Optimized realization of Bayesian networks in reduced normal form using latent variable model

Giovanni Di Gennaro² · Amedeo Buonanno^{1,2} · Francesco A. N. Palmieri²

Accepted: 29 January 2021 / Published online: 17 March 2021
© The Author(s) 2021

Abstract

Bayesian networks in their Factor Graph Reduced Normal Form are a powerful paradigm for implementing inference graphs. Unfortunately, the computational and memory costs of these networks may be considerable even for relatively small networks, and this is one of the main reasons why these structures have often been underused in practice. In this work, through a detailed algorithmic and structural analysis, various solutions for cost reduction are proposed. Moreover, an online version of the classic batch learning algorithm is also analysed, showing very similar results in an unsupervised context but with much better performance; which may be essential if multi-level structures are to be built. The solutions proposed, together with the possible online learning algorithm, are included in a C++ library that is quite efficient, especially if compared to the direct use of the well-known sum-product and Maximum Likelihood algorithms. The results obtained are discussed with particular reference to a Latent Variable Model structure.

Keywords Bayesian networks · Belief propagation · Factor graphs · Latent variable · Optimization

1 Introduction

The Factor Graph (FG) representation, and in particular the so-called Normal Form (FGn) (Forney 2001; Loeliger 2004), is a very appealing formulation to visualize and manipulate Bayesian graphs; representing their relative joint probability by assigning variables to arcs and functions (or factors) to nodes.

Furthermore, in the Factor Graph in Reduced Normal Form (FGrn), through the use of replicator units (or equal constraints), the graph is reduced to an architecture in which each variable is connected to two factors at most (Palmieri 2016); with belief messages that flow bidirectionally into the network. This paradigm has demonstrated its extensive

modularity and flexibility in managing variables of different types and cardinalities (Palmieri and Buonanno 2015), and can also be used to build multi-layer network (Palmieri and Buonanno 2014; Buonanno and Palmieri 2015b,c). For this reason, in a previous work a Simulink¹ library for the rapid prototyping of an FGn network was already implemented (Buonanno and Palmieri 2015a), but because of the limitations imposed by Simulink, it is not particularly suitable for treating large amounts of data and/or complex architectures with many variables.

Truthfully, despite the large presence of such arguments in the literature (Koller and Friedman 2009; Barber 2012; Murphy 2012), this type of structure always suffers from high computational and memory costs, due to the lack of attention given to the specific algorithmic implementation. Therefore, this work aims to improve complexity efficiency in both inference and learning (overcoming software limitations), and all the solutions obtained have been included in a C++ library (<https://github.com/mlunicampania/FGnLib>). The various problems, related to the propagation and learning of probabilities within the FGn paradigm, are addressed by focusing on the implementation of the Latent Variable Model

✉ Giovanni Di Gennaro
giovanni.digennaro@unicampania.it

Amedeo Buonanno
amedeo.buonanno@enea.it

Francesco A. N. Palmieri
francesco.palmieri@unicampania.it

¹ Department of Energy Technologies and Renewable Energy Sources, ENEA, P. Enrico Fermi 1, Portici, NA, Italy

² Dipartimento di Ingegneria, Università degli Studi della Campania “Luigi Vanvitelli”, via Roma 29, Aversa, CE, Italy

¹ Simulink® is a MATLAB-based graphical programming developed by MathWorks® (more information is available at <https://www.mathworks.com/products/simulink.html>).

Fig. 1 FGm components: **a** a Variable with the associated forward and backward messages; **b** a Diverter representing the replication of M variables; **c** a SISO block with the conditional distribution matrix of the connected variables; **d** a Source with the prior distribution

(LVM) (Bishop 1999; Murphy 2012), also called Autoclass (Cheeseman and Stutz 1996). LVMs can be used in a large number of applications and can be seen as a basic building block for more complex architectures.

After a brief introduction of the FGm paradigm in Sect. 2 and the LVM in Sect. 3, necessary to provide the fundamental elements for subsequent discussions, the C++ library project is described in Sect. 4. In Sect. 5 a detailed analysis of the computational complexity of the various algorithmic elements is presented for each building block. In Sect. 6 some simulation results that verify how the proposed algorithms produce indisputable advantages are presented. Finally, in Sect. 7, an incremental learning algorithm is introduced by modifying the ML recursions, that not only presents a significant decrease in terms of memory costs (allowing learning even in the presence of huge datasets), but shows in some cases better performance in avoiding local minima traps.

