

RESEARCH ARTICLE

A Fast Incremental Gaussian Mixture Model

Rafael Coimbra Pinto*, Paulo Martins Engel

Instituto de Informática, Universidade Federal do Rio Grande do Sul, Porto Alegre, RS, Brazil

* rcpinto@inf.ufrgs.br

Abstract

This work builds upon previous efforts in online incremental learning, namely the Incremental Gaussian Mixture Network (IGMN). The IGMN is capable of learning from data streams in a single-pass by improving its model after analyzing each data point and discarding it thereafter. Nevertheless, it suffers from the scalability point-of-view, due to its asymptotic time complexity of $O(NKD^3)$ for N data points, K Gaussian components and D dimensions, rendering it inadequate for high-dimensional data. In this work, we manage to reduce this complexity to $O(NKD^2)$ by deriving formulas for working directly with precision matrices instead of covariance matrices. The final result is a much faster and scalable algorithm which can be applied to high dimensional tasks. This is confirmed by applying the modified algorithm to high-dimensional classification datasets.



OPEN ACCESS

Citation: Pinto RC, Engel PM (2015) A Fast Incremental Gaussian Mixture Model. PLoS ONE 10 (10): e0139931. doi:10.1371/journal.pone.0139931

Editor: Thomas R. Ioerger, Texas A&M University, UNITED STATES

Received: May 22, 2015

Accepted: September 18, 2015

Published: October 7, 2015

Copyright: © 2015 Pinto, Engel. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: Data are available at Figshare (http://figshare.com/articles/A_Fast_Incremental_Gaussian_Mixture_Model/1552030). The MNIST data set is available at (<http://yann.lecun.com/exdb/mnist/>) and the CIFAR10 data set is available at (<http://www.cs.toronto.edu/~kriz/cifar.html>). The software binaries and source code are deposited to GitHub (<https://github.com/amtan/NeuralNetwork>).

Funding: The authors have no support or funding to report.

Competing Interests: The authors have declared that no competing interests exist.

1 Introduction

The Incremental Gaussian Mixture Network (IGMN) [1, 2] is a supervised algorithm which approximates the EM algorithm for Gaussian mixture models [3], as shown in [4]. It creates and continually adjusts a probabilistic model of the joint input-output space consistent to all sequentially presented data, after each data point presentation, and without the need to store any past data points. Its learning process is aggressive, meaning that only a single scan through the data is necessary to obtain a consistent model.

IGMN adopts a Gaussian mixture model of distribution components that can be expanded to accommodate new information from an input data point, or reduced if spurious components are identified along the learning process. Each data point assimilated by the model contributes to the sequential update of the model parameters based on the maximization of the likelihood of the data. The parameters are updated through the accumulation of relevant information extracted from each data point. New points are added directly to existing Gaussian components or new components are created when necessary, avoiding merge and split operations, much like what is seen in the Adaptive Resonance Theory (ART) algorithms [5]. It has been previously shown in [6] that the algorithm is robust even when data is presented in random order, having similar performance and producing similar number of clusters in any order. Also, [4] has shown that the resulting models are very similar to the ones produced by the batch EM algorithm.

The IGMN is capable of supervised learning, simply by assigning any of its input vector elements as outputs. In other words, any element can be used to predict any other element, like

auto-associative neural networks [7] or missing data imputation [8]. This feature is useful for simultaneous learning of forward and inverse kinematics [9], as well as for simultaneous learning of a value function and a policy in reinforcement learning [10].

Previous successful applications of the IGMN algorithm include time-series prediction [11–13], reinforcement learning [2, 14], mobile robot control and mapping [1, 15, 16] and outlier detection in data streams [17].

However, the IGMN suffers from cubic time complexity due to matrix inversion operations and determinant computations. Its time complexity is of $O(NKD^3)$, where N is the number of data points, K is the number of Gaussian components and D is the problem dimension. It makes the algorithm prohibitive for high-dimensional tasks (like visual tasks) and thus of limited use. One solution would be to use diagonal covariance matrices, but this decreases the quality of the results, as already reported in previous works [6, 11]. In this work, we propose the use of rank-one updates for both inverse matrices and determinants applied to full covariance matrices, thus reducing the time complexity to $O(NKD^2)$ for learning while keeping the quality of a full covariance matrix solution.

For the specific case of the IGMN algorithm, to the best of our knowledge, this has not been tried before, although we can find similar efforts for related algorithms. In [18], rank-one updates were applied to an iterated linear discriminant analysis algorithm in order to decrease the complexity of the algorithm. Rank-one updates were also used in [19], where Gaussian models are employed for feature selection. Finally, in [20], the same kind of optimization was applied to Maximum Likelihood Linear Transforms (MLLT).

The next Section describes the algorithm in more detail with the latest improvements to date. Section 3 describes our improvements to the algorithm. Section 4 shows the experiments and results obtained from both versions of the IGMN for comparison, and Section 5 finishes this work with concluding remarks.

2 Incremental Gaussian Mixture Network

In the next subsections we describe the current version of the IGMN algorithm, a slightly improved version of the one described in [15].

