The Shallow Gibbs Network, Double Backpropagation and Differential Machine learning

Nonvikan Karl-Augustt Alahassa, Ph. D Candidate

ALAHASSAN@DMS.UMONTREAL.CA

*Department of Mathematics and Statistics
University of Montreal
Montreal, Quebec, Canada*

Alejandro Murua, Full Professor

ALEJANDRO.MURUA@UMONTREAL.CA

*Department of Mathematics and Statistics
University of Montreal
Montreal, Quebec, Canada*

Editor: xxx

Abstract

We have built a Shallow Gibbs Network model as a Random Gibbs Network Forest to reach the performance of the Multilayer feedforward Neural Network in a few numbers of parameters, and fewer backpropagation iterations. To make it happens, we propose a novel optimization framework for our Bayesian Shallow Network, called the Double Backpropagation Scheme (DBS) that can also fit perfectly the data with appropriate learning rate, and which is convergent and universally applicable to any Bayesian neural network problem. The contribution of this model is broad. First, it integrates all the advantages of the Potts Model, which is a very rich random partitions model, that we have also modified to propose its Complete Shrinkage version using agglomerative clustering techniques. The model takes also an advantage of Gibbs Fields for its weights precision matrix structure, mainly through Markov Random Fields, and even has five (5) variants structures at the end: the Full-Gibbs, the Sparse-Gibbs, the Between layer Sparse Gibbs which is the B-Sparse Gibbs in a short, the Compound Symmetry Gibbs (CS-Gibbs in short), and the Sparse Compound Symmetry Gibbs (Sparse-CS-Gibbs) model. The Full-Gibbs is mainly to remind fully-connected models, and the other structures are useful to show how the model can be reduced in terms of complexity with sparsity and parsimony. All those models have been experimented with the Mulan project multivariate regression dataset, and the results arouse interest in those structures, in a sense that different structures help to reach different results in terms of Mean Squared Error (MSE) and Relative Root Mean Squared Error (RRMSE). For the Shallow Gibbs Network model, we have found the perfect learning framework : it is the $(l_1, \zeta, \epsilon_{dbs})$ – **DBS** configuration, which is a combination of the *Universal Approximation Theorem*, and the DBS optimization, coupled with the $(dist)$ -Nearest Neighbor-(h)-Taylor Series-Perfect Multivariate Interpolation ($dist$ -NN-(h)-TS-PMI) model [which in turn is a combination of the research of the Nearest Neighborhood for a good Train-Test association, the Taylor Approximation Theorem, and finally the Multivariate Interpolation Method]. It indicates that, with an appropriate number l_1 of neurons on the hidden layer, an optimal number ζ of DBS updates, an optimal DBS learning rate ϵ_{dbs} , an optimal distance $dist_{opt}$ in the research of the nearest neighbor in the training dataset for each test data x_i^{test} , an optimal order h_{opt} of the Taylor approximation for the Perfect Multivariate Interpolation ($dist$ -NN-(h)-TS-PMI) model once the **DBS** has overfitted the training dataset, the train and the test error converge to zero (0).

Keywords: Statistics, Neural Networks, Probability and stochastic processes, Graphical models, Structured models, Gibbs Fields, Sparse Models, Compound Symmetry, Perfect learning, Double Backpropagation, Multivariate Regression, Multivariate Interpolation, Taylor Theorem, Artificial intelligence.

Contents

1	Introduction	4
1.1	Efficiency of regression clustering	4
1.2	Combination of neural network regression and Potts clustering model	4
2	Shallow Gibbs networks	5
2.1	The sparse-Gibbs network	6
2.2	Compound symmetry Gibbs network	7
3	The random-Potts partition model	8
3.1	Some practical considerations	9
3.2	The Potts Clustering Model with Complete Shrinkage	9
4	The shallow Potts Gibbs-network mixture model	9
5	Bayesian variational inference	10
5.1	Regularization on the CS-Gibbs model	13
5.2	Keeping positive definiteness on the precision Matrix	14
6	Predictive Posterior	16
7	On Double backpropagation	18
7.1	A Proposed Double backpropagation scheme	18
8	On Differentiation Methods	20
8.1	Automatic differentiation.	20
8.2	Symbolic differentiation (SD).	20
8.3	Numerical differentiation.	21
9	Empirical differentiation and Differential Machine learning	21
9.1	Interpolation as a Machine <i>learner</i>	21
9.2	Multivariate Interpolation: Taylor Approximation is the future of Machine learning	22
9.2.1	Taylor Approximation is the future of Machine learning.	23
10	Data Augmentation for Empirical Differentiation (DAED)	24
11	Generalization Method	25
12	Experimental evaluation	26
12.1	The Results	27
	Concluding Remarks & Discussion notes	33
12.2	Note on the Shallow Gibbs Structure.	33
12.3	A Generalized Double Back-Propagation Scheme (GDBS) for any parametric model	35
12.4	The Infinite Zelda Stochastic Game.	36
	APPENDICES	39
A	Other Experiments Results with the Shallow Gibbs Models	39
B	Cholesky Decomposition	40

C	The learning rate	41
D	Stochastic gradient	42
	Bibliography	46

1. Introduction

We introduce a novel ensemble learning approach which combines random partitions models with a non-parametric predictor such as Multilayer feedforward networks. Neural networks are known as universal approximators (Hornik et al. (1989) and Cybenko (1989)), and are very well suited to explore other learning methods. We combine them with Potts clustering models to create a bagging-boosting-averaging-like learning framework where several estimates from each random partition are aggregated into one prediction. Our approach carries out the balance between overfitting and model stability in presence of higher dimensional data. More precisely, our model merges Potts model in a *multivariate multiple regression* task with Bayesian deep learning models to produce the ensemble learning method. Potts model clustering has been introduced to the statistical community by Murua et al. (2008a). The model is applied to the set of covariate values, formalizing a proximity co-clustering. The method has been used and proven to be effective (see Murua & Wicker (2014); Murua & Quintana (2017)).

The model called *Structured Potts Shallow Gibbs Neural Network* will be a hierarchical Bayesian model where we train individual neural nets to specialize on sub-groups (latent clusters components) while still being informed about representations of the overall population. Our Potts neural network model differs from those of Kanter (1988) and Philipson & Cluitmans (1993), which is a generalization of the Ising neural network. We call it a structured one, because we integrate the structured correlations among the weights (and offsets) of the network (Sun et al., 2017) through a Markov Random Fields (MRF) process. Bayesian learning allows the opportunity to quantify posterior uncertainty on neural networks (NNs) model parameters. We can specify priors to inform and constrain our models and get structured uncertainty estimation.

1.1 Efficiency of regression clustering

Regression clustering is a learning algorithm that allows multiple regression settings to be clustered where you have a dependent Y vector. And one or two autonomous (independent) variables, the X 's. Recovering the implicit partitioning of observations is the problem of regression clustering: the algorithm then divides the data into two or more clusters and performs an independent multiple regression within each cluster on the data. Many algorithms have emerged to run this approach in the past (see Späth (1979) and H.-J. (1985)). The ultimate aim is to discover classes without oversight before regression is implemented in each class. Nowadays, many variations have been proposed in machine learning; clustered linear regression¹ (CLR) is an example (Ari & Güvenir, 2002).

Regression clustering has shown a significant predictive advantage in several tasks: clustered regression trees (Torgo & da Costa, 2000), ridge regression clustering (Nagpal et al., 2013), and non-parametric Bayesian clustering regression (see Müller et al. (2008), Yang et al. (2014), and Murua & Quintana (2017)). The combination of clustering and regression approaches will potentially minimise the possible problems of predictive performance efficacy due to data heterogeneity (Chen et al., 2013). Using a clustered method, the data subsets are created with a degree of homogeneity that improves the accuracy of prediction.

1.2 Combination of neural network regression and Potts clustering model

Potts clustering has the ability to integrate multiple partitions from different cluster solutions to improve the robustness, stability, scalability of the clustering algorithms (Murua et al., 2008b). As shown by Murua & Quintana (2017), Potts clustering can be a good prior for statistical models such as regression. Our model is a *multivariate multiple regression clustering* one, aiming at combining the Potts clustering model with a non-linear multivariate multiple regression tool: neural networks. The first main advantage of the combination of neural networks regression and Potts clustering models, is that a few of the drawbacks of linear regression can be overcome by using artificial neural network (ANN), and *clustered* models can improve prediction accuracy. Their combination presents other key advantages that make them most suitable for certain problems and situations:

1. Clustered linear regression (CLR) is a modern machine learning algorithm that, by partitioning training space into subspaces, increases the precision of classical linear regression.

1. An artificial neural network is capable of fitting and modelling non-linear and complicated relationships, which is really important since many of the relationships between inputs and outputs are both non-linear and complex in real life.
2. An artificial neural network is a good predictor: it can also infer unseen relations on unseen data after the necessary fitting step from the initial inputs and their relationships, thereby allowing the model to predict on unseen data.
3. Extension of Potts model to various forecasting problems: in addition to exploring a graph-based consensus clustering² (GCC) to find cluster structures from heterogeneous data, the model proposed here offer a novel opportunity to couple neural network model with random partition models in diverse machine learning tasks, such as multivariate multiple regression (precisely: when Y response is a **vector**). In fact, Potts clustering as used here, is a *random partition model*, explicitly, a clustering model with prior distribution on partitions (Müller & Quintana, 2010a). Section 3 describes Potts clustering in detail.

However, ANN is a black box learning technique that can not interpret input-output relationships and can not cope with uncertainties. Using the Bayesian framework as done by Murua & Quintana (2017) can overcome uncertainties issue, and the results can be easily compared. The *Bayesian* technique is highly important, as traditional neural network training involves a lot of labelled data to monitor the possibility of overfitting. And when it comes to real-world regression assignments, the task gets more complicated. Such exercises (regression) also have less training data to use, which also makes it easy for neural networks to get stuck in overfitting.

A principled approach for solving this problem is Bayesian Neural Networks (see Vehtari & Lampinen (1999), Bishop (1997)). Prior distributions are placed on the neural network weights in Bayesian Neural Networks to consider the model uncertainty. One can fit a predictor by doing Bayesian inference on the weights, which both matches the training data and knows about the volatility of its own estimation on the test data (Blundell et al., 2015).

2. Shallow Gibbs networks

Let $\mathcal{D} = \{(y_i, x_i) : i = 1, \dots, n\}$ denote the complete data, where $\mathbf{x}^n = \{x_1, \dots, x_n\} \subset \mathbb{R}^q$ is the set of input vectors or covariables, and $\mathbf{y}^n = \{y_1, \dots, y_n\} \subset \mathbb{R}^p$, the set of associated responses. The matrix of covariables will be denote by $\mathbf{X} = (x_1|x_2|\dots|x_n)^T \in \mathbb{M}_{n \times q}$ (here and throughtout the chapter, the symbol T will indicate matrix transposition). Our shallow network is, as in the general case, a feedforward network given by the equations

$$h^{(k)} = b^{(k)} + g_{k-1}(h^{(k-1)})W^{(k)}, \quad k = 0, 1, 2, \quad (1)$$

where the layer 0 corresponds to the input data $h^0 = x, b^{(0)} = 0, g_0(x) = x$, the first layer corresponds to the hidden layer with vector output $h^{(1)}$, and the output layer corresponds to the network's predicted values $h^{(2)} = \hat{y}$. The parameters $\{(b^{(k)}, W^{(k)}) : k = 1, 2\}$ are, respectively, the matrices of offsets, also known as biases, and weights. The vector $g_{k-1}(h^{(k-1)})$ denotes a linear ($k=2$) or non-linear ($k=1$) function, such as the identity or a logistic sigmoid, that is applied element-wise. The top layer output $h^{(2)}$ is used for making a prediction and is combined with the supervised target y into a loss function $L(h^{(2)}, y)$, which is typically convex in $h^{(2)} = b^{(2)} + g_1(h^{(1)})W^{(2)}$. As in any regression model, the network's predicted value $h^{(2)}$ is an estimate of $f(y) = \mathbb{E}(y|x)$. Our model may be seen as a simplified feedforward network where the weight matrices are constrained to obey a Gibbs distribution whose neighborhood structure is explained below. Figure 1 sketches our network architecture.

2. Potts clustering is based on a consensus clustering approach (Blatt et al., 1996a).

Figure 1: Shallow feedforward neural network.

Let $\psi = (b^{(1)}, W^{(1)}, b^{(2)}, W^{(2)})$ be the parameters of the network. Suppose that the hidden layer contains l_1 nodes. Then the parameter dimensions are as follows: $b^{(1)} \in \mathbb{R}^{l_1}$, $W^{(1)} \in \mathbb{M}_{q \times l_1}$, $b^{(2)} \in \mathbb{R}^{p}$, and $W^{(2)} \in \mathbb{M}_{l_1 \times p}$. To ease the exposure of our model, we will use the notation, $l_0 = q$, $l_2 = p$. Let $\Sigma \in \mathbb{M}_{p \times p}$ be the variance-covariance matrix of y . We suppose that $y|x, \psi, \Sigma$ is distributed as a multivariate normal distribution with mean $f(y) = f_\psi(y) = \mathbb{E}(y|x, \psi)$, and variance Σ . That is, $p(y|x, \psi, \Sigma) = (2\pi)^{-p/2} |\Sigma|^{-1/2} \exp\{-1/2(y - f_\psi(x))' \Sigma^{-1} (y - f_\psi(x))\}$.

Let $w^{(1)} = \text{vec}(W^{(1)})$ be the $q \times l_1$ -dimensional vector of stacked columns of the weight matrix $W^{(1)}$. The vector $w^{(2)} = \text{vec}(W^{(2)}) \in \mathbb{R}^{l_1 \times p}$ is defined similarly. Let $w = (w^{(1)}, w^{(2)})$ be the vectorized version of the weight matrices $(W^{(1)}, W^{(2)})$. Our model is set into a Bayesian framework by setting a Gaussian Markov random field prior for ψ . The complete model is the similar to the one described in (Sun et al., 2017). But ours presents significant differences in the actual architecture, because we set up two independent fields: one on all the network weights $w = (w^{(1)}, w^{(2)})$, and another one on all the network biases $b = (b^{(1)}, b^{(2)})$. We assume that both random fields are zero-mean Gaussian fields.

Let $\Omega \in \mathbb{M}_{(q+p)l_1 \times (q+p)l_1}$ and $\Gamma \in \mathbb{M}_{(l_1+p) \times (l_1+p)}$ be the precision matrices associated with the random field of the weights, and the biases, respectively. Here we assume that the biases are independent, so that Γ is a block-diagonal matrix with blocks $\Gamma_1 \in \mathbb{M}_{l_1 \times l_1}$, and $\Gamma_2 \in \mathbb{M}_{p \times p}$. The corresponding prior densities are given by

$$p(w|\Omega) = (2\pi)^{-\frac{l_1(p+q)}{2}} \det(\Omega)^{\frac{1}{2}} \exp\left(-\frac{1}{2} w^T \Omega w\right), \quad \text{for the weights, and} \quad (2)$$

$$p(b|\Gamma) = (2\pi)^{-\frac{l_1+p}{2}} \det(\Gamma_1)^{\frac{1}{2}} \det(\Gamma_2)^{\frac{1}{2}} \exp\left(-\frac{1}{2} \{(b^{(1)})^T \Gamma_1 b^{(1)} + (b^{(2)})^T \Gamma_2 b^{(2)}\}\right), \quad (3)$$

for the biases. Our framework for this part is completed by assuming an Inverse-Wishart(Λ, ν) prior for the covariance matrix Σ . Here $\Lambda \in \mathbb{M}_{p \times p}$ is the prior scale matrix, and $\nu > 0$ is the prior degrees of freedom.

2.1 The sparse-Gibbs network

We define the neighborhood of the random field based on the nodes of the hidden layer. All weights coming or going out of the same node are considered neighbor weights. That is, weights $W_{ij}^{(1)}$ and $W_{st}^{(2)}$ are neighbors if and

only if $j = s$, $j, s \in \{1, \dots, l_1\}$, and weights $W_{ij}^{(1)}$ and $W_{st}^{(1)}$ are neighbors if and only if $j = t$, $j, t \in \{1, \dots, l_1\}$, and weights $W_{ij}^{(2)}$ and $W_{st}^{(2)}$ are neighbors if and only if $j = t$, $j, t \in \{1, \dots, p\}$. The matrix Ω is sparse, and composed of l_1 blocks of size $q \times q$, p blocks of size $l_1 \times l_1$, and l_1 blocks of size $q \times p$ as well as the associated transposed blocks of size $p \times q$. Note that the total number of non-zero elements in Ω is $l_1(q^2 + p^2 + 2pq) = l_1(q + p)^2$, which corresponds to a sparsity ratio of $l_1(p + q)^2 / (l_1(q + p))^2 = 1/l_1$. This rate is very significant even for small values of l_1 .

Let $V_{i\bullet} \in \mathbb{M}_{q \times q}$ be the block of the precision matrix associated with column i of $W^{(1)}$, $i = 1, \dots, l_1$, $V_{\bullet j} \in \mathbb{M}_{l_1 \times l_1}$ be the block of the precision matrix associated with column j of $W^{(2)}$, $j = 1, \dots, p$, and $V_{ii} \in \mathbb{M}_{q \times p}$ be the block of the precision matrix associated with column i of $W^{(1)}$, and row i of $W^{(2)}$, $i = 1, \dots, l_1$. Let also $e_{i;l_1} \in \mathbb{M}_{1 \times l_1}$ be the i th row of the $l_1 \times l_1$ identity matrix. Then

$$\Omega = \begin{pmatrix} V_{1\bullet} & 0 & \cdots & 0 & V_{11} \otimes e_{1;l_1} \\ 0 & V_{2\bullet} & \cdots & 0 & V_{22} \otimes e_{2;l_1} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & V_{l_1\bullet} & V_{l_1 l_1} \otimes e_{l_1;l_1} \\ V_{11}^T \otimes e_{1;l_1}^T & \cdots & V_{l_1 l_1}^T \otimes e_{l_1;l_1}^T & V_{\bullet 1} & 0 \cdots 0 \\ 0 & V_{\bullet 2} & \cdots & 0 & V_{\bullet p} \end{pmatrix} \quad (4)$$

where the symbol $\otimes$ stands for the Kronecker product. We will refer to this simplified network as the sparse-Gibbs network, as opposed to the full Gibbs network that imposed no constraints on Ω .

2.2 Compound symmetry Gibbs network

Note that the above matrix consists of four blocks, that is $\Omega = \begin{pmatrix} \Omega_{11} & \Omega_{12} \\ \Omega_{12}^T & \Omega_{22} \end{pmatrix}$, with $\Omega_{11} \in \mathbb{M}_{ql_1 \times ql_1}$, $\Omega_{22} \in \mathbb{M}_{pl_1 \times pl_1}$, and $\Omega_{12} \in \mathbb{M}_{ql_1 \times pl_1}$. Also, if all matrices $V_{i\bullet}$ were equal, then we could write $\Omega_{11} = I_{l_1} \otimes V_{1\bullet}$, where I_r denotes the identity matrix in $\mathbb{M}_{r \times r}$. Similarly, if all matrices $V_{\bullet j}$ were the same, we could write $\Omega_{22} = I_p \otimes V_{\bullet 1}$. We could also suppose that all matrices V_{ii} are equal. These simplifications would greatly reduce the number of parameters defining the network.

A less stringent Gaussian random field model consists of replacing I_{l_1} and I_p in the above Kronecker products for more general matrices $U_1 \in \mathbb{M}_{l_1 \times l_1}$, $U_2 \in \mathbb{M}_{p \times p}$. This is the second model we consider. However, we simplify it by constraining the inverse matrices of U_1 , U_2 , $V_{1\bullet}$, and $V_{\bullet 1}$ to have compound symmetry. That is, the inverse matrices of these matrices have to be of the form $\kappa_1(1 - \kappa_2)I_{l_1} + \kappa_1\kappa_2\mathbb{1}\mathbb{1}^T$, for $\kappa_1 > 0$, and $\kappa_2 \in [-1, 1]$, where $\mathbb{1}$ stands for the column matrix of all elements equal to 1. It can easily be shown that this constraint implies matrices of the form $aI_{l_1} + b\mathbb{1}\mathbb{1}^T$, with $(a, b) \in \mathbb{R}^+ \times \mathbb{R}$. The matrices V_{ii} are not constrained. We will refer to this network architecture as sparse compound symmetry Gibbs network or sparse-CS-Gibbs network for short.

Two less structured models. More complex models can be conceived by (a) changing the neighborhood in the Gibbs network, and by allowing the diagonal blocks Ω_{11} and Ω_{22} to be unconstrained. In the Gibbs network this corresponds to consider that all weights coming from the same layer are neighbors, while still keeping the neighborhood restriction on weights coming to and going out of the same node. We will refer to this model as the between-layer sparse-Gibbs network. The second extended model correspond to the compound symmetry network in which the off-diagonal block Ω_{12} is set to be of the form $U \otimes V$, for matrices $U \in \mathbb{M}_{l_1 \times l_1}$, and $V \in \mathbb{M}_{q \times p}$, or, as we have actually implemented in our model, for matrices $U \in \mathbb{M}_{l_1 \times p}$, and $V \in \mathbb{M}_{q \times l_1}$. Note that the latter matrix structure is always sparser than the former one when $l_1 \in [0, \min\{p, q\}) \cup (\max\{p, q\}, +\infty)$. We will refer to this architecture as the compound-symmetry Gibbs network or CS-Gibbs for short.

3. The random-Potts partition model

The purpose of creating the sparse networks of the preceding section is to be able to train the networks with a fraction of the data used by a classical network. As mentioned earlier, our prediction model is a finite but very large mixture of shallow and sparse networks. This is the result of using a non-parametric Bayesian model for prediction. The networks are trained on a subset of data that share similar characteristics in their features. Classical Bayesian non-parametric models rely on the Dirichlet process. However, these processes do not necessarily look at the data features to create clusters or partitions of the data. Recent work has overcome this limitation Müller & Quintana (2010b); Müller et al. (2011). In particular, the random-Potts partition model introduced in Murua & Quintana (2017) stresses the importance of data similarities in the resulting partitions by guiding the random partition process through feature characteristics. These are incorporated in the model through a kernel which in turns induces the partition probabilities. The random-Potts partition model may be seen as a stochastic version of the label propagation³ approach (Tibély & Kertész, 2008).

The random-Potts partition model is a probabilistic model defined on a graph. Keeping the notation introduced before, each data point x_i defines a vertex in a graph whose edges are pre-specified. For example, Murua & Quintana (2017) consider a k -nearest-neighbor graph. That is, a graph for which there is an edge between data points x_i and x_j if and only if either x_j is one of the k most similar data points to x_i , or x_i is one of the k most similar data points to x_j . Associated with any data point in the graph there is a color label $z_i \in \{1, \dots, r\}$. Usually r is chosen to be between 10 and 30 (Blatt et al., 1996b; Stanberry et al., 2008). The set of color labels $\mathbf{z}^n = (z_1, \dots, z_n)$ follows a Potts model, and thus it has probability mass function (pmf)

$$p(\mathbf{z}^n | \mathbf{x}^n, \beta, \sigma) \propto \exp\left\{-\beta \sum_{i \sim j} J_{ij}(\sigma)(1 - \mathbf{I}[z_i = z_j])\right\},$$

where $J_{ij}(\sigma) = J(x_i, x_j; \sigma)$ is the kernel evaluated at (x_i, x_j) , $\mathbf{I}[\cdot] = 1$ is the condition between brackets is satisfied, and is zero otherwise, and where the notation $i \sim j$ means that x_i and x_j are neighbors. In the Physics literature the parameter β is referred to as the inverse-temperature. In practice, one chooses the kernel so as to penalize color labelings that do not assign the same color to very similar data points. For example, one can work with $J_{ij} = J_{ij}(\sigma) = J(\|x_i - x_j\|/\sigma)$, where $\|\cdot\|$ denotes the distance (e.g., Euclidean), and $\sigma > 0$ is a bandwidth parameter. The partitions are generated by iteratively deleting edges and inserting deleted edges at random. This is done through a set of random binary variables $\mathbf{b} = \{b_{ij}\}$ named the *bonds*. Let $p_{ij} = 1 - \exp(-\beta J_{ij}(\sigma)\mathbf{I}[z_i = z_j, i \sim j])$. Each bond b_{ij} is generated independently of the other bonds as a Bernoulli(p_{ij}). If $i \sim j$, the edge between x_i and x_j remains in the graph with probability p_{ij} (i.e., $b_{ij} = 1$), and it is deleted from the graph with probability $1 - p_{ij}$ (i.e., $b_{ij} = 0$). A random partition ρ_n is generated as the (randomly generated) connected components $\rho_n = \{S_1, \dots, S_{k_n}\}$ of this new random graph. In what follows, we will refer to the components of a partition as clusters. It can be shown that this procedure is governed by the random cluster pmf (Sokal, 1997)

$$p(\mathbf{b} | \mathbf{x}^n) = \left[\prod_{b_{ij}=1, i \sim j} p_{ij} \right] \left[\prod_{b_{ij}=0, i \sim j} (1 - p_{ij}) \right] r^{k_n}.$$

Colors are also randomly assigned once the bonds have been chosen. Each connected component is assigned a color uniformly at random. This process of sampling bonds and colors alternatively is known as the Swendsen-Wang algorithm. It is a special case of the Markov chain Monte Carlo (MCMC) stochastic algorithm. Its goal is

3. A fast algorithm for finding communities in a graph is the Mark Propagation Algorithm (LPA). It defines certain communities as its guide using the network structure alone, and does not require a predefined objective feature or prior group knowledge. LPA is a relatively new algorithm, and was only proposed by Raghavan et al. (2007), in "Near linear time algorithm to detect community structures in large-scale networks". It functions by distributing labels across the network and building communities focused on this label dissemination mechanism. There are some references online : Xie & Szymanski (2011), Xie & Szymanski (2013), Gregory (2010). The paper of Tibély and Kertész (2008) demonstrates an equivalence of the label propagation method of community detection and Potts model approach.

to generate samples from the random-Potts partition model. This is achieved when the two sampling steps, bonds and colors, have been applied a sufficiently large number of times. The distribution of the random-Potts partitions is given by

$$p(\rho_n | \mathbf{x}^n) = \sum_{\mathbf{b} \Rightarrow \rho_n} p(\mathbf{b} | \mathbf{x}^n) \quad (5)$$

where the notation $\mathbf{b} \Rightarrow \rho_n$ means that the partition ρ_n arises from the connected components of the associated graph with edges given by the bonds $\mathbf{b}$ (see (Murua & Quintana, 2017) for more in-depth overview). The calculation in the above equation 5 might be intractable. Fortunately, we can sample partitions with MCMC methods without knowing these quantities exactly. These will be detailed below in Section 6.