2 Factor graphs in reduced normal form

A FGm requires only the combination of the elements shown in Figure 1:

- (a) The *Variable* V , which can take a single discrete value $v \in \mathcal{V} = \{v_1, \dots, v_{|\mathcal{V}|}\}$, is represented as an oriented edge with two discrete messages,² proportional ($\propto$) to the discrete distributions, that travel in both directions. Depending on the direction assigned to the variable, the two messages are respectively called *forward* $f_V(v)$ and *backward* $b_V(v)$, and can be also represented as $|\mathcal{V}|$ -dimensional column vectors: $\mathbf{f}_V$ and $\mathbf{b}_V$. Note that the

marginal distribution $p_V(v)$, which is proportional to the posterior given the observations anywhere else in the network, is proportional to the product

$$p_V(v) \propto f_V(v)b_V(v),$$

or in vector form

$$\mathbf{p}_V \propto \mathbf{f}_V \odot \mathbf{b}_V, \quad (1)$$

where $\odot$ denotes the Hadamard (element-by-element) product.

- (b) The *Replicator Block* (Diverter) represents the equality constraint of the connected variables. The constraint of equality between the variables is obtained by making sure that each replica can carry different forward and backward messages. A replicator acts like a bus with messages combined and diverted towards the connected branches. The combination rule (product rule) is such that outgoing messages are the product of all the incoming ones, except for the one that belongs to the same variable

$$b_{V(i)}(v) \propto \prod_{\substack{j=1 \\ j \neq i}}^m f_{V(j)}(v) \prod_{k=m+1}^M b_{V(k)}(v), \quad i = 1 : m$$

$$f_{V(i)}(v) \propto \prod_{j=1}^m f_{V(j)}(v) \prod_{\substack{k=m+1 \\ k \neq i}}^M b_{V(k)}(v), \quad i = m+1 : M,$$

or in vector form

$$\mathbf{b}_{V(i)} \propto \left(\bigodot_{\substack{j=1 \\ j \neq i}}^m \mathbf{f}_{V(j)} \right) \left(\bigodot_{k=m+1}^M \mathbf{b}_{V(k)} \right), \quad i = 1 : m$$

$$\mathbf{f}_{V(i)} \propto \left(\bigodot_{j=1}^m \mathbf{f}_{V(j)} \right) \left(\bigodot_{\substack{k=m+1 \\ k \neq i}}^M \mathbf{b}_{V(k)} \right), \quad i = m+1 : M \quad (2)$$

- (c) The *SISO block*, that is the core of the FGm paradigm, represents the conditional probability matrix $P(Y|V)$ of Y given V . Assuming that the output variable Y takes values in the alphabet $\mathcal{Y} = \{y_1, \dots, y_{|\mathcal{Y}|}\}$, this probability matrix is a $|\mathcal{V}| \times |\mathcal{Y}|$ row-stochastic matrix

$$P(Y|V) = [Pr\{Y = y_j | V = v_i\}]_{j=1:|\mathcal{Y}|}^{i=1:|\mathcal{V}|} = [\theta_{ij}]_{j=1:|\mathcal{Y}|}^{i=1:|\mathcal{V}|},$$

² Since the variable can only assume discrete values within the vocabulary, the two messages will be vectors of dimensions equal to the size of the vocabulary.

or more explicitly

$$P(Y|V) = \begin{bmatrix} P(Y = y_1|V = v_1) & \cdots & P(Y = y_{|\mathcal{Y}|}|V = v_1) \\ P(Y = y_1|V = v_2) & \cdots & P(Y = y_{|\mathcal{Y}|}|V = v_2) \\ \vdots & \ddots & \vdots \\ P(Y = y_1|V = v_{|\mathcal{V}|}) & \cdots & P(Y = y_{|\mathcal{Y}|}|V = v_{|\mathcal{V}|}) \end{bmatrix}$$

$$= \begin{bmatrix} \theta_{11} & \cdots & \theta_{1|\mathcal{Y}|} \\ \theta_{21} & \cdots & \theta_{2|\mathcal{Y}|} \\ \vdots & \ddots & \vdots \\ \theta_{|\mathcal{V}|1} & \cdots & \theta_{|\mathcal{V}||\mathcal{Y}|} \end{bmatrix}.$$

