



HAL
open science

Towards frugal unsupervised detection of subtle abnormalities in medical imaging

Geoffroy Oudoumanessah, Carole Lartizien, Michel Dojat, Florence Forbes

► **To cite this version:**

Geoffroy Oudoumanessah, Carole Lartizien, Michel Dojat, Florence Forbes. Towards frugal unsupervised detection of subtle abnormalities in medical imaging. 26th International Conference on Medical Image Computing and Computer Assisted Intervention, Oct 2023, Vancouver (BC), Canada. pp.1-13. hal-04192108

HAL Id: hal-04192108

<https://hal.science/hal-04192108>

Submitted on 31 Aug 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Towards frugal unsupervised detection of subtle abnormalities in medical imaging

Geoffroy Oudoumanessah^{1,2,3}, Carole Lartizien³, Michel Dojat², and Florence Forbes¹

¹ Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, 38000 Grenoble, France
{first.last}@inria.fr

² Univ. Grenoble Alpes, Inserm U1216, CHU Grenoble Alpes, Grenoble Institut des Neurosciences, 38000 Grenoble, France
{first.last}@univ-grenoble-alpes.fr

³ Univ. Lyon, CNRS, Inserm, INSA Lyon, UCBL, CREATIS, UMR5220, U1294, F-69621, Villeurbanne, France
{first.last}@creatis.insa-lyon.fr

Abstract. Anomaly detection in medical imaging is a challenging task in contexts where abnormalities are not annotated. This problem can be addressed through unsupervised anomaly detection (UAD) methods, which identify features that do not match with a reference model of normal profiles. Artificial neural networks have been extensively used for UAD but they do not generally achieve an optimal trade-off between accuracy and computational demand. As an alternative, we investigate mixtures of probability distributions whose versatility has been widely recognized for a variety of data and tasks, while not requiring excessive design effort or tuning. Their expressivity makes them good candidates to account for complex multivariate reference models. Their much smaller number of parameters makes them more amenable to interpretation and efficient learning. However, standard estimation procedures, such as the Expectation-Maximization algorithm, do not scale well to large data volumes as they require high memory usage. To address this issue, we propose to incrementally compute inferential quantities. This online approach is illustrated on the challenging detection of subtle abnormalities in MR brain scans for the follow-up of newly diagnosed Parkinsonian patients. The identified structural abnormalities are consistent with the disease progression, as accounted by the Hoehn and Yahr scale.

Keywords: Frugal computing · Online EM algorithm · Gaussian scale mixture · Unsupervised anomaly detection · Parkinson’s Disease

1 Introduction

Despite raising concerns about the environmental impact of artificial intelligence [35,37,38], the question of resource efficiency has not yet really reached medical imaging studies. The issue has multiple dimensions and the lack of clear metrics for a fair assessment of algorithms, in terms of resource and energy consumption,

contrasts with the obvious healthcare benefits of the ever growing performance of machine and statistical learning solutions.

In this work, we investigate the case of subtle abnormality detection in medical images, in an unsupervised context usually referred to as *Unsupervised Anomaly Detection* (UAD). This formalism requires only the identification of *normal* data to construct a normative model. *Anomalies* are then detected as outliers, *i.e.* as samples deviating from this normative model. Artificial neural networks (ANN) have been extensively used for UAD [21]. Either based on standard autoencoder (AE) architectures [3] or on more advanced architectures, *e.g.* combining a vector quantized AE with autoregressive transformers [33], ANN do not generally achieve an optimal trade-off between accuracy and computational demand. As an alternative, we show that more *frugal* approaches can be reached with traditional statistical models provided their cost in terms of memory usage can be addressed. Frugal solutions usually refer to strategies that can run with limited resources such as that of a single laptop. Frugal learning has been studied from several angles, in the form of constraints on the data acquired, on the algorithm deployed and on the nature of the proposed solution [9]. The angle we adopt is that of *online* or incremental learning, which refers to approaches that handle data in a sequential manner resulting in more efficient solutions in terms of memory usage and overall energy consumption. For UAD, we propose to investigate mixtures of probability distributions whose interpretability and versatility have been widely recognized for a variety of data and tasks, while not requiring excessive design effort or tuning. In particular, the use of multivariate Gaussian or generalized Student mixtures has been already demonstrated in many anomaly detection tasks, see [1,26,31] and references therein or [21] for a more general recent review. However, in their standard *batch* setting, mixtures are difficult to use with huge datasets due to the dramatic increase of time and memory consumption required by their estimation traditionally performed with an Expectation-Maximization (EM) algorithm [25]. Online more tractable versions of EM have been proposed and theoretically studied in the literature, *e.g.* [6,12], but with some restrictions on the class of mixtures that can be handled this way. A first natural approach is to consider Gaussian mixtures that belong to this class. We thus, present improvements regarding the implementation of an online EM for Gaussian mixtures. We then consider more general mixtures based on *multiple scale t-distributions* (MST) specifically adapted to outlier detection [10]. We show that these mixtures can be cast into the online EM framework and describe the resulting algorithm.

Our approach is illustrated with the MR imaging exploration of *de novo* (just diagnosed) Parkinson’s Disease (PD) patients, where brain anomalies are subtle and hardly visible in standard T1-weighted or diffusion MR images. The anomalies detected by our method are consistent with the Hoehn and Yahr (HY) scale [16], which describes how the symptoms of Parkinson’s disease progress. The results provide additional interesting clinical insights by pointing out the most impacted subcortical structures at both HY stages 1 and 2. The use of such an external scale appears to be an original and relevant indi-

rect validation, in the absence of ground truth at the voxel level. Energy and memory consumptions are also reported for batch and online EM to confirm the interesting performance/cost trade-off achieved. The code is available at <https://github.com/geoffroyO/onlineEM>.

2 UAD with mixture models

Recent studies have shown that, on subtle lesion detection tasks with limited data, alternative approaches to ANN, such as *one class support vector machine* or mixture models [1,26], were performing similarly [31,34]. We further investigate mixture-based models and show how the main UAD steps, *i.e.* the construction of a reference model and of a decision rule, can be designed.