3.1 Some practical considerations

For our experiments of Section 12, we use the Gaussian kernel $J_{ij}(\sigma) = \exp(-\frac{1}{2\sigma^2}\|x_i - x_j\|^2)$, which is the most popular kernel choice for the Potts or super-paramagnetic clustering model (Blatt et al., 1996a,b; Murua et al., 2008a). Although the scale parameter σ may be estimated through a Bayesian stochastic procedure (Murua & Wicker, 2014), we prefer to use its common estimator which is given by the average distances expressed as $\hat{\sigma}^2 = \sum_{i < j} \|x_i - x_j\|^2 / \binom{n}{2}$. Murua & Wicker (2014) show that the optimal scale is close to this simple estimator. To simplify the computations involved with the Potts model, we restricted the data neighborhood to a k -nearest-neighbors graph; that is, only the k -nearest-neighbors of a given point x_i are considered as neighbors of x_i . This implies that we set $J_{ij}(\sigma) = \exp(-\frac{1}{2\sigma^2}\|x_i - x_j\|^2)$, if x_j is one of the k -nearest-neighbors of x_i , or if x_i is one of the k -nearest-neighbors of x_j , and set $J_{ij}(\sigma) = 0$, otherwise.

One of the main issues encountered with the Potts model is the choice of the inverse temperature parameter β . This parameter controls the cluster sizes in the partitions. A value too low of β produces too many small clusters, while a value too high produces very few and large clusters. Because a shallow Gibbs network is going to be fitted for each cluster of the partition, we would rather not have too small size clusters. Therefore, a small value of β is preferred. In our experiments, for each dataset, we selected a value of β that produced large enough clusters. In general, Murua & Wicker (2014) gave a simple procedure to find a nearly optimal value of β . In our case, we just chose a value slightly smaller than the value suggested so as to ensure large cluster sizes.

3.2 The Potts Clustering Model with Complete Shrinkage

The drawback of the bonds approach is the increasing number of small clusters generated in a given partition. To deal with it, we apply a modified agglomerative clustering approach (Kurita, 1991) by merging all small clusters of size $\leq \mathbf{h}$ with their closest cluster in terms of minimal distance respectively, where $\mathbf{h}$ is an integer greater or equal to 2. The algorithm uses a technique in which distances of all pairs of observations are stored. Then the nearest cluster (with size $\geq \mathbf{h}$) is given by the cluster with the closest node in terms of minimal distance to the cluster to be merged. This is what we call Potts Clustering with Complete Shrinkage (PCCS).

4. The shallow Potts Gibbs-network mixture model

Combining the shallow Gibbs networks introduced in Section 2 with the random-Potts partition model described in the previous section we obtained the shallow Potts Gibbs-network mixture model. This model will be referred to as Potts Gibbs-network for short, since ‘‘Potts’’ already conveys the idea of mixture.

Keeping the notation introduced in the previous section, ρ_n denotes a partition of the data, $\psi_\sigma = (w, b, \Sigma)$ denotes the parameters of a shallow network, and $(\Omega, \Gamma, \Lambda, \nu)$ denotes the parameters defining the prior on ψ_σ . For a given partition $\rho_n = (S_1, \dots, S_{k_n})$, the variables $s_1, \dots, s_n$ will denote the cluster membership of the data points. That is, $s_i = \ell$ if and only if $x_i \in S_\ell$, $i = 1, \dots, n$, $\ell = 1, \dots, k_n$. Note that k_n , the number of clusters or connected

components of ρ_n , is a random variable. Proceeding as in (Murua & Quintana, 2017), we assume a hierarchical random partition model for the data

$$y_1, \dots, y_n | \mathbf{x}^n, \rho_n, \psi_{\sigma 1}^*, \dots, \psi_{\sigma k_n}^* \stackrel{\text{ind}}{\sim} p(y_i | x_i, \psi_{\sigma s_i}^*) \quad (6)$$

$$\psi_{\sigma 1}^*, \dots, \psi_{\sigma k_n}^* \stackrel{\text{ind}}{\sim} p(\psi_\sigma | \Omega, \Gamma, \Lambda, \nu), \text{ and } \rho_n | \mathbf{x}^n \sim p(\rho_n | \mathbf{x}^n), \quad (7)$$

where $(\psi_{\sigma 1}^*, \dots, \psi_{\sigma k_n}^*)$ denote the unique values of $(\psi_{\sigma 1}, \dots, \psi_{\sigma n})$ given the k_n -cluster partition ρ_n . Thus, $\psi_{\sigma i}^*$ are the parameters of the shallow Gibbs network associated with the i -th data cluster S_i of ρ_n . The notation $\stackrel{\text{ind}}{\sim}$ indicates that the variables are independent.

As suggested in (Murua & Quintana, 2017, Section 2.2), avoiding reversible jumps moves (Green, 1995), and following Besag (2004), we can sample from the posterior of ψ_σ by MCMC using these two moves:

(A) Sample a proposed partition $\rho'_n \sim p(\rho'_n | \mathbf{x}^n)$, and evaluate

$$\alpha = p(\mathbf{y}^n | \rho'_n, \mathbf{x}^n) / p(\mathbf{y}^n | \rho_n, \mathbf{x}^n); \quad (8)$$

if $\alpha \geq 1$, accept the proposed partition; otherwise, accept the proposed partition with probability α .

(B) If the proposed partition is accepted, sample (and accept) a new set of parameters $\psi_{\sigma'}^* = (\psi'_{\sigma 1}, \dots, \psi'_{\sigma k'_n})$ from the posterior given the partition $p(\psi_\sigma | \mathbf{y}^n, \mathbf{x}^n, \rho'_n)$, where $\rho'_n = (S'_1, \dots, S'_{k'_n})$. If the proposed partition is not accepted, then keep the current partition ρ_n and the corresponding parameters $\psi_\sigma^* = (\psi_{\sigma 1}, \dots, \psi_{\sigma k_n})$.

We will refer to the above steps as Besag's algorithm. To use this scheme, we need to be able (i) to sample from the posteriors $p(\psi_\sigma | \mathbf{y}^n, \mathbf{x}^n, \rho_n)$, and (ii) to compute the ratio α . In general, the probabilities $p(\mathbf{y}^n | \rho'_n, \mathbf{x}^n)$ are intractable. In fact, computing α is equivalent to computing Bayes factors. This algorithm works exactly only for certain problems such as the multivariate multiple regression with conjugate priors used in (Murua & Quintana, 2017). In fact, if $p(\psi_\sigma | \mathbf{y}^n, \mathbf{x}^n, \rho_n)$ is known, we can compute the marginal probabilities using the identity

$$p(\mathbf{y}^n | \rho_n, \mathbf{x}^n) = p(\mathbf{y}^n | \psi_\sigma, \rho_n, \mathbf{x}^n) p(\psi_\sigma | \rho_n) / p(\psi_\sigma | \mathbf{y}^n, \mathbf{x}^n, \rho_n). \quad (9)$$

In our case, the posteriors $p(\psi_{\sigma j} | \mathbf{y}^n, \mathbf{x}^n, \rho_n)$ are proportional to

$$p(T_j | S_j, \psi_{\sigma j}) p(w_j, b_j, \Sigma_j | \Omega, \Gamma, \Lambda, \nu) = \prod_{x_i \in S_j} p(y_i | S_j, \psi_{\sigma j}) p(w_j, b_j, \Sigma_j | \Omega, \Gamma, \Lambda, \nu),$$

where $T_j = \{y_i : x_i \in S_j\}$. The normalizing constants of these posteriors are not known because of the non-linearity in the neural network likelihood $p(y_i | S_j, \psi_{\sigma j})$. To avoid integrating over the posterior distribution so as to compute the normalizing constants, it is common to use Markov Chain Monte Carlo (MCMC) methods (Neal, 1996; Rasmussen, 1995). Another line of research for posterior inference uses stochastic gradient Markov Chain Monte Carlo (Chen et al., 2015). However, one major drawback of sampling, is that it is often very slow, especially for high-dimensional models (Robert & Casella, 2005). Faster and easy to parallelize methods may be drawn from Bayesian variational inference algorithms. These have been developed to be as flexible as MCMC (Zhai et al., 2012). Therefore, this is the route we follow to overcome the computational burden of posterior and Bayes factors calculation.

5. Bayesian variational inference

In this section we explain how we find a Bayesian variational approximation of the posterior distributions. We choose a family of distributions of the model parameters $\psi_\sigma = (w, b, \Sigma)$ indexed by a hyper-parameter λ , $Q(\psi_\sigma; \lambda)$. The idea is to approximate the posterior of ψ_σ with its closest distribution from this family. This corresponds to an optimization over the hyper-parameter λ . In this section, we suppose that we work with a given

cluster S of the partition ρ_n , and its associated set of response variables T . The core idea is to minimize over λ the Kullback-Leibler divergence (KL) between the posterior and $Q(\psi_\sigma; \lambda)$. That is, we need to minimize

$$\text{KL}(Q||p) = \int Q(\psi_\sigma; \lambda) \log\{Q(\psi_\sigma; \lambda)/p(\psi_\sigma|T, S, \rho_n)\}.$$

The optimal value λ_{opt} gives rise to the approximating *variational posterior distribution* $Q(\psi_\sigma; \lambda_{opt})$.

Approximated Bayes factors and Besag's algorithm. In practice, one obtains k_n such variational posteriors $Q(\psi_{\sigma j}; \lambda_j)$, $j = 1, \dots, k_n$ for a given partition ρ_n . These are the distributions that we use to replaced the intractable posteriors $p(\psi_\sigma|\mathbf{y}^n, \mathbf{x}^n, \rho_n)$ in:

- (i) STEP B of Besag's algorithm (see previous section) to sample the parameters ψ_σ ; and
- (ii) STEP A, to estimate the Bayes factor α . More specifically, let, respectively, $\hat{\psi}_\sigma = (\hat{\psi}_{\sigma 1}, \dots, \hat{\psi}_{\sigma k_n})$ and $\hat{\psi}'_\sigma = (\hat{\psi}'_{\sigma 1}, \dots, \hat{\psi}'_{\sigma k_n})$ be the maximum a posteriori (MAP) estimators of ψ_σ and ψ'_σ . We use the estimate

$$\hat{\alpha} = \frac{p(\mathbf{y}^n|\hat{\psi}'_\sigma, \rho'_n, \mathbf{x}^n)p(\hat{\psi}'_\sigma|\rho'_n) \prod_{j=1}^{k_n} Q(\hat{\psi}_{\sigma j}|\lambda_j)}{p(\mathbf{y}^n|\hat{\psi}_\sigma, \rho_n, \mathbf{x}^n)p(\hat{\psi}_\sigma|\rho_n) \prod_{j=1}^{k_n} Q(\hat{\psi}'_{\sigma j}|\lambda'_j)}.$$

The ELBO. Note that the optimal value is the solution of the optimization problem

$$\begin{aligned} \lambda_{opt} &= \arg \min_{\lambda} \text{KL}(Q(\psi_\sigma; \lambda)||p(\psi_\sigma|T, S, \rho_n)) \\ &= \arg \min_{\lambda} \int Q(\psi_\sigma; \lambda) \log \frac{Q(\psi_\sigma; \lambda)}{p(\psi_\sigma|\rho_n)p(T|\psi_\sigma, S, \rho_n)} d\psi_\sigma \\ &= \arg \max_{\lambda} \mathbb{E}_{Q(\psi_\sigma; \lambda)} (\log p(T|\psi_\sigma, S, \rho_n)) - \text{KL}(Q(\psi_\sigma; \lambda)||p(\psi_\sigma|\rho_n)). \end{aligned} \quad (10)$$

where the term to be maximized is known as the *evidence lower bound* or ELBO:

$$\text{ELBO}(\lambda) = \mathbb{E}_{Q(\psi_\sigma; \lambda)} (\log p(T|\psi_\sigma, S, \rho_n)), - \text{KL}(Q(\psi_\sigma; \lambda)||p(\psi_\sigma|\rho_n)) \quad (11)$$

The negative of this quantity is also known as the description length cost, or the variational free energy (see, for example, Friston et al. (2007)). The direct optimisation of this quantity instead of the Kullback-Leibler divergence was introduced by Hinton & van Camp (1993) (see also, Graves (2011)). In general, the ELBO cannot be computed exactly. The expectations are usually approximated through Monte Carlo methods (Blundell et al., 2015).

Choice of variational family. If the likelihood $p(T|\psi_\sigma, S, \rho_n)$ were linear in the weight and bias parameters, the posterior of the parameters would be conjugate to the prior (Murua & Quintana, 2017). Inspired by this observation, we consider $\lambda = (\lambda_w, \lambda_b, \lambda_\sigma)$ and the family of conjugate distributions to the prior

$$Q(\psi_\sigma; \lambda) = Q(w; \lambda_w)Q(b; \lambda_b)Q(\Sigma; \lambda_\sigma).$$

More specifically, we suppose that $Q(w; \lambda_w)$ is of the same form as the prior $p(w|\Omega)$. That is, $Q(w; \lambda_w)$ is a zero-mean Gaussian random field with covariance structure dictated by the Gibbs network induced by Ω . Thus, if the Gibbs field given by the prior is a compound symmetry Gibbs network, then so it is $Q(w; \lambda_w)$, etc. Similarly, $Q(b; \lambda_b)$ is a zero-mean Gaussian field with block diagonal covariance $\lambda_b = \text{diag}(\lambda_{b1}, \lambda_{b2})$ just as Γ . Also, just as the prior, the variational posterior approximation $Q(\Sigma; \lambda_\sigma)$ is assumed to be an Inverse-Wishart($\nu_\sigma, \Lambda_\sigma$). With this family of distributions, the ELBO becomes

$$\begin{aligned} \text{ELBO}(\lambda) &= \mathbb{E}_{Q(\psi_\sigma; \lambda)} (\log p(T|\psi_\sigma, S, \rho_n)) - \text{KL}(Q(w; \lambda_w)||p(w|\rho_n)) \\ &\quad - \text{KL}(Q(b; \lambda_b)||p(b|\rho_n)) - \text{KL}(Q(\Sigma; \lambda_\sigma)||p(\Sigma|\rho_n)) \end{aligned} \quad (12)$$

It is not difficult to show that the Kullback-Leibler divergences between the two sets of Gaussian distributions are

$$\begin{aligned}\text{KL}(Q(w; \lambda_w) \| p(w | \rho_n)) &= \frac{1}{2} \left\{ \text{trace}(\Omega^{-1} \lambda_w) - (q + p) l_1 + \log(\det(\Omega) / \det(\lambda_w)) \right\}, \\ \text{KL}(Q(b; \lambda_b) \| p(b | \rho_n)) &= \frac{1}{2} \left\{ \text{trace}(\Gamma^{-1} \lambda_b) - (l_1 + p) + \log(\det(\Gamma) / \det(\lambda_b)) \right\}.\end{aligned}$$

For the last divergence we have the following result.

Proposition 5.1

The Kullback-Leibler divergence $\text{KL}(Q(\Sigma; \lambda_\sigma) \| p(\Sigma | \rho_n))$ between the two inverse-Wisharts $Q(\Sigma; \lambda_\sigma)$ and $p(\Sigma | \rho_n)$, is given by

$$\frac{\nu}{2} \log \left(\frac{\det(\Lambda_\sigma)}{\det(\Lambda)} \right) + \frac{\nu_\sigma}{2} \text{trace}(\Lambda \Lambda_\sigma^{-1}) - (\nu - \nu_\sigma) \frac{\psi_p(\nu_\sigma/2)}{2} - \nu_\sigma \frac{p}{2} + \log \left(\frac{\Gamma_p(\nu/2)}{\Gamma_p(\nu_\sigma/2)} \right), \quad (13)$$

where $\Gamma_p(\cdot)$ stands for the p -multivariate gamma function, and $\psi_p(\cdot)$ denotes the derivative of the p -multivariate gamma function (i.e., the multivariate digamma function).

Proof: By definition $\text{KL}(Q(\Sigma; \lambda_\sigma) \| p(\Sigma | \rho_n))$ is equal to

$$\begin{aligned}& - \frac{(\nu_\sigma + p + 1)}{2} \mathbb{E}_{Q(\Sigma; \lambda_\sigma)}(\log \det(X)) - \frac{1}{2} \text{trace}(\Lambda_\sigma \mathbb{E}_{Q(\Sigma; \lambda_\sigma)}(X^{-1})) + \frac{\nu_\sigma}{2} \log \det(\Lambda_\sigma) \\& \quad - \frac{\nu_\sigma p}{2} \log(2) - \log(\Gamma_p(\nu_\sigma/2)) + \frac{(\nu + p + 1)}{2} \mathbb{E}_{Q(\Sigma; \lambda_\sigma)}(\log \det(X)) \\& \quad + \frac{1}{2} \text{trace}(\Lambda \mathbb{E}_{Q(\Sigma; \lambda_\sigma)}(X^{-1})) - \frac{\nu}{2} \log \det(\Lambda) + \frac{\nu p}{2} \log(2) + \log(\Gamma_p(\nu/2)) \\& = \frac{(\nu - \nu_\sigma)}{2} \mathbb{E}_{Q(\Sigma; \lambda_\sigma)}(\log \det(X)) + \frac{1}{2} \text{trace}([\Lambda - \Lambda_\sigma] \mathbb{E}_{Q(\Sigma; \lambda_\sigma)}(X^{-1})) \\& \quad + \frac{\nu_\sigma}{2} \log \det(\Lambda_\sigma) - \frac{\nu}{2} \log \det(\Lambda) + (\nu - \nu_\sigma) \frac{p \log(2)}{2} + \log \left(\frac{\Gamma_p(\nu/2)}{\Gamma_p(\nu_\sigma/2)} \right),\end{aligned}$$

where X denotes a positive definite matrix. Note that when $X \sim \text{Inverse-Wishart}(\nu_\sigma, \Lambda_\sigma)$, its inverse $X^{-1} \sim \text{Wishart}(\nu_\sigma, \Lambda_\sigma^{-1})$. Therefore,

$$\begin{aligned}\mathbb{E}_{Q(\Sigma; \lambda_\sigma)}(X^{-1}) &= \nu_\sigma \Lambda_\sigma^{-1}, \text{ and (Bishop, 2006, pp. 693)} \\ \mathbb{E}_{Q(\Sigma; \lambda_\sigma)}(\log \det(X)) &= -\psi_p(p/2) - p \log(2) - \log \det(\Lambda_\sigma^{-1}).\end{aligned}$$

Putting all together, we obtain the desired result (13). ■

The optimal solution is found using the stochastic gradient ascent variational Bayes algorithm (Paisley et al., 2012; Kucukelbir et al., 2014; Ye et al., 2020; Duchi et al., 2011). To obtain the gradient of $\text{ELBO}(\lambda)$, the gradients of the Kullback-Leibler divergence terms are needed. In practice, these derivatives are usually computed through automatic differentiation. That is, calculating numerically the value of the derivatives (e.g., using Chebyshev polynomial approximation (Press et al., 1996, Ch. 5.7)). These methods achieve very good accuracies (Bartholomew-Biggs et al., 2000).

The integrated log-likelihood, the first term in (12), is intractable. We use Monte Carlo sampling to estimate it. However, we only need to estimate the score (Kingma et al., 2015; Mohamed et al., 2019),

$$\frac{\partial}{\partial \lambda} \left(\mathbb{E}_{Q(\psi_\sigma; \lambda)}(\log p(T | \psi_\sigma, S, \rho_n)) \right) = \mathbb{E}_{Q(\psi_\sigma; \lambda)} \left(\log p(T | \psi_\sigma, S, \rho_n) \frac{\partial Q(\psi_\sigma; \lambda)}{\partial \lambda} \right).$$

The integral in the right-hand-side of the above equation is estimated by Monte Carlo sampling of ψ_σ from $Q(\psi_\sigma; \lambda)$, which is a multivariate Gaussian-inverse-Wishart distribution.

5.1 Regularization on the CS-Gibbs model

For the compound symmetry Gibbs network model described at the end of in Section 2.2, the off-diagonal block $\Omega_{12} = U \otimes V$, for $U \in \mathbb{M}_{l_1 \times p}$, and $V \in \mathbb{M}_{q \times l_1}$. These matrices are unconstrained. For the variational approximation, we suppose that the corresponding block of the precision matrix has the same structure. Let $U_{12} \in \mathbb{M}_{l_1 \times p}$, and $V_{12} \in \mathbb{M}_{q \times l_1}$ be the corresponding matrices in the variational approximation.

Note that the our prior for the weights correspond to a Gaussian graphical model (Giudici & Green, 1999; Rajaratnam et al., 2008). The usual procedure to find sparse covariance models for Gaussian random fields is the so-called graphical lasso (Friedman et al., 2008). This method is a penalized log-likelihood in which the size of the entries of Ω are penalized, so as to render this matrix sparse. Translating this method to our already sparse model leads to the penalization of the L_1 -norm of the matrices U_{12} and V_{12} . That is, we consider the maximization of the penalized log-likelihood

$$\log p(T|\psi_\sigma, S, \rho_n) - \beta_u \|U_{12}\|_1 - \beta_v \|V_{12}\|_1,$$

where β_u and β_v are the regularization constants. Introducing the variational approximation, we have

$$\begin{aligned} \mathbb{E}_{Q(\psi_\sigma; \lambda)} (\log p(T|\psi_\sigma, S, \rho_n) - \beta_u \|U_{12}\|_1 - \beta_v \|V_{12}\|_1) \\ = \text{ELBO}(\lambda) + \text{KL}(Q(\psi_\sigma; \lambda) \| p(\psi_\sigma | T, S, \rho_n)) - \beta_u \|U_{12}\|_1 - \beta_v \|V_{12}\|_1. \end{aligned}$$

Because $\text{KL}(Q(\psi_\sigma; \lambda) \| p(\psi_\sigma | T, S, \rho_n)) \geq 0$, we have

$$\text{ELBO}(\lambda) - \beta_u \|U_{12}\|_1 - \beta_v \|V_{12}\|_1 \leq \mathbb{E}_{Q(\psi_\sigma; \lambda)} (\log p(T|\psi_\sigma, S, \rho_n)) - \beta_u \|U_{12}\|_1 - \beta_v \|V_{12}\|_1,$$

and we can think of the right-hand-side of this expression as a lower bound for the expected graphical-lasso log-likelihood function. The goal is then to solve the maximization problem

$$\lambda^* = \arg \max_{\lambda} \{ \text{ELBO}(\lambda) - \beta_u \|U_{12}\|_1 - \beta_v \|V_{12}\|_1 \}$$

This result may also be embedded in a special case of sparse variational inference (Campbell & Beronov, 2019) or *Generalized ELBO with Constrained Optimization* (Rezende & Viola, 2018), where we build a regularized framework for our compound symmetry Gibbs network in order to set sparsity constraints.

The optimization problem is not straightforward to solve. First, the L_1 -norm is non differentiable. And more importantly, the solution for the covariance matrix must still be a positive definite matrix. There are several algorithms designed specially to solve the usual graphical lasso model (Mazumder & Hastie, 2015). But, they are not directly applicable to our model. We have addressed these two issues. In this section, we only explain how we have modified the maximization problem for this special case of our model. The problem dealing with the positive-definiteness of the matrix applies to all our models. It is addressed in the next section.

To deal with the non-differentiability of the L_1 norm, we have implemented two modifications to the original problem. The first one consists of replacing the absolute values of the elements of the two matrices $|u|$ by a smooth version of them. These are given by a smooth approximation (Schmidt et al., 2007) of the absolute value $|u|_\alpha = (1/\alpha) \{ \log(1 + \exp(-\alpha u)) + \log(1 + \exp(\alpha u)) \}$, for a given constant α . It can be easily seen that as the value of α increases $|u|_\alpha$ approaches $|u|$. In fact, Schmidt et al. (2007) show that $||u| - |u|_\alpha| \leq \frac{2 \log(2)}{\alpha}$. Therefore, the constant α must be fixed to a large enough value (a value of about 10^6 is large enough as suggested in (Schmidt et al., 2007), but smaller values also give good approximations). Our second modification is to replace the L_1 -norm by the Frobenius (that is, Euclidean) norm $\|\cdot\|$ of the matrices U_{12} and V_{12} . This modification can be thought of as a group-lasso version of the graphical lasso. Unfortunately, the Frobenius norm is also non-differentiable. Instead, we consider the squared of the Frobenius norm, which is a smooth function. In particular, it is easily seen that its derivative is given by $\nabla \|U\|_F^2 = \frac{\partial}{\partial U} \text{trace}(U^T U) = 2U$.