Outgoing messages are

$$f_Y(y_i) \propto \sum_{j=1}^{|\mathcal{V}|} \theta_{ij} f_V(v_j) \quad \text{and} \quad b_V(v_j) \propto \sum_{i=1}^{|\mathcal{Y}|} \theta_{ij} b_Y(y_i),$$

or in vector form

$$\mathbf{f}_Y \propto P(Y|V)^T \mathbf{f}_V \quad \text{and} \quad \mathbf{b}_V \propto P(Y|V) \mathbf{b}_Y. \quad (3)$$

- (d) The *Source block* defines an independent $|\mathcal{V}|$ -dimensional source variable V with its prior distribution π_V . Therefore, the outgoing message is

$$f_V(v_i) = \pi_V(v_i), \quad i = 1 : |\mathcal{V}|,$$

or in vector form

$$\mathbf{f}_V = \boldsymbol{\pi}_V.$$

It should be emphasized that the rules presented are a rigorous translation of the total probability theorem and Bayes' rule.

Note that the only parameters that need to be learned during the training phase are the matrices inside the SISO blocks and the priors inside the Sources. Although different variations are possible (Koller and Friedman 2009; Barber 2012; Palmieri 2016), training algorithms are derived mainly as maximization of the likelihood on the observed variables that can be anywhere in the network. Furthermore, within the FGn paradigm, learning takes place locally; that is, the parameters inside the SISO blocks and the Sources can be learned using only the backward and forward messages to that particular element. This also means that parameter learning in this representation can be addressed in a unified way because we can use a single rule to train any SISO or Source block in the system, simultaneously and independently of its position. Therefore, learning is done iterating over three simple steps:

1. present the observations to the network in the form of distributions; they can be anywhere in the network as backward or forward messages;

2. propagate the messages in the whole network, in accordance with the mathematical rules just described;
3. perform the update of SISO blocks and Sources using incoming messages.

The prior of a source is learned simply by calculating the new marginal probability (using Equation 1), due to the changes of the backward message to the Source. On the other hand, learning the matrices inside the SISO blocks is more complex. According to our experience (Palmieri 2016), the best algorithm for learning the conditional probability matrices inside the SISO blocks is the Maximum Likelihood (ML) algorithm, which uses the following update

$$\theta_{lm}^{(1)} \leftarrow \frac{\theta_{lm}^{(0)}}{\sum_{n=1}^N f_{V[n]}(l)} \sum_{n=1}^N \frac{f_{V[n]}(l) b_{Y[n]}(m)}{\mathbf{f}_{V[n]}^T \boldsymbol{\theta}^{(0)} \mathbf{b}_{Y[n]}}. \quad (4)$$

Equation 4 represents the heart of the ML algorithm and usually requires multiple cycles in order to achieve convergence. However, changing the matrices also changes the propagated messages. For this reason, the whole learning process (starting from the presentation of the evidence) also needs to be performed several times; namely for a fixed number of *epochs* (an hyperparameter).

3 Latent variable models

Factor Graphs and in particular FGn can be used to represent a joint probability distribution by appropriately using the dependencies/independence between the variables. In many cases, however, the probabilistic description of an event can be further simplified using a set of unobserved variables, which represent the unknown “causes” generating the event. In these models the variables are typically divided into *Visible* and *Hidden*: the first relating to the inputs observed during the inference, the latter belonging to the *internal representation*. Obviously, the main advantage of this type of model lies in the fact that it has fewer parameters than the complete model, producing a sort of compressed representation of the observed data (which are therefore easier to manage and analyse).

The simplest and most complete model that includes both visible and hidden variables is the bipartite graph of Fig. 2, named Latent Variable Model (LVM) (Murphy 2012; Bishop 1999) or Autoclass (Cheeseman and Stutz 1996); any other hidden variables model can ultimately be reduced to such a model. Within the figure, it is possible to distinguish H latent variables, $S_1, \dots, S_H$, and N observed variables, $Y_1, \dots, Y_N$; where typically $N \gg H$.

It should be noted that although the given nomenclature seems to subdivide the variables according to their position,

Fig. 2 The general structure for Latent Variable Models

Fig. 3 The general structure for the Latent Variable Models in Reduced Normal Form

where the variables below are known while those above need to be estimated, the bidirectional structure of the network remains quite general, including cases in which some of the above variables may be known and some of the bottom variables need to be estimated.

