1	_	0.793	0.690	0.660
	SA	_	0.799	0.711	0.590
	QA	_	0.790	0.729	0.656

Table 3: Comparison of clustering evaluation metrics for baseline segmentation and segmentation produced by our models. RBM+AHC-C and RBM+AHC-A, denote the RBM with AHC using complete and avarage linkage respectively.

The authors would like to thank Etos sp. z o.o. and the QBM4EO team for providing the QBM4EO code and support. We gratefully acknowledge the funding support by program "Excellence Initiative — Research University" for the AGH University of Kraków as well as the ARTIQ project ARTIQ/0004/2021.

References

- I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, et al., Generative adversarial networks, 2014.
- [2] A. Vaswani, N. Shazeer, N. Parmar, et al., Attention is all you need, 2023.
- R. Verdecchia, J. Sallou, and L. Cruz, "A systematic review of Green AI," WIREs Data Mining and Knowledge Discovery, vol. 13, no. 4, e1507, 2023.
 DOI: https://doi.org/10.1002/widm.1507. eprint: https://wires. onlinelibrary.wiley.com/doi/pdf/10.1002/widm.1507.
- [4] A. Decelle, B. Seoane, and L. Rosset, "Unsupervised hierarchical clustering using the learning dynamics of restricted Boltzmann machines," *Physical Review E*, vol. 108, 2023.
- [5] V. Dixit, R. Selvarajan, M. A. Alam, T. S. Humble, and S. Kais, "Training restricted Boltzmann machines with a D-Wave quantum annealer," *Frontiers in Physics*, vol. 9, 2021.

- [6] M. H. Amin, E. Andriyash, J. Rolfe, B. Kulchytskyy, and R. Melko, "Quantum boltzmann machine," *Phys. Rev. X*, vol. 8, p. 021050, 2 2018. DOI: 10.1103/PhysRevX.8.021050. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevX.8.021050.
- [7] Ł. Pawela and K. Jałowiecki, QBM4EO Supervised quantum machine learning system for Earth land cover understanding. 2022.
- [8] S. Kasi, P. A. Warburton, J. Kaewell, and K. Jamieson, A cost and power feasibility analysis of quantum annealing for NextG cellular wireless networks, 2022. arXiv: 2109.01465 [cs.NI]. [Online]. Available: https:// arxiv.org/abs/2109.01465.
- [9] P. Wang, S. Wang, Y. Zhang, and X. Duan, "Multispectral image under tissue classification algorithm in screening of cervical cancer," *Journal of Healthcare Engineering*, 2022.
- [10] J. Fajtl, V. Argyriou, D. Monekosso, and P. Remagnino, "Latent Bernoulli autoencoder," in *Proceedings of the 37th International Conference on Machine Learning*, PMLR, 2020.
- [11] S. Sifnaios, G. Arvanitakis, F. K. Konstantinidis, G. Tsimiklis, A. Amditis, and P. Frangos, A deep learning approach for pixel-level material classification via hyperspectral imaging, 2024.
- [12] D. Jaschke and S. Montangero, "Is quantum computing green? an estimate for an energy-efficiency quantum advantage," en, *Quantum Science and Technology*, vol. 8, no. 2, p. 025 001, Jan. 2023, ISSN: 2058-9565.
- [13] N. R. Pal and S. K. Pal, "A review on image segmentation techniques," *Pattern Recognition*, vol. 26, no. 9, 1993.
- [14] N. A. Hagen and M. W. Kudenov, "Review of snapshot spectral imaging technologies," *Optical Engineering*, 2013.
- [15] N. Bratchell, "Cluster analysis," Chemometrics and Intelligent Laboratory Systems, vol. 6, no. 2, 1989.
- [16] R. Garcia-Dias, S. Vieira, W. H. Lopez Pinaya, and A. Mechelli, "Chapter 13 — clustering analysis," in *Machine Learning*, A. Mechelli and S. Vieira, Eds., Academic Press, 2020, pp. 227–247, ISBN: 978-0-12-815739-8.
- [17] S. Morchhale, V. P. Pauca, R. J. Plemmons, and T. C. Torgersen, "Classification of pixel-level fused hyperspectral and lidar data using deep convolutional neural networks," in 2016 8th Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (WHISPERS), 2016.
- [18] X. Chen, "Hyperspectral image classification using pixel-level convolutional neural network based on attention mechanism," in *Proceedings of* the 5th International Conference on Computer Information and Big Data Applications, ser. CIBDA '24, Association for Computing Machinery, 2024, 819–824.
- [19] S. Raschka and V. Mirjalili, Python Machine Learning: Machine Learning and Deep Learning with Python, Scikit-learn, and TensorFlow 2, 3rd. Packt Publishing, 2019.
- [20] I. Goodfellow, Y. Bengio, and A. Courville, *Deep learning*. MIT Press, 2016.

