

Online Deterministic Annealing for Classification and Clustering

Christos N. Mavridis[✉], *Member, IEEE*, and John S. Baras[✉], *Life Fellow, IEEE*

Abstract—Inherent in virtually every iterative machine learning algorithm is the problem of hyperparameter tuning, which includes three major design parameters: 1) the complexity of the model, e.g., the number of neurons in a neural network; 2) the initial conditions, which heavily affect the behavior of the algorithm; and 3) the dissimilarity measure used to quantify its performance. We introduce an online prototype-based learning algorithm that can be viewed as a progressively growing competitive-learning neural network architecture for classification and clustering. The learning rule of the proposed approach is formulated as an online gradient-free stochastic approximation algorithm that solves a sequence of appropriately defined optimization problems, simulating an annealing process. The annealing nature of the algorithm contributes to avoiding poor local minima, offers robustness with respect to the initial conditions, and provides a means to progressively increase the complexity of the learning model, through an intuitive bifurcation phenomenon. The proposed approach is interpretable, requires minimal hyperparameter tuning, and allows online control over the performance-complexity tradeoff. Finally, we show that Bregman divergences appear naturally as a family of dissimilarity measures that play a central role in both the performance and the computational complexity of the learning algorithm.

Index Terms—Annealing optimization, Bregman divergences, classification, clustering, machine learning algorithms, progressive learning.

I. INTRODUCTION

LEARNING from data samples has become an important component of artificial intelligence. While virtually all learning problems can be formulated as constrained stochastic optimization problems, the optimization methods can be intractable, typically dealing with mixed constraints and very large, or even infinite-dimensional spaces [1]. For this reason, feature extraction, model selection and design, and analysis of optimization methods, have been the cornerstone of machine learning algorithms from their genesis until today.

Deep learning methods, currently dominating the field of machine learning due to their performance in multiple

applications, attempt to learn feature representations from data, using biologically inspired models in artificial neural networks [2], [3]. However, they typically use overly complex models of a great many parameters, which comes in the expense of time, energy, data, memory, and computational resources [4], [5]. Moreover, they are, by design, hard to interpret and vulnerable to small perturbations and adversarial attacks [6], [7]. The latter has led to an emerging hesitation in their implementation outside common benchmark datasets [8], and, especially, in security critical applications. On the other hand, it is understood that the tradeoff between model complexity and performance is closely related to overfitting, generalization, and robustness to input noise and attacks [9]. In this work, we introduce a learning model that progressively adjusts its complexity, offering online control over this tradeoff. The need for such approaches is reinforced by recent studies revealing that existing flaws in the current benchmark datasets may have inflated the need for overly complex models [10] and that overfitting to adversarial training examples may actually hurt generalization [11].

We focus on prototype-based models, mainly represented by vector quantization methods [12]–[14]. In vector quantization, originally introduced as a signal processing method for compression, a set of codevectors (or prototypes) $M := \{\mu_i\}$, is used to represent the data space in an optimal way according to an average distortion measure

$$\min_M J(M) := \mathbb{E} \left[\min_i d(X, \mu_i) \right]$$

where the proximity measure d defines the similarity between the random input X and a codevector μ_i . The codevectors can be viewed as a set of neurons, the weights of which live in the data space itself and constitute the model parameters. In this regard, vector quantization algorithms can be viewed as competitive-learning neural network architectures with a number of appealing properties: they are consistent, data-driven, interpretable, robust, topology-preserving [15], sparse in the sense of memory complexity, and fast to train and evaluate. In addition, they have recently shown impressive robustness against adversarial attacks, suggesting suitability in security-critical applications [16], while their representation of the input in terms of memorized exemplars is an intuitive approach, which parallels similar concepts from cognitive psychology and neuroscience. As iterative learning algorithms, however, their behavior heavily depends on three major design parameters: 1) the number of neurons/prototypes, which, defines the

Manuscript received 7 June 2021; revised 21 September 2021; accepted 18 December 2021. Date of publication 7 January 2022; date of current version 6 October 2023. This work was supported in part by the Defense Advanced Research Projects Agency under Agreement HR00111990027, in part by the Office of Naval Research under Grant N00014-17-1-2622, and in part by a grant from Northrop Grumman Corporation. (Corresponding author: Christos N. Mavridis.)

The authors are with the Department of Electrical and Computer Engineering and the Institute for Systems Research, University of Maryland, College Park, MD 20742 USA (e-mail: mavridis@umd.edu; baras@umd.edu).

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TNNLS.2021.3138676>.

Digital Object Identifier 10.1109/TNNLS.2021.3138676

2162-237X © 2022 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.

See <https://www.ieee.org/publications/rights/index.html> for more information.

complexity of the model; 2) the initial conditions, that affect the transient and steady-state behavior of the algorithm; and 3) the proximity measure d used to quantify the similarity between two vectors in the data space.

Inspired by the deterministic annealing (DA) approach [17], we propose a learning approach that resembles an annealing process, tending to avoid poor local minima, offering robustness with respect to the initial conditions, and providing a means to progressively increase the complexity of the learning model, allowing online control over the performance-complexity tradeoff. We relax the original problem to a soft-clustering problem, introducing the association probabilities $p(\mu_i|X)$ and replacing the cost function J by $D(M) := \mathbb{E}[\sum_i p(\mu_i|X)d(X, \mu_i)]$. This probabilistic framework (to be formally defined in Section III) allows us to define the Shannon entropy $H(M)$ that characterizes the “purity” of the clusters induced by the codevectors. We then replace the original problem by a *sequence of optimization problems*

$$\min_M F_T(M) := D(M) - TH(M)$$

parameterized by a temperature coefficient T , which acts as a Lagrange multiplier controlling the tradeoff between minimizing the distortion D and maximizing the entropy H . By successively solving the optimization problems $\min_M F_T(M)$ for decreasing values of T , the model undergoes a series of phase transitions that resemble an annealing process. Because of the nature of the entropy term, in high temperatures T , the effect of the initial conditions is greatly mitigated, while, as T decreases, the optimal codevectors of the last optimization problem are used as initial conditions to the next, which helps in avoiding poor local minima. Furthermore, as T decreases, the cardinality of the set of codevectors M increases, according to an intuitive bifurcation phenomenon.

Adopting the above-mentioned optimization framework, we introduce an online training rule based on *stochastic approximation* [18]. While stochastic approximation offers an online, adaptive, and computationally inexpensive optimization algorithm, it is also strongly connected to dynamical systems. This enables the study of the convergence of the learning algorithm through mathematical tools from dynamical systems and controls [18]. We take advantage of this property to prove the convergence of the proposed learning algorithm as a consistent density estimator (unsupervised learning), and a Bayes risk consistent classification rule (supervised learning). Finally, we show that the proposed stochastic approximation learning algorithm introduces *inherent regularization mechanisms* and is also *gradient-free*, provided that the proximity measure d belongs to the family of Bregman divergences. Bregman divergences are information-theoretic dissimilarity measures that have been shown to play an important role in learning applications [19], [20], including measures such as the widely used Euclidean distance and the Kullback–Leibler divergence. We believe that these results can potentially lead to new developments in learning with progressively growing models, including, but not limited to, communication, control, and reinforcement learning applications [21]–[23].

II. PROTOTYPE-BASED LEARNING

In this section, the mathematics and notation of prototype-based machine learning algorithms, which will be used as a base for our analysis, are briefly introduced. For more details see [14], [19], [20], [24].

A. Vector Quantization for Clustering

Unsupervised analysis can provide valuable insights into the nature of the dataset at hand, and it plays an important role in the context of visualization. Central to unsupervised learning is the representation of data in a vector space by typical representatives, which is formally defined in the following optimization problem.