The parameters β_u and β_v must be chosen adequately. We perform K -fold cross-validation to establish appropriate values for them. We stress that this is done during the fitting (training) stage of the model for each cluster of the partitions ρ_n . Because it is not possible to do cross-validation over the entire space of possible values, we set a sequence of possible values for $\beta = (\beta_u, \beta_v)$. Let $\{\beta^{(m)} : m = 1, \dots, M\}$ be this sequence. The procedure is as follows: Divide the training data (T, S) into K disjoint approximately equal size subsets $\mathcal{D}_\ell$, $\ell = 1, \dots, K$. For each $m \in \{1, \dots, M\}$, and $\ell \in \{1, \dots, K\}$, set $\mathcal{D}_{-\ell} = \cup_{j \neq \ell} \mathcal{D}_j$, and fit the model with the data in $\mathcal{D}_{-\ell}$ and parameters $\beta^{(m)}$. Let $\{\hat{y}_i^{(\ell, m)} : y_i \in T\}$ be the adjusted values of the model. Compute $\text{CV}(\beta^{(m)}) = \sum_{\ell=1}^K \sum_{y_i \in \mathcal{D}_\ell} \|y_i - \hat{y}_i^{(\ell, m)}\|^2$. Let $\hat{\beta} \in \{\beta^{(m)} : m = 1, \dots, M\}$ be the parameter values that minimizes $\text{CV}(\beta^{(m)})$. Fit the model with all data (T, S) and parameters $\hat{\beta}$. In our implementation we have use $M \geq 100$. The β values of the sequence were chosen automatically by the two-dimensional Golden-Section Search (GSS) algorithm (Chang, 2009). To initialize the search, several random values for each component of the vector β were drawn. Each draw is a value chosen uniformly at random from one of the intervals in the collection $\{[\epsilon + 100\ell, \epsilon + 100(\ell+1)] : \ell = 0, 1, \dots, \ell_{\max}\}$, where, in practice, $\ell_{\max} \leq 100$, and $\epsilon \leq 1/10000$. For each draw, the interval was also chosen uniformly at random. Technically, the GSS is based on the fact that the minimum lies within the interval defined by the two points adjacent to the last observed value. The method operates by successively narrowing the range of values on the specified search space, which makes it relatively slow, but very robust. It finds the best extrema $(\beta_u^{\text{opt}}, \beta_v^{\text{opt}})$ after going through all the regions. GSS presents an oracle complexity which converges to an ϵ -accurate solution at a linear rate, also known as geometric or exponential convergence in the numerical literature (Bertolazzi, 2008; Sebastien, 2013). In fact, GSS makes $O(\log(1/\epsilon))$ calls to an *Oracle Query Optimizer* (Wijesiriwardana & Firdhous, 2019) to compute the optimum.

5.2 Keeping positive definiteness on the precision Matrix

Let λ_w be the precision matrix associated with the variational density approximation to the posterior of the weight parameters w . The gradient of the ELBO with respect to the variational precision matrix λ_w , denoted $\nabla \text{ELBO}(\lambda_w)$, is computed component-wise or block-wise in all sparse structures. For example, in the Sparse-Gibbs network where the precision matrix λ_w has the form given by (4), we compute each gradient $\nabla \text{ELBO}(V_{i\bullet})$, $\nabla \text{ELBO}(V_{\bullet j})$, and $\nabla \text{ELBO}(V_{ii})$ separately in order to build the full update of λ_w . The same differentiation principle is applied for the sparse compound symmetry Gibbs model (sparse-CS-Gibbs network) where we constrain $U_1, U_2, V_{1\bullet}$, and $V_{\bullet 1}$ to have compound symmetry.

One of the issues that arised on the estimation of λ_w is how to ensure in practice that the matrix is kept positive definite during the gradient updates. Contrary to graphical lasso (Mazumder & Hastie, 2015), the positive definiteness property can be lost. There is no known algorithm that can ensure this property for our models while optimizing the matrix. To find a very good or optimal solution might require heavy exploration of the matrix space at each gradient update. Hence, it might be impractical to optimize the ELBO over λ_w with simultaneous constraints on both the sparse structure, and the positive definiteness property. Given the complexity of the model, and from a statistical viewpoint, it might be sufficient to obtain an approximate solution for λ_w that is both sparse as imposed, and positive definite. Therefore, betting on this observation, we based our matrix updates on the search for the nearest positive definite matrix (Higham, 1988). A simple procedure to find the nearest positive definite matrix $\hat{\lambda}_w$, where proximity is measure by the Frobenius norm, is based on the spectral decomposition of the matrix (Jewbali & Ore, 2009). We first determine the spectral decomposition of $\lambda_w = QDQ^T$, where Q is an unitary matrix, and D is a diagonal matrix with the eigenvalues of λ_w . An estimate of the nearest-positive definite matrix to λ_w is given by $\check{\lambda}_w = Q\check{D}Q^T$, where $\check{D}$ is the diagonal matrix D modified with all negative eigenvalues set to a small positive constant. However, $\check{\lambda}_w$ might not be the nearest positive definite matrix. The optimal small constant needs to be searched. One algorithm that does this search is the Iterative Spectral Method (Marée, 2012). Unfortunately, this and variants of this algorithm are costly, taking order $O(d^3)$ operations for a $d \times d$ matrix.

Two practical ways to verify if the above approximation is necessary, are, first to try a Cholesky decomposition of λ_w after each update, or to apply the well-known Sylvester’s criterion (Gilbert, 1991) to each update. The latter verifies if all leading principal minors of λ_w are positive. This criterion is doubly usefull, since it also helps to build the nearest symmetric positive definite matrix, with the imposed model structure. The minors can be used for non-positive definiteness correction. In fact, the k th principal minor is computed based on the $(k - 1)$ th principal minor augmented with the corresponding reduced k th column and row. If the $(k - 1)$ th principal minor is positive, and the k th principal minor is not, then one can make this latter positive with some slight modifications to the k th column and/or row. This can be done sequentially until all principal minors has been computed. However, one drawback of this method, besides its computational cost, is that it does not always guarantee to get the nearest positive definite matrix in terms of the Frobenius norm.

In our implementation, we use the *nearest Positive Definite Matrix computation*, which is mathematically equivalent to the gradient projection method known as *Projected gradient updates* (Cruz et al., 2011; Hassani et al., 2017). The combination of the gradient updates and the projection into the nearest positive definite matrix space is called *Iterative projected gradient* (IPG). This iterates between calculating the gradient and calculating the projection into the positive definite matrix space $\mathcal{S}_+$. More explicitly, at iteration k , the IPG computes

$$\lambda_w^{(k)} = \mathcal{P}_{\mathcal{S}_+} \left(\lambda_w^{(k-1)} + \alpha_k \nabla \text{ELBO} \left(\lambda_w^{(k-1)} \right) \right) \quad (14)$$

where, α_k is the step size, and $\mathcal{P}_{\mathcal{S}_+}$ denotes the Euclidean projection into the positive definite matrix space, that is,

$$\mathcal{P}_{\mathcal{S}_+}(\lambda_w) = \operatorname{argmin}_{A \in \mathcal{S}_+} \|A - \lambda_w\|_F$$

where, $\|\cdot\|_F$ stands for the Frobenius norm. To compute $\mathcal{P}_{\mathcal{S}_+}(\lambda_w)$, we have applied an approximation by matrices positive semi-definite on the subspace $\mathcal{S}_+$ as detailed in (Hayden & Wells, 1988). From properties of the Frobenius norm, we have that:

$$\|A - \lambda_w\|_F^2 = \|B - A + C\|_F^2 = \|B - \lambda_w\|_F^2 + \|C\|_F^2$$

with :

$$B = \frac{\lambda_w + \lambda_w^T}{2} \quad \text{and} \quad C = \frac{\lambda_w - \lambda_w^T}{2}$$

Thus we minimize $\|A - \lambda_w\|_F$ by minimizing $\|B - A\|_F$. Now, our approximation in the subspace $\mathcal{S}_+$, goes as it follows:

- Take $A \leftarrow B = \frac{\lambda_w + \lambda_w^T}{2}$. This is a transformation technique to force A to have the same properties as the symmetric matrix B .
- Compute the spectrale decomposition of A , let say $A = U \Lambda U^T$ where U is orthogonal and $\Lambda = \text{diag} [\lambda_1, \lambda_2, \dots, \lambda_n]$. Then the unique best approximation A^+ to A is given by:

$$A^+ = U \Lambda_+ U^T$$

where Λ_+ is obtained from Λ by replacing each negative eigenvalue by a number $a > 0$. We denote the new matrix as $A(a)$. So if there are no negative eigenvalues, A is taken as our approximation.

- If A as negative eigenvalues, we generate multiple values⁴ of a , to find a_{opt} , the optimal value of a .

$$\mathcal{P}_{\mathcal{S}_+}(\lambda_w) = \operatorname{argmin}_{A \in \mathcal{S}_+} \|A - \lambda_w\|_F = A(a_{opt})$$

4. In Python, it is implemented in Scipy library as the Golden Section Search technique.

However, we have also found that this projection was not robust. Our algorithm has then been reinforced with the **mbend** package from Nilforooshan (2020), which increase robustness by adding a weight matrix to A^+ (Jorjani et al., 2003), or by replacing each of the m negative eigenvalues (λ_i) with $\rho(s - \lambda_i)^2 / (100s^2 + 1)$, where ρ is the smallest positive eigenvalue and $s = 2 \sum_{i=1}^m \lambda_i$. This is called methode *LRS14* in their R package Nilforooshan & Nilforooshan (2020). To obtain the (ultimate) best solution $\mathcal{P}_{s+}(\lambda_w)$, those procedures are repeated until convergence is reached.

6. Predictive Posterior

In this section we describe how we estimate the predictive posterior distribution. That is, the estimation of $p(y^{(n+1)}|x^{(n+1)}, \mathbf{y}^n, \mathbf{x}^n)$ for a new data item $x^{(n+1)}$, which is the goal of our Potts Gibbs-network. Since the mixture is over all possible partitions ρ_{n+1} of the data $\mathbf{x}^n \cup \{x^{(n+1)}\}$, we have

$$p(y^{(n+1)}|x^{(n+1)}, \mathbf{y}^n, \mathbf{x}^n) = \sum_{\rho_{n+1}} p(y^{(n+1)}|x^{(n+1)}, \mathbf{y}^n, \mathbf{x}^n, \rho_{n+1}) p(\rho_{n+1}|x^{(n+1)}, \mathbf{y}^n, \mathbf{x}^n).$$

As in (Murua & Quintana, 2017), we consider the collection $\mathcal{A}(\rho_{n+1})$ of all partitions ρ_n of $\mathbf{x}^n$ giving rise to the partition ρ_{n+1} of $\mathbf{x}^n \cup \{x^{(n+1)}\}$ by generating appropriate bonds b_{n+1} in the extended graph with vertices $\mathbf{x}^n$ and $x^{(n+1)}$. The extended graph includes all edges of the the original graph plus the edges added to link the new data point. (e.g., if the graph is a k -nearest-neighbors graph, then all edges between $x^{(n+1)}$ and its k nearest neighbors are added). Elements in the collection $\mathcal{A}(\rho_{n+1})$ will be denoted as $[\rho_n, b_{n+1}]$ to make clear that ρ_{n+1} can be generated from ρ_n and the bonds b_{n+1} . Then, we can write

$$p(\rho_{n+1}|x^{(n+1)}, \mathbf{y}^n, \mathbf{x}^n) = \sum_{[\rho_n, b_{n+1}]} p(\rho_n|\mathbf{y}^n, \mathbf{x}^n) p(b_{n+1}|\rho_n, x^{(n+1)}, \mathbf{x}^n).$$

Therefore, the predictive posterior $p(y^{(n+1)}|x^{(n+1)}, \mathbf{y}^n, \mathbf{x}^n)$ is the mixture:

$$\sum_{[\rho_n, b_{n+1}]} p(y^{(n+1)}|x^{(n+1)}, \mathbf{y}^n, \mathbf{x}^n, [\rho_n, b_{n+1}]) p(\rho_n|\mathbf{y}^n, \mathbf{x}^n) p(b_{n+1}|\rho_n, x^{(n+1)}, \mathbf{x}^n). \quad (15)$$

This expression is intractable since the number of possible partitions of $\mathbf{x}^n$ is the well-known Bell number B_n . Fortunately, having sampled partitions and parameters from the posterior of ρ_n and ψ_σ , we can use these samples to estimate the predictive posterior as follows. Let $\{(\rho_{n,\ell}, \psi_{\sigma,\ell}) : \ell = 1, \dots, M\}$ be a sample from the posterior distribution. Also, let $\{b_{n+1,\ell} : \ell = 1, \dots, M\}$ be a sample of bonds linking the sampled partitions ρ_n to $x^{(n+1)}$. The estimate is

$$\hat{p}(y^{(n+1)}|x^{(n+1)}, \mathbf{y}^n, \mathbf{x}^n) = \frac{1}{M} \sum_{\ell=1}^M p(y^{(n+1)}|x^{(n+1)}, \mathbf{y}^n, \mathbf{x}^n, [\rho_{n,\ell}, b_{n+1,\ell}]). \quad (16)$$

The terms in the sum are

$$p(y^{(n+1)}|x^{(n+1)}, \mathbf{y}^n, \mathbf{x}^n, [\rho_n, b_{n+1}]) = \frac{p(y^{(n+1)}, \mathbf{y}^n|x^{(n+1)}, \mathbf{x}^n, [\rho_n, b_{n+1}])}{p(\mathbf{y}^n|x^{(n+1)}, \mathbf{x}^n, [\rho_n, b_{n+1}])} \quad (17)$$

$$= \frac{p(y^{(n+1)}, \mathbf{y}^n|x^{(n+1)}, \mathbf{x}^n, [\rho_n, b_{n+1}])}{\int_t p(t, \mathbf{y}^n|x^{(n+1)}, \mathbf{x}^n, [\rho_n, b_{n+1}]) dt}. \quad (18)$$

Let $\bar{j}$ be the cluster in which $x^{(n+1)}$ falls in the partition ρ_{n+1} , and suppose that this cluster is $\bar{S} = S_{\bar{j}} \cup \{x^{(n+1)}\}$, with $S_{\bar{j}} \subset \rho_n$. Let $\bar{T} = T_{\bar{j}} \cup \{y^{(n+1)}\}$. Note that

$$p(y^{(n+1)}, \mathbf{y}^n|x^{(n+1)}, \mathbf{x}^n, [\rho_n, b_{n+1}]) = p(\bar{T}|\bar{S}, \rho_{n+1}) \prod_{j \neq \bar{j}} p(T_j|S_j, \rho_n),$$

and similarly

$$\int_t p(t, \mathbf{y}^n | x^{(n+1)}, \mathbf{x}^n, [\rho_n, b_{n+1}]) = \left(\int_t p(\{T_{\bar{j}} \cup \{t\} | \bar{S}, \rho_{n+1}) dt \right) \left(\prod_{j \neq \bar{j}} p(T_j | S_j, \rho_n) \right).$$

Hence all the terms cancel in (17) except for the term associated with $x^{(n+1)}$. Therefore

$$p(y^{(n+1)} | x^{(n+1)}, \mathbf{y}^n, \mathbf{x}^n, [\rho_n, b_{n+1}]) = \frac{p(\bar{T} | \bar{S}, \rho_{n+1})}{\int_t p(\{T_{\bar{j}} \cup \{t\} | \bar{S}, \rho_{n+1}) dt}.$$

Using any value of $\psi_{\sigma_{\bar{j}}}$ in expression (9), we get

$$\begin{aligned} \hat{p}(y^{(n+1)} | x^{(n+1)}, \mathbf{y}^n, \mathbf{x}^n, [\rho_n, b_{n+1}]) &\approx \frac{p(\bar{T} | \psi_{\sigma_{\bar{j}}}, \bar{S}, \rho_{n+1}) p(\psi_{\sigma_{\bar{j}}} | \rho_{n+1})}{p(\psi_{\sigma_{\bar{j}}} | \bar{T}, \bar{S}, \rho_{n+1}) \int_t \frac{p(T_{\bar{j}} \cup \{t\} | \psi_{\sigma_{\bar{j}}}, \bar{S}, \rho_{n+1}) p(\psi_{\sigma_{\bar{j}}} | \rho_{n+1})}{p(\psi_{\sigma_{\bar{j}}} | T_{\bar{j}} \cup \{t\}, \bar{S}, \rho_{n+1})} dt} \\ &= \frac{p(\bar{T} | \psi_{\sigma_{\bar{j}}}, \bar{S}, \rho_{n+1})}{p(\psi_{\sigma_{\bar{j}}} | \bar{T}, \bar{S}, \rho_{n+1}) \int_t \frac{p(T_{\bar{j}} \cup \{t\} | \psi_{\sigma_{\bar{j}}}, \bar{S}, \rho_{n+1})}{p(\psi_{\sigma_{\bar{j}}} | T_{\bar{j}} \cup \{t\}, \bar{S}, \rho_{n+1})} dt}. \end{aligned}$$

Remark

In particular,

$$\begin{aligned} \hat{p}(y^{(n+1)} | x^{(n+1)}, \mathbf{y}^n, \mathbf{x}^n, [\rho_n, b_{n+1}]) &= \mathbb{E}_{p(\psi_{\sigma_{\bar{j}}} | \bar{T}, \bar{S}, \rho_{n+1})} \left\{ \frac{p(\bar{T} | \psi_{\sigma_{\bar{j}}}, \bar{S}, \rho_{n+1}) p(\psi_{\sigma_{\bar{j}}} | \rho_{n+1})}{p(\psi_{\sigma_{\bar{j}}} | \bar{T}, \bar{S}, \rho_{n+1})} \right\} \times \\ &\quad \left(\int_t \mathbb{E}_{p(\psi_{\sigma_{\bar{j}}} | \bar{T}, \bar{S}, \rho_{n+1})} \left[\frac{p(T_{\bar{j}} \cup \{t\} | \psi_{\sigma_{\bar{j}}}, \bar{S}, \rho_{n+1}) p(\psi_{\sigma_{\bar{j}}} | \rho_{n+1})}{p(\psi_{\sigma_{\bar{j}}} | T_{\bar{j}} \cup \{t\}, \bar{S}, \rho_{n+1})} dt \right] \right)^{-1}. \end{aligned}$$

In practice, we suppose that we can approximate $p(\psi_{\sigma_{\bar{j}}} | T_{\bar{j}} \cup \{t\}, \bar{S}, \rho_{n+1})$ by $p(\psi_{\sigma_{\bar{j}}} | T_{\bar{j}}, S_{\bar{j}}, \rho_n)$, and that, in turn, this last posterior is well approximated by the variational posterior approximation $Q(\psi_{\sigma_{\bar{j}}}; \lambda_{\bar{j}})$. Then each one of the expectations can be approximated by Monte Carlo, using our samples $\{\psi_{\sigma_{\bar{j}\ell}}, \ell=1, \dots, M\}$ from the variational posterior, where $\bar{j}_\ell$ denotes the cluster in $\rho_{n+1, \ell}$ in which $x^{(n+1)}$ falls, $\ell = 1, \dots, M$. This yields

$$\begin{aligned} \hat{p}(y^{(n+1)} | x^{(n+1)}, \mathbf{y}^n, \mathbf{x}^n, [\rho_n, b_{n+1}]) &= \left(\sum_{\ell=1}^M \frac{p(\bar{T}_\ell | \psi_{\sigma_{\bar{j}_\ell}}, \bar{S}_\ell, \rho_{n+1, \ell}) p(\psi_{\sigma_{\bar{j}_\ell}} | \rho_{n+1, \ell})}{Q(\psi_{\sigma_{\bar{j}_\ell}}; \lambda_{\bar{j}_\ell})} \right) \times \\ &\quad \left(\int_t \left[\sum_{\ell=1}^M \frac{p(T_{\bar{j}_\ell} \cup \{t\} | \psi_{\sigma_{\bar{j}_\ell}}, \bar{S}_\ell, \rho_{n+1, \ell}) p(\psi_{\sigma_{\bar{j}_\ell}} | \rho_{n+1, \ell})}{Q(\psi_{\sigma_{\bar{j}_\ell}}; \lambda_{\bar{j}_\ell})} dt \right] \right)^{-1}, \end{aligned}$$

with $\bar{S}_\ell = S_{\bar{j}_\ell, \ell} \cup \{x^{(n+1)}\}$, and $\bar{T}_\ell = T_{\bar{j}_\ell, \ell} \cup \{y^{(n+1)}\}$. The integral over t , which is a univariate integral, may be computed numerically, for example, using Romberg's algorithm (Li & Hu, 2017).

In particular, the value of the prediction $\mathbb{E}(y | x^{(n+1)}, \mathbf{y}^n, \mathbf{x}^n)$ may be estimated by the quantity:

$$\frac{1}{M} \sum_{\ell=1}^M \mathbb{E}(y | \psi_{\sigma_{\bar{j}_\ell}}, x^{(n+1)}) = \frac{1}{M} \sum_{\ell=1}^M [b_\ell^{(2)} + W_\ell^{(2)} g_1(b_\ell^{(1)} + W_\ell^{(1)} x^{(n+1)})].$$

The Shallow Potts as Random Gibbs Network Forest. There is an *intuitive* mathematical equivalence between *Random Forest* and the Shallow Gibbs. Random Forest is an ensemble learning : a machine learning method that combines many simple learners – base models– to construct an optimized predictive model (*powerful* model).

Random-forest does both row sampling and column sampling with Decision tree as a base learner (Liu et al. (2012); Ali et al. (2012); Rigatti (2017); Liu (2014); Schonlau & Zou (2020); Probst et al. (2019); Svetnik et al. (2003)). Analogically, the Shallow Gibbs does random clusters and rather than column sampling, it projects the data $\mathbf{x}$ into a fewer dimensional space on the shallow layer (one hidden layer with *very* few reduced number of neurons).

The variance will decrease in Random Forests, as you increase the number of base learners τ . When you reduce τ , there is a rise in variance. For the whole process, however, bias remains unchanged. It is *commonly* said: Random forest= Decision Trees (as a base learner)+ bagging (Row sampling with replacement)+ feature bagging (column sampling) + aggregation (mean/median, majority vote) [Skurichina & Duin (2001); Munson & Caruana (2009); Biau & Devroye (2010); Sutton et al. (2005); Strobl et al. (2009); Kotsiantis (2011); Panov & Džeroski (2007)]. In the Shallow Potts, when you increase the number of base learners (the number neurons one the hidden layer), the train and test errors increase. To reduce the errors in Shallow Gibbs, we *double* backpropagate it accross the hyper-parameters and parameters of the model [See 7.1].

The Shallow Potts as a Mixture Model (As a reminder of some practical computations) Computing expectation $\mathbb{E}(y^{(n+1)}|x^{(n+1)}, \mathbf{y}^n, \mathbf{x}^n)$ in a proper way would require an exact probability for each partition in 15. The value of the prediction $\mathbb{E}(y^{(n+1)}|x^{(n+1)}, \mathbf{y}^n, \mathbf{x}^n)$ need to adjusted with an estimate of $p(\rho_n|\mathbf{y}^n, \mathbf{x}^n)p(b_{n+1}|\rho_n, x^{(n+1)}, \mathbf{x}^n)$: it simply means that each partition has a probability that need to be taken into account in the final estimation, plus an estimation of the conditional probability of the new bonds (b_{n+1}). To solve this problem, we have approximated $\hat{p}(\rho_n|\mathbf{y}^n, \mathbf{x}^n) \approx 1/M$, and have introduced a smoothing parameter $0 \leq \delta \leq 1$ ($\delta = \hat{p}(b_{n+1}|\rho_n, x^{(n+1)}, \mathbf{x}^n)$) that will represent an estimate of this probability. This is not the optimal way, in a sense that $\hat{p}(b_{n+1}|\rho_n, x^{(n+1)}, \mathbf{x}^n)$ is different from a partition to another one. But, it will help out to transform the final prediction to:

$$\hat{y} \approx \frac{\delta}{M} \mathbb{E}(y^{(n+1)}|x^{(n+1)}, \mathbf{y}^n, \mathbf{x}^n)$$

where M is the number of partitions generated during the training step.

7. On Double backpropagation

Double backpropagation has various used cases. Under the name double backpropagation, the concept of penalizing the output gradient with respect to the input was introduced (Ross & Doshi-Velez, 2018), and later used to locate large connected areas of the error function called flat minima. The ultimate aim was the development of their models generalization, which is our goal here. It has effectively improved generalization performance [Drucker & Le Cun (1992), Drucker & Le Cun (1991)], and has also been applied to other adversarial models ((Seck et al., 2019), (Sun et al., 2020)), in character recognition (Kamruzzaman & Syed (1998), Kamruzzaman et al. (1997)), in robustness and saliency map interpretability (Etmann, 2019). We present in the following line our novel double backpropagation scheme (DBS).