Learning a reference model. We consider a set \mathbb{Y}_H of voxel-based features for a number of control (*e.g.* healthy) subjects, $\mathbb{Y}_H = \{\mathbf{y}_v, v \in \mathbb{V}_H\}$ where \mathbb{V}_H represents the voxels of all control subjects and $\mathbf{y}_v \in \mathbb{R}^M$ is typically deduced from image modality maps at voxel v or from abstract representation features provided by some ANN performing a pre-text task [22]. To account for the distribution of such normal feature vectors, we consider two types of mixture models, mixtures of Gaussian distributions with high tractability in multiple dimensions and mixtures of multiple scale t-distributions (MST) that are more appropriate when the data present elongated and strongly non-elliptical subgroups [10,1,26]. By fitting such a mixture model to the control data \mathbb{Y}_H , we build a reference model density f_H that depends on some parameter $\Theta_H = \{\theta_k, \pi_k, k = 1 : K_H\}$:

$$f_H(\mathbf{y}; \Theta_H) = \sum_{k=1}^{K_H} \pi_k f(\mathbf{y}; \theta_k), \quad (1)$$

with $\pi_k \in [0, 1]$, $\sum_{k=1:K_H} \pi_k = 1$ and K_H the number of components, each characterized by a distribution $f(\cdot; \theta_k)$. The EM algorithm is usually used to estimate Θ_H that best fits \mathbb{Y}_H while K_H can be estimated using *the slope heuristic* [2].

Designing a proximity measure. Given a reference model (1), a measure of proximity $r(\mathbf{y}_v; \Theta_H)$ of voxel v (with value \mathbf{y}_v) to f_H needs to be chosen. To make use of the mixture structure, we propose to consider distances to the respective mixture components through some weights acting as inverse Mahalanobis distances. We specify below this new proximity measure for MST mixtures. MST distributions are generalizations of the multivariate t-distribution that extend its Gaussian scale mixture representation [19]. The standard t-distribution univariate scale (weight) variable is replaced by a M -dimensional scale (weight) variable $\mathbf{W} = (W_m)_{m=1:M} \in \mathbb{R}^M$ with M the features dimension,

$$f_{MST}(\mathbf{y}; \theta) = \int_{[0, \infty]^M} \mathcal{N}_M(\mathbf{y}; \boldsymbol{\mu}, \mathbf{D} \boldsymbol{\Delta}_w \mathbf{A} \mathbf{D}^T) \prod_{m=1}^M \mathcal{G}(w_m; \frac{\nu_m}{2}) dw_1 \dots dw_M, \quad (2)$$

where $\mathcal{G}(\cdot, \frac{\nu_m}{2})$ denotes the gamma density with parameter $(\frac{\nu_m}{2}, \frac{\nu_m}{2}) \in \mathbb{R}^2$ and \mathcal{N}_M the multivariate normal distribution with mean parameter $\boldsymbol{\mu} \in \mathbb{R}^M$ and covariance matrix $\mathbf{D} \boldsymbol{\Delta}_w \mathbf{A} \mathbf{D}^T$ showing the scaling by the W_m 's through a diagonal

matrix $\mathbf{\Delta}_w = \text{diag}(w_1^{-1}, \dots, w_M^{-1})$. The MST parametrization uses the spectral decomposition of the scaling matrix $\mathbf{\Sigma} = \mathbf{D}\mathbf{A}\mathbf{D}^T$, with $\mathbf{D} \in \mathcal{O}(M) \subset \mathbb{R}^{M \times M}$ orthogonal and $\mathbf{A} = \text{diag}(A_1, \dots, A_M)$ diagonal. The whole set of parameters is $\boldsymbol{\theta} = \{\boldsymbol{\mu}, \mathbf{A}, \mathbf{D}, (\nu_m)_{m=1:M}\}$. The scale variable W_m for dimension m can be interpreted as accounting for the weight of this dimension and can be used to derive a measure of proximity. After fitting a mixture (1) with MST components to \mathbb{Y}_H , we set $r(\mathbf{y}_v; \boldsymbol{\theta}_H) = \max_{m=1:M} \bar{w}_m^{\mathbf{y}_v}$, with $\bar{w}_m^{\mathbf{y}} = \mathbb{E}[W_m | \mathbf{y}; \boldsymbol{\theta}_H]$. The proximity r is typically larger when at least one dimension of \mathbf{y}_v is well explained by the model. A similar proximity measure can also be derived for Gaussian mixtures, see details in the Supplementary Material Section 1.

Decision rule. For an effective detection, a threshold τ_α on proximity scores can be computed in a data-driven way by deciding on an acceptable false positive rate (FPR) α ; τ_α is the value such that $P(r(\mathbf{Y}; \boldsymbol{\theta}_H) < \tau_\alpha) = \alpha$, when \mathbf{Y} follows the f_H reference distribution. All voxels v whose proximity $r(\mathbf{y}_v; \boldsymbol{\theta}_H)$ is below τ_α are then labeled as abnormal. In practice, while f_H is known explicitly, the probability distribution of $r(\mathbf{Y}; \boldsymbol{\theta}_H)$ is not. However, it is easy to simulate this distribution or to estimate τ_α as an empirical α -quantile [1]. Unfortunately, learning f_H on huge datasets may not be possible due to the dramatic increase in time, memory and energy required by the EM algorithm. This issue often arises in medical imaging with the increased availability of multiple 3D modalities as well as the emergence of image-derived parametric maps such as radiomics [14] that should be analysed jointly, at the voxel level, and for a large number of subjects. A possible solution consists of employing powerful computers with graphics cards or grid-architectures in cloud computing. Here, we show that a more resource-friendly solution is possible using an online version of EM detailed in the next section.

3 Online mixture learning for large data volumes

Online learning refers to procedures able to deal with data acquired sequentially. Online variants of EM, among others, are described in [6,23,17,18,11,20,30]. As an archetype of such algorithms, we consider the online EM of [6] which belongs to the family of stochastic approximation algorithms [4]. This algorithm has been well theoretically studied and extended. However, it is designed only for distributions that admit a data augmentation scheme yielding a complete likelihood of the exponential family form, see (3) below. This case is already very broad, including Gaussian, gamma, t-distributions, etc. and mixtures of those. We recall below the main assumptions required and the online EM iteration.