2.1 Learning

The algorithm starts with no components, which are created as necessary (see subsection 2.2). Given input \mathbf{x} (a single instantaneous data point), the IGMN algorithm processing step is as follows. First, the squared Mahalanobis distance $d^2(\mathbf{x}, j)$ for each component j is computed:

$$d_M^2(\mathbf{x}, j) = (\mathbf{x} - \boldsymbol{\mu}_j)^T \mathbf{C}_j^{-1} (\mathbf{x} - \boldsymbol{\mu}_j) \tag{1}$$

where $\boldsymbol{\mu}_j$ is the j^{th} component mean, \mathbf{C}_j its full covariance matrix. If any $d^2(\mathbf{x}, j)$ is smaller than $\chi_{D, 1-\beta}^2$ (the $1 - \beta$ percentile of a chi-squared distribution with D degrees-of-freedom, where D is the input dimensionality and β is a user defined meta-parameter, e.g., 0.1), an update will occur, and posterior probabilities are calculated for each component as follows:

$$p(\mathbf{x}|j) = \frac{1}{(2\pi)^{D/2} \sqrt{|\mathbf{C}_j|}} \exp\left(-\frac{1}{2} d_M^2(\mathbf{x}, j)\right) \tag{2}$$

$$p(j|\mathbf{x}) = \frac{p(\mathbf{x}|j)p(j)}{\sum_{k=1}^K p(\mathbf{x}|k)p(k)} \quad \forall j \tag{3}$$

where K is the number of components. Now, parameters of the algorithm must be updated according to the following equations:

$$v_j(t) = v_j(t - 1) + 1 \tag{4}$$

$$sp_j(t) = sp_j(t - 1) + p(j|\mathbf{x}) \tag{5}$$

$$\mathbf{e}_j = \mathbf{x} - \boldsymbol{\mu}_j(t - 1) \tag{6}$$

$$\omega_j = \frac{p(j|\mathbf{x})}{sp_j} \tag{7}$$

$$\Delta\boldsymbol{\mu}_j = \omega_j\mathbf{e}_j \tag{8}$$

$$\boldsymbol{\mu}_j(t) = \boldsymbol{\mu}_j(t - 1) + \Delta\boldsymbol{\mu}_j \tag{9}$$

$$\mathbf{e}_j^* = \mathbf{x} - \boldsymbol{\mu}_j(t) \tag{10}$$

$$\mathbf{C}_j(t) = (1 - \omega_j)\mathbf{C}_j(t - 1) + \omega_j\mathbf{e}_j^*\mathbf{e}_j^{*T} - \Delta\boldsymbol{\mu}_j\Delta\boldsymbol{\mu}_j^T \tag{11}$$

$$p(j) = \frac{sp_j}{\sum_{q=1}^M sp_q} \tag{12}$$

where sp_j and v_j are the accumulator and the age of component j , respectively, and $p(j)$ is its prior probability. The equations are derived using the Robbins-Monro stochastic approximation [21] for maximizing the likelihood of the model. This derivation can be found in [4, 22].

2.2 Creating New Components

If the update condition in the previous subsection is not met, then a new component j is created and initialized as follows:

$$\boldsymbol{\mu}_j = \mathbf{x}; \quad sp_j = 1; \quad v_j = 1; \quad p(j) = \frac{1}{K}; \quad \mathbf{C}_j = \boldsymbol{\sigma}_{ini}^2 \mathbf{I}$$

$$\sum_{i=1}^K sp_i$$

where K already includes the new component and $\boldsymbol{\sigma}_{ini}$ can be obtained by:

$$\boldsymbol{\sigma}_{ini} = \delta std(\mathbf{x}) \tag{13}$$

where δ is a manually chosen scaling factor (e.g., 0.01) and std is the standard deviation of the dataset. Note that the IGMM is an online and incremental algorithm and therefore it may be the case that we do not have the entire dataset to extract descriptive statistics. In this case

the standard deviation can be just an estimation (e.g., based on sensor limits from a robotic platform), without impacting the algorithm.

2.3 Removing Spurious Components

Optionally, a component j is removed whenever $v_j > v_{min}$ and $sp_j < sp_{min}$, where v_{min} and sp_{min} are manually chosen (e.g., 5.0 and 3.0, respectively). In that case, also, $p(k)$ must be adjusted for all $k \in K, k \neq j$, using Eq (12). In other words, each component is given some time v_{min} to show its importance to the model in the form of an accumulation of its posterior probabilities sp_j . Those components are entirely removed from the model instead of merged with other components, because we assume they represent outliers. Since the removed components have small accumulated activations, it also implies that their removal has almost no negative impact on the model quality, often producing positive impact on generalization performance due to model simplification (a more throughout analysis of parameter sensibility for the IGMN algorithm can be found in [6]).