7.1 A Proposed Double backpropagation scheme

Similarly to our *Iterative projected gradient* (IPG) applied to our model variational parameter λ_w , we found that the Mean Squared Error (MSE) of our regression model can also be back-propagated (Rumelhart et al. (1986); Hecht-Nielsen (1992)) *w.r.t* to each of the model parameter, i.e $\psi = (b^{(1)}, W^{(1)}, b^{(2)}, W^{(2)})$, the main parameters of the network, with $b^{(1)} \in \mathbb{R}^{l_1}$, $W^{(1)} \in \mathbb{M}_{q \times l_1}$, $b^{(2)} \in \mathbb{R}^p$, and $W^{(2)} \in \mathbb{M}_{l_1 \times p}$, $l_0 = q$, $l_2 = p$, $\Sigma \in \mathbb{M}_{p \times p}$ being the variance-covariance matrix of y . First, we know that $y|x, \psi, \Sigma$ is distributed as a multivariate normal distribution with mean $f(x) = f_\psi(x) = \mathbb{E}(y|x, \psi)$, and variance Σ . That is, $p(y|x, \psi, \Sigma) = (2\pi)^{-p/2} |\Sigma|^{-1/2} \exp\{-1/2(y -$

$f_\psi(x))'\Sigma^{-1}(y - f_\psi(x))\}$. Then, by applying the sampling method described in the Cholesky decomposition section B, we have :

$$\hat{y}_{est} = f_\psi(x) + L \cdot u = [b^{(2)} + g_1(b^{(1)} + xW^{(1)})W^{(2)}] + L \cdot u \quad (19)$$

such that $L \in M_d(\mathbb{R})$ is a lower triangular matrix such that $\Sigma = LL^T$, and $u \sim N(0, I)$. The *double* back-propagation scheme for the Shallow Gibbs – **which is another main contribution of this research** – goes as follows:

1. Using the IPG (in 14), apply backpropagation method on hyper-parameter λ_w to reduce its Kullback-Leibler (KL) estimation error. Once done, generate an estimate $\hat{\psi}_0$ of $\psi = (b^{(1)}, W^{(1)}, b^{(2)}, W^{(2)})$, using Monte Carlo sampling method from the variational distribution of the parameters.
2. Use equation 19 to backpropagate the $MSE(y_i - \hat{y}_i) = \|y_i - \hat{y}_{est,i}\|^2$ to update $\hat{\psi}_0$ in the Potts cluster and **per observation** as follows:

$$\hat{\psi}_{1,i} \leftarrow \hat{\psi}_0 - \epsilon_{\psi,0} \frac{\partial MSE(y_i - \hat{y}_{est,i})}{\partial \psi} \quad (20)$$

$$\hat{\psi}_{t,i} \leftarrow \hat{\psi}_{t-1,i} - \epsilon_{\psi,t-1} \frac{\partial MSE(y_i - \hat{y}_{est,i})}{\partial \psi} \quad (21)$$

where $\epsilon_{\psi,t}$ is the learning rate schedule [See C.1] for this gradient update of ψ at step t , and $\hat{\psi}_{t,i}$ is the value of ψ at iteration t for observation i . Equivalently, for the i -th data x_i in the training data, it corresponds to the changes:

$$\begin{aligned} b_{i,t}^{(1)} &= b_{i,t-1}^{(1)} - \epsilon_{b_{i,t-1}^{(1)}} \cdot 2I_{l_1 \times l_1}^i g_1'(b_{i,t-1}^{(1)} + x_i W_{i,t-1}^{(1)}) W_{i,t-1}^{(2)} \star [f_{\psi_{i,t-1}}(x_i) - y_i]^T \\ b_{i,t}^{(2)} &= b_{i,t-1}^{(2)} - \epsilon_{b_{i,t-1}^{(2)}} \cdot 2I_{p \times p}^i [f_{\psi_{i,t-1}}(x_i) - y_i]^T \end{aligned} \quad (22)$$

$$W_{i,t}^{(1)} = W_{i,t-1}^{(1)} - \epsilon_{W_{i,t-1}^{(1)}} \cdot 2x_i^T [g_1'(b_{i,t-1}^{(1)} + x_i W_{i,t-1}^{(1)})] W_{i,t-1}^{(2)} [f_{\psi_{i,t-1}}(x_i) - y_i]^T \quad (23)$$

$$W_{i,t}^{(2)} = W_{i,t-1}^{(2)} - \epsilon_{W_{i,t-1}^{(2)}} \cdot 2[g_1(b_{i,t-1}^{(1)} + x_i W_{i,t-1}^{(1)})]^T [f_{\psi_{i,t-1}}(x_i) - y_i]^T \quad (24)$$

$$\Sigma_{i,t} = \Sigma_{i,t-1} - \epsilon_{\Sigma_{i,t-1}} \cdot 2[(L^T + L)^{-1} u_{i,t}] [f_{\psi_{i,t-1}}(x_i) - y_i]^T \quad (25)$$

$$\hat{y}_{est,(i,t)} = f_{\psi_{i,t}}(x) + L_{i,t} \cdot u_{i,t} = [b_{i,t}^{(2)} + g_1(b_{i,t}^{(1)} + xW_{i,t}^{(1)})W_{i,t}^{(2)}] + L_{i,t} \cdot u_{i,t} \quad (26)$$

where $[z]^T$ means the transpose of z , g_1' the derivative of g_1 , $\psi_{i,t} = (W_{i,t}^{(1)}, W_{i,t}^{(2)}, b_{i,t}^{(1)}, b_{i,t}^{(2)}, \Sigma_{i,t})$ represent the change for ψ for the i -th training observation at step t ; $\epsilon_{\psi_{i,t}} = (\epsilon_{W_{i,t}^{(1)}}, \epsilon_{W_{i,t}^{(2)}}, \epsilon_{b_{i,t}^{(1)}}, \epsilon_{b_{i,t}^{(2)}}, \epsilon_{\Sigma_{i,t}})$ is the learning rate of those gradients updates; $L_{i,t}$ the Choleski factor at step t from $\Sigma_{i,t}$; and finally $u_{i,t} \sim N(0, I)$. The changes for the i -th data in the testing set, is the change of its nearest neighbor in the current cluster *w.r.t* to the Euclidean distance ⁵.

The product symbol $\star$ means that the vector $[f_{\psi_{i,t-1}}(x_i) - y_i]^T$ will multiply each line of the final computation on his left. We will be calling the optimization scheme above as the Double Backpropagation Scheme – **DBS optimizer**. The DBS is an universal Bayesian neural network optimizer: given an initial $\psi_{i,0} = (W_{i,0}^{(1)}, W_{i,0}^{(2)}, b_{i,0}^{(1)}, b_{i,0}^{(2)}, \Sigma_{i,0})$,

5. We have updated this choice during our experiments.

and $\hat{y}_{est,(i,0)}$, it can make $\hat{y}_{est,(i,t)}$ converge to the right response value y_i with appropriate learning rate schedule. To find this appropriate learning schedule is an open work, and one can look for its convergence properties using the stochastic gradient learning convergence theorems, to show how $\hat{y}_{est,(i,t)}$ converges, and his speed of convergence.

8. On Differentiation Methods

Differentiation occurs everywhere, from the backpropagation algorithm in deep neural networks to the motion equations in physics to almost any region that needs to calculate a rate of change. When abording differentiation in machine learning, we encounter automatic differentiation, symbolic differentiation, numerical differentiation. We have introduced here in this paper empirical differentiation as a novel differential machine learning framework.

8.1 Automatic differentiation.

Automatic differentiation (also called *autodiff* in a short) is a set of techniques to automatically construct a procedure and numerically evaluate the derivative of a function, given the analytical expression of the function itself [Griewank et al. (1989), Baydin et al. (2018), Rall & Corliss (1996), Bückner et al. (2006), Bischof et al. (2008a)].

Automatic Differentiation can only calculate the partial derivative of an expression on a certain point, when we have assigned values to each of the variables . It has many applications such as numerical integration algorithms for ordinary differential equations (Bischof et al., 2008b), in Wave-function positivization (Torlai et al., 2020), in inverse design of photonic crystals (Minkov et al., 2020), for numerical solution of the Burgers' equation (Asaithambi, 2010), in nonlinear Mathematical Programming Language models Gay (1991), in periodic orbits and their bifurcations Guckenheimer & Meloon (2000), Runge-Kutta methods for optimal control Walther (2007), in engineering design Barthelemy & Hall (1995), in the computation of sparse Jacobian matrices Coleman & Verma (1998), in Fourier series and the radii polynomial Lessard et al. (2016); and has already been implemented in Matlab (Bischof et al. (2003), Forth (2006), Verma (1999)) and in Pytorch (Paszke et al. (2017)), as well as in The Stan Math library (Carpenter et al., 2015).

8.2 Symbolic differentiation (SD).

With respect to a defined variable, symbolic differentiation (Knott (2017)) is a program that seeks the derivative of a given formula, generating a new formula as its output (only applicable to closed-form mathematical functions). As for example, let say we are given function 27.

$$f(x) = \frac{x^2 - 1}{x\sqrt{x^2 + 1}} \quad (27)$$

Then the symbolic differentiation will compute directly the derivative solution:

$$f'(x) = \frac{2x^2\sqrt{x^2 + 1} - (x^2 - 1) \left(x \frac{x}{\sqrt{x^2 + 1}} + \sqrt{x^2 + 1} \right)}{x^2 (x^2 + 1)}. \quad (28)$$

Symbolic differentiation has been applied in graphics applications (Guenter, 2007), in serial kinematic chain (Bruyninckx & De Schutter, 1996), in multibody rigid systems (Villard & Arnaldi (1996), Shi & McPhee (2002)) , in machine and deep learning (Bergstra et al. (2011), Bastien et al. (2012), Zhabinski et al. (2017), Meurer et al. (2017), Clausen & Sokol (2020)), for fixed-Structure Controller Synthesis Osburn et al. (1996), for GPU

ShadersGunter et al. (2011), for web-based differentiation systems (Plaskura, 2017), for symbolic algebra (Kenny, Pantelides (1988)), and to solve Inverse Problems (Jégou (1997), Benghorbal (2007)).

There is also a link between Symbolic Differentiation (SD) and Automatic Differentiation (AD). In *The Simple Essence of Automatic Differentiation*, Conal Elliott wrote:

Automatic Differentiation is Symbolic Differentiation performed by a compiler Elliott (2018).

Laue (2019) argues an equivalence between both methods claiming that automatic forward mode differentiation and symbolic differentiation are similar in the sense that, when computing derivatives, they all perform the same operations.

8.3 Numerical differentiation.

This kind of differentiation is much more appreciated when we use the derivative formula in [29], with the analytical form of the function f itself. With numerical differentiation, the derivative of a function of $\nabla f, f : \mathbb{R}^n \rightarrow \mathbb{R}$ at $x = a$ (also called *Finite difference approximation*) is the limit:

$$\frac{\partial f(\mathbf{a})}{\partial x_i} \approx \frac{f(\mathbf{a} + h\mathbf{e}_i) - f(\mathbf{a})}{h}, \quad 0 < h \ll 1 \quad (29)$$

with $\mathbf{e}_i \in \mathbb{R}^n$ the i -th unit basis vector, and h is a small step size. This previous computation is very useful to approximate the gradient $\nabla f = \left(\frac{\partial f}{\partial x_1}, \dots, \frac{\partial f}{\partial x_n} \right)$. There are various forms of numerical differentiation in the literature, such as stable numerical differentiation, regularized numerical differentiation, noisy numerical differentiation, etc. [Cullum (1971), Ramm & Smirnova (2001), Anderssen & Bloomfield (1974), Chartrand (2011), Knowles & Renka (2014), Mboup et al. (2009), Lu & Pereverzev (2006)]

In the next section, we have introduced a novel differential machine learning framework based on an empirical differentiation procedure, which is a generalization of interpolation.

9. Empirical differentiation and Differential Machine learning

In chapter [??], equations [43] and [48] indicate the need for an adapted Train-test association function [43] for the prediction of $\mathbf{y}^{\text{test}}$ when we know $\mathbf{x}^{\text{test}}$. This form of data closest neighborhood association is one limitation of the **DBS** optimizer, since it *simply* implies that we are still searching for the right *learner*.

9.1 Interpolation as a Machine learner

You can notice with table [9], that in few iterations, the **DBS** can overfit the training data. The main question then is: knowing a perfect relation [46] for any train data, as $(x_j^{\text{train}}, \psi_j^{\text{train}})$:

$$x_j^{\text{train}} \xleftrightarrow{\iota} \psi_j^{\text{train}} = (W_j^{(1)}, W_j^{(2)}, b_j^{(1)}, b_j^{(2)}, \Sigma_j), y_j^{\text{train}} \quad (30)$$

where ι is our denoted *gate* function, how can we find the perfect ψ_i^{test} for the test data as:

$$x_i^{\text{test}} \xleftrightarrow{\iota} \psi_i^{\text{test}} = (W_i^{(1)}, W_i^{(2)}, b_i^{(1)}, b_i^{(2)}, \Sigma_i), y_i^{\text{test}} \quad (31)$$

Or simply said, is Nearest Neighborhood enough to find $\psi_i^{\text{test}}, y_i^{\text{test}}$, for a given x_i^{test} ?

This is where Multivariate Interpolation come into action, as the later can be taken also as a machine *learner*, because it can refine Nearest Neighborhood Train-Test association [43]. To perceive this, we present another perspective of machine learning problems with the following described properties about interpolation.

Linear interpolation usually requires two data points (u_a, v_a) and (v_b, v_b) , and at the point (u, v) , the interpolation equation is given by:

$$v = v_a + (v_b - v_a) \frac{u - u_a}{u_b - u_a} \quad (32)$$

Equation [32] specifies that the current line slope between (u_a, v_a) and (u, v) is equivalent to the line slope between (u_a, v_a) and (u_b, v_b) . The method have been used since a while in the literature [Joarder (2001), Dressler (2009), Otsuki et al. (2004), Guo-qiang et al. (2004), Kay (1983), Bourke (1999)]. When we know u , interpolation is useful to estimate v . But, it may be a poor estimator sometimes [Maeland (1988), Dumitru et al. (2013)].

To generalize equation [32] to a learning problem, remember the Taylor's theorem for a multivariate function ι in functions analysis theory [Giaquinta & Modica (2010), Browder (2012), Zorich (2016), Choquet-Bruhat et al. (1982), Pettersson (1992), Barbarossa & Manzonetto (2019), Cheney & Kincaid (2012), Canuto & Tabacco (2015)]. We already know the best linear approximation to ι . It involves the derivative $D\iota(\mathbf{a})$ such as:

$$\iota(\mathbf{x}) \approx \iota(\mathbf{a}) + D\iota(\mathbf{a}) \circ (\mathbf{x} - \mathbf{a}) \quad (33)$$

where $D\iota(\mathbf{a})$ is the matrix of partial derivatives of ι evaluated in the neighborhood of $\mathbf{a}$, and $\circ$ is the dot product between both vectors $D\iota(\mathbf{a})$ and $(\mathbf{x} - \mathbf{a})$. This approximation is linear and represents the first-order Taylor polynomial [Reimer (2012), Guessab et al. (2006), Phillips (2003), Dikumar (2016), Neumaier (2003), Cohen & Tan (2012), Smyth (2014)]. Proposition [9.1] is the generalization of approximation [33].

9.2 Multivariate Interpolation: Taylor Approximation is the future of Machine learning

Proposition 9.1 (Multivariate version of Taylor theorem (Qian (2011)))

Let $\iota : \mathbf{R}^n \rightarrow \mathbf{R}$ be a $m \geq \Xi \geq 2$ -times-differentiable function in a certain vicinity D of the point $\mathbf{u}_0 = (u_{01}, \dots, u_{0n}) \in \mathbf{R}^n$. Then, for any $v = (v_1, \dots, v_n)$, we have :

$$\begin{aligned} \iota(v_1, \dots, v_n) &= \iota(u_{01}, \dots, u_{0n}) + \sum_{i=1}^n \frac{\partial \iota}{\partial u_i} \Big|_{(u_{01}, \dots, u_{0n})} (v_i - u_{0i}) \\ &+ \frac{1}{2!} \sum_{i=1}^n \sum_{j=1}^n \frac{\partial^2 \iota}{\partial u_i \partial u_j} \Big|_{(u_{01}, \dots, u_{0n})} (v_i - u_{0i}) (v_j - u_{0j}) \\ &+ \dots + R_m(v - u_0) \end{aligned}$$

where $R_m(v - u_0)$ is a reminder function of $v - u_0$.

Remark

It is not a coincidence that Taylor Approximation theorem is only defined in a certain vicinity or a given neighborhood set! One of the main contribution of this chapter is to understand that : *Multivariate Interpolation using Taylor theorem is the refined generalization of simple Neighborhood Train-Test association*. Notice in equation [33], if $\mathbf{a} = \mathbf{x}^{\text{train}}$ and $\mathbf{x} = \mathbf{x}^{\text{test}}$, we have $\iota(\mathbf{x}^{\text{test}}) \approx \iota(\mathbf{x}^{\text{train}})$, when we suppose $\mathbf{x}^{\text{train}} \approx \mathbf{x}^{\text{test}}$. To avoid that simple approximation, we add more differential terms, as exposed in Proposition [9.1].

Can we overfit with proud? It is important to also understand the case $\iota : \mathbb{R}^q \rightarrow \mathbb{R}^m$ in [30], and how we expand $D\iota(\zeta)$ at $\zeta = x_0^{\text{train}}$. When ι is multivariate, we will use Proposition [9.1], or in a one order [33], to find [31] per component using the assumption $(\mathcal{F}_d)$ that :

There is an approximable differentiable function ι such as: $\iota(x_i^{\text{test}}) = (\psi_i^{\text{test}}, y_i^{\text{test}})$, where we already know in a perfect overfitting DBS configuration $(x_{j_1}^{\text{train}}, \psi_{j_1}^{\text{train}}, y_{j_1}^{\text{train}})$ [i.e $\iota(x_{j_1}^{\text{train}}) = (\psi_{j_1}^{\text{train}}, y_{j_1}^{\text{train}})$] and $(x_{j_2}^{\text{train}}, \psi_{j_2}^{\text{train}}, y_{j_2}^{\text{train}})$ [i.e $\iota(x_{j_2}^{\text{train}}) = (\psi_{j_2}^{\text{train}}, y_{j_2}^{\text{train}})$] for two training data j_1 and j_2 , With the precision that the training data point j_2 and the test data i are taken in a very closed neighborhood of j_1 .

In each Potts partition ρ_n clusters, this is applicable. Suppose $x_{j_1}^{\text{train}} = (x_{j_1}^{\text{train}}(1), \dots, x_{j_1}^{\text{train}}(q))$, and $x_{j_2}^{\text{train}} = (x_{j_2}^{\text{train}}(1), \dots, x_{j_2}^{\text{train}}(q))$, with the precision that the training data point j_2 is taken in a very closed neighborhood of j_1 (potential example is: the associate Potts cluster of j_1 given a partition ρ_n). In practice, for $x = x_j^{\text{test}}$ and $a = x_{j_1}^{\text{train}}$, we compute $D\iota(a) \circ (x - a)$ in approximation [33] as:

$$D\mathcal{C}_{\iota_j}(x_{j_1}^{\text{train}}) \circ (x_j^{\text{test}} - x_{j_1}^{\text{train}}) = \sum_{k=1}^q \frac{\mathcal{J}(x_{j_1}^{\text{train}}(k) \neq x_{j_2}^{\text{train}}(k))(x_j^{\text{test}}(k) - x_{j_2}^{\text{train}}(k))}{x_{j_1}^{\text{train}}(k) - x_{j_2}^{\text{train}}(k) + \mathcal{J}(x_{j_1}^{\text{train}}(k) = x_{j_2}^{\text{train}}(k))} [\mathcal{C}_{\iota_{j_1}}^{(1)} - \mathcal{C}_{\iota_{j_2}}^{(1)}] \quad (34)$$

where $\mathcal{C}_{\iota} \in \{W^{(1)}, W^{(2)}, b^{(1)}, b^{(2)}, \Sigma, y^{\text{train}}\}$, $\mathcal{J}(A)$ equal 1 if condition A is verified. Those constraints $\mathcal{J}(x_{j_1}^{\text{train}}(k) \neq x_{j_2}^{\text{train}}(k))$ and $\mathcal{J}(x_{j_1}^{\text{train}}(k) = x_{j_2}^{\text{train}}(k))$ are useful to indicate that we do not want the partial derivatives to be computed if there are no differentiation available for the k th dimension. Further, we shall reduced the Taylor series in [9.1] to integrate all those constraints, and the final computation will be called *reduced derivative*.

To use this technique of Reduced Taylor Series Multivariate Interpolation in [34] with the research of the nearest neighbor in the train dataset for a new x^{test} , is what we call the **(dist)-Nearest Neighbor-(h)-Taylor Series-Reduced Multivariate Interpolation (dist-NN-(h)-TS-RMI)** in reference to previous existent work [Rukundo & Cao (2012)], and it can reduce (possibly) test error, specially when we overfit sometimes. In *dist-NN-(h)-TS-RMI*, *dist* is set for the Nearest Neighbor distance applied [e.g 43], and *h* is the order of computed Taylor approximation in [34]. Assumption [9.2] denoted $(\mathcal{F}_d)$ may not be valid, but the method may be effective for some datasets; or another situation that could arise is that the closest data point j_2 to j_1 in a very small neighborhood set may not exist in the training data. Some experiments are required to confirm in which case the proposed *dist-NN-(h)-TS-RMI* model is a *good learner*.

Taylor Theorem can approximate any (differentiable) function. Two questions of interest may arise then:

- Q₁ What is the statistical distribution of any *dist*-Nearest Neighbor-(h)-Interpolated-Reduced response variable y , as the parameters of the model are shifted (by interpolation and reduction) ?**
- Q₂ Analogous to the *Universal Approximation Theorem* [see 12.1], can higher order Taylor Theorem solve any machine learning problem?**

9.2.1 TAYLOR APPROXIMATION IS THE FUTURE OF MACHINE LEARNING.

We can use Taylor Theorem [9.1] to develop higher order ($h \geq 2$) Multivariate Interpolation. In fact, Taylor Series can also be used to approximate any function. Formula [34] also opens the door to empirical derivatives computation. Empirical differentiation reduces or deletes the terms where differentiation is not applicable. This

is done to signify that the higher-order differentiated parts represent a *small* correction of the Nearest Neighbor associated values. This technique is valid because the reminder terms in proposition [9.1] get smaller when the differentiation order increases.

This is a great revolution in *Statistical Learning Theory*, as we can sufficiently (i.e $h \geq 2$) interpolate any machine learning regression problem, with the *reduced* differentiation form when we know the perfect parameters that fit the training data. By then, the famous **overfitting** problem will not appear to be an hindrance anymore, but an extraordinary gateway to another type of *learning* method, which is therefore the research of the perfect fit for both the training and the test data, using as e.g the proposed *dist*-NN-(h)-TS-RMI model.

10. Data Augmentation for Empirical Differentiation (DAED)

This section is a trial to answer question $\mathbf{Q}_2$, as Taylor Approximation can truly solve any machine learning problem. We illustrate this fact using data augmentation. To make the assumption $(\mathcal{F}_d)$ valid, we need to create more samples from existent training data.

To understand this intuition, remember that the partial derivative of a function $\iota(x_1, \dots, x_n)$ in the direction x_i at the point $(e_1, \dots, e_n)$ is defined by:

$$\frac{\partial \iota}{\partial x_i}(e_1, \dots, e_n) = \lim_{\delta_a \rightarrow 0} \frac{\iota(e_1, \dots, e_i + \delta_a, \dots, e_n) - \iota(e_1, \dots, e_i, \dots, e_n)}{\delta_a} \quad (35)$$

with $\delta_a \in \mathbb{R}$ has to be a very small real number. Because the closed neighborhood training data j_2 for j_1 in the differentiation computation [34] is not available in practice in the training data, we can create more data as follows, to compute derivative [35] and [34] with almost exact precision:

$$x_{j_1}^{\text{train}} = (x_{j_1}^{\text{train}}(1), \dots, x_{j_1}^{\text{train}}(q)) \longrightarrow \begin{cases} (x_{j_1}^{\text{train}}(1) + \delta_a, \dots, x_{j_1}^{\text{train}}(q)) \\ (x_{j_1}^{\text{train}}(1), x_{j_1}^{\text{train}}(2) + \delta_a, \dots, x_{j_1}^{\text{train}}(q)) \\ \dots\dots\dots \\ (x_{j_1}^{\text{train}}(1), x_{j_1}^{\text{train}}(2), \dots, x_{j_1}^{\text{train}}(q) + \delta_a) \end{cases} \quad (36)$$

where in [36], δ_a is added on each component of $x_{j_1}^{\text{train}}$ in each of his augmented version. When we have the model ι that overfit [e.g the **DBS** optimizer], we have to re-train the model with the augmented data [36] to find [37]:

$$\begin{cases} \iota(x_{j_1}^{\text{train}}(1), \dots, x_{j_1}^{\text{train}}(q)) \\ \iota(x_{j_1}^{\text{train}}(1) + \delta_a, \dots, x_{j_1}^{\text{train}}(q)) \\ \iota(x_{j_1}^{\text{train}}(1), x_{j_1}^{\text{train}}(2) + \delta_a, \dots, x_{j_1}^{\text{train}}(q)) \\ \dots\dots\dots \\ \iota(x_{j_1}^{\text{train}}(1), x_{j_1}^{\text{train}}(2), \dots, x_{j_1}^{\text{train}}(q) + \delta_a) \end{cases} \quad (37)$$

This data augmentation framework, the DAED, transforms the *dist*-NN-(h)-TS-RMI model in [9.2] into another one, which is the : **(dist)-Nearest Neighbor-(h)-Taylor Series-Perfect Multivariate Interpolation (dist-NN-(h)-TS-PMI)**. We have also find out that in practice, the reduced model *dist*-NN-(h)-TS-RMI was not precised, and not effective differentiation approach. We have run some experiments with Slump dataset, we still found the

DBS as a model that overfit the training data, but the reduced *dist*-NN-(h)-TS-RMI model can not generalize very well. The Perfect Multivariate Interpolation method, let say the *dist*-NN-(h)-TS-PMI) model, is recommended.

Using limit [35], and a re-trained model where we know [37], we can compute:

$$\frac{\partial \iota}{\partial x_j(k)} \left(x_j^{\text{train}}(1), \dots, x_j^{\text{train}}(q) \right) \approx \frac{\iota \left(x_j^{\text{train}}(1), \dots, x_j^{\text{train}}(k) + \delta_a, \dots, x_j^{\text{train}}(q) \right) - \iota \left(x_j^{\text{train}}(1), \dots, x_j^{\text{train}}(k), \dots, x_j^{\text{train}}(q) \right)}{\delta_a} \quad (38)$$

with a very small δ_a [e.g 10^{-10} , 10^{-50}], to get :

$$D\iota(x_j^{\text{train}}) = \left(\frac{\partial \iota}{\partial x_j(1)} \left(x_j^{\text{train}}(1), \dots, x_j^{\text{train}}(q) \right), \dots, \frac{\partial \iota}{\partial x_j(k)} \left(x_j^{\text{train}}(1), \dots, x_j^{\text{train}}(q) \right), \dots, \frac{\partial \iota}{\partial x_j(q)} \left(x_j^{\text{train}}(1), \dots, x_j^{\text{train}}(q) \right) \right) \quad (39)$$

$D\iota(x_j^{\text{train}})$ presented above is the best mathematical and empirical Multivariate function Derivative that you can compute for ι and evaluate at x_j^{train} .

We suppose that the augmented training data in [36], being denoted $x_j^{\text{aug-train}}$, are taken in very close neighborhood of $x_{j_1}^{\text{train}} = (x_{j_1}^{\text{train}}(1), \dots, x_{j_1}^{\text{train}}(q))$. So, **before the re-training step**, to compute the associated $y_j^{\text{aug-train}}$ for each $x_j^{\text{aug-train}}$, we may use $\psi_{j_1}^{\text{train}}$, as illustrated in [30].

Iterative Multivariate Interpolation Method. We want to highlight a strong link between Gradient descent and Taylor Approximation [HAMMER (2017)]. There is then possibility to refine the Perfect Multivariate Interpolation in [10], so that it can be iterated, most likely in a Gradient Descent Optimization Scheme, e.g illustrated by Lydia & Francis (2019).