Assume $(\mathbf{Y}_i)_{i=1}^n$ is a sequence of n independent and identically distributed replicates of a random variable $\mathbf{Y} \in \mathbb{Y} \subset \mathbb{R}^M$, observed one at a time. Extension to successive mini-batches of observations is straightforward [30]. In addition, \mathbf{Y} is assumed to be the visible part of the pair $\mathbf{X}^\top = (\mathbf{Y}^\top, \mathbf{Z}^\top) \in \mathbb{X}$, where $\mathbf{Z} \in \mathbb{R}^l$ is a latent variable, *e.g.* the unknown component label in a mixture model, and $l \in \mathbb{N}$. That is, each \mathbf{Y}_i is the visible part of a pair $\mathbf{X}_i^\top = (\mathbf{Y}_i^\top, \mathbf{Z}_i^\top)$. Suppose \mathbf{Y}

arises from some data generating process (DGP) characterised by a probability density function $f(\mathbf{y}; \boldsymbol{\theta}_0)$, with unknown parameters $\boldsymbol{\theta}_0 \in \mathbb{T} \subseteq \mathbb{R}^p$, for $p \in \mathbb{N}$.

Using the sequence $(\mathbf{Y}_i)_{i=1}^n$, the method of [6] sequentially estimates $\boldsymbol{\theta}_0$ provided the following assumptions are met:

(A1) The complete-data likelihood for \mathbf{X} is of the exponential family form:

$$f_c(\mathbf{x}; \boldsymbol{\theta}) = h(\mathbf{x}) \exp \left\{ [\mathbf{s}(\mathbf{x})]^\top \boldsymbol{\phi}(\boldsymbol{\theta}) - \psi(\boldsymbol{\theta}) \right\}, \quad (3)$$

with $h: \mathbb{R}^{M+l} \rightarrow [0, \infty)$, $\psi: \mathbb{R}^p \rightarrow \mathbb{R}$, $\mathbf{s}: \mathbb{R}^{M+l} \rightarrow \mathbb{R}^q$, $\boldsymbol{\phi}: \mathbb{R}^p \rightarrow \mathbb{R}^q$, for $q \in \mathbb{N}$.

(A2) The function

$$\bar{\mathbf{s}}(\mathbf{y}; \boldsymbol{\theta}) = \mathbb{E}[\mathbf{s}(\mathbf{X}) | \mathbf{Y} = \mathbf{y}; \boldsymbol{\theta}] \quad (4)$$

is well-defined for all \mathbf{y} and $\boldsymbol{\theta} \in \mathbb{T}$, where $\mathbb{E}[\cdot | \mathbf{Y} = \mathbf{y}; \boldsymbol{\theta}]$ is the conditional expectation when \mathbf{X} arises from the DGP characterised by $\boldsymbol{\theta}$.

(A3) There is a convex $\mathbb{S} \subseteq \mathbb{R}^q$, satisfying: (i) for all $\gamma \in (0, 1)$, $\mathbf{s} \in \mathbb{S}$, $\mathbf{y} \in \mathbb{Y}$, and $\boldsymbol{\theta} \in \mathbb{T}$, $(1 - \gamma)\mathbf{s} + \gamma\bar{\mathbf{s}}(\mathbf{y}; \boldsymbol{\theta}) \in \mathbb{S}$; and (ii) for any $\mathbf{s} \in \mathbb{S}$, the function $Q(\mathbf{s}; \boldsymbol{\theta}) = \mathbf{s}^\top \boldsymbol{\phi}(\boldsymbol{\theta}) - \psi(\boldsymbol{\theta})$ has a unique global maximizer on \mathbb{T} denoted by

$$\bar{\boldsymbol{\theta}}(\mathbf{s}) = \arg \max_{\boldsymbol{\theta} \in \mathbb{T}} Q(\mathbf{s}; \boldsymbol{\theta}). \quad (5)$$

Let $(\gamma_i)_{i=1}^n$ be a sequence of learning rates in $(0, 1)$ and let $\boldsymbol{\theta}^{(0)} \in \mathbb{T}$ be an initial estimate of $\boldsymbol{\theta}_0$. For each $i = 1 : n$, the online EM of [6] proceeds by computing

$$\mathbf{s}^{(i)} = \gamma_i \bar{\mathbf{s}}(\mathbf{y}_i; \boldsymbol{\theta}^{(i-1)}) + (1 - \gamma_i) \mathbf{s}^{(i-1)}, \quad (6)$$

and

$$\boldsymbol{\theta}^{(i)} = \bar{\boldsymbol{\theta}}(\mathbf{s}^{(i)}), \quad (7)$$

where $\mathbf{s}^{(0)} = \bar{\mathbf{s}}(\mathbf{y}_1; \boldsymbol{\theta}^{(0)})$. It is shown in Thm. 1 of [6] that when n tends to infinity, the sequence $(\boldsymbol{\theta}^{(i)})_{i=1:n}$ of estimators of $\boldsymbol{\theta}_0$ satisfies a convergence result to stationary points of the likelihood (cf. [6] for a more precise statement).

In practice, the algorithm implementation requires two quantities, $\bar{\mathbf{s}}$ in (4) and $\bar{\boldsymbol{\theta}}$ in (5). They are necessary to define the updating of sequences $(\mathbf{s}^{(i)})_{i=1:\infty}$ and $(\boldsymbol{\theta}^{(i)})_{i=1:\infty}$. We detail below these quantities for a MST mixture.

Online MST mixture EM. As shown in [29], the mixture case can be deduced from a single component case. The exponential form for a MST (2) writes:

$$\begin{aligned} f_c(\mathbf{x}; \boldsymbol{\theta}) &= \mathcal{N}_M(\mathbf{y}; \boldsymbol{\mu}, \mathbf{D} \boldsymbol{\Delta}_w \mathbf{A} \mathbf{D}^T) \prod_{m=1}^M \mathcal{G}\left(w_m; \frac{\nu_m}{2}\right), \quad \text{with } \mathbf{x} = (\mathbf{y}, \mathbf{w}) \quad (8) \\ &= h(\mathbf{y}, \mathbf{w}) \exp \left([\mathbf{s}(\mathbf{y}, \mathbf{w})]^\top \boldsymbol{\phi}(\boldsymbol{\mu}, \mathbf{D}, \mathbf{A}, \boldsymbol{\nu}) - \psi(\boldsymbol{\mu}, \mathbf{D}, \mathbf{A}, \boldsymbol{\nu}) \right) \end{aligned}$$