2.4 Inference

In the IGMN, any element can be predicted by any other element. In other words, inputs and targets are presented together as inputs during training. Thus, inference is done by reconstructing data from the target elements (\mathbf{x}_t , a slice of the entire input vector \mathbf{x}) by estimating the posterior probabilities using only the given elements (\mathbf{x}_p , also a slice of the entire input vector \mathbf{x}), as follows:

$$p(j|\mathbf{x}_i) = \frac{p(\mathbf{x}_i|j)p(j)}{\sum_{q=1}^M p(\mathbf{x}_i|q)p(q)} \quad \forall j \tag{14}$$

It is similar to Eq (3), except that it uses a modified input vector \mathbf{x}_i with the target elements \mathbf{x}_t removed from calculations. After that, \mathbf{x}_t can be reconstructed using the conditional mean equation:

$$\hat{\mathbf{x}}_t = \sum_{j=1}^M p(j|\mathbf{x}_i)(\boldsymbol{\mu}_{j,t} + \mathbf{C}_{j,tt}\mathbf{C}_{j,ii}^{-1}(\mathbf{x}_i - \boldsymbol{\mu}_{j,i})) \tag{15}$$

where $\mathbf{C}_{j,tt}$ is the sub-matrix of the j th component covariance matrix associating the unknown and known parts of the data, $\mathbf{C}_{j,ii}$ is the sub-matrix corresponding to the known part only and $\boldsymbol{\mu}_{j,i}$ is the j th's component mean without the element corresponding to the target element. This division can be seen below:

$$\mathbf{C}_j = \left(\begin{array}{c|c} \mathbf{C}_{j,ii} & \mathbf{C}_{j,it} \\ \hline \mathbf{C}_{j,ti} & \mathbf{C}_{j,tt} \end{array} \right)$$

3 Fast IGMN

One of the contributions of this work lies in the fact that Eq (1) (the squared Mahalanobis distance) requires a matrix inversion, which has a asymptotic time complexity of $O(D^3)$, for D dimensions ($O(D^{\log_2 7 + O(1)})$ for the Strassen algorithm or at best $O(D^{2.3728639})$ with the most recent algorithms to date [23]). This renders the entire IGMN algorithm as impractical for high-dimension tasks. Here we show how to work directly with the inverse of covariance

matrix (also called the precision or concentration matrix) for the entire procedure, therefore avoiding costly inversions.

Firstly, let us denote $\mathbf{C}^{-1} = \mathbf{\Lambda}$, the precision matrix. Our task is to adapt all equations involving \mathbf{C} to instead use $\mathbf{\Lambda}$.

We now proceed to adapt Eq (11) (covariance matrix update). This equation can be seen as a sequence of two rank-one updates to the \mathbf{C} matrix, as follows:

$$\bar{\mathbf{C}}_j(t) = (1 - \omega_j)\mathbf{C}_j(t - 1) + \omega_j\mathbf{e}_j^*\mathbf{e}_j^{*T} \tag{16}$$

$$\mathbf{C}_j(t) = \bar{\mathbf{C}}_j(t) - \Delta\boldsymbol{\mu}_j\Delta\boldsymbol{\mu}_j^T \tag{17}$$

This allows us to apply the Sherman-Morrison formula [24]:

$$(\mathbf{A} + \mathbf{u}\mathbf{v}^T)^{-1} = \mathbf{A}^{-1} - \frac{\mathbf{A}^{-1}\mathbf{u}\mathbf{v}^T\mathbf{A}^{-1}}{1 + \mathbf{v}^T\mathbf{A}^{-1}\mathbf{u}} \tag{18}$$

This formula shows how to update the inverse of a matrix plus a rank-one update. For the second update, which subtracts, the formula becomes:

$$(\mathbf{A} - \mathbf{u}\mathbf{v}^T)^{-1} = \mathbf{A}^{-1} + \frac{\mathbf{A}^{-1}\mathbf{u}\mathbf{v}^T\mathbf{A}^{-1}}{1 - \mathbf{v}^T\mathbf{A}^{-1}\mathbf{u}} \tag{19}$$

In the context of IGMM, we have $\mathbf{A} = (1 - \omega)\mathbf{C}_j(t - 1) = (1 - \omega)\mathbf{\Lambda}_j^{-1}(t - 1)$ and $\mathbf{u} = \mathbf{v} = \sqrt{\omega}\mathbf{e}^*$ for the first update, while for the second one we have $\mathbf{A} = \bar{\mathbf{C}}_j(t)$ and $\mathbf{u} = \mathbf{v} = \Delta\boldsymbol{\mu}_j$. Rewriting Eqs (18) and (19) we get (for the sake of compactness, assume all subscripts for $\mathbf{\Lambda}$ and $\Delta\boldsymbol{\mu}$ to be j):

$$\bar{\mathbf{\Lambda}}(t) = \frac{\mathbf{\Lambda}(t - 1)}{1 - \omega} - \frac{\frac{\omega}{(1-\omega)^2}\mathbf{\Lambda}(t - 1)\mathbf{e}^*\mathbf{e}^{*T}\mathbf{\Lambda}(t - 1)}{1 + \frac{\omega}{1-\omega}\mathbf{e}^{*T}\mathbf{\Lambda}(t - 1)\mathbf{e}^*} \tag{20}$$

$$\mathbf{\Lambda}(t) = \bar{\mathbf{\Lambda}}(t) + \frac{\bar{\mathbf{\Lambda}}(t)\Delta\boldsymbol{\mu}\Delta\boldsymbol{\mu}^T\bar{\mathbf{\Lambda}}(t)}{1 - \Delta\boldsymbol{\mu}^T\bar{\mathbf{\Lambda}}(t)\Delta\boldsymbol{\mu}} \tag{21}$$

These two equations allow us to update the precision matrix directly, eliminating the need for the covariance matrix \mathbf{C} . They have $O(N^2)$ complexity due to matrix-vector products.