11. Generalization Method

In the previous section, we see how data can be augmented. Now, it is time to reveal how we generalize the model in order to make any test data to have its associate neighbor in the training data. The secret is to understand that every augmented data can also be sufficiently differentiated, as much as we want to overlap all the possible data Space domain:

$$x_{j_1}^{\text{aug-train}} = (x_{j_1}^{\text{train}}(1) + \delta_a, \dots, x_{j_1}^{\text{train}}(q)) \longrightarrow \begin{cases} (x_{j_1}^{\text{train}}(1) + s_1 \cdot \delta_a, \dots, x_{j_1}^{\text{train}}(q)) \\ (x_{j_1}^{\text{train}}(1), x_{j_1}^{\text{train}}(2) + s_2 \cdot \delta_a, \dots, x_{j_1}^{\text{train}}(q)) \\ \dots\dots\dots \\ (x_{j_1}^{\text{train}}(1), x_{j_1}^{\text{train}}(2), \dots, x_{j_1}^{\text{train}}(q) + s_q \cdot \delta_a) \end{cases} \quad (40)$$

with $(s_1, s_2, \dots, s_q)$ is a q -tuple of integers or rationals. The re-training process has to be done progressively to obtain their perfect associated parameters. The last, but not the least, is that each novel augmented data can again be differentiated as done in [49], (again and again) with no end.

We also found that in practice, this generalisation method is effective for slump dataset, but it takes many data augmentation iterations over the training data to reach convergence for the test data.

12. Experimental evaluation

In this section, we study the performance of the Potts Gibbs-network mixture model, Potts-Gibbs-NM for short, on real datasets. As a baseline, we compare the performance of Potts-Gibbs-NM with that of the classical feed-forward neural network, and to the multi-layer multi-target regression (MMR) model of Zhen et al. (2017). The performance of the models is measured on datasets from the Mulan Project (Tsoumakas et al., 2011).

As described above, there are four variants of the Potss-Gibbs-NM model: the fully-connected Gibbs network or full Gibbs network (Full-Gibbs, for short), the between-layer sparse Gibbs network (B-Sparse-Gibbs), the sparse shallow Gibbs network (Sparse-Gibbs), the sparse compound symmetry Gibbs network (CS-Sparse-Gibbs), and the compound symmetry Gibbs network (CS-Gibbs). The activation function used for all the neural networks in this study is a smooth version of the Rectified Linear Unit or ReLU, namely, $\log(\exp(x) + 1) \approx \max\{0, x\}$ (Lee et al., 2019). All the code used in the experiments can be downloaded from the first author's `github` site (Alahassa, 2020).

To appreciate how we tweak each of our model, you need to pay attention on 13 factors:

- NumberPartitions (NP): an integer for the number of Potts model partitions generated;
- Minimum Potts Cluster Size (MPCS): an integer for the minimum size of each cluster in all partitions generated;
- NHLayer (NH): an integer for the number of neurons on the hidden layer (l_1);
- epoch_times (ET): an integer for the number of epochs training times;
- learning_rate (LR): an integer for the learning rate of parameters
- Number_EpLogLike (NE): Number of times we simulate λ_w to estimate the Expected-loglikelihood Score inside the KL to optimize.
- batch_psize (BS): an integer for the proportion of batch training data;
- Simulate_W_b-Pred (SWbP): After estimation of variational λ_w , Number of times We simulate W and b from to get a sample of (W,b) ready to be used for sampling the predicted test data from the model.
- Pred_Simulate_Ytest (PSY): Number of times we simulate Ytest given [each] W and b from our previous sample of (W, b). It means, for each W,b, we sample "Pred_Simulate_Ytest" times the Ytest predicted response data from the model.
- Simulate_Proba_Partition (SPP): After estimation of variational λ_w , Number of times We simulate W and b from to get a sample of (W, b) ready to be used to estimate the probability of acceptance each partition from the model.
- The smooth δ applied to smooth slightly every prediction [see 6]

- The number of times we backpropagate for the updates $\psi_{i,t} = (W_{i,t}^{(1)}, W_{i,t}^{(2)}, b_{i,t}^{(1)}, b_{i,t}^{(2)}, \Sigma_{i,t})$ for all the model parameters [see 7.1].
- The learning rate vector ϵ_{dbs} for all the parameters involved in the **DBS** optimizer.

The experiments were performed on 5 multivariate multiple regression datasets (Slump, EDM, Jura, Water Quality, and SCPF) taken from the multiple-output benchmark datasets available in the Mulan project website (Tsoumakas et al., 2020). The datasets are shown in Table 1.

Table 1: Summary of the five (5) datasets taken from the Mulan project.

Dataset	Domain	Number of Instances	Number of features	Response variable dimension (targets)
Slump	Concrete	103	7	3
EDM	Machining	154	16	2
Jura	Geology	359	15	3
Water quality	Biology	1 060	14	16
SCPF	Forecast	1 137	23	3

Evaluation metrics. To directly benchmark with state-of-the-art algorithms, we measure the performance of our Potts neural Gibbs networks with the Mean Square Error (MSE), and the commonly-used Relative Root Mean Squared Error (RRMSE)

$$\text{RRMSE} = \sqrt{\frac{\sum_{(x_i, y_i) \in \mathcal{D}} (\hat{y}_i - y_i)^2}{\sum_{(x_i, y_i) \in \mathcal{D}} (\hat{Y} - y_i)^2}}$$

where (x_i, y_i) , the i -th sample in the testing-set $\mathcal{D}$, is composed of features x_i and ground truth y_i ; $\hat{y}_i$ is the model prediction of y_i ; and $\hat{Y}$ is the average of the adjusted values $\hat{y}$ over the training-set. We take the average RRMSE across all the response dimensions (target variables) within the testing-set D_{test} . A lower aRRMSE indicates better performance. The MMR model has already substantially outperformed the best results from state-of-the-art algorithms on most of these eleven datasets. But, one counterfactual thing we have noticed is that the Relative Root Mean Squared Error (RRMSE) is not a measure of goodness of fit.

THE SIZES OF THE TESTING DATA IS 20% OF THE WHOLE DATASET EACH TIME.

12.1 The Results

The MMR Relative Root Mean Square Error (RRMSE) is beatable, but RRMSE is still not the best... Table 2 displays our multi-target prediction performance results on the 3 datasets from the Mulan project (Slump, EDM and SCPF). In the last column of the table you can compare our results with the multi-layer multi-target regression (MMR) RRMSE of Zhen et al. (2017). As you can see, when running the B-Sparse Gibbs model on Slump dataset, we have achieved a better RRMSE of 45.64 in comparison to the MMR model (58.70). This is true for the SCPF dataset, with a better RRMSE of 48.07 in comparison to the MMR model (81.20). Our goal with table 2 is to exhibit some datasets for which the RRMSE was low in comparison to the MMR model, but the MSE for the test data was not so much attractive. In other words, we have achieved better Relative Root Mean Square Error predictions, but the Mean Square Error of our predictions need more work or fine tuning to be competitive.

Table 2: Average root mean squared errors RRMSE (%) and mean squared error MSE Results with the shallow Gibbs Networks

	Shallow Gibbs settings								Measure of good fit S-Gibbs			MMR
Data sets	NH	ET	LR	NE	BS	SWbP	PSY	SPP	RMSE	RMSE	RRMSE	RRMSE
									Train	Test	-	-
Slump (FG)	1	1	$10e(-3)$	5	0.8	5	5	5	1433.99	482.73	45.64	58.70
EDM	1	1	$10e(-4)$	5	0.8	5	5	5	7.24	1.27	87.93	71.6
SCPF (BG)	1	1	$10e(-4)$	5	0.8	5	5	5	470.14	147.33	48.07	81.20

**We have generated 5 Potts partitions with a shrinkage constraint of 5 for each of the dataset, RMSE=Root Mean Squared Error, RRMSE=Relative Root Mean Squared Error; FG= Full Gibbs, BG = B-Sparse Gibbs, No DBS optimization and no δ -smoothing is applied to the data*

A broad view of the Model in terms of Mean Squared Error. We have performed many experiments before adapting and fixing our model parameters and hyper-parameters initialization included. You can check for some tables in appendix chapter [A] for some other results that may bring information about how the model behaves. As a reminder, our goal is to keep the model as simple as possible and propose an efficient estimation that can compete the Classical N_k Hidden Layer Multifeedforward Model in [??]. The proposed Potts-Gibbs-NM models has not achieved attractive Mean Square Error compare to the N_k Hidden Layer Multi-feedforward until we have added the **DBS** optimizer [See 12.1].

Some Comments about the model structures. We have designed the sparse and B-sparse Gibbs network for training with a small number of observations. As a side-effect, our sparse networks are also computationally fast to train. In fact, we observed in practice that in general, and in a relative comparison to the fully-connected Gibbs network (given the machine used), the sparse network structures speed up training to $2.5x \geq$ (sparse-Gibbs) and $1.5x \geq$ (B-sparse-Gibbs), all while doing a single training (epochs) at a time, specially for Slump dataset [see 4, 7, 5].

When we look at the sparse network structures, we also observe that the gradients oscillate or vanish less than in the fully-connected and compound symmetry Gibbs Networks. Moreover, we can argue (generally saying) that the sparse structured networks is the most effective at reducing the RRMSE than the fully-connected Gibbs Network, more effective than the compound symmetry Gibbs Network variants [4, 7, 5]. Of course, it depends on the dataset, and this finding is not consistent. However, it is impressive that sparse models can match and beat the performance of more dense networks with fewer weight parameters. Probably, one of the reasons for this behavior is that sparse neural network models are known to capture features useful to uncover broader and more general aspects of the data, resulting in better (learning) prediction (Liu, 2020). As many other works have shown (Srinivas et al., 2017), our work shows that sparse learning is possible.

The compound symmetry structure imposed on the precision matrix Ω is conceived so as to simplify the model in terms of parsimony. As opposed to the fully-connected Gibbs network for which the unstructured precision matrix specifies no patterns in the weight spatial connection (that is, the precision matrix is completely general), the compound symmetry structure specifies that weights on the same layer have homogeneous (or nearly equal) covariances, and homogeneous (or nearly equal) variances. In practice, this saves a lot in terms of number of parameters, leading to training times up to more faster than the fully-connected network in some cases. Obviously, these results cannot be over-generalized too. There is an open door for research about conceiving a good Bayesian tests (Mulder & Fox, 2013), or other model selection mechanisms, to choose the best suitable model among the shallow Gibbs networks for a given dataset.

Table 3: Average root mean squared errors RRMSE (%) and and mean squared error MSE Results with the shallow Gibbs Networks: Sparse-Gibbs

	Shallow Gibbs settings								Measure of fitness S-Gibbs			Time (sec.)
Data sets	NH	ET	LR	NE	BS	SWbP	PSY	SPP	RMSE	RMSE	RRMSE	Time
									Train	Test	-	
Slump	1	1	$10e(-3)$	5	0.8	5	5	5	67.36	67.36	99.34	116.12
Jura	1	3	$10e(-3)$	2	0.3	5	5	5	32.94	35.99	100.0	1385.67
SCPF	1	1	$10e(-4)$	5	0.8	5	5	5	111.82	41.30	116.57	633.94

**We have generated 5 Potts partitions with a shrinkage constraint of 5, RMSE=Root Mean Squared Error, RRMSE=Relative Root Mean Squared Error; DBS epoch times =3, DBS learning rate = $10e - 3$, and no smoother δ is applied*

Table 4: Average root mean squared errors RRMSE (%) and and mean squared error MSE Results with the shallow Gibbs Networks: B-Sparse Gibbs

	Shallow Gibbs settings								Measure of fitness S-Gibbs			Time (sec.)
Data sets	NH	ET	LR	NE	BS	SWbP	PSY	SPP	RMSE	RMSE	RRMSE	Time
									Train	Test	-	
Slump	1	1	$10e(-3)$	2	0.8	5	5	5	66.32	67.23	99.65	825.62
Jura	1	3	$10e(-2)$	2	0.3	5	5	5	32.57	35.69	103.05	2133.79
SCPF	1	1	$10e(-4)$	5	0.8	5	5	5	111.82	40.25	136.69	950.19

**We have generated 5 Potts partitions with a shrinkage constraint of 5, RMSE=Root Mean Squared Error, RRMSE=Relative Root Mean Squared Error; DBS epoch times =3, DBS learning rate = $10e - 5$, and no smoother δ is applied*

Table 5: Average root mean squared errors RRMSE (%) and and mean squared error MSE Results with the shallow Gibbs Networks: Full-Gibbs

	Shallow Gibbs settings								Measure of fitness S-Gibbs			Time (sec.)
Data sets	NH	ET	LR	NE	BS	SWbP	PSY	SPP	RMSE	RMSE	RRMSE	Time
									Train	Test	-	
Slump	1	1	$10e(-5)$	2	0.8	5	5	5	70.06	70.68	98.46	2156.10
Jura	1	3	$10e(-2)$	5	0.8	5	5	5	32.70	35.70	110.26	5349.34
SCPF	1	1	$10e(-4)$	5	0.8	5	5	5	111.82	31.30	186.87	428.97

**We have generated 5 Potts partitions with a shrinkage constraint of 5, RMSE=Root Mean Squared Error, RRMSE=Relative Root Mean Squared Error; DBS epoch times =3, DBS learning rate = $10e - 5$, and no smoother δ is applied*

Table 6: Average root mean squared errors RRMSE (%) and and mean squared error MSE Results with the shallow Gibbs Networks: Sparse-CS-Gibbs

	Shallow Gibbs settings								Measure of fitness S-Gibbs			Time (sec.)
Data sets	NH	ET	LR	NE	BS	SWbP	PSY	SPP	RMSE	RMSE	RRMSE	Time
									Train	Test	-	
Slump	1	1	$10e(-5)$	2	0.8	5	5	5	66.90	68.21	100.64	1102.48
Jura	1	3	$10e(-2)$	5	0.8	5	5	5	32.94	38.39	190.26	2929.34
SCPF	1	1	$10e(-4)$	5	0.8	5	5	5	111.82	51.89	171.82	430.84

**We have generated 5 Potts partitions with a shrinkage constraint of 5, RMSE=Root Mean Squared Error, RRMSE=Relative Root Mean Squared Error; DBS epoch times =3, DBS learning rate = $10e - 5$, and no smoother δ is applied*

Table 7: Average root mean squared errors RRMSE (%) and mean squared error MSE Results with the shallow Gibbs Networks: CS-Gibbs

	Shallow Gibbs settings								Measure of fitness S-Gibbs			Time (sec.)
Data sets	NH	ET	LR	NE	BS	SWbP	PSY	SPP	RMSE	RMSE	RRMSE	Time
									Train	Test	-	
Slump	1	1	$10e(-5)$	2	0.8	5	5	5	70.06	70.68	98.46	456.10
Jura	1	3	$10e(-2)$	5	0.8	5	5	5	32.70	45.70	110.26	1349.34
SCPF	1	1	$10e(-4)$	5	0.8	5	5	5	111.82	41.30	116.57	798.09

*We have generated 5 Potts partitions with a shrinkage constraint of 5, RMSE=Root Mean Squared Error, RRMSE=Relative Root Mean Squared Error; DBS epoch times =3, DBS learning rate = $10e - 5$, and no smoother δ is applied

Convergence of the DBS optimizer. We have found experimentally that the **DBS** optimizer is convergent. We have choosen **Slump** dataset to illustrate this finding, look for table [8] to be convinced.

To reach this convergence, we have identify the need to modify this optimizer with a slight correction. Mainly $\frac{\partial MSE(y - \hat{y}_{est})}{\partial \hat{y}_{est}}$ has to integrate the updates system as follows:

$$\hat{y}_{est,(i,t)} \leftarrow f_{\psi_{i,t}}(x) + L_{i,t} \cdot u_{i,t} - \epsilon_{\hat{y}_{est,(i,t)}} \frac{\partial MSE(y - \hat{y}_{est})}{\partial \hat{y}_{est}} \quad (41)$$

where $\frac{\partial MSE(y - \hat{y}_{est,(i,t)})}{\partial \hat{y}_{est}} = 2 * (y_i - \hat{y}_{est,(i,t)})$.

So each training data has its own learning rate which is here set as $\epsilon_{\hat{y}_{est,(i,t)}}$. This is valuable for each test data as follows:

$$\hat{y}_{est,(i,t)}^{test} \leftarrow f_{\psi_{i,t}}(x^{test}) + L_{i,t} \cdot u_{i,t} - \epsilon_{\hat{y}_{est,(i,t)}}^{test} \frac{\partial MSE(y - \hat{y}_{est})}{\partial \hat{y}_{est}} \quad (42)$$

where for the k -th test data x_k^{test} the changes $f_{\psi_{i,t}}, L_{i,t}, u_{i,t}$ are taken from the j -th training data y_j which verify:

$$j^{chosen} = \arg \min_{x_j \in \text{Training Set}} \text{Mean}(x_j - x_k^{test}) \quad (43)$$

where the operation $\text{Mean}(u)$ for vector u is taken upon all dimension of u . In our experiments with **Slump** dataset, we have choosen $\epsilon_{\hat{y}_{est,(i,t)}}^{test}$ and $\epsilon_{\hat{y}_{est,(i,t)}}$ to be double time larger than:

$$\epsilon_{\psi_{i,t}} = (\epsilon_{W_{i,t}^{(1)}}, \epsilon_{W_{i,t}^{(2)}}, \epsilon_{b_{i,t}^{(1)}}, \epsilon_{b_{i,t}^{(2)}}, \epsilon_{\Sigma_{i,t}})$$

In equation [45], $\frac{\partial MSE(y - \hat{y}_{est})}{\partial \hat{y}_{est}}$ has to be approximated as :

$$\frac{\partial MSE(y - \hat{y}_{est})}{\partial \hat{y}_{est}} \approx 2 * (y_{j^{chosen}}^{train} - \hat{y}_{est,(j^{chosen},t)}^{train}). \quad (44)$$

When the training error was getting lower, with a not decreasing test error, we double again $\epsilon_{\hat{y}_{est,(i,t)}}^{\text{test}}$ compared to $\epsilon_{\hat{y}_{est,(i,t)}}^{\text{train}}$ specifically. The convergence property of the **DBS** optimizer in equation [45] for the training data is not surprising due to the convergence of gradient descent properties [Lee et al. (2016); Ruder (2016)] for the Mean Square Error $MSE(\hat{y}_{est}) = \|y - \hat{y}_{est}\|_F^2$ which is convex in $\hat{y}_{est}$. However, because of the approximation [44], this convergence is not reached with the same *speed* for the test data, but the model still do converge. In their case, to reduce test errors to (zero), a last ingredient needs to be applied, called *Universal Approximation Theorem*:

Proposition 12.1 (Guilhoto (2018))

A Neural Network with a single hidden layer and sufficiently many nodes is capable of approximating any continuous function.

Intuitively, it means that we need to increase the number of neurons within the hidden layer, to reach a good prediction for the test data. This is also valid if the model was multi-layer architecture. To understand this intuition, look for equation [45], and beware that when convergence is reached for the training data, we have :

$$\frac{\partial MSE(y - \hat{y}_{est,(i,t)})}{\partial \hat{y}_{est}} = 2 * (y_i - \hat{y}_{est,(i,t)}) = 0$$

and equation [45] for the test data x^{test} becomes:

$$\hat{y}_{est,(i,t)}^{\text{test}} \longleftarrow f_{\psi_{i,t}}(x^{\text{test}}) + L_{i,t} \cdot u_{i,t} \quad (45)$$

We simply need then to find the right (neural) machine that will associate with perfection x^{test} to y^{test} , as:

$$\hat{y}_{est,(i)}^{\text{test}} = f_{\psi_i}(x) + L_i \cdot u_i = [b_i^{(2)} + g_1(b_i^{(1)} + x_i^{\text{test}} W_i^{(1)}) W_i^{(2)}] + L_i \cdot u_i$$

Or simply, very like an encoder or neural cryptography machine [Ruttor et al. (2007), Ruttor et al. (2004), Volná (2000), Kanter & Kinzel (2003), Pattanayak & Ludwig (2017), Hadke & Kale (2016), Sharma et al. (2019), Kinzel & Kanter (2002a), Kinzel & Kanter (2002b), Dong & Huang (2019), Volna et al. (2012), Klimov et al. (2002), Godhavari et al. (2005), Blackledge et al. (2015), Baird et al. (2005), Behrmann et al. (2019)]:

$$x_i^{\text{test}} \xleftrightarrow[y_i^{\text{test}}]{\psi_i^{\text{test}}} = (W_i^{(1)}, W_i^{(2)}, b_i^{(1)}, b_i^{(2)}, \Sigma_i) \quad (46)$$

To obtain this machine ψ_i^{test} which has to encode as much as information to be efficient, we need to extract it from its associate training data in equation⁶ [43], and to make it robust, the *Universal Approximation Theorem* [see 12.1] requires $[W_i^{(1)}, W_i^{(2)}, b_i^{(1)}]$ to be high dimensional⁷, i.e. an appropriate number of neurons on the hidden layer. This last requirement transforms the **DBS** optimizer into another **DBS** optimizer [the combination of *Universal Approximation Theorem* and **DBS** optimization] with an optimal number l_1 of neurons on the hidden layer [see 2] of the Shallow net, an adapted number ζ of **DBS** optimization, an optimal DBS learning rate ϵ_{dbs} , called the $(l_1, \zeta, \epsilon_{dbs}) - \mathbf{DBS}$, and which, combined with findings of section [9], i.e. the *(dist)-Nearest Neighbor-(h)-Taylor Series-Perfect Multivariate Interpolation (dist-NN-(h)-TS-PMI)* presented in [10], is the Perfect fit⁸ (or the Perfect learning) for the Shallow Gibbs Network, summarized in equation [47]:

6. This equation can be improved too for better Train-Test association.

7. Where the model can also be multi-layer as well.

8. Chapter [9] is added to reinforce the $(l_1, \zeta, \epsilon_{dbs}) - \mathbf{DBS}$ model.

$$\lim_{l_1, \text{opt}, \zeta, \text{opt}, \epsilon_{\text{dbs}}, \text{opt}, \text{dist}_{\text{opt}}, h_{\text{opt}}} (MSE^{\text{Train}}, MSE^{\text{Test}}) = (0, 0) \quad (47)$$

where $MSE^{\text{Train}}, MSE^{\text{Test}}$ are the Mean Squared Error of the train and test data respectively, dist_{opt} is the optimal distance for the research of the nearest neighbor in the training dataset for each test data x_i^{test} , h_{opt} is the optimal order of the Taylor approximation for the Perfect Multivariate Interpolation ($\text{dist-NN-(h)-TS-PMI}$) model once the $(l_1, \zeta, \epsilon_{\text{dbs}}) - \mathbf{DBS}$ has overfitted the training dataset. $l_1, \text{opt}, \zeta, \text{opt}$ are respectively the optimal number of hidden neurons, and the optimal number of $\mathbf{DBS}$ updates. ϵ_{dbs} integrates simultaneously the DBS learning rate vector for all the model parameters, the DBS learning rate for the training data, and the DBS learning rate for the test data. $\epsilon_{\text{dbs}, \text{opt}}$ is the optimal one. We have performed two more experiments to confirm this fact, summarized in table [9]. Table [9] may be treated as over-fitting, but the advantage of this finding is a confirmation of the convergence ($\downarrow 0$) of the with $(l_1, \zeta, \epsilon_{\text{dbs}}) - \mathbf{DBS}$, where we may applied a simple optimization –[Grid Search or Golden Section Search method]– to find the perfect parametrization of (l_1, ζ) and the required learning rate ϵ_{dbs} .

We have also found in practice, that when l_1 and ζ are too large, the model may diverge. So, it is important to find the right ratio ζ/l_1 (with appropriate $\mathbf{DBS}$ learning rate).

When there is enough training data available, another way to increase the convergence speed for the test data is to modify the criteria used in optimization [43], and use a distance dist for which each test data x_k^{test} is ensured to find an associate x_i^{train} in the training data with :

$$\text{dist}(x_k^{\text{test}}, x_i^{\text{train}}) \leq \varepsilon \quad (48)$$

where ε is a very small number.

Table 8: Average root mean squared errors RRMSE (%) and and mean squared error MSE Results with the shallow Gibbs Networks on Slump dataset with DBS optimization

Shallow Gibbs settings with DBS applied on Slump dataset					Measure of fitness Sparse-Gibbs with Slump dataset			Time (sec.)
NH	ET	LR	DBS	DBS learning rate	RMSE	RMSE	RRMSE	Time
					Train	Test	-	
1	1	$10e(-5)$	3	$10e(-3)$	66.60	67.36	99.34	116.12
1	1	$10e(-5)$	10	$10e(-3)$	46.35	50.06	99.15	100.02
1	1	$10e(-5)$	30	$10e(-3)$	43.56	49.13	102.85	4239.41
1	1	$10e(-5)$	30	$10e(-3)$	43.56	49.13	102.85	4239.41
1	1	$10e(-5)$	30	$10e(-6)$	47.66	53.13	102.51	4300.36
1	1	$10e(-5)$	50	$10e(-3)$	14.20	30.24	130.59	9358.63
1	1	$10e(-5)$	70	$10e(-3)$	9.20	27.19(*)	128.11	10609.91

*We have generated 5 Potts partitions with a shrinkage constraint of 5, RMSE=Root Mean Squared Error, RRMSE=Relative Root Mean Squared Error; DBS epoch times =3, NE = 2, BS = 0.8, SWbP = PSY = SPP = 5, DBS learning rate = $10e - 2$, and no smoother δ is applied; test data DBS learning rate = $2 * (10e - 2)$, (*) the DBS learning rate of the test data is twice the learning rate of the training data

Model Execution Time. The results we have presented above are relatively comparable in terms of execution times, simply because during all the simulations and model trials, we have used different machines with different capacities, sometimes not comparable in terms of computation and degree of availability to run large or huge multiple simulations at a time [look in table 10]. Simply to say, our time comparison should be taken as relative for it require precaution and regards about the machine used to run the model, not absolute. Also, whatever the

Table 9: Average root mean squared errors RRMSE (%) and and mean squared error MSE Results with the shallow Gibbs Networks on Slump dataset with DBS optimization

Shallow Gibbs settings with DBS applied on Slump dataset					Measure of fitness Sparse-Gibbs with Slump dataset			Time (sec.)
NH	ET	LR	DBS	DBS learning rate	RMSE	RMSE	RRMSE	Time
					Train	Test	-	
6	1	$10e(-5)$	3	$10e(-2)$	2.05	26.30(*)	126.20	6814.34
10	1	$10e(-5)$	40	$10e(-2)$	3.11	25.89(*)	127.26	5894.04

*We have generated 5 Potts partitions with a shrinkage constraint of 5, $RMSE=$ Root Mean Squared Error, $RRMSE=$ Relative Root Mean Squared Error; DBS epoch times =3, $NE = 2$, $BS = 0.8$, $SWbP = PSY = SPP = 5$, DBS learning rate = $10e - 2$, and no smoother δ is applied; test data DBS learning rate = $2 * (10e - 2)$, (*) the DBS learning rate of the test data is twice the learning rate of the training data

machine being used, increasing the training times (ET==epoch.times) for the neural networks in each cluster of the Shallow Gibbs model is costly in CPUs [Figure 2]. The CPUs available are overfull in less than 20 minutes of execution. To solve this problem, we have applied **Reusable Process Pool Executor** which is a novel framework that combines threading and multiprocessing primitives for robust concurrent futures (McCullagh (2017)).