with $\mathbf{s}(\mathbf{y}, \mathbf{w}) = [w_1 \mathbf{y}, w_1 \text{vec}(\mathbf{y} \mathbf{y}^\top), w_1, \log w_1, \dots, w_M \mathbf{y}, w_M \text{vec}(\mathbf{y} \mathbf{y}^\top), w_M, \log w_M]^\top$, $\boldsymbol{\phi}(\boldsymbol{\mu}, \mathbf{D}, \mathbf{A}, \boldsymbol{\nu}) = [\boldsymbol{\phi}_1, \dots, \boldsymbol{\phi}_M]^\top$ with $\boldsymbol{\phi}_m$ equal to:

$$\boldsymbol{\phi}_m = \left[\frac{\mathbf{d}_m \mathbf{d}_m^\top \boldsymbol{\mu}}{A_m}, -\frac{\text{vec}(\mathbf{d}_m \mathbf{d}_m^\top)}{2A_m}, -\frac{\text{vec}(\mathbf{d}_m \mathbf{d}_m^\top)^\top \text{vec}(\boldsymbol{\mu} \boldsymbol{\mu}^\top)}{2A_m} - \frac{\nu_m}{2}, \frac{1 + \nu_m}{2} \right]$$

$$\text{and } \psi(\boldsymbol{\mu}, \mathbf{D}, \mathbf{A}, \boldsymbol{\nu}) = \sum_{m=1}^M \left(\frac{\log A_m}{2} + \log \Gamma\left(\frac{\nu_m}{2}\right) - \frac{\nu_m}{2} \log\left(\frac{\nu_m}{2}\right) \right),$$

where \mathbf{d}_m denotes the m^{th} column of \mathbf{D} and $\text{vec}(\cdot)$ the vectorisation operator, which converts a matrix to a column vector. The exact form of h is not important for the algorithm. It follows that $\bar{\boldsymbol{\theta}}(\mathbf{s})$ is defined as the unique maximizer of function $Q(\mathbf{s}, \boldsymbol{\theta}) = \mathbf{s}^T \boldsymbol{\phi}(\boldsymbol{\theta}) - \psi(\boldsymbol{\theta})$ where \mathbf{s} is a vector that matches the definition and dimension of $\boldsymbol{\phi}(\boldsymbol{\theta})$ and can be conveniently written as $\mathbf{s} = [\mathbf{s}_{11}, \text{vec}(\mathbf{S}_{21}), s_{31}, s_{41}, \dots, \mathbf{s}_{1M}, \text{vec}(\mathbf{S}_{2M}), s_{3M}, s_{4M}]^T$, with for each m , \mathbf{s}_{1m} is a M -dimensional vector, \mathbf{S}_{2m} is a $M \times M$ matrix, s_{3m} and s_{4m} are scalars. Solving for the roots of the Q gradients leads to $\bar{\boldsymbol{\theta}}(\mathbf{s}) = (\bar{\boldsymbol{\mu}}(\mathbf{s}), \bar{\mathbf{A}}(\mathbf{s}), \bar{\mathbf{D}}(\mathbf{s}), \bar{\nu}(\mathbf{s}))$ whose expressions are detailed in Supplementary Material Section 2.

A second important quantity is $\bar{\mathbf{s}}(\mathbf{y}, \boldsymbol{\theta}) = \mathbb{E}[\mathbf{s}(\mathbf{X}) | \mathbf{Y} = \mathbf{y}; \boldsymbol{\theta}]$. This quantity requires to compute the following expectations for all m , $\mathbb{E}[W_m | \mathbf{Y} = \mathbf{y}; \boldsymbol{\theta}]$ and $\mathbb{E}[\log W_m | \mathbf{Y} = \mathbf{y}; \boldsymbol{\theta}]$. More specifically in the update equation (6), these expectations need to be computed for $\mathbf{y} = \mathbf{y}_i$ the observation at iteration i . We therefore denote these expectations respectively by

$$w_{im}^{(i-1)} = \mathbb{E}[W_m | \mathbf{Y} = \mathbf{y}_i; \boldsymbol{\theta}^{(i-1)}] = \alpha_m^{(i-1)} / \beta_m^{(i-1)} \quad (9)$$

and $\tilde{u}_{im}^{(i-1)} = \mathbb{E}[\log W_m | \mathbf{Y} = \mathbf{y}_i; \boldsymbol{\theta}^{(i-1)}] = \Psi^{(0)}(\alpha_m^{(i-1)}) - \log \beta_m^{(i-1)}$, where $\alpha_m^{(i-1)} = \frac{\nu_m^{(i-1)} + 1}{2}$ and $\beta_m^{(i-1)} = \frac{\nu_m^{(i-1)}}{2} + \frac{(\mathbf{d}_m^{(i-1)T}(\mathbf{y}_i - \boldsymbol{\mu}^{(i-1)}))^2}{2A_m^{(i-1)}}$. The update of $\mathbf{s}^{(i)}$

in (6) follows from the update for each m . From this single MST iteration, the mixture case is easily derived, see [29] or Supplementary Material Section 2.

Online Gaussian mixture EM. This case can be found in previous work *e.g.* [6,30] but to our knowledge, implementation optimizations are never really addressed. We propose an original version that saves computations, especially in a multivariate case where $\bar{\boldsymbol{\theta}}(\mathbf{s})$ involves large matrix inverses and determinants. Such inversions are avoided using results detailed in Supplementary Section 3.

4 Brain abnormality exploration in *de novo* PD patients

Data description and preprocessing. The Parkinson’s Progression Markers Initiative (PPMI) [24] is an open-access database dedicated to PD. It includes MR images of *de novo* PD patients, as well as of healthy subjects (HC), all acquired on the same 3T Siemens Trio Tim scanner. For our illustration, we use 108 HC and 419 PD samples, each composed of a 3D T1-weighted image (T1w), Fractional Anisotropy (FA) and Mean Diffusivity (MD) volumes. The two latter are extracted from diffusion imaging using the DiPy package [13], registered onto T1w and interpolated to the same spatial resolution with SPM12. Standard T1w preprocessing steps, comprising non-local mean denoising, skull stripping and tissue segmentation are also performed with SPM12. HC and PD groups are age-matched (median age: 64 y.) with the male-female ratio equal to 6:4. We focus on some subcortical structures, which are mostly impacted at the early stage of the disease [7], Globus Pallidus external and internal (GPe and GPi), Nucleus Accumbens (NAC), Substantia Nigra reticulata (SNr), Putamen (Pu), Caudate (Ca) and Extended Amygdala (EXA). Their position is determined by projecting the CIT168 atlas [32] onto each individual image.