Following on the adaptation of the IGMM equations, Eq (1) (the squared Mahalanobis distance) allows for a direct substitution, yielding the following new equation:

$$d_M^2(\mathbf{x}, j) = (\mathbf{x} - \boldsymbol{\mu}_j)^T \mathbf{\Lambda}_j (\mathbf{x} - \boldsymbol{\mu}_j) \tag{22}$$

which now has a $O(N^2)$ complexity, since there is no matrix inversion as the original equation. Note that the Sherman-Morrison identity is exact, thus the Mahalanobis computation yields exactly the same result, as will be shown in the experiments. After removing the cubic complexity from this step, the determinant computation will be dealt with next.

Since the determinant of the inverse of a matrix is simply the inverse of the determinant, it is sufficient to invert the result. But computing the determinant itself is also a $O(D^3)$ operation, so we will instead perform rank-one updates using the Matrix Determinant Lemma [25],

which states the following:

$$|\mathbf{A} + \mathbf{u}\mathbf{v}^T| = |\mathbf{A}|(1 + \mathbf{v}^T \mathbf{A}^{-1} \mathbf{u}) \tag{23}$$

$$|\mathbf{A} - \mathbf{u}\mathbf{v}^T| = |\mathbf{A}|(1 - \mathbf{v}^T \mathbf{A}^{-1} \mathbf{u}) \tag{24}$$

Since the IGMM covariance matrix update involves a rank-two update, adding a term and then subtracting one, both rules must be applied in sequence, similar to what has been done with the Λ equations. Eqs (16) and (17) may be reused here, together with the same substitutions previously showed, leaving us with the following new equations for updating the determinant (again, j subscripts were dropped):

$$|\bar{\mathbf{C}}(t)| = (1 - \omega)^D |\mathbf{C}(t-1)| \left(1 + \frac{\omega}{1 - \omega} \mathbf{e}^{*T} \mathbf{\Lambda}(t-1) \mathbf{e}^*\right) \tag{25}$$

$$|\mathbf{C}(t)| = |\bar{\mathbf{C}}(t)| (1 - \Delta \boldsymbol{\mu}^T \bar{\mathbf{\Lambda}}(t) \Delta \boldsymbol{\mu}) \tag{26}$$

This was the last source of cubic complexity, which is now quadratic.

Finishing the adaptation in the learning part of the algorithm, we just need to define the initialization for Λ for each component. What previously was $\mathbf{C}_j = \sigma_{ini}^2 \mathbf{I}$ now becomes $\Lambda_j = \sigma_{ini}^{-2} \mathbf{I}$, the inverse of the variances of the dataset. Since this matrix is diagonal, there are no costly inversions involved. And for initializing the determinant $|\mathbf{C}|$, just set it to $\prod \sigma_{ini}^2$, which again takes advantage of the initial diagonal matrix to avoid costly operations. Note that we keep the precision matrix Λ , but the determinant of the covariance matrix \mathbf{C} instead. See algorithms 1 to 3 for a summary of the new learning algorithm (excluding pruning, for brevity).

Algorithm 1 Fast IGMM Learning

```

Input:  $\delta, \beta, \mathbf{X}$ 
 $K > 0, \sigma_{ini}^{-1} = (\delta std(\mathbf{X}))^{-1}, M = \emptyset$ 
for all input data vector  $\mathbf{x} \in \mathbf{X}$  do
  if  $K = 0$  or  $\exists j, d_M^2(\mathbf{x}, j) < \chi_{D, 1-\beta}^2$  then
    update  $(\mathbf{x})$ 
  else
     $M \leftarrow MU\ create(\mathbf{x})$ 
  end if
end for

```

Finally, the inference Eq (15) must also be updated in order to allow the IGMM to work in supervised mode. This can be accomplished by the use of a block matrix decomposition (note that here \mathbf{C} is just another sub-matrix, not the covariance matrix as used before):

$$\Lambda = \begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{bmatrix}^{-1} = \begin{bmatrix} \mathbf{X} & \mathbf{Y} \\ \mathbf{Z} & \mathbf{W} \end{bmatrix} = \begin{bmatrix} (\mathbf{A} - \mathbf{B}\mathbf{D}^{-1}\mathbf{C})^{-1} & -\mathbf{A}^{-1}\mathbf{B}(\mathbf{D} - \mathbf{C}\mathbf{A}^{-1}\mathbf{B})^{-1} \\ -\mathbf{D}^{-1}\mathbf{C}(\mathbf{A} - \mathbf{B}\mathbf{D}^{-1}\mathbf{C})^{-1} & (\mathbf{D} - \mathbf{C}\mathbf{A}^{-1}\mathbf{B})^{-1} \end{bmatrix}$$