Problem 1: Let $X : \Omega \rightarrow S \subseteq \mathbb{R}^d$ be a random variable defined in a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, and $d : S \times \text{ri}(S) \rightarrow [0, \infty)$ be a divergence measure, where $\text{ri}(S)$ represents the relative interior of S . Let $V := \{S_h\}_{h=1}^K$ be a partition of S and $M := \{\mu_h\}_{h=1}^K$ a set of codevectors, such that $\mu_h \in \text{ri}(S_h)$, for all $h = 1, \dots, K$. A quantizer $Q : S \rightarrow M$ is defined as the random variable $Q(X) = \sum_{h=1}^K \mu_h \mathbb{1}_{[X \in S_h]}$ and the vector quantization problem is formulated as

$$\min_{M, V} J(Q) := \mathbb{E}[d(X, Q)].$$

Vector quantization is a hard-clustering algorithm, and, as such, assumes that the quantizer Q assigns an input vector X to a unique codevector $\mu_h \in M$ with probability one. As a result, Problem 1 becomes equivalent to

$$\min_{\{\mu_h\}_{h=1}^K} \sum_{h=1}^K \mathbb{E}[d(X, \mu_h) \mathbb{1}_{[X \in S_h]}] \quad (1)$$

for V being a Voronoi partition, i.e., for

$$S_h = \left\{ x \in S : h = \arg \min_{\tau=1, \dots, K} d(x, \mu_\tau) \right\}, \quad h = 1, \dots, K.$$

It is typically the case that the actual distribution of $X \in S$ is unknown, and a set of independent realizations $\{X_i\}_{i=1}^n := \{X(\omega_i)\}_{i=1}^n$, for $\omega_i \in \Omega$, are available. In case the observations $\{X_i\}_{i=1}^n$ are available *a priori*, the solution of the VQ problem is traditionally approached with variants of the LBG algorithm [25], a generalization of the Lloyd algorithm [26], which includes the widely used k -means algorithm [27].

When the training data are not available *a priori* but are being observed online, or when the processing of the entire dataset in every optimization iteration is computationally infeasible, a stochastic vector quantization (sVQ) algorithm can be defined as a recursive asynchronous stochastic approximation algorithm based on gradient descent [14]:

Definition 1 Stochastic Vector Quantization Algorithm: Repeat

$$\begin{cases} \mu_h^{t+1} = \mu_h^t - \alpha(v(h, t)) \mathbb{1}_{[X_{t+1} \in S_h^{t+1}]} \nabla_{\mu_h} d(X_{t+1}, \mu_h^t) \\ S_h^{t+1} = \left\{ X \in S : h = \arg \min_{\tau=1, \dots, k} d(X, \mu_\tau^t) \right\}, \quad h \in K \end{cases}$$

for $t \geq 0$ until convergence, where μ_h^0 is given during initialization and $v(h, t)$ represents the number of times the component μ_h has been updated up until time t .

B. Learning Vector Quantization for Classification

The supervised counterpart of vector quantization is the particularly attractive and intuitive approach of the competitive-learning learning vector quantization (LVQ) algorithm, initially proposed by Kohonen [12]. LVQ for binary classification is formulated in the following optimization problem (and generalized to any type of classification task, see [28]).

Problem 2: Let the pair of random variables $\{X, c\} \in S \times \{0, 1\}$ defined in a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, with c representing the class of X and $S \subseteq \mathbb{R}^d$. Let $M := \{\mu_h\}_{h=1}^K$, where $\mu_h \in \text{ri}(S_h)$ represent codevectors and define the set $C_\mu := \{c_{\mu_h}\}_{h=1}^K$, such that $c_{\mu_h} \in \{0, 1\}$ represents the class of μ_h for all $h \in \{1, \dots, K\}$. The quantizer $Q^c : S \rightarrow \{0, 1\}$ is defined such that $Q^c(X) = \sum_{h=1}^K c_{\mu_h} \mathbb{1}_{[X \in S_h]}$. Then, the minimum-error classification problem is formulated as

$$\min_{\{c_{\mu_h}, S_h\}_{h=1}^K} J_B(Q^c) := \pi_1 \sum_{H_0} \mathbb{P}_1[X \in S_h] + \pi_0 \sum_{H_1} \mathbb{P}_0[X \in S_h]$$

where $\pi_i := \mathbb{P}[c = i]$, $\mathbb{P}_i\{\cdot\} := \mathbb{P}\{\cdot | c = i\}$, and H_i is defined as $H_i := \{h \in \{1, \dots, K\} : Q^c = i\}$, $i \in \{0, 1\}$.

LVQ algorithms that solve Problem 2 are similar in structure with the sVQ algorithm of Definition 1 and make use of a modified distortion measure, which in the case of the original LVQ1 algorithm [12] takes the form

$$d^l(x, c_x, \mu, c_\mu) = \begin{cases} d(x, \mu), & c_x = c_\mu \\ -d(x, \mu), & c_x \neq c_\mu. \end{cases}$$

Generalizations of this definition based on similar principles have also been proposed [29], [30].

C. Bregman Divergences as Dissimilarity Measures

Prototype-based algorithms rely on measuring the proximity between different vector representations. In most cases, the Euclidean distance or another convex metric is used, but this can be generalized to alternative dissimilarity measures inspired by information theory and statistical analysis, such as the Bregman divergences.

Definition 2 (Bregman Divergence): Let $\phi : H \rightarrow \mathbb{R}$, be a strictly convex function defined on a vector space H such that ϕ is twice F-differentiable on H . The Bregman divergence $d_\phi : H \times H \rightarrow [0, \infty)$ is defined as

$$d_\phi(x, \mu) = \phi(x) - \phi(\mu) - \frac{\partial \phi}{\partial \mu}(\mu)(x - \mu)$$

where $x, \mu \in H$, and the continuous linear map $(\partial \phi / \partial \mu)(\mu) : H \rightarrow \mathbb{R}$ is the Fréchet derivative of ϕ at μ .

Notice that, as a divergence measure, the Bregman divergence can be used to measure the dissimilarity of one probability distribution to another on a statistical manifold and is a notion weaker than that of the distance. In particular, it does not need to be symmetric or satisfy the triangle inequality. In this work, we will concentrate on nonempty, compact convex sets $S \subseteq \mathbb{R}^d$ so that the derivative of d_ϕ with respect to the second argument can be written as

$$\begin{aligned} \frac{\partial d_\phi}{\partial \mu}(x, \mu) &= \frac{\partial \phi(x)}{\partial \mu} - \frac{\partial \phi(\mu)}{\partial \mu} - \frac{\partial^2 \phi(\mu)}{\partial \mu^2}(x - \mu) + \frac{\partial \phi(\mu)}{\partial \mu} \\ &= -\frac{\partial^2 \phi(\mu)}{\partial \mu^2}(x - \mu) = -\langle \nabla^2 \phi(\mu), (x - \mu) \rangle \end{aligned}$$

where $x, \mu \in S$, $(\partial / \partial \mu)$ represents differentiation with respect to the second argument of d_ϕ , and $\nabla^2 \phi(\mu)$ represents the Hessian matrix of ϕ at μ .

Example 1: As a first example, $\phi(x) = \langle x, x \rangle$, $x \in \mathbb{R}^d$, gives the squared Euclidean distance

$$d_\phi(x, \mu) = \|x - \mu\|^2$$

for which $(\partial d_\phi / \partial \mu)(x, \mu) = -2(x - \mu)$.

Example 2: A second interesting Bregman divergence that shows the connection to information theory is the generalized I-divergence which results from $\phi(x) = \langle x, \log x \rangle$, $x \in \mathbb{R}_{++}^d$ such that

$$d_\phi(x, y) = \langle x, \log x - \log y \rangle - \langle \mathbb{1}, x - y \rangle$$

for which $(d_\phi) / (\mu)(x, \mu) = -\text{diag}^{-1}(\mu)(x - \mu)$, where $\mathbb{1} \in \mathbb{R}^d$ is the vector of ones and $\text{diag}^{-1}(\mu) \in \mathbb{R}_{++}^{d \times d}$ is the diagonal matrix with diagonal elements the inverse elements of μ . It is easy to see that $\phi(x)$ reduces to the Kullback–Leibler divergence if $\langle \mathbb{1}, x \rangle = 1$.

The family of Bregman divergences provides proximity measures that have been shown to enhance the performance of a learning algorithm [31]. In addition, the following theorem shows that the use of Bregman divergences is both necessary and sufficient for the optimizer μ_h of (1) to be analytically computed as the expected value of the data inside S_h , which is implicitly used by many “centroid” algorithms, such as k -means [27].