Pipeline and results. We follow Sections 2 and 3 using T1w, FA and MD volumes as features ($M = 3$) and a FPR $\alpha = 0.02$. The pipeline is repeated 10 times for cross-validation. Each fold is composed of 64 randomly selected HC images for training (about 70M voxels), the remaining 44 HC and all the PD samples for testing. For the reference model, we test Gaussian and MST mixtures, with respectively $K_H = 14$ and $K_H = 8$, estimated with the slope heuristic. Abnormal voxels are then detected for all test subjects, on the basis of their proximity to the learned reference model, as detailed in Section 2.

The PPMI does not provide ground truth information at the voxel level. This is a recurring issue in UAD, which limits validations to mainly qualitative ones. For a more quantitative evaluation, we propose to resort to an auxiliary task whose success is likely to be correlated with a good anomaly detection. We consider the classification of test subjects into healthy and Parkinsonian subjects based on their global (over all brain) percentages of abnormal voxels. We exploit the availability of HY values to divide the patients into two HY=1 and HY=2 groups, representing the two early stages of the disease’s progression. Classification results yield a median g-mean, for stage 1 vs stage 2, respectively of 0.59 vs 0.63 for the Gaussian mixtures model and 0.63 vs 0.65 for the MST mixture. The ability of both mixtures to better differentiate stage 2 than stage 1 patients from HC is consistent with the progression of the disease. Note that the structural differences between these two PD stages remain subtle and difficult to detect, demonstrating the efficiency of the models. The MST mixture model appears better in identifying stage 2 PD patients based on their abnormal voxels.

To gain further insights, we report, in Fig. 1, the percentages of anomalies detected in each subcortical structure, for control, stage 1 and stage 2 groups. For each structure and both mixture models, the number of anomalies increases from control to stage 1 and stage 2 groups. As expected the MST mixture shows a better ability to detect outliers with significant differences between HC and PD groups, while for the Gaussian model, percentages do not depart much from that in the control group. Overall, in line with the know pathophysiology [7], MST results suggest clearly that all structures are potential good markers of the disease progression at these early stages, with GPe, GPi, EXA and SNr showing the largest impact.

Regarding efficiency, energy consumption in kilojoules (kJ) is measured using the PowerAPI library [5]. In Table 1, we report the energy consumption for the training and testing of one random fold, comparing our online mixtures with AE-supported methods for UAD [3], namely the patch-based reconstruction error [3] and FastFlow [39]. We implemented both methods with two different AE architectures: a lightweight AE already used for *de novo* PD detection [34], and a larger one, ResNet-18 [15]. The global g-mean (not taking HY stages into account) is also reported for the chosen fold. The experiments were run on a CPU with Intel Cascade Lake 6248@2.5GHz (20 cores), and a GPU Nvidia V100-32GB. Online mixtures exhibit significantly lower energy consumption, both for training and inference. In terms of memory cost, DRAM peak results, as measured by the *tracemalloc* Python library, also show lower costs for online




Fig. 1. Left: Median, over 10 folds, percentages of anomalies (0 to 22%) in each subcortical structure (see text for full names) for control subjects (green), stage 1 (blue) and stage 2 (red) patients. Plain and dotted lines indicate respectively results obtained with the MST and Gaussian mixtures. Structure sizes in voxels are indicated in parenthesis. SC refers to the combination of all structures. Right: 3D rendering of the subcortical structures colored according to MST percentages from 0% (green) to 22% (red), for healthy controls (HC), stage 1 and stage 2 groups.

mixtures, which by design deal with batches of voxels of smaller sizes than the batches of patches used in AE solutions. These results highlight the advantage of online mixtures, which compared to other hardware-demanding methods, can be run on a minimal configuration while maintaining good performance.

Table 1. UAD methods comparison for one fold: online Gaussian (OGMM) and Student (OMMST) mixtures, Lightweight AE and ResNet-18 architectures with reconstruction error (RE) and FastFlow (FF) based detection. Best values in bold font.

Method	Backend	Training			Inference			Gmean	Parameters
		Time	Consumption	DRAM peak	Time	Consumption	DRAM peak		
Online Mixtures (ours)									
OGMM	CPU	50s	85 kJ	494 MB	17min	23 kJ	92 MB	0.65	140
OMMST	CPU	1min20	153 kJ	958 MB	18min	32 kJ	96 MB	0.67	128
Lightweight AE									
RE	GPU	1h26	5040 kJ	26 GB	3h30	8350 kJ	22 GB	0.61	5266
FF	GPU	4h	6854 kJ	27 GB	3h53	13158 kJ	27 GB	0.55	1520
Resnet-18									
RE	GPU	17h40	53213 kJ	26 GB	59h	108593 kJ	28 GB	0.64	23730218
FF	GPU	4h10	7234 kJ	28 GB	19h45	18481 kJ	28GB	0.61	1520

5 Conclusion and perspectives

Despite a challenging medical problematic of PD progression at early stages, we have observed that energy and memory efficient methods could yield interesting and comparable results with other studies performed on the same database [34,28] and with similar MR modalities [36,8,26,27]. An interesting future work

would be to investigate the possibility to use more structured observations, such as patch-based features [28] or latent representations from a preliminary pretext task, provided the task cost is reasonable. Overall, we have illustrated that the constraints of Green AI [35] could be considered in medical imaging by producing innovative results without increasing computational cost or even reducing it. We have investigated statistical mixture models for an UAD task and shown that their expressivity could account for multivariate reference models, and their much simpler structure made them more amenable to efficient learning than most ANN solutions. Although very preliminary, we hope this attempt will open the way to the development of more methods that can balance the environmental impact of growing energy cost with the obtained healthcare benefits.

6 Data use declaration and acknowledgement

G. Oudoumanessah was financially supported by the AURA region. This work has been partially supported by MIAI@Grenoble Alpes (ANR-19-P3IA-0003), and was granted access to the HPC resources of IDRIS under the allocation 2022-AD011013867 made by GENCI. The data used in the preparation of this article were obtained from the Parkinson’s Progression Markers Initiative database www.ppmi-info.org/access-data-specimens/download-data openly available for researchers.