Table 10: Machines used during our simulation and their main characteristics

Machine Host	Cpus	C.L.	Mem	Mem L	GPU	G. Nom	CPU S	CPU M
simulation7	24	10	251	133	-	-	13030	312720
simulation8	24	12	251	111	-	-	13030	312720
venice	12	8	15	8	0	GTX 1050	15971	191652
jupiter	8	5	11	6	-	-	5200	41600
fox	12	7	15	5	0	GTX 1050	15971	191652
acapulco	12	3	15	2	0	GTX 1050	15971	191652
panthere	12	8	15	2	0	GTX 1050	15971	191652
lion	12	0	15	1	0	GTX 1050	15971	191652
jaguar	12	7	15	1	0	GTX 1050	15971	191652
Cpus =>	Cores				C.L. =>	Free Cpu		
Mem =>	Total Memory				Mem L =>	Available Memory (Go)		
GPU =>	GPUs used				G. Nom =>	GPU model		
CPU S =>	Cpu Mark single thread				CPU M =>	Cpu Mark total multi threads		

Concluding Remarks & Discussion notes

12.2 Note on the Shallow Gibbs Structure.

We present in chapter ?? a framework to build a new structured *Potts clustered Gibbs Multivariate Regression model* which is properly say is Random Gibbs Neural Network Forest (Barber & Bishop (1997), Barber & Bishop (1998)), by combining structured precision matrixes, Potts Neural Network Regression (SPNNR), variational learning and backpropagation. You may call the model under many appellations: Shallow Gibbs Network, Random Gibbs Network Forest, Shallow Potts Neural Network, the Potts-Gibbs-NM model etc. The model has four variants: the fully-connected Gibbs network or full Gibbs network (Full-Gibbs, for short), the between-layer sparse Gibbs network (B-Sparse-Gibbs), the sparse shallow Gibbs network (Sparse-Gibbs), the sparse compound symmetry Gibbs network (CS-Sparse-Gibbs), and the compound symmetry Gibbs network (CS-Gibbs). As literature references to many alike sparse-structured or similar models, you may look for the work from: Ionescu et al.

Figure 2: Percentage of CPUs charged by the Shallow Gibbs during execution for a training times ≥ 10 . The right axis is RAM memory in Mo.

(2015), Kepner & Robinett (2019), Wainwright (2014), Louizos & Welling (2016), Paul et al. (2016), Schwing & Urtasun (2015), Ardakani et al. (2016), Mazza-Anthony (2019), with an exception of the Compound symmetry Gibbs network, which is a particular *innovation* of our research. We can even say that one main innovation of this research is the introduction of *compound symmetry* precision matrix on the weights on our Bayesian Shallow neural networks. The effectiveness of our results will definitely increase interests in those types of matrices.

The neural network (our base learner) weights are Gaussian Markov Random Field distributed. To offer flexibility in terms of model structure, the precision matrix has been modified to infer three variants of analysis : sparse analysis, compound symmetry analysis and fully-connected framework. Those models have the properties to adapt themselves easily to the data in a few iterations during the learning process. The model is mounted in three stages: a set of data $\mathbf{y}^n = \{y_1, \dots, y_n\} \subset \mathbb{R}^p$ with associate covariables ; each y_i following a Gaussian distribution $p(y|x, \psi, \Sigma) = (2\pi)^{-p/2} |\Sigma|^{-1/2} \exp\{-1/2(y - f_\psi(x))' \Sigma^{-1} (y - f_\psi(x))\}$, with covariance Σ and mean $f_\psi(x)$ which represent specifically a neural network function of the covariable x . The weights and the biases of the model vectorized as $\psi = (b^{(1)}, W^{(1)}, b^{(2)}, W^{(2)})$ are high dimensional and their size can be increased at will. But, keep in mind that the goal of the Shallow Gibbs Network is to build a simplified model with few neurons and better mean *squared* error in a *very* short training time (reduced number of epoch ≤ 10).

First, we define the Kullback-Leibler of the model variational distribution [See 5] that will help to simulate variational parameters. This approach is very rich in terms of properties and convergence (da Rocha et al. (2011), Wiegierinck & Kappen (2000), Challis & Barber (2013)); even though it suffers from many drawbacks as the positive definiteness of your precision/covariances matrices during gradients updates [see 5.2]. Gradients updates have good advantages such as a fact approximations in a few iterative steps, mainly if you tweak appropriately the learning rate. When you applied an approximation for your parameters or hyper-parameters during the updates, you step forward the **Iterative Projected Gradient** (IPG) updates approach (Cruz et al. (2011)) which converge also if the projection space is an optimal subspace. As the case of PSD - Positive Semi Definite - approximation is current in literature, we can index this paper "**Approximation by matrices positive semidefinite on a subspace**" from Hayden & Wells (1988) that was a great help in building our model.

It is no need to say that the *double* backpropagation scheme in [see 7.1] is truly a requirement to reach an outstanding result. The secret behind is that the second backpropagation system represents itself another *neural network* by its own. In other words, given a set of initial predictions, for the data (x, y) and an estimated/or random values for $\psi = (b^{(1)}, W^{(1)}, b^{(2)}, W^{(2)})$ and Σ , it reduces naturally the model error by backpropagating under a suitable

learning rate schedule. Our experiments and results were not only efficient in comparison to the MMR model Zhen et al. (2017), but also was impressive against Murua & Quintana (2017) model.

The four (5) models developed in this framework include a clustered neural network regression via a Potts model. Monte Carlo simulations are limited to run posterior simulation for these models, and variational inference was the last resort. This novel representation (framework) aims to capture also more complex pattern in small datasets with high dimensional features. We have increased relative prediction power compare to simple prediction, and offered to the novice an adapted regression model to heterogeneous data, that can surpass any Multilayer Feed Forward Neural Network (FFNN) and shall over-become the next revolution of neural networks. And finally, we have found for the Shallow Gibbs Network model [2], the perfect learning configuration is: the *dist*-NN-(h)-TS-PMI)-(l_1, ζ) – **DBS**, which is a combination of the *Universal Approximation Theorem*, and the DBS optimization, all coupled with the (*dist*)-Nearest Neighbor-(h)-Taylor Series-Perfect Multivariate Interpolation (*dist*-NN-(h)-TS-PMI). It indicates that, with an optimal number l_1 of neurons on the hidden layer, and an adapted ζ of DBS updates, an optimal distance $dist_{opt}$ in the research of the nearest neighbor in the training dataset for each test data x_i^{test} , an optimal order h_{opt} of the Taylor approximation for the Perfect Multivariate Interpolation (*dist*-NN-(h)-TS-PMI) model once the **DBS** has overfitted the training dataset, the train and the test error converge to zero (0). **This finding is a great revolution in all fields and subfields of Statistical Learning Theory, from the now to the forever.** The generalization power of this model is infinite, as the secret is to understand that every augmented data can also be sufficiently differentiated, as much as we want to overlap all the possible data Space domain:

$$x_{j_1}^{\text{aug-train}} = (x_{j_1}^{\text{train}}(1) + \delta_a, \dots, x_{j_1}^{\text{train}}(q)) \longrightarrow \begin{cases} (x_{j_1}^{\text{train}}(1) + s_1 \cdot \delta_a, \dots, x_{j_1}^{\text{train}}(q)) \\ (x_{j_1}^{\text{train}}(1), x_{j_1}^{\text{train}}(2) + s_2 \cdot \delta_a, \dots, x_{j_1}^{\text{train}}(q)) \\ \dots\dots\dots \\ (x_{j_1}^{\text{train}}(1), x_{j_1}^{\text{train}}(2), \dots, x_{j_1}^{\text{train}}(q) + s_q \cdot \delta_a) \end{cases} \quad (49)$$

with $(s_1, s_2, \dots, s_q) \in \mathbb{N}^q$ is a q -tuple of integers or rationals, where $x_{j_1}^{\text{aug-train}}$ is our augmented data in the training set, and $\delta_a \in \mathbb{R}$ a very small number. The re-training process has to be done progressively to obtain their perfect associated parameters [see 11].

12.3 A Generalized Double Back-Propagation Scheme (GDBS) for any parametric model

We propose an effective Generalized Double Back-Propagation for any parametric model, augmented with a differential and local neighborhood machine learning framework for almost sure convergence. As an extension of section [7.1], we propose a general double back-propagation scheme (GDBS) using the Mean Squared Error (MSE), and for any (parametric) model with parameter ψ as :

$$\hat{y}_{est,i} = f_{\psi}(x_i) \quad (50)$$

1. Using the equation (in 50), apply any suitable optimization framework to obtain a general estimate $\hat{\psi}_0$ of ψ for all the model (in a first step).
2. Use again equation 50 to backpropagate the $MSE(y_i - \hat{y}_i) = \|y_i - \hat{y}_{est,i}\|^2$ to update $\hat{\psi}_0$ per observation as follows:

$$\hat{\psi}_{1,i} \longleftarrow \hat{\psi}_0 - \epsilon_{\psi,0} \frac{\partial MSE(y_i - \hat{y}_{est,i})}{\partial \psi} \quad (51)$$

$$\hat{\psi}_{t,i} \leftarrow \hat{\psi}_{t-1,i} - \epsilon_{\psi,t-1} \frac{\partial MSE(y_i - \hat{y}_{est,i})}{\partial \psi} \quad (52)$$

where $\epsilon_{\psi,t}$ is the learning rate schedule [See C.1] for this gradient update of ψ at step t , and $\hat{\psi}_{t,i}$ is the value of ψ at iteration t for observation i .

To reach convergence, above assignments have to integrate updates for y as follows:

$$\hat{y}_{est,(i,t)} \leftarrow f_{\psi_{i,t}}(x) + L_{i,t} \cdot u_{i,t} - \epsilon_{\hat{y}_{est,(i,t)}} \frac{\partial MSE(y_i - \hat{y}_{est,i})}{\partial \hat{y}_{est,i}} \quad (53)$$

where $\frac{\partial MSE(y_i - \hat{y}_{est,(i,t)})}{\partial \hat{y}_{est,i}} = 2 * (y_i - \hat{y}_{est,(i,t)})$.

So each training data has its own learning rate which is here set as $\epsilon_{\hat{y}_{est,(i,t)}}$. This is valuable for each test data as follows:

$$\hat{y}_{est,(i,t)}^{\text{test}} \leftarrow f_{\psi_{i,t}}(x^{\text{test}}) + L_{i,t} \cdot u_{i,t} - \epsilon_{\hat{y}_{est,(i,t)}}^{\text{test}} \frac{\partial MSE(y_i - \hat{y}_{est,i})}{\partial \hat{y}_{est,i}} \quad (54)$$

where for the k -th test data x_k^{test} the changes $f_{\psi_{i,t}}$ are taken from the j -th training data y_j which verify:

$$j^{\text{choosen}} = \arg \min_{x_j \in \text{Training Set}} \text{Mean}(x_j - x_k^{\text{test}}) \quad (55)$$

where the operation $\text{Mean}(u)$ for vector u is taken upon all dimension of u . The criteria used in optimization [43] can be modify for a distance $dist$ for which each test data x_k^{test} is ensured to find an associate x_i^{train} in the training data with :

$$dist(x_k^{\text{test}}, x_i^{\text{train}}) \leq \varepsilon \quad (56)$$

where ε is a very small number. This presented framework will be called $(\zeta, \epsilon_{dbs}) - \mathbf{GDBS}$, and augmented with the data Augmentation for Empirical Differentiation (DAED) framework in section 10, shall be called the $dist\text{-NN(h)-TS-PMI}-(l_1, \zeta, \epsilon_{dbs}) - \mathbf{GDBS}$. When the model is truly differentiable, with assumption $(\mathcal{F}_d)$ being valid, the learning with this model is almost surely perfect. This simply means in one equation [57]:

$$\lim_{l_1, \text{opt}, \zeta, \text{opt}, \epsilon_{dbs}, \text{opt}, \text{dist}, \text{opt}, h, \text{opt}} (MSE^{\text{Train}}, MSE^{\text{Test}}) = (0, 0) \quad (57)$$

12.4 The Infinite Zelda Stochastic Game.

This is a stochastic game framework we have derived for a direct application scheme of the Potts Shrinkage model coupled with the Shallow Potts Gibbs Models. We all know that the next generation of games will come from artificial intelligence (AI) networks (Hsu (2004), Hassabis (2017), DeCoste (1997), Granter et al. (2017)). The

core idea of this part is inspired from Bowling & Veloso (2002) that have described a scalable learning algorithm for stochastic games, and mainly the equivalence between Potts Model and percolation (Essam (1979), Ding et al. (2012), Kemppainen & Smirnov (2019)). We present here the Infinite Zelda Game (IZG) that rely on the learning process of the Shallow Gibbs Network (SGN) model.

The random partitions can be seen as random graphs. Each partition from the random bond Potts models used for its simulation presents some connected components and isolated elements. With a given subset of data $\{(x_i, y_i)\}_{1 \leq i \leq N}$, the graph generated at each step is equivalent to a network graph built from pair interactions between the neighbors. Regardless of the connectivity structure, the graph has v^N state (s) if each data point is characterized by v possible states (or spins). Each of the v states are represented graphically by a color. At each step (a new distinct *generated* partition from a previous one), y_i can move from state v_1 to v_2 , thanks to the clustering process, and its estimate from the Shallow Gibbs Network (SGN) is denoted $\hat{y}_i$.

Figure 3: A 3-component Potts graph of size 8. Shrinkage constraint is reduced to 1.

The green colored circle shall indicate isolated points, and the same colored squared indicates a set of points all connected forming a component. Given the size of the hidden layer of the Shallow Gibbs Network (SGN), there two possibilities (denoted pos_1 and pos_2 and detailed in the following) [see figure 3]:

- i. Estimated $\hat{y}_i$ is close (in distance metric) to real y_i more than its assigned cluster S_i estimated mean $\frac{1}{s_i} \sum_{j \in S_i} y_j$, with $|S_i| = s_i$. This possibility is called pos_1 .
- ii. Estimated $\hat{y}_i$ is not close (in distance metric) to y_i more than its its assigned cluster S_i mean $\frac{1}{s_i} \sum_{j \in S_i} y_j$, with $|S_i| = s_i$. This possibility is called pos_2 .

We set each available dataset $\mathcal{D}_s$ to represent a Mansion (the training part). In each Mansion, the player has to maximize its coins collection to be eligible to move out to another Mansion. Each generated partition proposes a conditional estimate $\hat{y}_i$ for $y_i \in \mathbf{Y} = \{y_i, 1 \leq i \leq q\}$ through the Shallow Gibbs Network (SGN). The player has to find *by mouse or finger click* in the network graph, those colored points in the position pos_2 . Obviously, the isolated points (in white color) are always in position pos_2 .

To facilitate the game, we always declare the number b of data points in position pos_2 , and ask the player to find them in b_e trials with $b + 5 = b_e$. Because the partitions are generated by Swenden-Wang cuts (Barbu & Zhu (2003)), the game is set for the player to accumulate coins as many as he detects the right squares in position pos_2 , in few trials, and in a given recorded time t_e .

The Goal of the Infinite Zelda Game (IZG): “Win or Learn More” The goal of the Infinite Zelda Game (IZG) is to maximize your coins through the mansions. The more you win, the quicker you can move to another mansion. One more trick is to be done each time before generating a partition (a new network graph): the player has to choose the number of $dist\text{-}NN\text{-}(h)\text{-}TS\text{-}PMI\text{-}(l_1, \zeta, \epsilon_{dbs}) - \mathbf{DBS}$ updates for the Shallow Gibbs Network (SGN) characterising its *power level*. By default, the power is set to **5**, and he has to buy more coins to set more power.

The more $dist\text{-}NN\text{-}(h)\text{-}TS\text{-}PMI\text{-}(l_1, \zeta, \epsilon_{dbs}) - \mathbf{DBS}$ updates you set, the more you may win at each partition configuration, because it reduces the number of data points (y_i) in position pos_2 with a Potts Shrinkage constraint $\geq m \geq 5$: the more *power level* you set, the more you reduce the number of positions pos_2 to find. The Infinite Zelda Game (IZG) for the Shallow Gibbs Network (SGN) can be summarized in the WoLM principle (“**Win or Learn More**”) which encourages convergence.

Appendix A. Other Experiments Results with the Shallow Gibbs Models

Table 11: RRMSE (%) and MSE Results with the shallow Gibbs Networks: B-sparse-Gibbs

	Shallow Gibbs settings								Measure of fitness S-Gibbs			Time (sec.)
Data sets	NH	ET	LR	NE	BS	SWbP	PSY	SPP	RMSE	RMSE	RRMSE	Time
									Train	Test	-	
Slump	1	1	$10e(-3)$	10	0.8	5	5	5	1542.67	422.17	399.11	2633.44
Slump	1	1	$10e(-6)$	10	0.8	5	5	5	1351.71	527.27	104.17	2622.03
Slump	1	1	$10e(-6)$	10	0.8	5	5	5	1815.48	778.58	117.12	2616.32
Slump	1	1	$10e(-6)$	10	0.8	5	5	5	3457.77	1479.01	175.50	716.58
EDM	1	1	$10e(-4)$	5	0.8	5	5	5	7.24	1.27	87.93	1160.06
EDM	1	1	$10e(-3)$	5	0.8	5	5	5	6.82	1.55	94.77	1143.09
EDM	1	1	$10e(-3)$	5	0.8	5	5	5	5.33	3.37	129.76	2414.18
EDM	1	1	$10e(-5)$	5	0.8	5	5	5	5.55	2.50	112.12	1155.88
EDM	2	1	$10e(-5)$	5	0.8	5	5	5	10.23	3.07	87.36	1549.87
EDM	10	1	$10e(-5)$	2	0.8	5	5	5	24.17	8.92	192.35	1299.76
Jura	1	10	$10e(-3)$	10	0.8	2	2	2	476.95	174.69	650.35	9833.41
SCPF	1	1	$10e(-4)$	5	0.8	5	5	5	470.14	147.33	48.07	2101.72
SCPF	1	1	$10e(-3)$	5	0.8	5	5	5	324.71	116.63	120.83	2083.83
SCPF	1	1	$10e(-4)$	5	0.8	5	5	5	312.39	130.77	91.36	4432.99
SCPF	1	2	$10e(-3)$	5	0.8	5	5	5	380.89	137.21	60.00	4560.25

**We have generated 5 Potts partitions with a shrinkage constraint of 5, RMSE=Root Mean Squared Error, RRMSE=Relative Root Mean Squared Error; all without the DBS optimizer and the smoother δ*

Table 12: RRMSE (%) and MSE Results with the shallow Gibbs Networks: B-sparse-Gibbs, with different smoother δ applied

	Shallow Gibbs settings								Measure of fitness S-Gibbs			Time (sec.)
Data sets	NH	ET	LR	NE	BS	SWbP	PSY	SPP	RMSE	RMSE	RRMSE	Time
									Train	Test	-	
Slump	1	1	$10e(-3)$	10	0.8	5	5	5	82.47	50.24	83.74	493.65
Slump	1	1	$10e(-6)$	10	0.8	5	5	5	66.09	67.86	100.63	497.35
Slump	1	1	$10e(-6)$	10	0.8	5	5	5	66.87	68.13	100.00	491.90
Jura	1	1	$10e(-3)$	10	0.8	2	2	2	148.55	88.09	90.24	9312.22
Jura	1	1	$10e(-2)$	5	0.8	2	2	2	64.16	52.02	520.49	4305.73
Jura	10	1	$10e(-2)$	5	0.8	5	5	5	48.97	41.33	284.72	14395.23
SCPF	1	1	$10e(-3)$	5	0.8	5	5	5	142.82	40.06	87.93	1997.34
SCPF	1	1	$10e(-4)$	5	0.8	5	5	5	111.84	41.15	97.97	1980.34
SCPF	1	2	$10e(-3)$	5	0.8	5	5	5	380.89	137.21	60.00	4560.25

**We have generated 5 Potts partitions with a shrinkage constraint of 5, RMSE=Root Mean Squared Error, RRMSE=Relative Root Mean Squared Error; with different smoother δ , without the DBS optimizer*

Table 13: RRMSE (%) and MSE Results with the shallow Gibbs Networks: Full-Gibbs

	Shallow Gibbs settings								Measure of fitness S-Gibbs			Time (sec.)
Data sets	NH	ET	LR	NE	BS	SWbP	PSY	SPP	RMSE	RMSE	RRMSE	Time
									Train	Test	-	
Slump	1	1	$10e(-3)$	5	0.8	5	5	5	1433.99	482.73	45.64	1782.12
Slump	10	1	$10e(-3)$	5	0.8	5	5	5	3178.67	1330.54	934.75	2122.28
EDM	1	1	$10e(-4)$	5	0.8	5	5	5	7.24	1.27	87.93	1160.06
EDM	1	1	$10e(-3)$	5	0.8	5	5	5	6.82	1.55	94.77	1143.09
EDM	1	1	$10e(-3)$	5	0.8	5	5	5	5.33	3.37	129.76	2414.18
EDM	1	1	$10e(-5)$	5	0.8	5	5	5	5.55	2.50	112.12	1155.88
EDM	2	1	$10e(-5)$	5	0.8	5	5	5	10.23	3.07	87.36	1549.87
EDM	10	1	$10e(-5)$	2	0.8	5	5	5	24.17	8.92	192.35	1299.76
Jura	1	3	$10e(-3)$	2	0.8	5	5	5	75.12	50.64	102.82	5360.25
Jura	1	3	$10e(-3)$	2	0.3	5	5	5	32.94	35.99	100.0	1385.67
SCPF	1	1	$10e(-4)$	5	0.8	5	5	5	470.14	147.33	48.07	

*We have generated 5 Potts partitions with a shrinkage constraint of 5, RMSE=Root Mean Squared Error, RRMSE=Relative Root Mean Squared Error; all without the DBS optimizer and the smoother δ

Table 14: RRMSE (%) and MSE Results with the shallow Gibbs Networks: Sparse-Gibbs

	Shallow Gibbs settings								Measure of fitness S-Gibbs			Time (sec.)
Data sets	NH	ET	LR	NE	BS	SWbP	PSY	SPP	RMSE	RMSE	RRMSE	Time
									Train	Test	-	
Slump												
EDM	1	1	$10e(-4)$	5	0.8	5	5	5	7.24	1.27	87.93	1160.06
EDM	1	1	$10e(-3)$	5	0.8	5	5	5	6.82	1.55	94.77	1143.09
EDM	1	1	$10e(-3)$	5	0.8	5	5	5	5.33	3.37	129.76	2414.18
EDM	1	1	$10e(-5)$	5	0.8	5	5	5	5.55	2.50	112.12	1155.88
EDM	2	1	$10e(-5)$	5	0.8	5	5	5	10.23	3.07	87.36	1549.87
EDM	10	1	$10e(-5)$	2	0.8	5	5	5	24.17	8.92	192.35	1299.76
SCPF	1	1	$10e(-4)$	5	0.8	5	5	5	470.14	147.33	48.07	

*We have generated 5 Potts partitions with a shrinkage constraint of 5, RMSE=Root Mean Squared Error, RRMSE=Relative Root Mean Squared Error; all without the DBS optimizer and the smoother δ

Appendix B. Cholesky Decomposition

For symmetric, positive definite matrices, the Cholesky decomposition or Cholesky factorization is useful in practice to generate samples from a Gaussian-like distribution.

Theorem B.1 (Cholesky Decomposition)

A symmetric, positive definite, matrix A can be factorized into a product $A = \mathcal{Z}\mathcal{Z}^\top$, described as:

$$\begin{bmatrix} a_{11} & \cdots & a_{1n} \\ \vdots & \ddots & \vdots \\ a_{n1} & \cdots & a_{nn} \end{bmatrix} = \begin{bmatrix} l_{11} & \cdots & 0 \\ \vdots & \ddots & \vdots \\ l_{n1} & \cdots & l_{nn} \end{bmatrix} \begin{bmatrix} l_{11} & \cdots & l_{n1} \\ \vdots & \ddots & \vdots \\ 0 & \cdots & l_{nn} \end{bmatrix}$$

where $\mathcal{Z}$ is a lower triangular matrix with positive diagonal elements, $\mathcal{Z}$ is called the Cholesky factor of A , and $\mathcal{Z}$ is unique.

Application to Gaussian distribution sampling. The following example is from Williams & Rasmussen (2006). For a Gaussian variable $y \sim N(\eta, \kappa_0)$, the multivariate normal distribution has a joint probability density given

by:

$$p(y \mid \eta, \kappa_0) = (2\pi)^{-d/2} |\kappa_0|^{-1/2} \exp \left(-\frac{1}{2} (y - \eta)^T \kappa_0^{-1} (y - \eta) \right)$$

where $\eta \in \mathbb{R}^d$ is the mean vector and $\kappa_0 \in M_d(\mathbb{R})$ is the (symmetric, positive definite) covariance matrix. To get some samples of y from this distribution, apply the following steps:

- Compute the Cholesky decomposition: we want to compute the Cholesky decomposition of the covariance matrix κ_0 . That is, we want to find a lower triangular matrix $\mathcal{Z} \in M_d(\mathbb{R})$ such that $\kappa_0 = \mathcal{Z}\mathcal{Z}^T$. Matrix $\mathcal{Z}$ will be useful in a further step.
- Generate Independent Samples $v \sim N(0, I)$;
- Compute $y = \eta + \mathcal{Z}v$. The variable $y = \eta + \mathcal{Z}v$ has a multivariate normal distribution since is a linear combination of independent normally distributed variables. Moreover,

$$E[y] = E[\eta + \mathcal{Z}v] = \eta + \mathcal{Z}E[v] = \eta$$

and

$$E[(y - \eta)^T (y - \eta)] = \kappa_0$$

Appendix C. The learning rate

Definition C.1 (The learning rate schedule). This function $\epsilon(s) : \mathbb{N} \rightarrow \mathbb{R}$ is called the learning rate schedule.