Theorem 1: Let $X : \Omega \rightarrow S$ be a random variable defined in the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ such that $\mathbb{E}[X] \in \text{ri}(S)$, and let a distortion measure $d : S \times \text{ri}(S) \rightarrow [0, \infty)$, where $\text{ri}(S)$ denotes the relative interior of S . Then, $\mu := \mathbb{E}[X]$ is the unique minimizer of $\mathbb{E}[d(X, s)]$ in $\text{ri}(S)$, if and only if d is a Bregman divergence for any function ϕ that satisfies the definition.

Proof: For necessity, identical arguments, as in [19, Appendix B], are followed. For sufficiency

$$\begin{aligned} &\mathbb{E}[d_\phi(X, s)] - \mathbb{E}[d_\phi(X, \mu)] \\ &= \phi(\mu) + \frac{\partial \phi}{\partial \mu}(\mu)(\mathbb{E}[X] - \mu) - \phi(s) - \frac{\partial \phi}{\partial s}(s)(\mathbb{E}[X] - s) \\ &= \phi(\mu) - \phi(s) - \frac{\partial \phi}{\partial s}(s)(\mu - s) = d_\phi(\mu, s) \geq 0 \quad \forall s \in S \end{aligned}$$

with equality holding only when $s = \mu$ by the strict convexity of ϕ , which completes the proof. \square

In Section III, we will show a similar result for the proposed algorithm that uses a soft-partition approach.

III. ONLINE DETERMINISTIC ANNEALING FOR UNSUPERVISED AND SUPERVISED LEARNING

Online vector quantization algorithms are proven to converge to locally optimal configurations [14]. However, as iterative machine learning algorithms, their convergence properties and final configuration depend heavily on two design parameters: the number of neurons/clusters K , and their initial configuration. Inspired by the DA framework [17], we relax the original optimization problem (1) to a soft-clustering problem and replace it by a sequence of deterministic optimization

problems, parameterized by a temperature coefficient, that are progressively solved at successively reducing temperature levels. As will be shown, the annealing nature of this algorithm will contribute to avoiding poor local minima, provide robustness with respect to the initial conditions, and induce a progressive increase in the cardinality of the set of clusters needed to be used, via an intuitive bifurcation phenomenon.

A. Soft-Clustering and Annealing Optimization

In the clustering problem (Problem 1), the distortion function J is typically nonconvex and riddled with poor local minima. To partially deal with this phenomenon, soft-clustering approaches have been proposed as a probabilistic framework for clustering. In this case, an input vector X is assigned, through the quantizer Q , to all codevectors $\mu_h \in M$ with probabilities $p(\mu_h|X)$, where $\sum_{h=1}^K p(\mu_h|X) = 1$. In this regard, the quantizer $Q : S \rightarrow M$ becomes a discrete random variable, with the set M being its image and can be fully described by the values of $M = \{\mu_h\}_{h=1}^K$ and the probability functions $\{p(\mu_h|x)\}_{h=1}^K$. In contrast, hard clustering assumes that Q is a simple random variable that can be described fully by M and $V = \{S_h\}_{h=1}^K$, since $p(\mu_h|X) = \mathbb{1}_{[X \in S_h]}$ (see Problem 1).

For the randomized partition, we can rewrite the expected distortion as

$$\begin{aligned} D &= \mathbb{E}[d_\phi(X, Q)] \\ &= \mathbb{E}[\mathbb{E}[d_\phi(X, Q)|X]] \\ &= \int p(x) \sum_{\mu} p(\mu|x) d_\phi(x, \mu) dx \end{aligned}$$

where $p(\mu|x)$ is the association probability relating the input vector x with the codevector μ . We note that at the limit, where each input vector is assigned to a unique codevector with probability one, this reduces to the hard clustering distortion. The main idea in DA is to seek the distribution that minimizes D subject to a specified level of randomness measured by the Shannon entropy

$$\begin{aligned} H(X, M) &= \mathbb{E}[-\log p(X, Q)] \\ &= H(X) + H(Q|X) \\ &= H(X) - \int p(x) \sum_{\mu} p(\mu|x) \log p(\mu|x) dx \end{aligned}$$

by appealing to Jaynes' maximum entropy principle¹ [32]. This multi-objective optimization is conveniently formulated as the minimization of the Lagrangian

$$F = D - TH \quad (2)$$

where T is the temperature parameter that acts as a Lagrange multiplier. Clearly, for large values of T , we maximize the entropy, and, as T is lowered, we trade entropy for reduction in distortion. Equation (2) also represents the scalarization method for tradeoff analysis between two performance metrics [33]. As T varies, we essentially transition from one

Pareto point to another, and the sequence of the solutions will correspond to a Pareto curve of the multi-objective optimization (2) that resembles annealing processes in chemical engineering. In this regard, the entropy H , which is closely related to the "purity" of the clusters, acts as a regularization term which is given progressively less weight as T decreases.

As in the case of vector quantization, we form a coordinate block optimization algorithm to minimize F , by successively minimizing it with respect to the association probabilities $p(\mu|x)$ and the codevector locations μ . Minimizing F with respect to the association probabilities $p(\mu|x)$ is straightforward and yields the Gibbs distribution

$$p(\mu|x) = \frac{e^{-\frac{d(x,\mu)}{T}}}{\sum_{\mu} e^{-\frac{d(x,\mu)}{T}}} \quad \forall x \in S \quad (3)$$

while in order to minimize F with respect to the codevector locations μ , we set the gradients to zero

$$\begin{aligned} \frac{d}{d\mu} D = 0 &\implies \frac{d}{d\mu} \mathbb{E}[\mathbb{E}[d(X, \mu)|X]] = 0 \\ &\implies \int p(x) p(\mu|x) \frac{d}{d\mu} d(x, \mu) dx = 0. \end{aligned} \quad (4)$$

In the following theorem, we show that we can have an analytical solution to the last optimization step (4) in a convenient centroid form if d is a Bregman divergence. This is a similar result to Theorem 1 for vector quantization.

Theorem 2: Assuming the conditional probabilities $p(\mu|x)$ are fixed, the Lagrangian F in (2) is minimized with respect to the codevector locations μ by

$$\mu^* = \mathbb{E}[X|\mu] = \frac{\int x p(x) p(\mu|x) dx}{p(\mu)} \quad (5)$$

if $d := d_\phi$ is a Bregman divergence for some function ϕ that satisfies Definition 2.

Proof: If $d := d_\phi$ is a Bregman divergence, then, by Definition 2, it follows that:

$$\frac{d}{d\mu} d_\phi(x, \mu) = -\frac{\partial^2 \phi(\mu)}{\partial \mu^2} (x - \mu).$$

Therefore, (4) becomes

$$\int (x - \mu) p(x) p(\mu|x) dx = 0 \quad (6)$$

which is equivalent to (5) since $\int p(x) p(\mu|x) dx = p(\mu)$. \square

B. Bifurcation Phenomena

This optimization procedure takes place for decreasing values of the temperature coefficient T such that the solution maintains minimum free energy (thermal equilibrium) while gradually lowering the temperature. Adding to the physical analogy, it is significant that, as the temperature is lowered, the system undergoes a sequence of "phase transitions," which consists of natural cluster splits where the cardinality of the codebook (number of clusters) increases. This is a bifurcation phenomenon and provides a useful tool for controlling the size of the clustering model relating it to the scale of the solution.

¹Informally, Jaynes' principle states: of all the probability distributions that satisfy a given set of constraints, choose the one that maximizes the entropy.