References

1. Arnaud, A., Forbes, F., Coquery, N., Collomb, N., Lemasson, B., Barbier, E.: Fully Automatic Lesion Localization and Characterization: Application to Brain Tumors Using Multiparametric Quantitative MRI Data. *IEEE Transactions on Medical Imaging* **37**(7), 1678–1689 (2018)
2. Baudry, J.P., Maugis, C., Michel, B.: Slope heuristic: overview and implementation. *Stat. Comp.* **22**, 455–470 (2012)
3. Baur, C., Denner, S., Wiestler, B., Navab, N., Albarqouni, S.: Autoencoders for unsupervised anomaly segmentation in brain MR images: A comparative study. *Medical Image Analysis* **69**, 101952 (2021)
4. Borkar, V.: *Stochastic approximation: a dynamical view point*. Cambridge University Press (2008)
5. Bourdon, A., Noureddine, A., Rouvoy, R., Seinturier, L.: PowerAPI: A software library to monitor the energy consumed at the process-level. *ERCIM News* (2013)
6. Cappé, O., Moulines, E.: On-line Expectation-Maximization algorithm for latent data models. *Journal of the Royal Statistical Society B* **71**, 593–613 (2009)
7. Dexter, D.T., Wells, F.R., Agid, F., Agid, Y., Lees, A.J., Jenner, P., Marsden, C.D.: Increased nigral iron content in postmortem Parkinsonian brain. *Lancet* (1987)
8. Du, G., Lewis, M.M., Styner, M., Shaffer, M.L., Sen, S., Yang, Q.X., Huang, X.: Combined R2* and Diffusion Tensor Imaging Changes in the Substantia Nigra in Parkinson’s Disease. *Movement Disorders* **26**(9), 1627–1632 (2011)
9. Evchenko, M., Vanschoren, J., Hoos, H.H., Schoenauer, M., Sebag, M.: Frugal machine learning. *ArXiv abs/2111.03731* (2021)
10. Forbes, F., Wraith, D.: A new family of multivariate heavy-tailed distributions with variable marginal amounts of tailweights: Application to robust clustering. *Statistics and Computing* **24**(6), 971–984 (2014)
11. Fort, G., Moulines, E., Wai, H.T.: A stochastic path-integrated differential estimator expectation maximization algorithm. In: 34th Conference on Neural Information Processing Systems (NeurIPS) (2020)
12. Fort, G., Gach, P., Moulines, E.: Fast Incremental Expectation Maximization for finite-sum optimization: nonasymptotic convergence. *Stat. Comp.* **31**(48) (2021)
13. Garyfallidis, E., Brett, M., Amirkhanyan, B., Rokem, A., Van Der Walt, S., Descoteaux, M., Nimmo-Smith, I., Contributors, D.: Dipy, a library for the analysis of diffusion MRI data. *Frontiers in neuroinformatics* **8** (2014)
14. Gillies, R.J., Kinahan, P.E., Hricak, H.: Radiomics: Images are more than pictures, they are data. *Radiology* **278**(2), 563–77 (2016)
15. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *IEEE Conf. Comp. Vis. Patt. Rec.* pp. 770–778 (2016)
16. Hoehn, M.M., Yahr, M.D.: Parkinsonism: onset, progression, and mortality. *Neurology* **50**(2), 318–318 (1998)
17. Karimi, B., Miasojedow, B., Moulines, E., Wai, H.T.: Non-asymptotic analysis of biased stochastic approximation scheme. *Proc. Mach. Learn. Res.* **99**, 1–31 (2019)
18. Karimi, B., Wai, H.T., Moulines, E., Lavielle, M.: On the global convergence of (fast) incremental Expectation Maximization methods. In: 33rd Conference on Neural Information Processing Systems (NeurIPS) (2019)
19. Kotz, S., Nadarajah, S.: *Multivariate t Distributions And Their Applications*. Cambridge University Press (2004)
20. Kuhn, E., Matias, C., Rebafka, T.: Properties of the stochastic approximation EM algorithm with mini-batch sampling. *Stat. Comp.* **30**, 1725–1739 (2020)