The FGn of a Bayesian network represented by the general LVM structure of Fig. 2 is shown in Fig. 3. The system is intrinsically represented as a generative model; that is, choosing to direct the variables downwards and positioning the marginally independent sources $S_1, \dots, S_H$ at the top. Moreover, it should be noted that the supposed independence of all known variables given the hidden variables (provided by the latent variable model) allows to greatly simplify the analysis of the total joint probability, which can now be represented by the factorization

$$P(Y_1, \dots, Y_N, S_1, \dots, S_H) \\ = P(Y_1|S_1, \dots, S_H) \cdots P(Y_N|S_1, \dots, S_H) P(S_1) \cdots P(S_H).$$

As said, the structure allows to manage totally heterogeneous variables, but it is good to clarify that (being a representation of a Bayesian network) the variables must be presented as probability vectors; that is, through vectors containing the “degree of similarity” of the variable to each element of its discrete alphabet.

The source variables, which have prior distributions $\pi_{S_1}, \dots, \pi_{S_H}$, are mapped to the product space $\mathcal{P}$, of dimensions $|\mathcal{P}| = |S_1| \times \cdots \times |S_H|$, via the fixed row-stochastic matrices (shaded blocks in Fig. 3)

$$P((S_1 S_2 \dots S_H)^{(1)} | S_1) = \frac{|S_1|}{\prod_{i=1}^H |S_i|} I_{|S_1|} \otimes 1_{|S_2|}^T \otimes \cdots \otimes 1_{|S_H|}^T, \\ \vdots \\ P((S_1 S_2 \dots S_H)^{(H)} | S_H) = \frac{|S_H|}{\prod_{i=1}^H |S_i|} 1_{|S_1|}^T \otimes \cdots \otimes 1_{|S_{H-1}|}^T \otimes I_{|S_H|}, \quad (5)$$

where $\otimes$ denotes the Kronecker product, 1_K is a K -dimensional column vector with all ones, and I_K is the $K \times K$ identity matrix (Palmieri 2016).

The conditional probability matrix is such that each variable contributes to the product space with its value, and it is uniform on the components that compete to the other source variables. This is the FGn counterpart of the Junction Tree reduction procedure because it is equivalent to “marry the parents” in Bayesian Graphs (Koller and Friedman 2009), but here there are explicit branches for the product space variable. For this reason, although the messages traveling bi-directionally and the initial Bayesian network present many loops (Fig. 2), the FGn architecture will not show any convergence problem because the LVM has been reduced to a tree.

Finally, the j -th SISO block at the bottom of Fig. 3, with $j = 1, \dots, N$, represents the conditional probability matrices $P(Y_j|S_1 S_2 \dots S_H)$, which then will have dimensions $|\mathcal{P}| \times |\mathcal{Y}_j|$.

3.1 The LVM with one hidden variable

When $H > 1$ we have a Many-To-Many LVM model, which was already discussed in a previous work (Palmieri and Buonanno 2015); calling it Discrete Independent Component Analysis (DICA) because it uses the same generative model of the Independent Component Analysis (ICA), but on discrete variables.

Vice versa, when $H = 1$, we obtain a One-To-Many LVM model, in which there is just one latent factor (parent) that conditions $Y_1, \dots, Y_N$ (children) and where obviously fixed matrices (previously represented by the shaded blocks) are no longer necessary. Although the general paradigm has the great advantage of allowing the presence of Sources of dif-

Fig. 4 A N -tuple with the Latent Variable S and a Class Variable L drawn as **a** Bayesian Network, **b** Factor Graph in Reduced Normal Form

ferent types, for simplicity in the tests performed we have preferred to focus exclusively on this more manageable architecture (Fig. 4). The figure shows the most general case (used in the final tests), in which a Class Variable L is also added to the bottom variables. This configuration is used in the case of supervised learning, allowing (after learning) to perform various tasks, including:

- Pattern classification*, achievable by injecting the observations as delta distributions on the backwards of $Y_1, \dots, Y_N$ and leaving the backward on L uniform. The classification will be returned to L through its forward message f_L .
- Pattern completion*, where only some of the observed variables $Y_1, \dots, Y_N$ are available in inference (and then injected into the network through delta distributions on their backwards) while the others are unknown (and therefore represented by uniform backward distributions).
Also L may be unknown (uniform distribution), partially known (generic distribution), or known (delta distribution). The forward distributions at the unknown variables complete the pattern, and at L provide the best inference in case it is not perfectly known. The posterior on L is obtained by multiplying the two messages on it (Equation 1); avoidable step if L is not known at the beginning (because in this case the backward is uniform).
- Prototype inspection*, obtainable by injecting only a delta distribution at L on the j th label. The forward distributions $f_{Y_1}, \dots, f_{Y_N}$ will represent the separable prototype distribution of that class.