At very high temperature ($T \rightarrow \infty$) the optimization yields uniform association probabilities

$$p(\mu|x) = \lim_{T \rightarrow \infty} \frac{e^{-\frac{d(x,\mu)}{T}}}{\sum_{\mu} e^{-\frac{d(x,\mu)}{T}}} = \frac{1}{K}$$

and provided $d := d_{\phi}$ is a Bregman divergence, all the codevectors are located at the same point

$$\mu = \mathbb{E}[X]$$

which is the expected value of X (Theorem 1). This is true regardless of the number of codevectors available. We refer to the number of different codevectors resulting from the optimization process as *effective codevectors*. These define the cardinality of the codebook, which changes as we lower the temperature. The bifurcation occurs when the solution above a critical temperature T_c is no longer the minimum of the free energy F for $T < T_c$. A set of coincident codevectors then splits into separate subsets. These critical temperatures T_c can be traced when the Hessian of F loses its positive definite property, and are, in some cases, computable (see [17, Th. 1]). In other words, an algorithmic implementation needs only as many codevectors as the number of effective codevectors, which depends only on the temperature parameter, i.e., the Lagrange multiplier of the multi-objective minimization problem in (2). As will be shown in Section III-E, we can detect the bifurcation points by maintaining and perturbing pairs of codevectors at each effective cluster so that they separate only when a critical temperature is reached.

C. Online Deterministic Annealing for Clustering

The conditional expectation $\mathbb{E}[X|\mu]$ in (5) can be approximated by the sample mean of the data points weighted by their association probabilities $p(\mu|x)$, that is,

$$\hat{\mathbb{E}}[X|\mu] = \frac{\sum x p(\mu|x)}{p(\mu)}.$$

This approach, however, defines an offline (batch) optimization algorithm and requires the entire dataset to be available *a priori*, subtly assuming that it is possible to store and also quickly access the entire dataset at each iteration. This is rarely the case in practical applications and results in computationally costly iterations that are slow to converge. We propose an online DA (ODA) algorithm that dynamically updates its estimate of the effective codevectors with every observation. This results in a significant reduction in complexity that comes in two levels. The first refers to huge reduction in memory complexity, since we bypass the need to store the entire dataset and the association probabilities $\{p(\mu|x), \forall x\}$ that map each data point in the dataset to each cluster. The second level refers to the nature of the optimization iterations. In the online approach, the optimization iterations increase in number but become much faster, and practical convergence is often after a smaller number of observations.

To define an online training rule for the above-mentioned optimization framework, we formulate a stochastic approximation algorithm to *recursively estimate* $\mathbb{E}[X|\mu]$ *directly*.

Stochastic approximation, first introduced in [34], was originally conceived as a tool for statistical computation, and, since then, has become a central tool in a number of different disciplines, oftentimes unbeknownst to the users, researchers, and practitioners. It offers an online, adaptive, and computationally inexpensive optimization framework, properties that make it an ideal optimization method for machine learning algorithms. In addition to its connection with optimization and learning algorithms, however, the stochastic approximation is strongly connected to dynamical systems, as well, a property that allows the study of its convergence through the analysis of an ordinary differential equation (ODE), as illustrated in the following theorem.

Theorem 3 [18, Ch. 2]: Almost surely, the sequence $\{x_n\} \in S \subseteq \mathbb{R}^d$ generated by the following stochastic approximation scheme:

$$x_{n+1} = x_n + \alpha(n)[h(x_n) + M_{n+1}], \quad n \geq 0 \quad (7)$$

with prescribed x_0 , *converges* to a (possibly sample path dependent) compact, connected, internally chain transitive, and invariant set of the ODE

$$\dot{x}(t) = h(x(t)), \quad t \geq 0 \quad (8)$$

where $x : \mathbb{R}_+ \rightarrow \mathbb{R}^d$ and $x(0) = x_0$, provided the following assumptions hold.

- (A1) The map $h : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is Lipschitz in S , i.e., $\exists L$ with $0 < L < \infty$ such that $\|h(x) - h(y)\| \leq L\|x - y\|$, $x, y \in S$.
- (A2) The stepsizes $\{\alpha(n) \in \mathbb{R}_{++}, n \geq 0\}$ satisfy $\sum_n \alpha(n) = \infty$, and $\sum_n \alpha^2(n) < \infty$.
- (A3) $\{M_n\}$ is a martingale difference sequence with respect to the increasing family of σ -fields $\mathcal{F}_n := \sigma(x_m, M_m, m \leq n)$, $n \geq 0$, i.e., $\mathbb{E}[M_{n+1}|\mathcal{F}_n] = 0$ a.s., for all $n \geq 0$, and, furthermore, $\{M_n\}$ are square-integrable with $\mathbb{E}[\|M_{n+1}\|^2|\mathcal{F}_n] \leq K(1 + \|x_n\|^2)$, a.s., where $n \geq 0$ for some $K > 0$.
- (A4) The iterates $\{x_n\}$ remain bounded a.s., i.e., $\sup_n \|x_n\| < \infty$ a.s.

As an immediate result, the following corollary also holds.

Corollary 1: If the only internally chain transitive invariant sets for (8) are isolated equilibrium points, then, almost surely, $\{x_n\}$ converges to a, possibly sample dependent, equilibrium point of (8).

Now, we are in place to prove the following theorem.

Theorem 4: Let S a vector space, $\mu \in S$, and $X : \Omega \rightarrow S$ be a random variable defined in a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Let $\{x_n\}$ be a sequence of independent realizations of X , and $\{\alpha(n) > 0\}$ a sequence of stepsizes such that $\sum_n \alpha(n) = \infty$, and $\sum_n \alpha^2(n) < \infty$. Then, the random variable $m_n = \sigma_n/\rho_n$, where (ρ_n, σ_n) are sequences defined by

$$\begin{aligned} \rho_{n+1} &= \rho_n + \alpha(n)[p(\mu|x_n) - \rho_n] \\ \sigma_{n+1} &= \sigma_n + \alpha(n)[x_n p(\mu|x_n) - \sigma_n] \end{aligned} \quad (9)$$

converges to $\mathbb{E}[X|\mu]$ almost surely, i.e., $m_n \xrightarrow{a.s.} \mathbb{E}[X|\mu]$.

Proof: We will use the facts that $p(\mu) = \mathbb{E}[p(\mu|x)]$ and $\mathbb{E}[\mathbb{1}_{\mu}X] = \mathbb{E}[xp(\mu|x)]$. The recursive equations (9) are

stochastic approximation algorithms of the form

$$\begin{aligned}\rho_{n+1} &= \rho_n + \alpha(n)[(p(\mu) - \rho_n) + (p(\mu|x_n) - \mathbb{E}[p(\mu|X)])] \\ \sigma_{n+1} &= \sigma_n + \alpha(n) \\ &\quad \times [(\mathbb{E}[\mathbb{1}_{[\mu]}X] - \sigma_n) + (x_n p(\mu|x_n) - \mathbb{E}[x_n p(\mu|X)])].\end{aligned}\quad (10)$$

It is obvious that both stochastic approximation algorithms satisfy the conditions of Theorem 3 and Corollary 1. As a result, they converge to the asymptotic solution of the differential equations

$$\begin{aligned}\dot{\rho} &= p(\mu) - \rho \\ \dot{\sigma} &= \mathbb{E}[\mathbb{1}_{[\mu]}X] - \sigma\end{aligned}$$

which can be trivially derived through standard ODE analysis to be $(p(\mu), \mathbb{E}[\mathbb{1}_{[\mu]}X])$. In other words, we have shown that

$$(\rho_n, \sigma_n) \xrightarrow{a.s.} (p(\mu), \mathbb{E}[\mathbb{1}_{[\mu]}X]). \quad (11)$$

The convergence of m_n follows from the fact that $\mathbb{E}[X|\mu] = \mathbb{E}[\mathbb{1}_{[\mu]}X]/p(\mu)$ and standard results on the convergence of the product of two random variables. \square

As a direct consequence of this theorem, the following corollary provides an online learning rule that solves the optimization problem of the DA algorithm.

Corollary 2: The online training rule

$$\begin{cases} \rho_i(n+1) = \rho_i(n) + \alpha(n)[\hat{p}(\mu_i|x_n) - \rho_i(n)] \\ \sigma_i(n+1) = \sigma_i(n) + \alpha(n)[x_n \hat{p}(\mu_i|x_n) - \sigma_i(n)] \end{cases} \quad (12)$$

where the quantities $\hat{p}(\mu_i|x_n)$ and $\mu_i(n)$ are recursively updated as follows:

$$\begin{aligned}\hat{p}(\mu_i|x_n) &= \frac{\rho_i(n)e^{-\frac{d(x_n, \mu_i(n))}{T}}}{\sum_i \rho_i(n)e^{-\frac{d(x_n, \mu_i(n))}{T}}} \\ \mu_i(n) &= \frac{\sigma_i(n)}{\rho_i(n)}\end{aligned} \quad (13)$$

converges almost surely to a possibly sample path dependent solution of the block optimization (3) and (5).