21. Lagogiannis, I., Meissen, F., Kaissis, G., Rueckert, D.: Unsupervised pathology detection: A deep dive into the state of the art. ArXiv **abs:2303.00609** (2023)
22. Li, C., Sohn, K., Yoon, J., Pfister, T.: Cutpaste: Self-supervised learning for anomaly detection and localization. In: IEEE Conf. Comp. Vis. Patt. Rec. pp. 9664–9674 (2021)
23. Maire, F., Moulines, E., Lefebvre, S.: Online EM for functional data. *Computational Statistics and Data Analysis* **111**, 27–47 (2017)
24. Marek, K., Chowdhury, S., Siderowf, A., Lasch, S., Coffey, C., Caspell-Garcia, C., Simuni, T., Jennings, D., M. Tanner, C., Q. Trojanowski, J., Shaw, L., Seibyl, J., Schuff, N., Singleton, A., Kieburtz, K., W. Toga, A., Mollenhauer, B., Galasko, D., M. Chahine, L.: The Parkinson’s progression markers initiative - establishing a PD biomarker cohort. *Ann. Clinical Translational Neurology* p. 1460–1477 (2018)
25. McLachlan, G.J., Krishnan, T.: The EM algorithm and extensions. John Wiley & Sons (2007)
26. Munoz-Ramirez, V., Forbes, F., Arbel, J., Arnaud, A., Dojat, M.: Quantitative MRI characterization of brain abnormalities in de novo Parkinsonian patients. In: IEEE Int. Symp. on Bio. Im. (2019)
27. Muñoz-Ramírez, V., Kmetzsch, V., Forbes, F., Meoni, S., Moro, E., Dojat, M.: Subtle anomaly detection in MRI brain scans: Application to biomarkers extraction in patients with de novo parkinson’s disease. *Artificial Intelligence Medicine* (2021)
28. Muñoz-Ramírez, V., Pinon, N., Forbes, F., Lartizen, C., Dojat, M.: Patch vs. Global Image-Based Unsupervised Anomaly Detection in MR Brain Scans of Early Parkinsonian Patients. In: Machine Learning in Clinical Neuroimaging (2021)
29. Nguyen, H.D., Forbes, F.: Global implicit function theorems and the online expectation-maximisation algorithm. *Australian New Zealand J. Stat.* (2022)
30. Nguyen, H.D., Forbes, F., McLachlan, G.J.: Mini-batch learning of exponential family finite mixture models. *Stat. Comp.* **30**, 731–748 (2020)
31. Oluwasegun, A., Jung, J.C.: A multivariate Gaussian mixture model for anomaly detection in transient current signature of control element drive mechanism. *Nuclear Engineering and Design* **402**, 112098 (2023)
32. Pauli, W.M., Nili, A.N., Tyszka, J.M.: A high-resolution probabilistic in vivo atlas of human subcortical brain nuclei. *Scientific data* **5**(1), 1–13 (2018)
33. Pinaya, W.H., Tudosiu, P.D., Gray, R., Rees, G., Nachev, P., Ourselin, S., Cardoso, M.J.: Unsupervised brain imaging 3D anomaly detection and segmentation with transformers. *Medical Image Analysis* **79**, 102475 (2022)
34. Pinon, N., Oudoumanessah, G., Trombetta, R., Dojat, M., Forbes, F., Lartizen, C.: Brain subtle anomaly detection based on auto-encoders latent space analysis : application to de novo Parkinson patients. In: IEEE Int. Symp. on Bio. Im. (2023)
35. Schwartz, R., Dodge, J., Smith, N., Etzioni, O.: Green AI. *Commun. ACM* **63**(12), 54–63 (2020)
36. Schwarz, S.T., Abaei, M., Gontu, V., Morgan, P.S., Bajaj, N., Auer, D.P.: Diffusion tensor imaging of nigral degeneration in Parkinson’s disease: A region-of-interest and voxel-based study at 3T and systematic review with meta-analysis. *NeuroImage: Clinical* **3**, 481–488 (2013)
37. Strubell, E., Ganesh, A., McCallum, A.: Energy and Policy Considerations for Deep Learning in NLP. In: 57th Meeting of Assoc. Computational Linguistics (2019)
38. Thompson, N.C., Greenewald, K., Lee, K., Manso, G.F.: The computational limits of deep learning. ArXiv **abs/2007.05558** (2022)
39. Yu, J., Zheng, Y., Wang, X., Li, W., Wu, Y., Zhao, R., Wu, L.: Fastflow: Unsupervised anomaly detection and localization via 2D normalizing flows. ArXiv **abs/2111.07677** (2021)

Supplementary Material: Towards frugal unsupervised detection of subtle abnormalities in medical imaging

1 Weight-based proximity measure for Gaussian mixtures

For a Gaussian distribution with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma} = \mathbf{D}^T \mathbf{A} \mathbf{D}$, the Mahalanobis distance of a sample \mathbf{y} is $(\mathbf{y} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{y} - \boldsymbol{\mu})$ which can be decomposed as $\sum_{m=1:M} (\mathbf{d}_m^T (\mathbf{y} - \boldsymbol{\mu}))^2 / A_m$. Each term in the sum can be interpreted as a dimension-wise distance and its inverse as a weight. This inverse corresponds to expression u_{im} in Section 3 of the manuscript with $\nu_m = 0$. Using this similarity we define a proximity for Gaussian mixtures as the maximum over m of these weights whose expressions are the same than for MST mixtures with $\nu_m = 0$.

2 Online MST mixture EM updates

We first specify the expression of $\bar{\boldsymbol{\theta}}(\mathbf{s})$ for a single MST distribution. We can define $\bar{\boldsymbol{\theta}}(\mathbf{s})$ as the root of $\mathbf{J}_\phi(\boldsymbol{\theta}) \mathbf{s} - \frac{\partial \psi}{\partial \boldsymbol{\theta}}(\boldsymbol{\theta}) = \mathbf{0}$, where $\mathbf{J}_\phi(\boldsymbol{\theta}) = \partial \phi / \partial \boldsymbol{\theta}$ is the Jacobian of ϕ , with respect to $\boldsymbol{\theta}$, as a function of $\boldsymbol{\theta}$. Parameters $(\boldsymbol{\mu}, \mathbf{A}, \mathbf{D})$ and $\boldsymbol{\nu}$ can be optimized separately. It follows that,

$\boldsymbol{\mu} = \left(\sum_{m=1}^M \frac{s_{3m}}{A_m} (\mathbf{d}_m \mathbf{d}_m^T) \right)^{-1} \left(\sum_{m=1}^M \frac{\mathbf{d}_m \mathbf{d}_m^T}{A_m} \mathbf{s}_{1m} \right) = \mathbf{D} \mathbf{S}_3^{-1} \mathbf{v}$ where $\mathbf{S}_3 = \text{diag}(s_{31}, \dots, s_{3M})$ and \mathbf{v} is defined as $\mathbf{v}^T = (\mathbf{d}_1^T \mathbf{s}_{11}, \dots, \mathbf{d}_M^T \mathbf{s}_{1M})$. For matrix \mathbf{A} , we get for each m ,

$$\begin{aligned} A_m &= \text{vec}(\mathbf{S}_{2m})^T \text{vec}(\mathbf{d}_m \mathbf{d}_m^T) + s_{3m} \text{vec}(\mathbf{d}_m \mathbf{d}_m^T)^T \text{vec}(\boldsymbol{\mu} \boldsymbol{\mu}^T) - 2 \mathbf{s}_{1m}^T \mathbf{d}_m \mathbf{d}_m^T \boldsymbol{\mu} \\ &= \mathbf{d}_m^T \mathbf{S}_{2m} \mathbf{d}_m - \frac{(\mathbf{d}_m^T \mathbf{s}_{1m})^2}{s_{3m}}, \end{aligned} \quad (1)$$

where the last equality is obtained by replacing $\boldsymbol{\mu}$ by its expression above. Then, plugging-in the expressions of $\boldsymbol{\mu}$ and \mathbf{A} above and omitting parts that depend on $\boldsymbol{\nu}$ only, it comes,

$$\mathbf{D} = \arg \min_{\mathbf{D}, \mathbf{D}^T \mathbf{D} = \mathbf{I}d} \sum_{m=1}^M \log \left(\mathbf{d}_m^T \left(\mathbf{S}_{2m} - \frac{\mathbf{s}_{1m} \mathbf{s}_{1m}^T}{s_{3m}} \right) \mathbf{d}_m \right) \quad (2)$$

With the orthogonality constraint on \mathbf{D} , this can be solved using a version of the conjugate gradient algorithm designed for manifolds with the Polak-Ribiere line search computing the Stiefel gradient. See [1, 5] for details. For $\boldsymbol{\nu}$, as for the standard t -distribution, we have to solve the following equation in ν_m for each m , $s_{4m} - s_{3m} - \Psi^{(0)}(\nu_m/2) + \log(\nu_m/2) + 1 = 0$, where $\Psi^{(0)}$ is the digamma function.