Another way to use the previous network is to make the class variable coincide with the hidden variable $S = L$, forcing the corresponding SISO block matrix to be diagonal. This constraint will create a so-called “Naive Bayes Classifier” (Barber 2012), further simplifying the factorization into

$$P(Y_1, \dots, Y_N, L) = P(Y_1|L) \cdots P(Y_N|L) P(L). \quad (6)$$

In this case, usually all the variables are observed during training, and the typical use in inference is to obtain L from observed $Y_1, \dots, Y_N$.

Note that the case related to unsupervised learning can be obtained from the general model presented simply by eliminating the variable L . In this case, after learning, the elements of the alphabet $\mathcal{S} = \{s_1, \dots, s_{|\mathcal{S}|}\}$ of the hidden variable S represent “Bayesian clusters” of the data, which follows the prior distribution π_S (learned blindly). The network can be used in inference both for the pattern completion, in the case where (as seen previously) only some of the underlying variables are known and we try to estimate the others through the corresponding forward messages, and to create a so-called embedding representation, in which the backward message becomes a different representation of the underlying known variables. In the latter case, in order to understand the representation that the network has created, we can look at the j -th centroid of the bayesian clusters injecting as $\mathbf{f}_S$ a delta distribution $\mathbf{f}_j = [0 \dots 1 \dots 0]^T$, where the 1 is at the j -th position. The set of forward distributions $f_{Y_1}, \dots, f_{Y_N}$ generated by the network will represent the marginal distributions around the centroid of the j th cluster.

4 Design of FGrnLib

There are several software packages, known in the literature, that can be used to design and/or simulate Bayesian networks (an updated list can be found in Murphy (2014)). Unfortunately, many of them are in closed packages and/or run only on private servers, preventing proper performance analysis. Others either have limitations on the number of variables and the size of the network, or do not use the FG architecture. Therefore, the main purpose of this work (ie the reduction of complexity) has merged with the design of an optimized library, called *FGrnLib*, for the realization of a Bayesian network through the use of the FGrn model; that is open and contains an efficient implementation of the elements in Fig. 1.

Fig. 5 The basic class diagram for FGnLib, that shows the dependencies between the classes

The *FGnLib* library has been written in C++, following the classic object-oriented paradigm, and it has been adapted for parallel computing (on multiprocessor systems with shared memory) through the use of the OpenMP application interface. The various algorithmic operations have been implemented to limit as much as possible the computational complexity without significantly affecting memory requirements.

4.1 Data structures

Before starting the analysis of individual operations, it is necessary to focus on the structure of the main classes (Fig. 5) and their individual roles. These classes correspond to the main elements presented in Fig. 1 (and mathematically described in Sect. 2), but also have subtle design choices that need to be clarified (especially in reference to what was previously done in Buonanno and Palmieri (2015a)). The main classes are:

- The **Link** class, which represents a single discrete variable of the model. This class contains the two forward and backward messages for the variable, and is designed to ensure that each message at each time step is represented by a single vector. This means that in every instant of time every single link takes on only two messages (in the two different directions), providing better control of information traveling on the network but preventing the possibility of learning through the simultaneous presentation of all the evidence.
- The **Diverter** class, which imposes the constraint of equality among the variables. This class has been created to be as general as possible, in the sense that it can automatically adapt the parameters of the net to the number of variables. For this reason, it includes not only the process of replication of variables but also the creation and control of product space matrices (Equation 5). For space

complexity, these (sparse and row-stochastic) matrices are stored in column vectors, whose elements represent the index of the active column in that particular row of the matrix

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \\ \vdots & & \end{bmatrix} \rightarrow \begin{bmatrix} 0 \\ 1 \\ 2 \\ 0 \\ \vdots \end{bmatrix}.$$

- The **SISOBlock** class, which represents the probability of the output variable Y given an input variable V , contains the row-stochastic conditional probability matrix $P(Y|V)$. Due to the fact that the **Link** class allows only two vector at a time, the SISO blocks are realized to permit the storage and retrieval of all messages that reached the blocks during the batch learning phase; avoiding transmitting the evidences all at the same time but reusing the single **Link** memory units for each step.
- The **Source** class, which represents the independent variables inside the model, with their prior probabilities.

4.2 Data flow management

Since the FGn paradigm allows to work only through local information, the flow of messages within the network can be parallelized. However, parallelizing the message flow in the network imposes essential changes to the **Diverter** class.