Finally, the learning rule (12) and (13) can be used to define a consistent (histogram) density estimator at the limit $T \rightarrow 0$. This follows from the fact that as $T \rightarrow 0$, the number of clusters K goes to infinity, $p(\mu_h|X) \rightarrow \mathbb{1}_{[X \in S_h]}$, and, as a result, $F \rightarrow J$, i.e., the consistency of Algorithm 1 can be studied with similar arguments to the stochastic divergence-based vector quantization algorithm (1) (see [13], [14]).

D. Online Deterministic Annealing for Classification

We can extend the proposed learning algorithm to be used for classification as well. In this case, we can rewrite the expected distortion as

$$D = \mathbb{E}[d^b(c_X, Q^c)]$$

where $d^b(c_X, c_\mu) = \mathbb{1}_{[c_X \neq c_\mu]}$. Because d^b is not differentiable, using similar principles as in the case of LVQ, we can instead approximate the optimal solution by solving the minimization problem for the following distortion measure:

$$d^c(x, c_X, \mu, c_\mu) = \begin{cases} d(x, \mu), & c_X = c_\mu \\ 0, & c_X \neq c_\mu. \end{cases} \quad (14)$$

This particular choice for the distortion measure d^c will lead to some interesting regularization properties of the proposed online approach (see Section III-E).

It is easy to show that the coordinate block optimization steps (3) and (5), in this case become

$$\begin{aligned}p(\mu, c_\mu|x, c_X) &= \frac{e^{-\frac{d^c(x, c_X, \mu, c_\mu)}{T}}}{\sum_{\mu, c_\mu} e^{-\frac{d^c(x, c_X, \mu, c_\mu)}{T}}} \\ \text{and} \\ \mu^* &= \frac{\sum_{c_X=c_\mu} x p(x, c_X) p(\mu, c_\mu|x, c_X)}{\sum_{c_X=c_\mu} p(x, c_X) p(\mu, c_\mu|x, c_X)}\end{aligned}$$

respectively. In the last step, we have assumed that the class c_μ of each centroid μ is given and cannot be changed dynamically by the algorithm, which results to the minimization with respect to μ only. In a similar fashion, it can be shown that the online learning rule that solves the optimization problem of the DA algorithm for classification, based on the distortion measure (14), is given by

$$\begin{aligned}\rho_i(n+1) &= \rho_i(n) + \alpha(n)\mathbb{1}_{[c_{X_j}=c_{\mu_i}]} \\ &\quad \times [\hat{p}(\mu_i, c_{\mu_i}|x_n, c_{X_n}) - \rho_i(n)] \\ \sigma_i(n+1) &= \sigma_i(n) + \alpha(n)\mathbb{1}_{[c_{X_j}=c_{\mu_i}]} \\ &\quad \times [x_n \hat{p}(\mu_i, c_{\mu_i}|x_n, c_{X_n}) - \sigma_i(n)]\end{aligned} \quad (15)$$

where

$$\begin{aligned}\hat{p}(\mu_i, c_{\mu_i}|x_n, c_{X_n}) &= \frac{\rho_i(n)e^{-\frac{d^c(x_n, c_{X_n}, \mu_i(n), c_{\mu_i}(n))}{T}}}{\sum_i \rho_i(n)e^{-\frac{d^c(x_n, c_{X_n}, \mu_i(n), c_{\mu_i}(n))}{T}}} \\ \mu_i(n) &= \frac{\sigma_i(n)}{\rho_i(n)}.\end{aligned} \quad (16)$$

At the limit $T \rightarrow 0$, the quantization scheme described earlier equipped with a majority-vote classification rule is strongly Bayes risk consistent, i.e., converges to the optimal (Bayes) probability of error (see [13, Ch. 21]). However, due to the choice of the distortion measure d^c in (14) used in ODA for classification, the algorithm can be used to estimate consistent class-conditional density estimators, which define the natural classification rule

$$\hat{c}(x) = c_{\mu_{h^*}} \quad (17)$$

where $h^* = \arg \max_{\tau=1, \dots, K} p(\mu_\tau|x)$, $h \in \{1, \dots, K\}$.

E. Algorithm

The proposed ODA algorithm (Algorithm 1) is based on (15) and (16) and can be used for both clustering and classification alike, depending on whether the data belong to a single (clustering) or several classes (classification).

1) *Temperature Schedule:* The temperature schedule T_i plays an important role in the behavior of the algorithm. Starting at high temperature T_{\max} ensures the correct operation of the algorithm. The value of T_{\max} depends on the domain of the data and should be large enough such that there is only one effective codevector at $T = T_{\max}$. When the range of the domain of the data is not known *a priori*, overestimation is recommended. The stopping temperature T_{\min} can be set

a priori or be decided online depending on the performance of the model at each temperature level. The temperature step $dT_i = T_{i-1} - T_i$ should be small enough such that no critical temperature is missed. On the other hand, the smaller the step dT_i , the more optimization problems need to be solved. It is common practice to use the geometric series $T_{i+1} = \gamma T_i$.

2) *Stochastic Approximation*: Regarding the stochastic approximation step sizes, simple time-based learning rates, e.g., of the form $\alpha_n = 1/a + bn$, have been sufficient for fast convergence in all our experiments so far. Convergence is checked with the condition $d_\phi(\mu_i^n, \mu_i^{n-1}) < \epsilon_c$ for a given threshold ϵ_c that can depend on the domain of X . Exploring adaptive learning rates would be an interesting research direction for the future.

3) *Bifurcation and Perturbations*: To every temperature level T_i , corresponds a set of effective codevectors $\{\mu_j\}_{j=1}^{K_i}$, which consist of the different solutions of the optimization problem (2) at T_i . Bifurcation, at T_i , is detected by maintaining a pair of perturbed codevectors $\{\mu_j + \delta, \mu_j - \delta\}$ for each effective codevector μ_j generated at T_{i-1} , i.e. for $j = 1, \dots, K_{i-1}$. Using arguments from variational calculus [17], it is easy to see that, upon convergence, the perturbed codevectors will merge if a critical temperature has not been reached, and will get separated otherwise. In case of a merge, one of the perturbed codevectors is removed from the model. Therefore, the cardinality of the model is at most doubled at every temperature level. For classification, a perturbed codevector for each distinct class is generated.

4) *Regularization*: Merging is detected by the condition $d_\phi(\mu_j, \mu_i) < \epsilon_n$, where ϵ_n is a design parameter that acts as a regularization term for the model. Large values for ϵ_n (compared with the support of the data X) lead to fewer effective codevectors, while small ϵ_n values lead to a fast growth in the model size, which is connected to overfitting. It is observed that, for practical convergence, the perturbation noise δ is best to not exceed ϵ_n . An additional regularization mechanism that comes as a natural consequence of the stochastic approximation learning rule is the detection of idle codevectors. To see that, notice that the sequence $\rho_i(n)$ resembles an approximation of the probability $p(\mu_i, c_{\mu_i})$. In the updates (12) and (13), $\rho_i(n)$ becomes negligible ($\rho_i(n) < \epsilon_r$) if not updated by any nearby observed data, which is a natural criterion for removing the codevector μ_i . This happens if all observed data samples x_n are largely dissimilar to μ_i . In classification, because of the choice of d^c in (14), codevectors μ_i that are not assigned the same class as the data in their vicinity, will end up to be removed, as well. The threshold ϵ_r is a parameter that usually takes values near zero.

5) *Complexity*: The worst case complexity of Algorithm 1 behaves as $O(\sigma_{\max} N K_{\max}^2 d)$, where the following holds.