Here, according to Eq (15), we need \mathbf{C} and \mathbf{A}^{-1} . But since the terms that constitute these sub-matrices are relative to the original covariance matrix (which we do not have), they must be extracted from the precision matrix directly. Looking at the decomposition, it is clear that $\mathbf{Y}\mathbf{W}^{-1} = -\mathbf{A}^{-1}\mathbf{B} = -\mathbf{C}\mathbf{A}^{-1}$ (the terms between parenthesis in \mathbf{Y} and \mathbf{W} cancel each other, while

$\mathbf{B} = \mathbf{C}^T$ due to symmetry). So Eq (15) can be rewritten as:

$$\hat{\mathbf{x}}_i = \sum_{j=1}^M p(j|\mathbf{x}_i)(\boldsymbol{\mu}_{j,t} - \mathbf{Y}\mathbf{W}^{-1}(\mathbf{x}_i - \boldsymbol{\mu}_{j,i})) \tag{27}$$

where \mathbf{Y} and \mathbf{W} can be extracted directly from $\boldsymbol{\Lambda}$. However, we still need to compute the inverse of \mathbf{W} . So we can say that this particular implementation has $O(NKD^2)$ complexity for learning and $O(NKD^3)$ for inference. The reason for us to not worry about that is that $d = i + o$, where i is the number of inputs and o is the number of outputs. The inverse computation acts only upon the output portion of the matrix. Since, in general, $o \ll i$ (in many cases even $o = 1$), the impact is minimal, and the same applies to the $\mathbf{Y}\mathbf{W}^{-1}$ product. In fact, Weka (the data mining platform used in this work [26]) allows for only 1 output, leaving us with just scalar operations.

Algorithm 2 update

```

Input:  $\mathbf{x}$ 
for all Gaussian components  $j \in M$  do
     $d_M^2(\mathbf{x}, j) = (\mathbf{x} - \boldsymbol{\mu}_j)^T \boldsymbol{\Lambda}_j (\mathbf{x} - \boldsymbol{\mu}_j)$ 
     $p(\mathbf{x} | j) = \frac{1}{(2\pi)^{D/2} \sqrt{|\mathbf{C}_j|}} \exp(-\frac{1}{2} d_M^2(\mathbf{x}, j))$ 
     $p(j | \mathbf{x}) = \frac{p(\mathbf{x}|j)p(j)}{\sum_{k=1}^K p(\mathbf{x}|k)p(k)} \quad \forall j$ 
     $v_j(t) = v_j(t-1) + 1$ 
     $sp_j(t) = sp_j(t-1) + p(j|\mathbf{x})$ 
     $\mathbf{e}_j = \mathbf{x} - \boldsymbol{\mu}_j(t-1)$ 
     $\omega_j = \frac{p(j|\mathbf{x})}{sp_j}$ 
     $\Delta \boldsymbol{\mu}_j = \omega_j \mathbf{e}_j$ 
     $\boldsymbol{\mu}_j(t) = \boldsymbol{\mu}_j(t-1) + \Delta \boldsymbol{\mu}_j$ 
     $\mathbf{e}_j^* = \mathbf{x} - \boldsymbol{\mu}_j(t)$ 
     $\bar{\boldsymbol{\Lambda}}(t) = \frac{\boldsymbol{\Lambda}(t-1)}{1-\omega} - \frac{\frac{\omega}{(1-\omega)^2} \boldsymbol{\Lambda}(t-1) \mathbf{e}^* \mathbf{e}^{*T} \boldsymbol{\Lambda}(t-1)}{1 + \frac{\omega}{1-\omega} \mathbf{e}^{*T} \boldsymbol{\Lambda}(t-1) \mathbf{e}^*}$ 
     $\boldsymbol{\Lambda}(t) = \bar{\boldsymbol{\Lambda}}(t) + \frac{\bar{\boldsymbol{\Lambda}}(t) \Delta \boldsymbol{\mu} \Delta \boldsymbol{\mu}^T \bar{\boldsymbol{\Lambda}}(t)}{1 - \Delta \boldsymbol{\mu}^T \bar{\boldsymbol{\Lambda}}(t) \Delta \boldsymbol{\mu}}$ 
     $p(j) = \frac{sp_j}{\sum_{q=1}^M sp_q}$ 
     $|\bar{\mathbf{C}}(t)| = (1 - \omega)^D |\mathbf{C}(t-1)| (1 + \frac{\omega}{1-\omega} \mathbf{e}^{*T} \boldsymbol{\Lambda}(t-1) \mathbf{e}^*)$ 
     $|\mathbf{C}(t)| = |\bar{\mathbf{C}}(t)| (1 - \Delta \boldsymbol{\mu}^T \bar{\boldsymbol{\Lambda}}(t) \Delta \boldsymbol{\mu})$ 
end for

```

Algorithm 3 create

```

Input:  $\mathbf{x}$ 
     $K \leftarrow K + 1$ 
    return new Gaussian component  $K$  with  $\boldsymbol{\mu}_K = \mathbf{x}$ ,  $\boldsymbol{\Lambda}_K = \boldsymbol{\sigma}_{ini}^{-1} \mathbf{I}$ ,  $|\mathbf{C}_K| = |\boldsymbol{\Lambda}_K|^{-1}$ ,  $sp_j = 1$ ,  $v_j = 1$ ,  $p(j) = \frac{1}{\sum_{k=1}^K sp_k}$ 

```

4 Experiments

The first experiment was meant to verify that both IGMM implementations produce exactly the same results. They were both applied to 7 standard datasets distributed with the Weka software (Table 1). Parameters were set to $\delta = 0.5$ (chosen by 2-fold cross-validation) and $\beta = 4.9E - 324$, the smallest possible double precision number available for the Java Virtual Machine (and also the default value for this implementation of the algorithm), such that Gaussian components are created only when strictly necessary. The same parameters were used for all

Table 1. Datasets.