Then, the online EM for a K -component mixture \mathcal{M} can be derived from the online EM for a single component. It consists of computing $\bar{\boldsymbol{\theta}}_{\mathcal{M}}(\mathbf{s}_{\mathcal{M}})$ where $\mathbf{s}_{\mathcal{M}} = (s_{01}, \mathbf{s}_{\mathcal{M}1}^\top, \dots, s_{0K}, \mathbf{s}_{\mathcal{M}K}^\top)$ and $\mathbf{s}_{\mathcal{M}k}$ has the structure given in the manuscript for one MST distribution. We can show [2], that $\bar{\boldsymbol{\theta}}_{\mathcal{M}}(\mathbf{s}_{\mathcal{M}})^\top = (\bar{\pi}_1(\mathbf{s}_{\mathcal{M}}), \boldsymbol{\theta}(\mathbf{s}_{\mathcal{M}1}/s_{01})^\top, \dots, \bar{\pi}_K(\mathbf{s}_{\mathcal{M}}), \boldsymbol{\theta}(\mathbf{s}_{\mathcal{M}K}/s_{0K})^\top)$ where $\bar{\pi}_k(\mathbf{s}_{\mathcal{M}}) = s_{0k} / (\sum_{z=1:K} s_{0z})$ and $\bar{\boldsymbol{\theta}}$ is the expression found for one single MST component, detailed in the manuscript.

3 Improved implementation of online Gaussian mixture EM

We cannot detail the online EM for K-Gaussian mixtures but the algorithm can be found in previous work. We focus here on our proposed improvements. We use the notation in [3] just changing the indices to match notation in our previous section, that is for each $k = 1 : K$, we set $(s_{0k}, \mathbf{s}_{1k}, \mathbf{S}_{2k}) = (s_{1k}, \mathbf{s}_{2k}, \mathbf{S}_{3k})$. At iteration i , the update of $\boldsymbol{\theta}^{(i)}$ involves the update of K covariance matrices, for $k = 1 : K$, $\boldsymbol{\Sigma}_k^{(i)} = \mathbf{S}_{2k}^{(i)} / s_{0k}^{(i)} - \mathbf{s}_{1k}^{(i)} \mathbf{s}_{1k}^{(i)T} / s_{0k}^{(i)2}$. The subsequent update of $\mathbf{s}^{(i+1)}$ then requires the inverse covariances (i.e. precisions) and their determinants whose computation is costly. We thus propose to work directly with precision matrices. Applying the Sherman-Morrison formula [4] to the equation above gives,

$$\boldsymbol{\Sigma}_k^{(i)-1} = s_{0k}^{(i)} \mathbf{S}_{2k}^{(i)-1} + \frac{\mathbf{S}_{2k}^{(i)-1} \mathbf{s}_{1k}^{(i)} \mathbf{s}_{1k}^{(i)T} \mathbf{S}_{2k}^{(i)-1}}{1 - \frac{1}{s_{0k}^{(i)}} \mathbf{s}_{1k}^{(i)T} \mathbf{S}_{2k}^{(i)-1} \mathbf{s}_{1k}^{(i)}}.$$

To avoid any matrix inversion, we also apply the Sherman-Morrison formula to the update of \mathbf{S}_2 ,

$$\mathbf{S}_{2k}^{(i)-1} = \frac{1}{1 - \gamma_i} \mathbf{S}_{2k}^{(i-1)-1} - \frac{\gamma_i \tau_{ki}}{(1 - \gamma_i)^2} \frac{\mathbf{S}_{2k}^{(i-1)-1} \mathbf{y}_i \mathbf{y}_i^T \mathbf{S}_{2k}^{(i-1)-1}}{1 + \frac{\gamma_i \tau_{ki}}{(1 - \gamma_i)} \mathbf{y}_i^T \mathbf{S}_{2k}^{(i-1)-1} \mathbf{y}_i}.$$

Similarly, the determinant version of the Sherman-Morrison formula gives our new updates for the precision determinants and the \mathbf{S}_2 statistics:

$$\det(\boldsymbol{\Sigma}_k^{(i)-1}) = \frac{s_{0k}^{(i)M}}{1 - \frac{\mathbf{s}_{1k}^{(i)T} \mathbf{S}_{2k}^{(i-1)-1} \mathbf{s}_{1k}^{(i)}}{s_{0k}^{(i)}}} \det(\mathbf{S}_{2k}^{(i)-1})$$

$$\det(\mathbf{S}_{2k}^{(i)-1}) = \left(1 + \frac{\gamma_i \tau_{ki}}{1 - \gamma_i} \mathbf{y}_i^T \mathbf{S}_{2k}^{(i-1)-1} \mathbf{y}_i \right)^{-1} (1 - \gamma_i)^{-M} \det(\mathbf{S}_{2k}^{(i-1)-1}).$$

A similar improvement can be derived for a mini-batch online EM [3] by applying both Sherman-Morrison formulas recursively.

References

- [1] Absil, P.A., Mahony, R., Sepulchre, R.: Optimization Algorithms on Matrix Manifolds. Princeton University Press, Princeton, NJ (2008)
- [2] Nguyen, H.D., Forbes, F.: Global implicit function theorems and the online expectation-maximisation algorithm. Australian New Zealand J. Stat. (2022)
- [3] Nguyen, H.D., Forbes, F., McLachlan, G.J.: Mini-batch learning of exponential family finite mixture models. Stat. Comp. **30**, 731–748 (2020)
- [4] Sherman, J., Morrison, W.J.: Adjustment of an Inverse Matrix Corresponding to a Change in One Element of a Given Matrix. The Annals of Mathematical Statistics **21**(1), 124 – 127 (1950)
- [5] Townsend, J., Koep, N., Weichwald, S.: Pymanopt: A python toolbox for optimization on manifolds using automatic differentiation. Journal of Machine Learning Research **17**(137), 1–5 (2016)