- 1) N is an upper bound of the number of data samples observed, which should be large enough to overestimate the iterations needed for convergence.
- 2) d is the dimension of the input vectors, i.e., $x \in \mathbb{R}^d$.
- 3) K_{\max} is the maximum number of codevectors allowed.
- 4) $\sigma_{\max} = \{\sigma_1, \sigma_2, \dots, \sigma_{K_{\max}}\}$, where σ_i is the number of temperature values in our temperature schedule that lie between two critical temperatures T_i and

Algorithm 1 ODA

```

Select Bregman divergence  $d_\phi$ 
Set temperature schedule:  $T_{\max}, T_{\min}, \gamma$ 
Decide maximum number of codevectors  $K_{\max}$ 
Set convergence parameters:  $\{\alpha_n\}, \epsilon_c, \epsilon_n, \epsilon_r, \delta$ 
Select initial configuration  $\{\mu^i\} : c_{\mu^i} = c, \forall c \in \mathcal{C}$ 
Initialize:  $K = 1, T = T_{\max}$ 
Initialize:  $p(\mu^i) = 1, \sigma(\mu^i) = \mu^i p(\mu_i), \forall i$ 
while  $K < K_{\max}$  and  $T > T_{\min}$  do
  Perturb  $\mu^i \leftarrow \{\mu^i + \delta, \mu^i - \delta\}, \forall i$ 
  Increment  $K \leftarrow 2K$ 
  Update  $p(\mu^i), \sigma(\mu^i) \leftarrow \mu^i p(\mu^i), \forall i$ 
  Set  $n \leftarrow 0$ 
  repeat
    Observe data point  $x$  and class label  $c$ 
    for  $i = 1, \dots, K$  do
      Compute membership  $s^i = \mathbb{1}_{[c_{\mu^i}=c]}$ 
      Update:
        
$$p(\mu^i|x) \leftarrow \frac{p(\mu^i)e^{-\frac{d_\phi(x, \mu^i)}{T}}}{\sum_i p(\mu^i)e^{-\frac{d_\phi(x, \mu^i)}{T}}}$$

        
$$p(\mu^i) \leftarrow p(\mu^i) + \alpha_n[s^i p(\mu^i|x) - p(\mu^i)]$$

        
$$\sigma(\mu^i) \leftarrow \sigma(\mu^i) + \alpha_n[s^i x p(\mu^i|x) - \sigma(\mu^i)]$$

        
$$\mu^i \leftarrow \frac{\sigma(\mu^i)}{p(\mu^i)}$$

      Increment  $n \leftarrow n + 1$ 
    end for
  until  $d_\phi(\mu_n^i, \mu_{n-1}^i) < \epsilon_c, \forall i$ 
  Keep effective codevectors:
    discard  $\mu^i$  if  $d_\phi(\mu^j, \mu^i) < \epsilon_n, \forall j, j, i \neq j$ 
  Remove idle codevectors:
    discard  $\mu^i$  if  $p(\mu^i) < \epsilon_r, \forall i$ 
  Update  $K, p(\mu^i), \sigma(\mu^i), \forall i$ 
  Lower temperature  $T \leftarrow \gamma T$ 
end while

```

T_{i+1} , with the understanding that at T_i there are i distinct effective codevectors present. Here, we have assumed that K_{\max} is achievable within our temperature schedule.

6) *Fine-Tuning*: In practice, because the convergence to the Bayes decision surface comes at the limit $(K, T) \rightarrow (\infty, 0)$, a fine-tuning mechanism should be designed to run on top of the proposed algorithm after T_{\min} . This can be either an LVQ algorithm (Section II-B) or some other local model.

IV. EXPERIMENTAL EVALUATION AND DISCUSSION

We illustrate the properties and evaluate the performance of the proposed algorithm is widely used artificial and real datasets for clustering and classification.²

A. Toy Examples

We first showcase how Algorithm 1 works in three simple, but illustrative, classification problems in two

²Code and Reproducibility: The source code is publicly available online at <https://github.com/MavridisChristos/OnlineDeterministicAnnealing>.




Fig. 1. (a)–(c) Illustration of the evolution of Algorithm 1 for decreasing temperature T in binary classification in 2-D. (d) Showcasing robustness with respect to bad initial conditions. (a) Concentric circles. (b) Half-moons. (c) Gaussians. (d) Poor initial conditions.

dimensions (Fig. 1). The first two are binary classification problems with the underlying class distributions shaped as concentric circles [Fig. 1(a)], and half-moons [Fig. 1(b)], respectively. The third is a multiclass classification problem with Gaussian mixture class distributions [Fig. 1(c)]. All datasets consist of 1500 samples. Since the objective is to give a geometric illustration of how the algorithm works in the 2-D plane, the Euclidean distance is used. The algorithm starts at high temperature with a single codevector for each class. As the temperature coefficient gradually decreases (Fig. 1, from left to right), the number of codevectors progressively increases. The accuracy of the algorithm typically increases as well. As the temperature goes to zero, the complexity of the model, i.e., the number of codevectors, rapidly increases (Fig. 1, rightmost pictures). This may, or may not, translate to a corresponding performance boost. A single parameter—the temperature T —offers online control on this complexity-accuracy tradeoff. Finally, Fig. 1(d) showcases the robustness of the proposed algorithm with respect to the initial configuration. Here, the codevectors are poorly initialized outside the support of the data, which is not assumed known *a priori* (e.g., online observations of unknown domain). In this example, the LVQ algorithm has been shown to fail [35]. In contrast, the entropy term H in the optimization objective of Algorithm 1, allows for the online adaptation to the domain of the dataset and helps to prevent poor local minima.

B. Real Datasets

1) *Clustering*: For clustering, we consider the following datasets: 1) the dataset of Fig. 1(c) (Gaussians); 2) the WBCD dataset [36]; 3) the PIMA dataset [37]; and 4) the Adult dataset³ [36]. In Fig. 2, we compare Algorithm 1 with the

³15 000 samples randomly selected. Nonnumerical features removed.

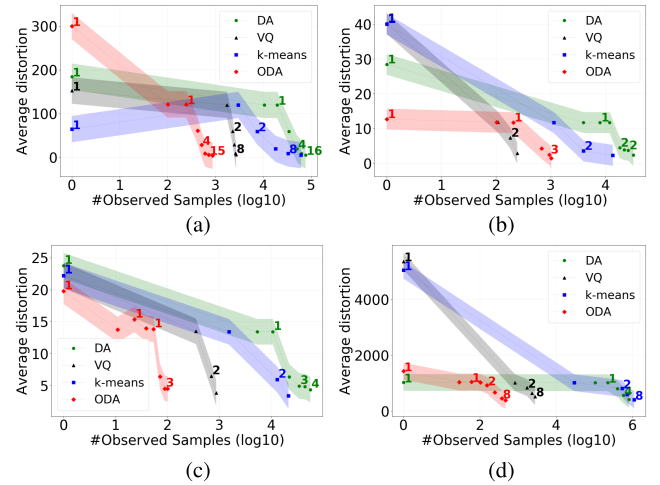


Fig. 2. Algorithm comparison for clustering. (a) Gaussians. (b) WBCD. (c) PIMA. (d) Adult.

TABLE I
CLASSIFICATION ACCURACIES IN FIVE-FOLD CROSS VALIDATION

DATA SET	ODA	SVM	NN	RF
GAUSSIAN	98.9 \pm 0.0	79.5 \pm 0.0	98.6 \pm 0.0	98.7 \pm 0.0
WBCD	90.7 \pm 0.0	85.6 \pm 0.0	92.7 \pm 0.0	94.6 \pm 0.0
CREDIT (F1)	95.6 \pm 0.0	69.1 \pm 0.2	58.9 \pm 0.1	62.8 \pm 0.1
PIMA	70.5 \pm 0.0	62.9 \pm 0.0	76.3 \pm 0.0	74.4 \pm 0.0

online sVQ algorithm (Definition 1), and two offline algorithms, namely, k -means [27], and the original DA algorithm [17]. The algorithms are compared in terms of the minimum average distortion achieved, as a function of the number of samples they observed, and the number of clusters they used (floating numbers inside Fig. 2). The Euclidean distance is used for fair comparison. Since there is no criterion to decide the number of clusters K for k -means and sVQ, we run them sequentially for the K values estimated by DA and add up the computational time. All algorithms are able to achieve comparable average distortion values given good initial conditions and appropriate size K . Therefore, the progressive estimation of K and the robustness with respect to the initial conditions are key features of both annealing algorithms, i.e., DA and ODA (Algorithm 1). Compared with the offline algorithms, i.e., k -means and DA, ODA and sVQ achieve practical convergence with significantly smaller number of observations, which corresponds to reduced computational time, as argued earlier. Notice the substantial difference in running time between the original DA algorithm and the proposed ODA algorithm in Fig. 3. Compared with the online sVQ (and LVQ), the probabilistic approach of ODA introduces additional computational cost: all neurons are now updated in every iteration, instead of only the winner neuron. However, the updates can still be computed fast when using Bregman divergences (Theorem 2), and the aforementioned benefits of the annealing nature of ODA, outweigh this additional cost in many real-life problems.