Dataset	Instances (N)	Attributes (D)	Classes
breast-cancer	286	9	2
pima-diabetes	768	8	2
Glass	214	9	7
ionosphere	351	34	2
iris	150	4	3
labor-neg-data	57	16	2
soybean	683	35	19
MNIST [28]	70000	784	10
CIFAR-10 [29]	60000	3072	10

doi:10.1371/journal.pone.0139931.t001

Table 2. Accuracy of different algorithms on standard datasets.

Dataset	RF	NN	Lin. SVM	RBF SVM	IGMN	FIGMN
breast-cancer	69.6 ± 9.1	75.2 ± 6.5	69.3 ± 7.5	70.6 ± 1.5	71.4 ± 7.4	71.4 ± 7.4
pima-diabetes	75.8 ± 3.5	74.2 ± 4.9	77.5 ± 4.4	65.1 ± 0.4•	73.0 ± 4.5	73.0 ± 4.5
Glass	79.9 ± 5.0	53.8 ± 7.4•	62.7 ± 7.8•	68.8 ± 8.7•	65.4 ± 4.9•	65.4 ± 4.9•
ionosphere	92.9 ± 3.6	92.6 ± 2.4	88.0 ± 3.5	93.5 ± 3.0	92.6 ± 3.8	92.6 ± 3.8
iris	95.3 ± 4.5	95.3 ± 5.5	96.7 ± 4.7	96.7 ± 3.5	97.3 ± 3.4	97.3 ± 3.4
labor-neg-data	89.7 ± 14.3	89.7 ± 14.3	93.3 ± 11.7	93.3 ± 8.6	94.7 ± 8.6	94.7 ± 8.6
soybean	93.0 ± 3.1	93.0 ± 2.4	94.0 ± 2.2	88.7 ± 3.0•	91.5 ± 5.4	91.5 ± 5.4
Average	85.2	82.0	83.1	82.4	83.7	83.7

• statistically significant degradation

doi:10.1371/journal.pone.0139931.t002

datasets. Results were obtained from 10-fold cross-validation (resulting in training sets with 90% of the data and test sets with the remaining 10%) and statistical significances came from paired t-tests with $p = 0.05$. As can be seen in Table 2, both IGMN and FIGMN algorithms produced exactly the same results, confirming our expectations. The number of clusters created by them was also the same, and the exact quantity for each dataset is shown in Table 3. The Weka packages for both variations of the IGMN algorithm, as well as the datasets used in the experiments can be found at [27]. The MNIST dataset can be found at <http://yann.lecun.com/exdb/mnist/>, while the CIFAR10 dataset is available at <http://www.cs.toronto.edu/~kriz/cifar.html>.

Table 3. Number of Gaussian components created.

Dataset	# of Components
breast-cancer	14.2 ± 1.9
pima-diabetes	19.4 ± 1.3
Glass	15.9 ± 1.1
ionosphere	74.4 ± 1.4
iris	2.7 ± 0.7
labor-neg-data	12.0 ± 1.2
soybean	42.6 ± 2.2

doi:10.1371/journal.pone.0139931.t003

Table 4. Training and testing running times (in seconds).

Dataset	IGMN Training	FIGMN Training	IGMN Testing	FIGMN Testing
MNIST	32,544.69	1,629.81	3,836.06	230.92
CIFAR-10	2,758,252*	15,545.05	-	795.98

* estimated time projected from 100 data points

doi:10.1371/journal.pone.0139931.t004

Besides the confirmation we wanted, we could also compare the IGMN/FIGMN classification accuracy for the referred datasets against other 4 algorithms: Random Forest (RF), Neural Network (NN), Linear SVM and RBF SVM. The neural network is a parallel implementation of a state-of-the-art Dropout Neural Network [30] with 100 hidden neurons, 50% dropout for the hidden layer and 20% dropout for the input layer (this specific implementation can be found at <https://github.com/amenten/NeuralNetwork>). The 4 algorithms were kept with their default parameters. The IGMN algorithms produced competitive results, with just one of them (Glass) being statistically significant below the accuracy produced by the Random Forest algorithm. This value was significantly inferior for all other algorithms too. On average, the IGMN algorithms were the second best from the set, losing only to the Random Forest. Note, however, that the Random Forest is a batch algorithm, while the IGMN learns incrementally from each data point. Also, the resulting Random Forest model used 6 times more memory than the IGMN model. We also tested the FIGMN accuracy on the MNIST dataset, but even after parameter tuning, the results were not on par with the state-of-the-art (above 99%), reaching a maximum of around 93% accuracy.