Appendix D. Stochastic gradient

The Stochastic gradient can be also ascent or descent. We will illustrate the standard properties of this algorithm using the descent configuration, because both (ascent & descent) have similar characteristics (or identical features).

The stochastic descent of the gradient (often shortened as SGD) is a stochastic approximation of the method of gradient descent to minimise an objective descent. A functionality that is written as a sum of distinguishable functions. The term stochastic here applies to the fact that we understand that we do not exactly know the gradient, but rather know a chaotic estimate of it instead.

This paragraph may require further readings about the Evidence Lower Bound (ELBO) in section 5 where more details have been proposed. Following Bottou (2012), the stochastic gradient descent algorithm replaces the gradient by an estimate :

$$\lambda^{n+1} = \lambda^n - \epsilon_n Z(\lambda^n; \xi^n)$$

where we describe the estimate of the gradient by $Z(\lambda^n; \xi^n)$, with the optimized parameter λ , emphasizing the stochastic nature through the random vector ξ^n . A class of possibilities are given by

$$Z(\lambda^n; \xi^n) = \left[\frac{1}{n_t} \sum_{i \in \mathcal{S}_t} \nabla ELBO_i(\lambda^n) \right]$$

where $\mathcal{S}_t \subset \{1, \dots, n_o\}$ and $n_t = |\mathcal{S}_t|$ gives the number of observations to base the estimate of the gradient on. In our case, we have set $n_t \leq 2$ enabling high-speed computation, but requiring many iterations ≥ 2 . We will consider the following assumptions on $\{\epsilon_n\}$:

$$\begin{aligned} \epsilon_n &> 0 & (A-1) \\ \sum_{n=1}^{\infty} \epsilon_n &= \infty & (A-2) \\ \sum_{n=1}^{\infty} \epsilon_n^2 &< \infty & (A-3) \end{aligned}$$

Theorem D.1 (Quasimartingale convergence theorem (Robbins & Siegmund, 1971) and (Fisk, 1965))

If $(X_n)_{n=1}^{\infty}$ is a positive stochastic process, and

$$\sum_{n=1}^{\infty} \mathbb{E} \left[(\mathbb{E}[X_{n+1} | \mathcal{F}_n] - X_n) \mathbb{1}_{\{\mathbb{E}[X_{n+1} | \mathcal{F}_n] - X_n > 0\}} \right] < \infty$$

then $X_n \rightarrow X_{\infty}$ almost surely on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_n)_{n=0}^{\infty}, \mathbb{P})$, with

$$\mathcal{F}_n = \sigma(X_m \mid m \leq \hat{n})$$

Now, let $C : \mathbb{R}^k \rightarrow \mathbb{R}$ be differentiable. How do solutions $s : \mathbb{R} \rightarrow \mathbb{R}^k$ to the following Ordinary differential Equation (ODE) behave:

$$\frac{d}{dt} s(t) = -\nabla C(s(t)) \text{ or in a discretized form: } \frac{s_{n+1} - s_n}{\epsilon_n} = -\nabla C(s_n)$$

The discretization uses Forward Euler discretisation method (Villatoro & Ramos, 1999).

Proposition D.1 ((Robbins & Monro, 1951) and (Bach, 2018))

If $s_{n+1} = s_n - \epsilon_n H_n(s_n)$, with $H_n(s_n)$ an unbiased estimator for $\nabla C(s_n)$, and C satisfies:

1. C has a unique minimiser x^*
2. $\forall \epsilon > 0, \inf_{\|x - x^*\|_2^2 > \epsilon} \langle x - x^*, \nabla C(x) \rangle > 0$,
3. $\mathbb{E} [\|H_n(x)\|_2^2] \leq A + B \|x - x^*\|_2^2$ for some $A, B \geq 0$ independent of n ,

then subject to (A-2) and (A-3) we have $s_n \rightarrow x^*$.

Proof. The proof structure broadly follows Bottou (1998) paper.

Step 1 **Define Lyapunov sequence:** $h_n = \|s_n - x^*\|_2^2$. h_n is not guaranteed to be decreasing. The main Idea from (Bottou (1998)) is: h_n may fluctuate, but if we can show that the cumulative 'up' movements aren't too big, we can still prove convergence of h_n .

Step 2 **Consider the h_n variations:** $h_{n+1} - h_n$.

$$\begin{aligned}
 h_{n+1} - h_n &= \langle s_{n+1} - x^*, s_{n+1} - x^* \rangle - \langle s_n - x^*, s_n - x^* \rangle \\
 &= \langle s_{n+1}, s_{n+1} \rangle - \langle s_n, s_n \rangle - 2 \langle s_{n+1} - s_n, x^* \rangle \\
 &= \langle s_n - \epsilon_n H_n(s_n), s_n - \epsilon_n H_n(s_n) \rangle - \langle s_n, s_n \rangle + 2\epsilon_n \langle H_n(s_n), x^* \rangle \\
 &= -2\epsilon_n \langle s_n - x^*, H_n(s_n) \rangle + \epsilon_n^2 \|H_n(s_n)\|_2^2
 \end{aligned}$$

So

$$\mathbb{E} [h_{n+1} - h_n \mid \mathcal{F}_n] = -2\epsilon_n \langle s_n - x^*, \nabla C(s_n) \rangle + \epsilon_n^2 \mathbb{E} [\|H_n(s_n)\|_2^2 \mid \mathcal{F}_n]$$

Step 3 **Show that h_n converge almost surely:**

Assuming $\mathbb{E} [\|H_n(x)\|_2^2] \leq A + B \|x - x^*\|_2^2$, we get:

$$\begin{aligned}
 h_{n+1} - h_n &\leq -2\epsilon_n \langle s_n - x^*, H_n(s_n) \rangle + \epsilon_n^2 (A + B h_n) \\
 \implies h_{n+1} - (1 + \epsilon_n^2 B) h_n &\leq -2\epsilon_n \langle s_n - x^*, H_n(s_n) \rangle + \epsilon_n^2 A \\
 &\leq \epsilon_n^2 A
 \end{aligned}$$

Condition (2) simply states that the opposite of the gradient $-\nabla_x C(x)$ always points towards the minimum x^* . This is also a convexity criterion that ensures that the term $\langle s_n - x^*, H_n(s_n) \rangle$ is always negative. We have:

$$\forall \epsilon > 0, \inf_{\|x - x^*\|_2^2 > \epsilon} \langle x - x^*, \nabla C(x) \rangle > 0 \implies \langle s_n - x^*, H_n(s_n) \rangle > 0$$

So, we get:

$$h_{n+1} - (1 + \epsilon_n^2 B) h_n \leq \epsilon_n^2 A$$

Introduce the series $\mu_n = \prod_{i=1}^{n-1} \frac{1}{1 + \epsilon_i^2 B} \xrightarrow{n \rightarrow \infty} \mu_\infty$, and $h'_n = \mu_n h_n$ Get:

$$\begin{aligned} \mathbb{E} [h'_{n+1} - h'_n \mid \mathcal{F}_n] &\leq \epsilon_n^2 \mu_n A \\ \implies \mathbb{E} [(h'_{n+1} - h'_n) \mathbb{1}_{\mathbb{E}[h'_{n+1} - h'_n \mid \mathcal{F}_n] > 0} \mid \mathcal{F}_n] &\leq \epsilon_n^2 \mu_n A \end{aligned}$$

$$\begin{aligned} \left\{ \begin{array}{l} \mu_n \leq \frac{1}{1+\epsilon_n^2 B} \\ \sum_{n=1}^{\infty} \epsilon_n^2 < \infty \end{array} \right. &\implies [\text{Quasimartingale convergence}] \textbf{Theorem 1.} \implies (h'_n)_{n=1}^{\infty} \text{ converges a.s.} \\ &\implies (h_n)_{n=1}^{\infty} \text{ converges a.s.} \end{aligned}$$

Step 4 **Show that h_n must converge to 0:**

From previous calculations:

$$\mathbb{E} [h_{n+1} - (1 + \epsilon_n^2 B) h_n \mid \mathcal{F}_n] = -2\epsilon_n \langle s_n - x^*, \nabla C(s_n) \rangle + \epsilon_n^2 A$$

$(h_n)_{n=1}^{\infty}$ converges, so the sequence in the first member of the equation above is summable a.s. . Because $\sum_{n=1}^{\infty} \epsilon_n^2 < \infty$, so right term ($\epsilon_n^2 A$) is summable a.s., so the left term side is also summable a.s. :

$$\sum_{n=1}^{\infty} \epsilon_n \langle s_n - x^*, \nabla C(s_n) \rangle < \infty \text{ almost surely, and } \epsilon_n \langle s_n - x^*, \nabla C(s_n) \rangle \longrightarrow 0 \text{ almost surely}$$

We can conclude that $(h_n)_{n=1}^{\infty}$ converge to zero, and this forces also to have: $\langle s_n - x^*, \nabla C(s_n) \rangle \rightarrow 0$ almost surely .

We can reach this conclusion because we know that h_n converges. Reasoning by contradiction: let us assume that $h_n = \|s_n - x^*\|_2^2$ converges to a value greater than zero and therefore, after a certain time, remains greater than some $\epsilon > 0$. Assumption (2) implies that $\langle s_n - x^*, \nabla C(s_n) \rangle > 0$, and then it remains greater than a strictly positive quantity. Since this would cause the sum ($\sum_{n=1}^{\infty} \epsilon_n \langle s_n - x^*, \nabla C(s_n) \rangle$) to diverge (but, this is not the case: $\sum_{n=1}^{\infty} \epsilon_n \langle s_n - x^*, \nabla C(s_n) \rangle < \infty$), we can conclude on that h_n converges to zero. We get by then: $s_n \rightarrow x^*$ (and simultaneously $\langle s_n - x^*, \nabla C(s_n) \rangle \rightarrow 0$), which ends the proof. ■

■

Application on our stochastic gradient algorithm: If C is the ELBO function [Section 5] averaged across a data set, the true gradient is of the form:

$$\nabla C(\lambda^t) = \frac{1}{N} \sum_{i \in \mathcal{S}_t} \nabla ELBO_i(\lambda^n)$$

and an approximation is formed by subsampling:

$$\widehat{\nabla C}(\lambda^n) = \frac{1}{n_t} \sum_{i \in \mathcal{S}_t} \nabla ELBO_i(\lambda^n) \quad (S_t \sim \text{Uniforme}(\text{subsets of size } n_t))$$

The only difference between the Stochastic Gradient Ascent (SGA) algorithm versus the Stochastic Gradient Descent (SGD), is that we want to maximize the ELBO function. For that purpose, we simply reformat maximizing the ELBO as minimizing its negative to apply the convergence theorem. In our case:

- i. The learning rate vector $\alpha_t = (\alpha_t^1, \alpha_t^2, \alpha_t^3)$ defined in section 7.1 (α_t^1 for the weights, α_t^2 for the biases and α_t^3 for the covariance matrix Σ) is set to be decreasing to make conditions [(A – 1), (A – 2), and (A – 3)] hold in our experiments :

$$\alpha_t^1 = \alpha_t^2 = \alpha_t^3 = \frac{10^{-50}}{t}$$

- ii. The Evidence Lower Bound (ELBO) is continuous and differentiable in every point, and is bounded by $\log p(y)$ (Yang, 2017). So in Proposition [see D.1], the condition (3) always holds for the ELBO. But, it is not the case for condition (1) and (2) because of non-convexity of the ELBO. In their seminal paper from 1951, Robbins and Monro showed that the stochastic optimization will converge to a local optimum in our case (Robbins & Monro (1951)). So, it requires good initialization.

In practice, the quality of the approximation depends on the variance of the estimator of $Z(\lambda^n; \xi^n)$ (Johnson & Zhang, 2013). The main advantage of this algorithm is that one can even estimate the model with a cluster of small size, and still get good estimations. The Robbins-Siegmund theorem (Robbins & Siegmund, 1971) provides the means to establish almost sure convergence under conditions including (A – 2), also in cases where the loss function is non-smooth (Saad, 1998).

The size of the subset used to measure the gradient can be considered in the same manner as we think of *sample size* in simple estimation problem. Big mini-batch sizes can have reliable gradient forecasts, reducing the parameters update variances. Small mini-batches, by comparison, are easy to estimate.

References

- Alahassa, K.-A. (2020).
URL <http://my-github.com/kgalahassa>
- Ali, J., Khan, R., Ahmad, N., & Maqsood, I. (2012). Random forests and decision trees. *International Journal of Computer Science Issues (IJCSI)*, 9(5), 272.
- Anderssen, R. S., & Bloomfield, P. (1974). Numerical differentiation procedures for non-exact data. *Numerische Mathematik*, 22(3), 157–182.
- Ardakani, A., Condo, C., & Gross, W. J. (2016). Sparsely-connected neural networks: towards efficient vlsi implementation of deep neural networks. *arXiv preprint arXiv:1611.01427*.
- Ari, B., & Güvenir, H. (2002). Clustered linear regression. *Knowledge-Based Systems*, 15(3), 169 – 175.
URL <http://www.sciencedirect.com/science/article/pii/S095070510100154X>
- Asaithambi, A. (2010). Numerical solution of the burgers' equation by automatic differentiation. *Applied Mathematics and Computation*, 216(9), 2700–2708.
- Bach, F. (2018). Statistical machine learning and convex optimization.
- Baird, L., Smalenberger, D., & Ingkiriwang, S. (2005). One-step neural network inversion with pdf learning and emulation. In *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.*, vol. 2, (pp. 966–971). IEEE.
- Barbarossa, D., & Manzonetto, G. (2019). About the power of taylor expansion. In *3rd International Workshop on Trends in Linear Logic and Applications*.
- Barber, D., & Bishop, C. (1997). Ensemble learning for multi-layer networks. *Advances in neural information processing systems*, 10, 395–401.
- Barber, D., & Bishop, C. M. (1998). Ensemble learning in bayesian neural networks. *Nato ASI Series F Computer and Systems Sciences*, 168, 215–238.
- Barbu, A., & Zhu, S.-C. (2003). Graph partition by swendsen-wang cuts. In *null*, (p. 320). IEEE.
- Barthelemy, J.-F., & Hall, L. E. (1995). Automatic differentiation as a tool in engineering design. *Structural optimization*, 9(2), 76–82.
- Bartholomew-Biggs, M., Brown, S., Christianson, B., & Dixon, L. (2000). Automatic differentiation of algorithms. *Journal of Computational and Applied Mathematics*, 124(1), 171 – 190. Numerical Analysis 2000. Vol. IV: Optimization and Nonlinear Equations.
URL <http://www.sciencedirect.com/science/article/pii/S0377042700004222>
- Bastien, F., Lamblin, P., Pascanu, R., Bergstra, J., Goodfellow, I., Bergeron, A., Bouchard, N., Warde-Farley, D., & Bengio, Y. (2012). Theano: new features and speed improvements. *arXiv preprint arXiv:1211.5590*.
- Baydin, A. G., Pearlmutter, B. A., Radul, A. A., & Siskind, J. M. (2018). Automatic differentiation in machine learning: a survey. *Journal of machine learning research*, 18.
- Behrmann, J., Grathwohl, W., Chen, R. T., Duvenaud, D., & Jacobsen, J.-H. (2019). Invertible residual networks. In *International Conference on Machine Learning*, (pp. 573–582). PMLR.
- Benghorbal, M. M. (2007). Unified formulas for integer and fractional order symbolic derivatives and integrals of the power-inverse trigonometric class i. *International Journal of Pure and Applied Mathematics*, 40(1), 77.

- Bergstra, J., Bastien, F., Breuleux, O., Lamblin, P., Pascanu, R., Delalleau, O., Desjardins, G., Warde-Farley, D., Goodfellow, I., Bergeron, A., et al. (2011). Theano: Deep learning on gpus with python. In *NIPS 2011, BigLearning Workshop, Granada, Spain*, vol. 3, (pp. 1–48). Citeseer.
- Bertolazzi, E. (2008). One-dimensional minimization. Lecture Notes, Università di Trento.
- Besag, J. (2004). Markov Chain Monte Carlo Methods for Statistical Inference. Tech. rep., Department of Statistics, University of Washington, Seattle, USA.
- Biau, G., & Devroye, L. (2010). On the layered nearest neighbour estimate, the bagged nearest neighbour estimate and the random forest method in regression and classification. *Journal of Multivariate Analysis*, 101(10), 2499–2518.
- Bischof, C., Lang, B., & Vehreschild, A. (2003). Automatic differentiation for matlab programs. In *PAMM: Proceedings in Applied Mathematics and Mechanics*, vol. 2, (pp. 50–53). Wiley Online Library.
- Bischof, C. H., Bücker, H. M., Hovland, P., Naumann, U., & Utke, J. (2008a). Advances in automatic differentiation.
- Bischof, C. H., Hovland, P. D., & Norris, B. (2008b). On the implementation of automatic differentiation tools. *Higher-Order and Symbolic Computation*, 21(3), 311–331.
- Bishop, C. M. (1997). Bayesian Neural Networks. *Journal of the Brazilian Computer Society*, 4.
URL http://www.scielo.br/scielo.php?script=sci_arttext&pid=S0104-65001997000200006&nrm=iso
- Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. Springer.
- Blackledge, J., Bezobrazov, S., & Tobin, P. (2015). Cryptography using artificial intelligence. In *2015 International Joint Conference on Neural Networks (IJCNN)*, (pp. 1–6). IEEE.
- Blatt, M., Wiseman, S., & Domany, E. (1996a). Clustering data through an analogy to the potts model. In *Advances in Neural Information Processing Systems*, (pp. 416–422).
- Blatt, M., Wiseman, S., & Domany, E. (1996b). Superparamagnetic clustering of data. *Physical Review Letters*, 76, 3251–3254.
- Blundell, C., Cornebise, J., Kavukcuoglu, K., & Wierstra, D. (2015). Weight uncertainty in neural networks. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37, ICML'15*, (pp. 1613–1622). JMLR.org.
URL <http://dl.acm.org/citation.cfm?id=3045118.3045290>
- Bottou, L. (1998). Online learning and stochastic approximations. *On-line learning in neural networks*, 17(9), 142.
- Bottou, L. (2012). Stochastic gradient descent tricks. In *Neural networks: Tricks of the trade*, (pp. 421–436). Springer.
- Bourke, P. (1999). Interpolation methods. *Miscellaneous: projection, modelling, rendering, I*.
- Bowling, M., & Veloso, M. (2002). Scalable learning in stochastic games. In *AAAI Workshop on Game Theoretic and Decision Theoretic Agents*, (pp. 11–18).
- Browder, A. (2012). *Mathematical analysis: an introduction*. Springer Science & Business Media.

- Bruyninckx, H., & De Schutter, J. (1996). Symbolic differentiation of the velocity mapping for a serial kinematic chain. *Mechanism and machine theory*, 31(2), 135–148.
- Bücker, H. M., Corliss, G., Hovland, P., Naumann, U., & Norris, B. (2006). *Automatic differentiation: applications, theory, and implementations*, vol. 50. Springer Science & Business Media.
- Campbell, T., & Beronov, B. (2019). Sparse variational inference: Bayesian coresets from scratch. In *Advances in Neural Information Processing Systems*, (pp. 11461–11472).
- Canuto, C., & Tabacco, A. (2015). *Mathematical analysis II*, vol. 85. Springer.
- Carpenter, B., Hoffman, M. D., Brubaker, M., Lee, D., Li, P., & Betancourt, M. (2015). The stan math library: Reverse-mode automatic differentiation in c++. *arXiv preprint arXiv:1509.07164*.
- Challis, E., & Barber, D. (2013). Gaussian kullback-leibler approximate inference. *The Journal of Machine Learning Research*, 14(1), 2239–2286.
- Chang, Y.-C. (2009). N-dimension golden section search: Its variants and limitations. In *2009 2nd International Conference on Biomedical Engineering and Informatics*, (pp. 1–6). IEEE.
- Chartrand, R. (2011). Numerical differentiation of noisy, nonsmooth data. *International Scholarly Research Notices*, 2011.
- Chen, C., Ding, N., & Carin, L. (2015). On the convergence of stochastic gradient mcmc algorithms with high-order integrators. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2, NIPS'15*, (pp. 2278–2286). Cambridge, MA, USA: MIT Press.
URL <http://dl.acm.org/citation.cfm?id=2969442.2969494>
- Chen, H. H., Wang, P. F., Sung, C. T., Yeh, Y. R., & Lee, Y. J. (2013). Energy disaggregation via clustered regression models: A case study in the convenience store. In *2013 Conference on Technologies and Applications of Artificial Intelligence*, (pp. 37–42).
- Cheney, E. W., & Kincaid, D. R. (2012). *Numerical mathematics and computing*. Cengage Learning.
- Choquet-Bruhat, Y., DeWitt-Morette, C., de Witt, C., Bleick, M. D., & Dillard-Bleick, M. (1982). *Analysis*. Gulf Professional Publishing.
- Clausen, A., & Sokol, S. (2020). Deriv: R-based symbolic differentiation.
- Cohen, M. A., & Tan, C. O. (2012). A polynomial approximation for arbitrary functions. *Applied Mathematics Letters*, 25(11), 1947–1952.
- Coleman, T. F., & Verma, A. (1998). The efficient computation of sparse jacobian matrices using automatic differentiation. *SIAM Journal on Scientific Computing*, 19(4), 1210–1233.
- Cruz, J. B., Pérez, L. L., & Melo, J. (2011). Convergence of the projected gradient method for quasiconvex multiobjective optimization. *Nonlinear Analysis: Theory, Methods & Applications*, 74(16), 5268–5273.
- Cullum, J. (1971). Numerical differentiation and regularization. *SIAM Journal on numerical analysis*, 8(2), 254–265.
- Cybenko, G. (1989). Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4), 303–314.
- da Rocha, J. C. F., Guimarães, A. M., & Kozłowski Jr, V. (2011). Convergence of iterative algorithms for learning bayesian networks. *Iberoamerican Journal of Applied Computing*, 1(2).

- DeCoste, D. (1997). The future of chess-playing technologies and the significance of kasparov versus deep blue. In *Deep Blue Versus Kasparov: The Significance for Artificial Intelligence*, (pp. 9–13).
- Dikusar, N. (2016). Higher-order polynomial approximation. *Mathematical Models and Computer Simulations*, 8(2), 183–200.
- Ding, C., Wang, Y., & Li, Y. (2012). Potts and percolation models on bowtie lattices. *Physical Review E*, 86(2), 021125.
- Dong, T., & Huang, T. (2019). Neural cryptography based on complex-valued neural network. *IEEE Transactions on Neural Networks and Learning Systems*.
- Dressler, M. (2009). Art of surface interpolation. *Technical University of Liberec Faculty of Mechatronics and Interdisciplinary Engineering Studies: Czech Republic*.
- Drucker, H., & Le Cun, Y. (1991). Double backpropagation increasing generalization performance. In *IJCNN-91-Seattle International Joint Conference on Neural Networks*, vol. 2, (pp. 145–150). IEEE.
- Drucker, H., & Le Cun, Y. (1992). Improving generalization performance using double backpropagation. *IEEE Transactions on Neural Networks*, 3(6), 991–997.
- Duchi, J., Hazan, E., & Singer, Y. (2011). Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research*, 12(Jul), 2121–2159.
- Dumitru, P. D., Plopeanu, M., & Badea, D. (2013). Comparative study regarding the methods of interpolation. *Recent advances in geodesy and Geomatics engineering*, 1, 45–52.
- Elliott, C. (2018). The simple essence of automatic differentiation. *Proceedings of the ACM on Programming Languages*, 2(ICFP), 1–29.
- Essam, J. (1979). Potts models, percolation, and duality. *Journal of Mathematical Physics*, 20(8), 1769–1773.
- Etmann, C. (2019). *Double Backpropagation with Applications to Robustness and Saliency Map Interpretability*. Ph.D. thesis, Universität Bremen.
- Fisk, D. L. (1965). Quasi-martingales. *Transactions of the American Mathematical Society*, 120(3), 369–389.
- Forth, S. A. (2006). An efficient overloaded implementation of forward mode automatic differentiation in matlab. *ACM Transactions on Mathematical Software (TOMS)*, 32(2), 195–222.
- Friedman, J., Hastie, T., & Tibshirani, R. (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3), 432–441.
- Friston, K. J., Mattout, J., Trujillo-Barreto, N. J., Ashburner, J., & Penny, W. D. (2007). Variational free energy and the laplace approximation. *NeuroImage*, 34 1, 220–34.
- Gay, D. M. (1991). Automatic differentiation of nonlinear ampl models. *Automatic Differentiation of Algorithms: Theory, Implementation, and Application*, (pp. 61–73).
- Giaquinta, M., & Modica, G. (2010). *Mathematical analysis: An introduction to functions of several variables*. Springer Science & Business Media.
- Gilbert, G. T. (1991). Positive definite matrices and sylvester’s criterion. *The American Mathematical Monthly*, 98(1), 44–46.