In fact, the multiplication of the messages internally to the **Diverter** can take place only after all the messages of the variables relating thereto have been updated. Being the only responsible for the combination of messages coming from different directions, the **Diverter** must necessarily act as a “barrier” (with the meaning that this term assumes in parallel programming); solving the synchronization problems

Fig. 6 Messages propagation: **a** inference phase; **b** batch learning phase

simply by not updating the output values until it receives the activation signal from a *supervisor*. Although within the FGrnLib a *Supervisor* class has been defined to more easily manage some predefined network, the meaning of a supervisor is quite general here because it refers to any class that possesses all the references to the single elements of the network realized.

In Fig. 6 it is shown how the supervisor handles the message scheduling, relatively to both the inference mode and the batch learning mode. Note that, in order for messages that travel in parallel to be propagated anywhere on the network, a number of steps equal to the diameter of the graph is required (Pearl 1988).

Recalling therefore that, in graph theory, the diameter of the network is the longest possible path (in terms of number of crossed arcs) among the smallest existing paths able to connect each vertex inside the graph, it is easy to understand that a simple one-layer LVM network (like in Fig. 3 or in Fig. 4a) will only need three propagation steps.

4.2.1 Inference

The inference phase, depicted in Fig. 6a, is relatively simple: the various messages proceed in parallel until they reach the Diverter. All that is necessary is to block the messages at the Diverter, to prevent the start of the multiplication process before all the messages have become available. The supervisor must therefore perform only two phases: parallelize the input variables and then activate the Diverter. It should be noted that typically all messages are initialized to uniform vectors, and updated according to the rules described in Sect. 2.

For what has been said, in the first step of Fig. 6a the supervisor performs the parallelization of the input variables $Y_1, \dots, Y_N$, placing in the backward messages appropriate distributions for the known (or partially known) variables and uniform distributions for the unknown ones. Hence, in

the second step, the messages start to be propagated in the network (using the second in Equation 3). At this point, when all the variables have performed the second step, the supervisor must activate the Diverter to allow the propagation of the messages to continue. The last message propagation step allows to obtain the desired output values (using the first in Equation 3).

4.2.2 Batch learning

Every single epoch of the batch learning phase (Fig. 6b) is not so different from what we have just seen for the inference phase, since the supervisor basically performs only two more operations. The first one, which only applies at the beginning of the whole learning phase (that is, it is not reiterated at each epoch), consists in activating the batch learning mode inside the SISO blocks and the Sources. This step enables storage within the SISO blocks, which consequently memorize the incoming vectors (from both sides) as soon as they are updated, and the Sources, which will begin to add together all incoming backward messages within a temporary vector. It should be noted that, in the case of multi-layer structures, the supervisor will also have to worry about preventing the SISO blocks of the next level from propagating incoming messages, thus transforming the latter into “barriers” as well. In fact, in order to get the classic layer-by-layer approach, the information should not propagate above the Diverter until the underlying batch learning phase is complete, but only if the top layer does not provide Sources. So as not to complicate the Diverter too much, forcing it to know the connected elements, the SISO blocks also provide a pause mode, which if enabled prevents forward propagation of messages.

After activation of the batch learning mode the messages propagate within the network through the same previous steps, being blocked again to the Diverter. However, as we can see in Fig. 6b, the descending phase will not include the production of the final messages, which will only be stored

in the SISO blocks in order to avoid performing the related mathematical operations.

The second different operation, performed by the supervisor at the end of each epoch, consists in the activation of the actual learning procedure, which will execute the ML algorithm using the stored vectors (up to a sufficient convergence or at most for a fixed number of times) and change the prior of the Source. After this last phase, the procedure can be repeated for another epoch, presenting the evidence to the network again. At the end of the learning, the last message from the supervisor will also modify the operating modes of the SISO blocks and the Sources, thus making them ready for the subsequent inference phases.

5 Complexity and efficient algorithms

Having to pay particular attention to the computational and memory costs, in the creation of the library we worked on the details of each individual element. This, together with the probabilistic nature of the Bayesian networks, has led to the preliminary definition of particular basic data structures that it is perhaps necessary to analyse quickly. In fact, the library also defines classes that represent probability vectors and row-stochastic matrices, to facilitate the interpretation and definition of the variables and to easily manage all the algebraic operations.