A second experiment was performed in order to evaluate the speed performance of the proposed algorithm, both the original and improved IGMN algorithms, using the parameters $\delta = 1$ and $\beta = 0$, such that a single component was created and we could focus on speedups due only to dimensionality (this also made the algorithm highly insensitive to the δ parameter). They were applied to the 2 highest dimensional datasets in Table 1, namely, the MNIST and CIFAR-10 datasets. The MNIST dataset was split into a training set with 60000 data points and a testing set containing 10000 data points, the standard procedure in the machine learning community [28]. Similarly, the CIFAR-10 dataset was split into 50000 training data points and 10000 testing data points, also a standard procedure for this dataset [29].

Results can be seen in Table 4. Training time for the MNIST dataset was 20 times smaller for the fast version while the testing time was 16 times smaller. It makes sense that the testing time has shown a bit less improvement, since inference only takes advantage from the incremental determinant computation but not from the incremental inverse computation. For the CIFAR-10 dataset, it was impractical to run the original IGMN algorithm on the entire dataset, requiring us to estimate the total time, linearly projecting it from 100 data points (note that, since the model always uses only 1 Gaussian component during the entire training, the computation time per data point does not increase over time). It resulted in 32 days of CPU time estimated for the original algorithm against 15545s ($\sim 4h$) for the improved algorithm, a speedup above 2 orders of magnitude. Testing time is not available for the original algorithm on this dataset, since the training could not be concluded. Additionally, we compared a pure clustering version of the FIGMN algorithm on the MNIST training set against batch EM (the implementation found in the Weka software). While the FIGMN algorithm took $\sim 7.5h$ hours to finish,

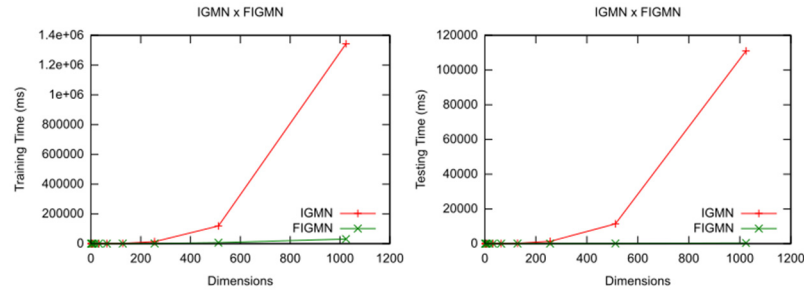


Fig 1. Training and testing times for both versions of the IGMN algorithm with growing number of dimensions.

doi:10.1371/journal.pone.0139931.g001

using 208 Gaussian components, the batch EM algorithm took $\sim 1.3h$ to complete a *single iteration* (we set the fixed number of components to 208 too) using 4 CPU cores. Besides generally requiring more than one iteration to achieve best results, the batch algorithm required the entire dataset in RAM. The FIGMN memory requirements were much lower.

Finally, both versions of the IGMN algorithm with $\delta = 1$ and $\beta = 0$ were compared on 11 synthetic datasets generated by Weka. All datasets have 1000 data points drawn from a single Gaussian distribution (90% training, 10% testing) and an exponentially growing number of dimensions: 1, 2, 4, 8, 16, 32, 64, 128, 256, 512 and 1024. This experiment was performed in order to compare the scalability of both algorithms. Results for training and testing can be seen in [Fig 1](#):

5 Conclusion

We have shown how to work directly with precision matrices in the IGMN algorithm, avoiding costly matrix inversions by performing rank-one updates. The determinant computations were also avoided using a similar method, effectively eliminating any source of cubic complexity for the learning algorithm. This resulted in substantial speedups for high-dimensional datasets, turning the IGMN into a good option for this kind of tasks. The inference operation still has cubic complexity, but we argue that it has a much smaller impact on the total runtime of the algorithm, since the number of outputs is usually much smaller than the number of inputs. This was confirmed in the experiments.

In general, we could see that the fast IGMN is a good option for supervised learning, with low runtimes and good accuracy. It should be noted that this is achieved with a single-pass through the data, making it also a valid option for data streams.

Author Contributions

Conceived and designed the experiments: RCP PME. Performed the experiments: RCP. Analyzed the data: RCP PME. Wrote the paper: RCP PME.

References

1. Heinen MR, Engel PM, Pinto RC. Using a Gaussian mixture neural network for incremental learning and robotics. In: Neural Networks (IJCNN), The 2012 International Joint Conference on. IEEE; 2012. p. 1–8.