Fig. 3. Running time of the algorithms in Figs. 2(a) and 4(a).

2) *Classification*: For classification, we consider the Gaussian [Fig. 1(c)], WBCD, PIMA, and credit card⁴ [38] datasets. We compare Algorithm 1 against an SVM model with a linear kernel [39], a feedforward fully connected neural network with a single-hidden layer of n_{NN} neurons (NN), and the random forests (RFs) algorithm with t_{RF} estimators [40]. These algorithms have been selected to represent today's standards in simple classification tasks, i.e., when no sophisticated feature extraction is required. The SVM classifier represents the class of linear classification models, the neural network represents the class of nonlinear approximation models, and the random forests algorithm represents the class of partition-based methods with bootstrap aggregating. Table I shows the results of a five-fold cross validation (80%/20%), and Fig. 4 shows the performance of the algorithms during a random test. The evolution of the complexity of the ODA model is shown as a function of the observed samples and the classification accuracy achieved. We use the generalized I divergence (Example 2) in the WBCD dataset and the Euclidean distance in the rest. ODA (Algorithm 1) outperforms the linear SVM classifier and can achieve comparable performance with the NN and the RF algorithms, which are today's standards in classification tasks where no feature extraction is required. In the greatly unbalanced credit card dataset, all algorithms achieved accuracy close to 100%, but the $F1$ scores dropped significantly [Fig. 4(d)]. Notably, this was not the case with the ODA algorithm. This may be due to the generative nature of the algorithm and might also be an instance of the robustness expected by vector quantization algorithms [16]. Justifying and quantifying this robustness is beyond the scope of this article.

3) *Parameters*: The parameters $n_{NN} \in [10, 100]$ and $t_{RF} \in [10, 100]$ were selected through extensive grid search. In contrast, the parameters of the ODA algorithm for all the experiments were set as follows: $T_{\max} = 100\Delta_S d$, $T_{\min} = 0.001\Delta_S d$, $K_{\max} = 100$, $\gamma = 0.8$, $\epsilon_c = 0.0001\Delta_S d$, $\epsilon_n = 0.001\Delta_S d$, $\epsilon_r = 10^{-7}$, $\delta = 0.01\Delta_S d$, and $\alpha_n = 1/1 + 0.9n$, where d is the number of dimensions of the input $X \in S \subseteq \mathbb{R}^d$, and Δ_S represents the length of the largest edge of the smallest d -orthotope that contains S . We stress that no parameter tuning has taken place for the proposed algorithm.

4) *Limitations*: Finally, we note that both NN and RF outperform Algorithm 1 in some datasets (Table I). A fine-tuning mechanism, as discussed in Section III-E, could alleviate these differences and is currently not used in our experiments. Regarding the running time of the ODA algorithm, Fig. 3 shows the execution time of the learning algorithms used in Figs. 2(a) and 4(a). All experiments were implemented

⁴15 000 samples randomly selected.

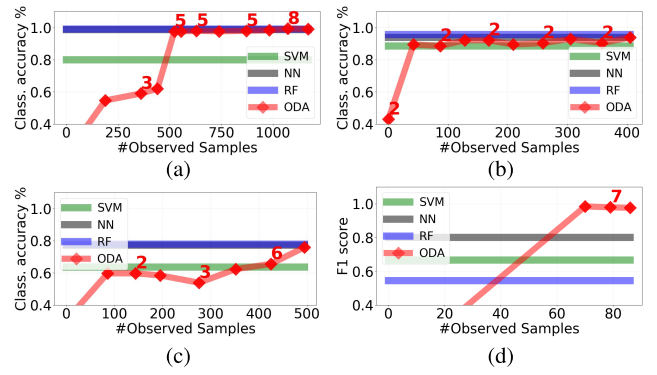


Fig. 4. Algorithm comparison for classification. (a) Gaussians. (b) WBCD. (c) PIMA. (d) Credit card ($F1$ score).

in a personal computer. We note that in contrast to the commercial, and, therefore, optimized versions of the k -means, SVM, NN, and RF algorithms, the algorithmic implementation of the proposed algorithm is not yet optimized and substantial speedup is expected through appropriate software development.

V. CONCLUSION

It is understood that the tradeoff between model complexity and performance in machine learning algorithms is closely related to overfitting, generalization, and robustness to input perturbations and adversarial attacks. We investigate the properties of learning with progressively growing models and propose an online annealing optimization approach as a learning algorithm that progressively adjusts its complexity with respect to new observations, offering online control over the performance-complexity tradeoff. The proposed algorithm can be viewed as a neural network with inherent regularization mechanisms, the learning rule of which is formulated as an online gradient-free stochastic approximation algorithm. As a prototype-based learning algorithm, it offers a progressively growing knowledge base that can be interpreted as a memory unit that parallels similar concepts from cognitive psychology and neuroscience. The annealing nature of the algorithm prevents poor local minima, offers robustness to initial conditions, and provides a means to progressively increase the complexity of the learning model as needed. To the best of our knowledge, this is the first time such a progressive approach has been proposed for machine learning applications. We believe that these results can lead to new developments in learning with progressively growing models, especially in communication, control, and reinforcement learning applications.

REFERENCES

- [1] K. P. Bennett and E. Parrado-Hernández, "The interplay of optimization and machine learning research," *J. Mach. Learn. Res.*, vol. 7, pp. 1265–1281, Jul. 2006.
- [2] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2012, pp. 1097–1105.
- [3] H. Lee, R. Grosse, R. Ranganath, and A. Y. Ng, "Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations," in *Proc. 26th Annu. Int. Conf. Mach. Learn.*, 2009, pp. 609–616.