5.1 Probability vector

A probability vector is a vector with elements in $[0, 1]$ that sum to one. Although the network does not use only probability vectors, the execution of every operation that uses them must then provide normalization. Every normalization consists of $d - 1$ sums and d divisions, so the computational cost will be $\mathcal{O}(d)$

$$\left[\begin{array}{c} v_1 \\ v_2 \\ \vdots \\ v_d \end{array} \right]_{\text{norm}} = \left[\begin{array}{c} \zeta_1 \\ \zeta_2 \\ \vdots \\ \zeta_d \end{array} \right] \rightarrow \zeta_i = \frac{v_i}{\sum_{j=1}^d v_j}.$$

5.2 Row-stochastic matrix multiplication

In a row-stochastic matrix P each row is a probability vector. It is important to observe that the premultiplication or postmultiplication of a row-stochastic matrix for a vector (with the appropriate dimensions)

Fig. 7 A diverter with $M = 4$ (one input and three output variables)

$$[\xi_1 \dots \xi_l] \begin{bmatrix} p_{11} & \dots & p_{1d} \\ \vdots & \ddots & \vdots \\ p_{l1} & \dots & p_{ld} \end{bmatrix} = [\zeta_1 \dots \zeta_l] \rightarrow \zeta_j = \sum_{i=1}^l p_{ij} \xi_i$$

$$\begin{bmatrix} p_{11} & \dots & p_{1d} \\ \vdots & \ddots & \vdots \\ p_{l1} & \dots & p_{ld} \end{bmatrix} \begin{bmatrix} \xi_1 \\ \vdots \\ \xi_d \end{bmatrix} = \begin{bmatrix} \zeta_1 \\ \vdots \\ \zeta_l \end{bmatrix} \rightarrow \zeta_i = \sum_{j=1}^d p_{ij} \xi_j$$

will consist of ld multiplications and $(l - 1)d$ or $l(d - 1)$ sums respectively, producing the same computational cost $\mathcal{O}(ld)$.

5.3 Diverter

From a computational point of view, the most critical structure in the implementation of a Bayesian network using FGm is represented by the Diverter, where, obviously, the greatest criticality is in the efficient implementation of the internal multiplication process (Equation 2). In fact, the possible presence of zeros, in different positions of the single vectors, obliges to perform the calculation of the outgoing vectors individually. In the general case of Fig. 3, indicating the total value of all the variables connected to the diverter with $M = H + N$ and assuming that the input vectors have dimensions equal to d , bare application of the multiplication rule would require $M(M - 2)d$ multiplications. Regardless of the size of the input vectors, the computational cost is therefore polynomial and equal to $\mathcal{O}(M^2)$ vectorial multiplications. Considering as an example the Diverter of Fig. 7 (with $M = 4$), the simple application of the product rule in the production of the outgoing messages ($\mathbf{b}_{V(0)}$, $\mathbf{f}_{V(1)}$, $\mathbf{f}_{V(2)}$, $\mathbf{f}_{V(3)}$) would require eight vectorial multiplications:

$$\mathbf{b}_{V(0)} = \mathbf{b}_{V(1)} \odot \mathbf{b}_{V(2)} \odot \mathbf{b}_{V(3)};$$

$$\mathbf{f}_{V(1)} = \mathbf{f}_{V(0)} \odot \mathbf{b}_{V(2)} \odot \mathbf{b}_{V(3)};$$

$$\mathbf{f}_{V(2)} = \mathbf{f}_{V(0)} \odot \mathbf{b}_{V(1)} \odot \mathbf{b}_{V(3)};$$

$$\mathbf{f}_{V(3)} = \mathbf{f}_{V(0)} \odot \mathbf{b}_{V(1)} \odot \mathbf{b}_{V(2)}.$$

This process can be performed more efficiently by defining an order among the variables connected to the Diverter and by performing a double cascade process in which each variable is responsible only for passing the correct value to

Fig. 8 Details of the efficient implementation of the products inside the Diverter, with in red the *input messages* and in blue the *output messages*. For each output message two contributions are used: one derived from the left part of the computational graph and the other one derived from the right part

the neighboring variable. In this way, the variables at the ends of the chain will perform no multiplication while each variable inside the chain will perform only three multiplications, relative to the passage of the two temporary vectors along the chain and to the output of the outgoing message.

With reference to the example of Fig. 7 we have the data flow represented in the Fig. 8. In other words, the proposed solution exploits the presence of the same multiplication groups through a round-trip process. This reduces the computational complexity from quadratic to linear, $\mathcal{O}(M)$, finally requiring only $3(M - 2)$ vector multiplications. Although there is obviously an increase in the memory required, due to temporary vectors along the chain, it remains linear with M , and has been further optimized (requiring only $M - 1$ temporary vectors altogether) by choosing to reuse the same vectors ($a_i = a'_i$) when changing direction.