2. Heinen MR, Engel PM, Pinto RC. IGMM: An incremental gaussian mixture network that learns instantaneously from data flows. Proc VIII Encontro Nacional de Inteligência Artificial (ENIA2011). 2011;.
3. Dempster AP, Laird NM, Rubin DB, et al. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society Series B (Methodological)*. 1977; 39(1):1–38.
4. Engel P, Heinen M. Incremental learning of multivariate gaussian mixture models. *Advances in Artificial Intelligence—SBIA 2010*. 2011;p. 82–91.
5. Grossberg S. Competitive learning: From interactive activation to adaptive resonance. *Cognitive science*. 1987; 11(1):23–63. doi: [10.1111/j.1551-6708.1987.tb00862.x](https://doi.org/10.1111/j.1551-6708.1987.tb00862.x)
6. Heinen MR. A connectionist approach for incremental function approximation and on-line tasks. Universidade Federal do Rio Grande do Sul; 2011.
7. Rumelhart DE, McClelland JL. *Parallel distributed processing*. MIT Pr.; 1986.
8. Ghahramani Z, Jordan MI. Supervised learning from incomplete data via an EM approach. In: *Advances in neural information processing systems 6*. Citeseer; 1994.
9. Damas B, Santos-Victor J. An online algorithm for simultaneously learning forward and inverse kinematics. In: *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*. IEEE; 2012. p. 1499–1506.
10. Heinen MR, Engel PM. IGMM: An incremental connectionist approach for concept formation, reinforcement learning and robotics. *Journal of Applied Computing Research*. 2011; 1(1):2–19. doi: [10.4013/jacr.2011.11.01](https://doi.org/10.4013/jacr.2011.11.01)
11. Pinto RC, Engel PM, Heinen MR. Echo State Incremental Gaussian Mixture Network for Spatio-Temporal Pattern Processing. In: *Proceedings of the IX ENIA-Brazilian Meeting on Artificial Intelligence, Natal (RN)*; 2011.
12. Pinto RC, Engel PM, Heinen MR. Recursive incremental gaussian mixture network for spatio-temporal pattern processing. Proc 10th Brazilian Congr Computational Intelligence (CBIC) Fortaleza, CE, Brazil (Nov 2011);.
13. Flores JHF, Engel PM, Pinto RC. Autocorrelation and partial autocorrelation functions to improve neural networks models on univariate time series forecasting. In: *Neural Networks (IJCNN), The 2012 International Joint Conference on*. IEEE; 2012. p. 1–8.
14. Heinen MR, Bazzan AL, Engel PM. Dealing with continuous-state reinforcement learning for intelligent control of traffic signals. In: *Intelligent Transportation Systems (ITSC), 2011 14th International IEEE Conference on*. IEEE; 2011. p. 890–895.
15. Pinto RC, Engel PM, Heinen MR. One-shot learning in the road sign problem. In: *Neural Networks (IJCNN), The 2012 International Joint Conference on*. IEEE; 2012. p. 1–6.
16. de Pontes Pereira R, Engel PM, Pinto RC. Learning Abstract Behaviors with the Hierarchical Incremental Gaussian Mixture Network. In: *Neural Networks (SBRN), 2012 Brazilian Symposium on*. IEEE; 2012. p. 131–135.
17. Santos ADPd, Wives LK, Alvares LO. Location-Based Events Detection on Micro-Blogs. arXiv preprint arXiv:12104008. 2012;.
18. Salmen J, Schlipfing M, Igel C. Efficient update of the covariance matrix inverse in iterated linear discriminant analysis. *Pattern Recognition Letters*. 2010; 31(13):1903–1907. Meta-heuristic Intelligence Based Image Processing. doi: [10.1016/j.patrec.2010.03.001](https://doi.org/10.1016/j.patrec.2010.03.001)
19. Lefakis L, Fleuret F. Jointly Informative Feature Selection. In: *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics*; 2014. p. 567–575.
20. Olsen PA, Gopinath RA. Extended mllt for gaussian mixture models. *Transactions in Speech and Audio Processing*. 2001;.
21. Robbins H, Monro S. A stochastic approximation method. *The annals of mathematical statistics*. 1951; p. 400–407. doi: [10.1214/aoms/1177729586](https://doi.org/10.1214/aoms/1177729586)
22. Engel PM. INBC: An incremental algorithm for dataflow segmentation based on a probabilistic approach. INBC: an incremental algorithm for dataflow segmentation based on a probabilistic approach. 2009;.
23. Gall FL. Powers of tensors and fast matrix multiplication. arXiv preprint arXiv:14017714. 2014;.
24. Sherman J, Morrison WJ. Adjustment of an Inverse Matrix Corresponding to a Change in One Element of a Given Matrix. *The Annals of Mathematical Statistics*. 1950 03; 21(1):124–127. Available from: <http://dx.doi.org/10.1214/aoms/1177729893> doi: [10.1214/aoms/1177729893](https://doi.org/10.1214/aoms/1177729893)
25. Harville DA. *Matrix algebra from a statistician's perspective*. Springer; 2008.
26. Hall M, Frank E, Holmes G, Pfahringer B, Reutemann P, Witten IH. The WEKA data mining software: an update. *ACM SIGKDD explorations newsletter*. 2009; 11(1):10–18. doi: [10.1145/1656274.1656278](https://doi.org/10.1145/1656274.1656278)

27. Pinto, RC. Experiment Data for "A Fast Incremental Gaussian Mixture Model". <http://dx.doi.org/10.6084/m9.figshare.1552030>
28. LeCun Y, Bottou L, Bengio Y, Haffner P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*. 1998; 86(11):2278–2324. doi: [10.1109/5.726791](https://doi.org/10.1109/5.726791)
29. Krizhevsky A, Hinton G. Learning multiple layers of features from tiny images. Computer Science Department, University of Toronto, Tech Rep. 2009;.
30. Hinton GE, Srivastava N, Krizhevsky A, Sutskever I, Salakhutdinov RR. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv preprint arXiv:12070580*. 2012;.