- [4] N. C. Thompson, K. Greenewald, K. Lee, and G. F. Manso, "The computational limits of deep learning," 2020, *arXiv:2007.05558*.
- [5] E. Strubell, A. Ganesh, and A. McCallum, "Energy and policy considerations for deep learning in NLP," 2019, *arXiv:1906.02243*.
- [6] C. Szegedy *et al.*, "Intriguing properties of neural networks," 2013, *arXiv:1312.6199*.
- [7] N. Carlini and D. Wagner, "Towards evaluating the robustness of neural networks," in *Proc. IEEE Symp. Secur. Privacy*, May 2017, pp. 39–57.
- [8] V. Schwag *et al.*, "Analyzing the robustness of open-world machine learning," in *Proc. 12th ACM Workshop Artif. Intell. Secur.*, 2019, pp. 105–116.
- [9] H. Xu and S. Mannor, "Robustness and generalization," *Mach. Learn.*, vol. 86, no. 3, pp. 391–423, 2012.
- [10] C. G. Northcutt, A. Athalye, and J. Mueller, "Pervasive label errors in test sets destabilize machine learning benchmarks," 2021, *arXiv:2103.14749*.
- [11] A. Raghunathan, S. M. Xie, F. Yang, J. C. Duchi, and P. Liang, "Adversarial training can hurt generalization," 2019, *arXiv:1906.06032*.
- [12] T. Kohonen, *Learning Vector Quantization*. Berlin, Germany: Springer, 1995, pp. 175–189.
- [13] L. Devroye, L. Györfi, and G. Lugosi, *A Probabilistic Theory of Pattern Recognition*, vol. 31. Springer, 2013.
- [14] C. N. Mavridis and J. S. Baras, "Convergence of stochastic vector quantization and learning vector quantization with Bregman divergences," in *Proc. 21st IFAC World Congr.*, 2020, pp. 2244–2249.
- [15] E. A. Uriarte and F. D. Martín, "Topology preservation in SOM," *Int. J. Appl. Math. Comput. Sci.*, vol. 1, no. 1, pp. 19–22, 2005.
- [16] S. Saralajew, L. Holdijk, M. Rees, and T. Villmann, "Robustness of generalized learning vector quantization models against adversarial attacks," in *Proc. Int. Workshop Self-Organizing Maps*. Springer, 2019, pp. 189–199.
- [17] K. Rose, "Deterministic annealing for clustering, compression, classification, regression, and related optimization problems," *Proc. IEEE*, vol. 86, no. 11, pp. 2210–2239, Nov. 1998.
- [18] V. S. Borkar, *Stochastic Approximation: A Dynamical Systems Viewpoint*, vol. 48. Springer, 2009.
- [19] A. Banerjee, S. Merugu, I. S. Dhillon, and J. Ghosh, "Clustering with Bregman divergences," *J. Mach. Learn. Res.*, vol. 6, pp. 1705–1749, Oct. 2005.
- [20] T. Villmann, S. Haase, F.-M. Schleif, B. Hammer, and M. Biehl, "The mathematics of divergence based online learning in vector quantization," in *Proc. IAPR Workshop Artif. Neural Netw. Pattern Recognit.* Springer, 2010, pp. 108–119.
- [21] C. N. Mavridis and J. S. Baras, "Progressive graph partitioning based on information diffusion," in *Proc. Conf. Decis. Control (CDC)*, 2021, pp. 37–42.
- [22] C. N. Mavridis, N. Suriyarachchi, and J. S. Baras, "Detection of dynamically changing leaders in complex swarms from observed dynamic data," in *Proc. Conf. Decis. Game Theory Secur. (GameSec)*, 2020, pp. 223–240.
- [23] C. N. Mavridis, N. Suriyarachchi, and J. S. Baras, "Maximum-entropy progressive state aggregation for reinforcement learning," in *Proc. Conf. Decis. Control (CDC)*, 2021, pp. 5144–5149.
- [24] M. Biehl, B. Hammer, and T. Villmann, "Prototype-based models in machine learning," *Wiley Interdiscipl. Rev. Cognit. Sci.*, vol. 7, pp. 92–111, Jan. 2016.
- [25] Y. Linde, A. Buzo, and R. Gray, "An algorithm for vector quantizer design," *IEEE Trans. Commun.*, vol. COM-28, no. 1, pp. 84–95, Jan. 1980.
- [26] M. Sabin and R. Gray, "Global convergence and empirical consistency of the generalized Lloyd algorithm," *IEEE Trans. Inf. Theory*, vol. IT-32, no. 2, pp. 148–155, Mar. 1986.
- [27] L. Bottou and Y. Bengio, "Convergence properties of the K-means algorithms," in *Proc. Adv. Neural Inf. Process. Syst.*, 1995, pp. 585–592.
- [28] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*. Hoboken, NJ, USA: Wiley, 2012.
- [29] A. Sato and K. Yamada, "Generalized learning vector quantization," in *Proc. Adv. Neural Inf. Process. Syst.*, 1996, pp. 423–429.
- [30] B. Hammer and T. Villmann, "Generalized relevance learning vector quantization," *Neural Netw.*, vol. 15, nos. 8–9, pp. 1059–1068, Oct. 2002.
- [31] H. K. B. Babiker and R. Goebel, "Using KL-divergence to focus deep visual explanation," 2017, *arXiv:1711.06431*.
- [32] E. T. Jaynes, "Information theory and statistical mechanics," *Phys. Rev.*, vol. 106, no. 4, p. 620, 1957.
- [33] K. Miettinen, *Nonlinear Multiobjective Optimization*, vol. 12. Springer, 2012.
- [34] H. Robbins and S. Monro, "A stochastic approximation method," *Ann. Math. Statist.*, vol. 22, no. 3, pp. 400–407, 1951.
- [35] J. S. Baras and A. LaVigna, "Convergence of a neural network classifier," in *Proc. Adv. Neural Inf. Process. Syst.*, 1991, pp. 1–6.
- [36] D. Dua and C. Graff, (2017). *UCI Machine Learning Repository*. [Online]. Available: <http://archive.ics.uci.edu/ml>
- [37] J. W. Smith, J. Everhart, W. Dickson, W. Knowler, and R. Johannes, "Using the ADAP learning algorithm to forecast the onset of diabetes mellitus," in *Proc. Annu. Symp. Comput. Appl. Med. Care*. Bethesda, MD, USA: AMIA, 1988, p. 261.
- [38] F. Carcillo, Y.-A. Le Borgne, O. Caelen, Y. Kessaci, F. Oblé, and G. Bontempi, "Combining unsupervised and supervised learning in credit card fraud detection," *Inf. Sci.*, vol. 557, pp. 317–331, May 2021.
- [39] M. A. Hearst, S. T. Dumais, E. Osman, J. Platt, and B. Scholkopf, "Support vector machines," *IEEE Intell. Syst. Appl.*, vol. 13, no. 4, pp. 18–28, Jul./Aug. 2008.
- [40] L. Breiman, "Random forests," *Mach. Learn.*, vol. 45, pp. 5–32, Oct. 2001.

Christos N. Mavridis (Member, IEEE) received the Diploma degree in electrical and computer engineering from the National Technical University of Athens, Athens, Greece, in 2017, and the M.S. and Ph.D. degrees in electrical and computer engineering from the University of Maryland, College Park, MD, USA, in 2021.



He was a Researcher with the Department of Electrical and Computer Engineering, University of Maryland, and a Research Intern with the Math and Algorithms Research Group, Nokia Bell Labs, Murray Hill, NJ, USA, and the System Sciences Laboratory, PARC, Palo Alto, CA, USA. He is currently a member of the Institute for Systems Research, Maryland, and the Autonomy, Robotics, and Cognition Laboratory, Maryland. His research interests include learning theory, stochastic optimization, systems and control theory, multi-agent systems, and robotics.

Dr. Mavridis received the A. James Clark School of Engineering Distinguished Graduate Fellowship in 2017, the Outstanding Graduate Research Assistant Award in 2020, the Future Faculty Fellowship in 2021, and the Ann G. Wylie Dissertation Fellowship in 2021. He has been a Finalist in the Qualcomm Innovation Fellowship USA, San Diego, CA, USA, 2018. He has received the Best Student Paper Award (first place) in the IEEE International Conference on Intelligent Transportation Systems in 2021.



John S. Baras (Life Fellow, IEEE) received the Diploma degree in electrical and mechanical engineering from the National Technical University of Athens, Athens, Greece, in 1970, and the M.S. and Ph.D. degrees in applied mathematics from Harvard University, Cambridge, MA, USA, in 1971 and 1973, respectively.

From 1985 to 1991, he was the Founding Director of the Institute for Systems Research (ISR), University of Maryland, College Park, MD, USA. He is currently a Distinguished University Professor and the Lockheed Martin Chair in systems engineering with the Department of Electrical and Computer Engineering, ISR. Since 1992, he has been the Director of the Maryland Center for Hybrid Networks, which he was a Co-Founder. His research interests include systems and control, optimization, communication networks, applied mathematics, machine learning, artificial intelligence, signal processing, robotics, computing systems, security, trust, systems biology, healthcare systems, and model-based systems engineering.

Dr. Baras is a fellow of SIAM, AAAS, NAI, IFAC, AMS, and AIAA, a member of the National Academy of Inventors, and a Foreign Member of the Royal Swedish Academy of Engineering Sciences. He received the 1980 George Axelby Award from the IEEE Control Systems Society, the 2006 Leonard Abraham Prize from the IEEE Communications Society, the 2017 IEEE Simon Ramo Medal, the 2017 AACC Richard E. Bellman Control Heritage Award, and the 2018 AIAA Aerospace Communications Award. In 2018, he was received the Doctorate Honoris Causa from the National Technical University of Athens, Greece. In 2016, he was inducted in the A. J. Clark School of Engineering Innovation Hall of Fame.