- Giudici, P., & Green, P. (1999). Decomposable graphical Gaussian model determination. *Biometrika*, 86(4), 785–801.
- Godhavari, T., Alamelu, N., & Soundararajan, R. (2005). Cryptography using neural network. In *2005 Annual IEEE India Conference-Indicon*, (pp. 258–261). IEEE.
- Granter, S. R., Beck, A. H., & Papke Jr, D. J. (2017). Alphago, deep learning, and the future of the human microscopist. *Archives of pathology & laboratory medicine*, 141(5), 619–621.
- Graves, A. (2011). Practical variational inference for neural networks. In J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, & K. Q. Weinberger (Eds.) *Advances in Neural Information Processing Systems 24*, (pp. 2348–2356). Curran Associates, Inc.
URL <http://papers.nips.cc/paper/4329-practical-variational-inference-for-neural-net.pdf>
- Green, P. J. (1995). Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika*, 82(4), 711–732.
- Gregory, S. (2010). Finding overlapping communities in networks by label propagation. *New journal of Physics*, 12(10), 103018.
- Griewank, A., et al. (1989). On automatic differentiation. *Mathematical Programming: recent developments and applications*, 6(6), 83–107.
- Guckenheimer, J., & Meloon, B. (2000). Computing periodic orbits and their bifurcations with automatic differentiation. *SIAM Journal on Scientific Computing*, 22(3), 951–985.
- Guenter, B. (2007). Efficient symbolic differentiation for graphics applications. In *ACM SIGGRAPH 2007 papers*, (pp. 108–es).
- Guenter, B., Rapp, J., & Finch, M. (2011). Symbolic differentiation in gpu shaders. Tech. rep., Citeseer.
- Guessab, A., Nouisser, O., & Schmeisser, G. (2006). Multivariate approximation by a combination of modified taylor polynomials. *Journal of Computational and Applied Mathematics*, 196(1), 162–179.
- Guilhoto, L. F. (2018). An overview of artificial neural networks for mathematicians.
- Guo-qiang, Y., Wan-jin, H., Wen-cai, L., et al. (2004). Linear interpolation method for processing the test data of five-hole probes [j]. *Journal of Engineering for Thermal Energy and Power*, 5.
- Hadke, P. P., & Kale, S. G. (2016). Use of neural networks in cryptography: A review. In *2016 World Conference on Futuristic Trends in Research and Innovation for Social Welfare (Startup Conclave)*, (pp. 1–4). IEEE.
- HAMMER, F. G. H. (2017). Taylor expansion, gradient descent and newton’s method in machine learning.
- Hassabis, D. (2017). Artificial intelligence: Chess match of the century. *Nature*, 544(7651), 413–414.
- Hassani, H., Soltanolkotabi, M., & Karbasi, A. (2017). Gradient methods for submodular maximization. In *Advances in Neural Information Processing Systems*, (pp. 5841–5851).
- Hayden, T. L., & Wells, J. (1988). Approximation by matrices positive semidefinite on a subspace. *Linear Algebra and its Applications*, 109, 115–130.
- Hecht-Nielsen, R. (1992). Theory of the backpropagation neural network. In *Neural networks for perception*, (pp. 65–93). Elsevier.

- Higham, N. J. (1988). Computing a nearest symmetric positive semidefinite matrix. *Linear algebra and its applications*, 103, 103–118.
- Hinton, G. E., & van Camp, D. (1993). Keeping the neural networks simple by minimizing the description length of the weights. In *Proceedings of the Sixth Annual Conference on Computational Learning Theory, COLT '93*, (pp. 5–13). New York, NY, USA: ACM.
URL <http://doi.acm.org/10.1145/168304.168306>
- Hornik, K., Stinchcombe, M., & White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5), 359 – 366.
URL <http://www.sciencedirect.com/science/article/pii/0893608089900208>
- Hsu, F.-H. (2004). *Behind Deep Blue: Building the computer that defeated the world chess champion*. Princeton University Press.
- H.-J., M. (1985). Spath, h.: Cluster dissection and analysis: theory, fortran programs, examples. (translator: Johannes goldschmidt.) ellis horwood ltd wiley, chichester 1985. 226 pp. £25. *Biometrical Journal*, 28(2), 182–182.
- Ionescu, C., Vantzou, O., & Sminchisescu, C. (2015). Training deep networks with structured layers by matrix backpropagation. *arXiv preprint arXiv:1509.07838*.
- Jégou, S. (1997). Using maple for symbolic differentiation to solve inverse problems. *MapleTech*, (pp. 32–40).
- Jewbali, A., & Ore, R. T. I. (2009). Finding the nearest positive definite matrix for input to semi-automatic variogram fitting (varfit_lm). *Centre for Computational Geostatistics*, 11, 402.
- Joarder, A. H. (2001). Six ways to look at linear interpolation. *International Journal of Mathematical Education in Science and Technology*, 32(6), 932–937.
- Johnson, R., & Zhang, T. (2013). Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in neural information processing systems*, (pp. 315–323).
- Jorjani, H., Klei, L., & Emanuelson, U. (2003). A simple method for weighted bending of genetic (co) variance matrices. *Journal of dairy science*, 86(2), 677–679.
- Kamruzzaman, J., Kumagai, Y., & Aziz, S. M. (1997). Character recognition by double backpropagation neural network. In *TENCON'97 Brisbane-Australia. Proceedings of IEEE TENCON'97. IEEE Region 10 Annual Conference. Speech and Image Technologies for Computing and Telecommunications (Cat. No. 97CH36162)*, vol. 1, (pp. 411–414). IEEE.
- Kamruzzaman, J., & Syed, A., Mahfuzul (1998). A neural network based character recognition system using double backpropagation. *Malaysian Journal of Computer Science*, 11(1), 58–64.
- Kanter, I. (1988). Potts-glass models of neural networks. *Phys. Rev. A*, 37, 2739–2742.
URL <https://link.aps.org/doi/10.1103/PhysRevA.37.2739>
- Kanter, I., & Kinzel, W. (2003). The theory of neural networks and cryptography. In *The Physics of Communication*, (pp. 631–642). World Scientific.
- Kay, S. (1983). Some results in linear interpolation theory. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 31(3), 746–749.
- Kemppainen, A., & Smirnov, S. (2019). Conformal invariance of boundary touching loops of fk ising model. *Communications in Mathematical Physics*, 369(1), 49–98.

- Kenny, K. B. (????). Symbolic differentiation in tcl: reusing the tcl parser for symbolic algebra.
- Kepner, J., & Robinett, R. (2019). Radix-net: Structured sparse matrices for deep neural networks. In *2019 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, (pp. 268–274). IEEE.
- Kingma, D. P., Salimans, T., & Welling, M. (2015). Variational dropout and the local reparameterization trick. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, & R. Garnett (Eds.) *Advances in Neural Information Processing Systems 28*, (pp. 2575–2583). Curran Associates, Inc.
URL <http://papers.nips.cc/paper/5666-variational-dropout-and-the-local-reparameterization-trick.pdf>
- Kinzel, W., & Kanter, I. (2002a). Interacting neural networks and cryptography. In *Advances in solid state physics*, (pp. 383–391). Springer.
- Kinzel, W., & Kanter, I. (2002b). Neural cryptography. In *Proceedings of the 9th International Conference on Neural Information Processing, 2002. ICONIP'02.*, vol. 3, (pp. 1351–1354). IEEE.
- Klimov, A., Mityagin, A., & Shamir, A. (2002). Analysis of neural cryptography. In *International Conference on the Theory and Application of Cryptology and Information Security*, (pp. 288–298). Springer.
- Knott, G. D. (2017). Symbolic differentiation. In *Interpreting LISP*, (pp. 93–99). Springer.
- Knowles, I., & Renka, R. J. (2014). Methods for numerical differentiation of noisy data. *Electron. J. Differ. Equ*, *21*, 235–246.
- Kotsiantis, S. (2011). Combining bagging, boosting, rotation forest and random subspace methods. *Artificial intelligence review*, *35*(3), 223–240.
- Kucukelbir, A., Ranganath, R., Gelman, A., & Blei, D. (2014). Fully automatic variational inference of differentiable probability models. In *NIPS Workshop on Probabilistic Programming*.
- Kurita, T. (1991). An efficient agglomerative clustering algorithm using a heap. *Pattern Recognition*, *24*(3), 205–209.
- Laue, S. (2019). On the equivalence of forward mode automatic differentiation and symbolic differentiation. *arXiv preprint arXiv:1904.02990*.
- Lee, J., Shridhar, K., Hayashi, H., Iwana, B. K., Kang, S., & Uchida, S. (2019). Probact: A probabilistic activation function for deep neural networks. *arXiv preprint arXiv:1905.10761*.
- Lee, J. D., Simchowitz, M., Jordan, M. I., & Recht, B. (2016). Gradient descent only converges to minimizers. In *Conference on learning theory*, (pp. 1246–1257).
- Lessard, J.-P., James, J. M., & Ransford, J. (2016). Automatic differentiation for fourier series and the radii polynomial approach. *Physica D: Nonlinear Phenomena*, *334*, 174–186.
- Li, Z.-H., & Hu, W.-Z. (2017). A high-precision digital integrator based on the romberg algorithm. *Review of Scientific Instruments*, *88*(4), 045111.
- Liu, S. (2020). Learning sparse neural networks for better generalization. In *29th International Joint Conference on Artificial Intelligence-17th Pacific Rim International Conference on Artificial Intelligence.*
- Liu, Y. (2014). Random forest algorithm in big data environment. *Computer Modelling & New Technologies*, *18*(12A), 147–151.

- Liu, Y., Wang, Y., & Zhang, J. (2012). New machine learning algorithm: Random forest. In *International Conference on Information Computing and Applications*, (pp. 246–252). Springer.
- Louizos, C., & Welling, M. (2016). Structured and efficient variational deep learning with matrix gaussian posteriors. In *International Conference on Machine Learning*, (pp. 1708–1716).
- Lu, S., & Pereverzev, S. (2006). Numerical differentiation from a viewpoint of regularization theory. *Mathematics of computation*, 75(256), 1853–1870.
- Lydia, A., & Francis, S. (2019). Adagrad—an optimizer for stochastic gradient descent. *Int. J. Inf. Comput. Sci.*, 6(5).
- Maelland, E. (1988). On the comparison of interpolation methods. *IEEE transactions on medical imaging*, 7(3), 213–217.
- Marée, S. (2012). *Correcting non positive definite correlation matrices*. BSc Thesis Applied Mathematics, TU Delft.
- Mazumder, R., & Hastie, T. (2015). The graphical lasso: new insights and alternatives. *Electronic Journal of Statistics*, 6, 2125–2149.
- Mazza-Anthony, C. A. (2019). *Structured Sparsity and Precision Matrix Estimation*. Ph.D. thesis, McGill University.
- Mboup, M., Join, C., & Fliess, M. (2009). Numerical differentiation with annihilators in noisy environment. *Numerical algorithms*, 50(4), 439–467.
- McCullagh, P. (2017). Robustifying concurrent.futures with loky.
URL <https://tommorai.github.io/talks/pyparis17/#41>
- Meurer, A., Smith, C. P., Paprocki, M., Čertík, O., Kirpichev, S. B., Rocklin, M., Kumar, A., Ivanov, S., Moore, J. K., Singh, S., et al. (2017). Sympy: symbolic computing in python. *PeerJ Computer Science*, 3, e103.
- Minkov, M., Williamson, I. A., Andreani, L. C., Gerace, D., Lou, B., Song, A. Y., Hughes, T. W., & Fan, S. (2020). Inverse design of photonic crystals through automatic differentiation. *ACS Photonics*, 7(7), 1729–1741.
- Mohamed, S., Rosca, M., Figurnov, M., & Mnih, A. (2019). Monte carlo gradient estimation in machine learning. *arXiv preprint arXiv:1906.10652*.
- Mulder, J., & Fox, J.-P. (2013). Bayesian tests on components of the compound symmetry covariance matrix. *Statistics and Computing*, 23(1), 109–122.
- Müller, P., & Quintana, F. (2010a). Random partition models with regression on covariates. *Journal of statistical planning and inference*, 140(10), 2801–2808.
- Müller, P., & Quintana, F. (2010b). Random partition models with regression on covariates. *Journal of Statistical Planning and Inference*, 140(10).
- Müller, P., Quintana, F., & Rosner, G. (2008). Bayesian clustering with regression. Tech. rep., Working paper series - European Central Bank.
- Müller, P., Quintana, F. A., & Rosner, G. L. (2011). A Product Partition Model with Regression on Covariates. *Journal of Computational and Graphical Statistics*, 20(1), 260–278.
- Munson, M. A., & Caruana, R. (2009). On feature selection, bias-variance, and bagging. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, (pp. 144–159). Springer.

- Murua, A., & Quintana, F. A. (2017). Semiparametric bayesian regression via potts model. *Journal of Computational and Graphical Statistics*, 26(2), 265–274.
URL <https://doi.org/10.1080/10618600.2016.1172015>
- Murua, A., Stanberry, L., & Stuetzle, W. (2008a). On potts model clustering, kernel k-means, and density estimation. *Journal of Computational and Graphical Statistics*, 17(3), 629–658.
- Murua, A., Stanberry, L., & Stuetzle, W. (2008b). On potts model clustering, kernel k means and density estimation. *Journal of Computational and Graphical Statistics*, 17, 629–658.
- Murua, A., & Wicker, N. (2014). The conditional-potts clustering model. *Journal of Computational and Graphical Statistics*, 23(3), 717–739.
URL <https://doi.org/10.1080/10618600.2013.837828>
- Nagpal, G., Uddin, M., & Kaur, A. (2013). Estimating project development effort using clustered regression approach. In *Computer Science & Information Technology*, vol. 3, (pp. 493–507).
- Neal, R. M. (1996). *Bayesian Learning for Neural Networks*. Secaucus, NJ, USA: Springer-Verlag New York, Inc.
- Neumaier, A. (2003). Taylor forms—use and limits. *Reliable computing*, 9(1), 43–79.
- Nilforooshan, M. A. (2020). mbend: an r package for bending non-positive-definite symmetric matrices to positive-definite. *BMC genetics*, 21(1), 1–8.
- Nilforooshan, M. A., & Nilforooshan, M. M. A. (2020). Package ‘mbend’.
- Osburn, S. L., Ahmed, J., Coppola, V., & Bernstein, D. S. (1996). Symbolic differentiation for fixed-structure controller synthesis. *IFAC Proceedings Volumes*, 29(1), 1697–1702.
- Otsuki, T., Ogino, H., Ide, S., & Chiba, T. (2004). Curve interpolation method. US Patent 6,823,234.
- Paisley, J., Blei, D., & Jordan, M. (2012). Variational bayesian inference with stochastic search. *arXiv preprint arXiv:1206.6430*.
- Panov, P., & Džeroski, S. (2007). Combining bagging and random subspaces to create better ensembles. In *International Symposium on Intelligent Data Analysis*, (pp. 118–129). Springer.
- Pantelides, C. C. (1988). Symbolic and numerical techniques for the solution of large systems of algebraic equations.
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., & Lerer, A. (2017). Automatic differentiation in pytorch.
- Pattanayak, S., & Ludwig, S. A. (2017). Encryption based on neural cryptography. In *International Conference on Health Information Science*, (pp. 321–330). Springer.
- Paul, D., Wang, L., et al. (2016). Discussion of “estimating structured high-dimensional covariance and precision matrices: Optimal rates and adaptive estimation”. *Electronic Journal of Statistics*, 10(1), 74–80.
- Pettersson, R. (1992). Stratonovich–taylor expansion and numerical methods. *Stochastic Analysis and Applications*, 10(5), 603–612.
- Philipsen, W., & Cluitmans, L. (1993). Using a genetic algorithm to tune potts neural networks. In *Artificial Neural Nets and Genetic Algorithms*, (pp. 650–657). Springer.

- Phillips, G. M. (2003). *Interpolation and approximation by polynomials*, vol. 14. Springer Science & Business Media.
- Plaskura, P. (2017). Derivwww-web-based symbolic differentiation system. , (60, . 4), 254–263.
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. (1996). *Numerical recipes in C*. Cambridge University Press.
- Probst, P., Wright, M. N., & Boulesteix, A.-L. (2019). Hyperparameters and tuning strategies for random forest. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 9(3), e1301.
- Qian, Z. (2011). High order directional derivative and the simple form of multivariate taylor theorem. *Journal of Heze University*, (2), 4.
- Raghavan, U. N., Albert, R., & Kumara, S. (2007). Near linear time algorithm to detect community structures in large-scale networks. *Physical review E*, 76(3), 036106.
- Rajaratnam, B., Massam, H., & Carvalho, C. M. (2008). Flexible covariance estimation in graphical Gaussian models. *The Annals of Statistics*, 36(6), 2818–2849.
- Rall, L. B., & Corliss, G. F. (1996). An introduction to automatic differentiation. *Computational Differentiation: Techniques, Applications, and Tools*, 89.
- Ramm, A., & Smirnova, A. (2001). On stable numerical differentiation. *Mathematics of computation*, 70(235), 1131–1153.
- Rasmussen, C. E. (1995). A practical monte carlo implementation of bayesian learning. In *Proceedings of the 8th International Conference on Neural Information Processing Systems*, NIPS’95, (pp. 598–604). Cambridge, MA, USA: MIT Press.
URL <http://dl.acm.org/citation.cfm?id=2998828.2998913>
- Reimer, M. (2012). *Multivariate polynomial approximation*, vol. 144. Birkhäuser.
- Rezende, D. J., & Viola, F. (2018). Generalized elbo with constrained optimization, geco. In *Workshop on Bayesian Deep Learning, NeurIPS*.
- Rigatti, S. J. (2017). Random forest. *Journal of Insurance Medicine*, 47(1), 31–39.
- Robbins, H., & Monroe, S. (1951). A stochastic approximation method. *The annals of mathematical statistics*, (pp. 400–407).
- Robbins, H., & Siegmund, D. (1971). A convergence theorem for non negative almost supermartingales and some applications. In *Optimizing methods in statistics*, (pp. 233–257). Elsevier.
- Robert, C. P., & Casella, G. (2005). *Monte Carlo Statistical Methods (Springer Texts in Statistics)*. Secaucus, NJ, USA: Springer-Verlag New York, Inc.
- Ross, A., & Doshi-Velez, F. (2018). Improving the adversarial robustness and interpretability of deep neural networks by regularizing their input gradients. In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32.
- Ruder, S. (2016). An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747*.
- Rukundo, O., & Cao, H. (2012). Nearest neighbor value interpolation. *arXiv preprint arXiv:1211.1768*.

- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *nature*, 323(6088), 533–536.
- Ruttor, A., Kinzel, W., & Kanter, I. (2007). Dynamics of neural cryptography. *Physical Review E*, 75(5), 056104.
- Ruttor, A., Kinzel, W., Shacham, L., & Kanter, I. (2004). Neural cryptography with feedback. *Physical Review E*, 69(4), 046110.
- Saad, D. (1998). Online algorithms and stochastic approximations. *Online Learning*, 5, 6–3.
- Schmidt, M., Fung, G., & Rosales, R. (2007). Fast optimization methods for l_1 regularization: A comparative study and two new approaches. In J. N. Kok, J. Koronacki, R. L. d. Mantaras, S. Matwin, D. Mladenič, & A. Skowron (Eds.) *Machine Learning: ECML 2007*, (pp. 286–297). Berlin, Heidelberg: Springer Berlin Heidelberg.
- Schonlau, M., & Zou, R. Y. (2020). The random forest algorithm for statistical learning. *The Stata Journal*, 20(1), 3–29.
- Schwing, A. G., & Urtasun, R. (2015). Fully connected deep structured networks. *arXiv preprint arXiv:1503.02351*.
- Sebastien, B. (2013). Orf523: Oracle complexity, large-scale optimization.
URL <https://blogs.princeton.edu/imabandit/2013/03/15/orf523-oracle-complexity-large-scale-optimization/>
- Seck, I., Loosli, G., & Canu, S. (2019). L 1-norm double backpropagation adversarial defense. *arXiv preprint arXiv:1903.01715*.
- Sharma, K., Aggarwal, A., Singhania, T., Gupta, D., & Khanna, A. (2019). Hiding data in images using cryptography and deep neural network. *arXiv preprint arXiv:1912.10413*.
- Shi, P., & McPhee, J. (2002). Symbolic programming of a graph-theoretic approach to flexible multibody dynamics.
- Skurichina, M., & Duin, R. P. (2001). Bagging and the random subspace method for redundant feature spaces. In *International Workshop on Multiple Classifier Systems*, (pp. 1–10). Springer.
- Smyth, G. K. (2014). Polynomial approximation. *Wiley StatsRef: Statistics Reference Online*.
- Sokal, A. (1997). Monte Carlo methods in statistical mechanics: foundations and new algorithms. In *Functional integration (Cargèse, 1996)*, vol. 361 of *NATO Adv. Sci. Inst. Ser. B Phys.*, (pp. 131–192). Plenum, New York.
- Späth, H. (1979). Algorithm 39 clusterwise linear regression. *Computing*, 22(4), 367–373.
URL <https://doi.org/10.1007/BF02265317>
- Srinivas, S., Subramanya, A., & Venkatesh Babu, R. (2017). Training sparse neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, (pp. 138–145).
- Stanberry, L., Murua, A., & Cordes, D. (2008). Functional connectivity mapping using the ferromagnetic Potts spin model. *Human Brain Mapping*, 29, 422–440.
- Strobl, C., Malley, J., & Tutz, G. (2009). An introduction to recursive partitioning: rationale, application, and characteristics of classification and regression trees, bagging, and random forests. *Psychological methods*, 14(4), 323.

- Sun, C., Chen, S., & Huang, X. (2020). Double backpropagation for training autoencoders against adversarial attack. *arXiv preprint arXiv:2003.01895*.
- Sun, S., Chen, C., & Carin, L. (2017). Learning Structured Weight Uncertainty in Bayesian Neural Networks. In A. Singh, & J. Zhu (Eds.) *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, vol. 54 of *Proceedings of Machine Learning Research*, (pp. 1283–1292). Fort Lauderdale, FL, USA: PMLR.
URL <http://proceedings.mlr.press/v54/sun17b.html>
- Sutton, C., Sindelar, M., & McCallum, A. (2005). Feature bagging: Preventing weight undertraining in structured discriminative learning. *Center for Intelligent Information Retrieval, U. of Massachusetts*.
- Svetnik, V., Liaw, A., Tong, C., Culberson, J. C., Sheridan, R. P., & Feuston, B. P. (2003). Random forest: a classification and regression tool for compound classification and qsar modeling. *Journal of chemical information and computer sciences*, 43(6), 1947–1958.
- Tibély, G., & Kertész, J. (2008). On the equivalence of the label propagation method of community detection and a potts model approach. *Physica A: Statistical Mechanics and its Applications*, 387(19), 4982–4984.
- Torgo, L., & da Costa, J. P. (2000). Clustered multiple regression. In H. A. L. Kiers, J.-P. Rasson, P. J. F. Groenen, & M. Schader (Eds.) *Data Analysis, Classification, and Related Methods*, (pp. 217–222). Berlin, Heidelberg: Springer Berlin Heidelberg.
- Torlai, G., Carrasquilla, J., Fishman, M. T., Melko, R. G., & Fisher, M. P. (2020). Wave-function positivization via automatic differentiation. *Physical Review Research*, 2(3), 032060.
- Tsoumakas, G., Spyromitros-Xioufis, E., Vilcek, J., & Vlahavas, I. (2011). Mulan: A java library for multi-label learning. *Journal of Machine Learning Research*, 12(Jul), 2411–2414.
- Tsoumakas, G., Spyromitros-Xioufis, E., Vilcek, J., & Vlahavas, I. (2020). Datasets from mulan: A java library for multi-label learning.
URL <http://mulan.sourceforge.net/datasets-mtr.html>
- Vehtari, A., & Lampinen, J. (1999). Bayesian neural networks for industrial applications. In *SMCia/99 Proceedings of the 1999 IEEE Midnight - Sun Workshop on Soft Computing Methods in Industrial Applications (Cat. No.99EX269)*, (pp. 63–68).
- Verma, A. (1999). Admat: Automatic differentiation in matlab using object oriented methods. In *SIAM Interdisciplinary Workshop on Object Oriented Methods for Interoperability*, (pp. 174–183). Citeseer.
- Villard, D., & Arnaldi, B. (1996). Symbolic differentiation library for simulation of multibody rigid systems. *Mathematics and Computers in Simulation*, 42(4-6), 659–673.
- Villatoro, F. R., & Ramos, J. I. (1999). On the method of modified equations. i: Asymptotic analysis of the euler forward difference method. *Applied mathematics and computation*, 103(2-3), 111–139.
- Volná, E. (2000). Using neural network in cryptography. In *The State of the Art in Computational Intelligence*, (pp. 262–267). Springer.
- Volna, E., Kotyrba, M., Kocian, V., & Janosek, M. (2012). Cryptography based on neural network. In *ECMS*, (pp. 386–391).
- Wainwright, M. J. (2014). Structured regularizers for high-dimensional problems: Statistical and computational issues. *Annual Review of Statistics and Its Application*, 1, 233–253.

- Walther, A. (2007). Automatic differentiation of explicit runge-kutta methods for optimal control. *Computational Optimization and Applications*, 36(1), 83–108.
- Wiegerinck, W., & Kappen, B. (2000). Approximations of bayesian networks through kl minimisation. *New Generation Computing*, 18(2), 167–175.
- Wijesiriwardana, C., & Firdhous, M. (2019). An innovative query tuning scheme for large databases. In *2019 International Conference on Data Science and Engineering (ICDSE)*, (pp. 154–159). IEEE.
- Williams, C. K., & Rasmussen, C. E. (2006). *Gaussian processes for machine learning*, vol. 2. MIT press Cambridge, MA.
- Xie, J., & Szymanski, B. K. (2011). Community detection using a neighborhood strength driven label propagation algorithm. In *2011 IEEE Network Science Workshop*, (pp. 188–195). IEEE.
- Xie, J., & Szymanski, B. K. (2013). Labelrank: A stabilized label propagation algorithm for community detection in networks. In *2013 IEEE 2nd Network Science Workshop (NSW)*, (pp. 138–143). IEEE.
- Yang, S., Cui, X., & Fang, Z. (2014). Bcrgt: a bayesian cluster regression-based genotyping algorithm for the samples with copy number alterations. *BMC bioinformatics*, 15(1), 74.
- Yang, X. (2017). Understanding the variational lower bound.
- Ye, L., Beskos, A., De Iorio, M., & Hao, J. (2020). Monte carlo co-ordinate ascent variational inference. *Statistics and Computing*, (pp. 1–19).
- Zhabinski, A., Zhabinskii, S., & Adzinets, D. (2017). Symbolic tensor differentiation for applications in machine learning.
- Zhai, K., Boyd-Graber, J., Asadi, N., & Alkhouja, M. L. (2012). Mr. lda: A flexible large scale topic modeling package using variational inference in mapreduce. In *Proceedings of the 21st International Conference on World Wide Web, WWW '12*, (pp. 879–888). New York, NY, USA: ACM.
URL <http://doi.acm.org/10.1145/2187836.2187955>
- Zhen, X., Yu, M., He, X., & Li, S. (2017). Multi-target regression via robust low-rank learning. *IEEE transactions on pattern analysis and machine intelligence*, 40(2), 497–504.
- Zorich, V. A. (2016). *Mathematical analysis II*. Springer.