DOI: 10.1007/978-3-031-97570-7_16

14

15

- [21] S. Naeem, A. Ali, S. Anam, and M. Ahmed, "An unsupervised machine learning algorithms: Comprehensive review," *IJCDS Journal*, 2023.
- [22] K. Bindra and A. Mishra, "A detailed study of clustering algorithms," in 2017 6th International Conference on Reliability, Infocom Technologies and Optimization (Trends and Future Directions) (ICRITO), 2017.
- [23] W. Wang, Y. Huang, Y. Wang, and L. Wang, Generalized autoencoder: A neural network framework for dimensionality reduction, 2014.
- [24] J. A. Koziol, E. M. Tan, L. Dai, P. Ren, and J.-Y. Zhang, "Restricted boltzmann machines for classification of hepatocellular carcinoma," *Computational Biology Journal*, 2014.
- [25] J. Li, X. Li, D. He, and Y. Qu, "A novel method for early gear pitting fault diagnosis using stacked sae and gbrbm," *Sensors*, 2019.
- [26] R. Benedetti, M. M. Dickson, G. Espa, F. Pantalone, and F. Piersimoni, "A simulated annealing-based algorithm for selecting balanced samples," *Computational Statistics*, 2021.
- [27] T. Kadowaki and H. Nishimori, "Quantum annealing in the transverse Ising model," *Physical review. E, Statistical physics, plasmas, fluids, and related interdisciplinary topics*, vol. 58, 1998.
- [28] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, Quantum computation by adiabatic evolution, 2000.
- [29] W. Vinci and D. A. Lidar, "Non-stoquastic hamiltonians in quantum annealing via geometric phases," *npj Quantum Information*, vol. 3, no. 1, 2017.
- [30] A. Datta, S. T. Flammia, and C. M. Caves, "Entanglement and the Power of One Qubit," *Physical Review A*, 2005.
- [31] B. B. Sinaice, N. Owada, H. Ikeda, et al., "Spectral angle mapping and ai methods applied in automatic identification of placer deposit magnetite using multispectral camera mounted on uav," *Minerals*, vol. 12, no. 2, 2022.
- [32] M. Romaszewski, P. Głomb, A. Sochan, and M. Cholewa, "A dataset for evaluating blood detection in hyperspectral images," *Forensic Science International*, vol. 320, 2021.
- [33] P. Głomb, P. Sadowski, O. Faucoz, *et al.*, "Hyper-spectral image classification using adiabatic quantum computation," 2023.
- [34] M. Murugan, M. Mani, and G. M. Arasakumar, "Application of pareto's principle on deep learning research output: A scientometric analysis," 2021.
- [35] S. W. Emmons, S. G. Kobourov, M. Gallant, and K. Börner, Analysis of network clustering algorithms and cluster quality metrics at scale, 2016.
- [36] A. Rosenberg and J. Hirschberg, "V-measure: A conditional entropy-based external cluster evaluation measure," in *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, Association for Computational Linguistics, 2007.
- [37] D-Wave documentation, Accessed: 6th October 2024, D-Wave Systems Inc., 2024. [Online]. Available: https://docs.dwavesys.com/.