5.4 Unknown variables

A very attractive property of using a probabilistic paradigm, which makes it preferable in certain contexts, is represented by its ability to manage unknown inputs. Even in the case of maximum uncertainty, that is when nothing is known about the particular variable in that particular observation set, a process of inference or learning can still be performed by making the corresponding message of that variable a uniformly distributed probability vector

$$\bar{\mathbf{b}}_{Y_i} = \begin{bmatrix} \frac{1}{|\mathcal{Y}_i|} \\ \vdots \\ \frac{1}{|\mathcal{Y}_i|} \end{bmatrix}.$$

In this particular circumstance, the message propagation process can be optimized by avoiding the multiplication of backward vectors with the matrices inside the SISO blocks, noting that

$$\begin{aligned} \mathbf{b}_{S(i)} &= P(Y_i | S) \bar{\mathbf{b}}_{Y_i} \\ &= \begin{bmatrix} P(Y_i = \xi_1 | S = \sigma_1) & \cdots & P(Y_i = \xi_{|\mathcal{Y}_i|} | S = \sigma_1) \\ \vdots & & \vdots \\ P(Y_i = \xi_1 | S = \sigma_{|S|}) & \cdots & P(Y_i = \xi_{|\mathcal{Y}_i|} | S = \sigma_{|S|}) \end{bmatrix} \begin{bmatrix} \frac{1}{|\mathcal{Y}_i|} \\ \vdots \\ \frac{1}{|\mathcal{Y}_i|} \end{bmatrix} \\ &= \begin{bmatrix} \sum_{j=1}^{|\mathcal{Y}_i|} \theta_{1j} \bar{b}_{Y_i}(\xi_j) \\ \vdots \\ \sum_{j=1}^{|\mathcal{Y}_i|} \theta_{|S|j} \bar{b}_{Y_i}(\xi_j) \end{bmatrix} \\ &= \begin{bmatrix} \frac{1}{|\mathcal{Y}_i|} \sum_{j=1}^{|\mathcal{Y}_i|} \theta_{1j} \\ \vdots \\ \frac{1}{|\mathcal{Y}_i|} \sum_{j=1}^{|\mathcal{Y}_i|} \theta_{|S|j} \end{bmatrix} = \begin{bmatrix} \frac{1}{|\mathcal{Y}_i|} \\ \vdots \\ \frac{1}{|\mathcal{Y}_i|} \end{bmatrix}. \end{aligned}$$

By not propagating the unknown variable (setting $\mathbf{b}_{S(i)}$ as an expanded/reduced version of $\bar{\mathbf{b}}_{Y_i}$), for every single unknown variable present in input during the inferential and the learning process we can save $|S|$ vector multiplications, improving overall network performance.

5.5 Efficient ML implementation

Regarding the learning phase, particular attention has been given to the realization of an efficient implementation of the ML algorithm. First of all, since the matrix inside the SISO blocks is set to be row-stochastic by construction, it has been noted that the first divisor in the Equation 4 becomes unnecessary. For this reason, the equation can be rewritten as follows

$$\theta_{lm}^{(1)} \leftarrow \theta_{lm}^{(0)} \sum_{n=1}^N \frac{f_{V[n]}(l) b_{Y[n]}(m)}{\mathbf{f}_{V[n]}^T \boldsymbol{\theta}^{(0)} \mathbf{b}_{Y[n]}},$$

or in vector form

$$\boldsymbol{\theta}^{(1)} \leftarrow \boldsymbol{\theta}^{(0)} \odot \sum_{n=1}^N \frac{\mathbf{f}_{V[n]} \mathbf{b}_{Y[n]}^T}{\mathbf{f}_{V[n]}^T \boldsymbol{\theta}^{(0)} \mathbf{b}_{Y[n]}}.$$

Furthermore, it can be observed that the value obtained through the vector multiplications $\mathbf{f}_{V[n]}^T \boldsymbol{\theta}^{(0)} \mathbf{b}_{Y[n]}$ is actually equal to the sum of all elements of the matrix $\boldsymbol{\theta}^{(0)} \odot \mathbf{f}_{V[n]} \mathbf{b}_{Y[n]}^T$. This assertion is provable by observing that

$$\begin{aligned} \boldsymbol{\theta}^{(0)} \odot \mathbf{f}_{V[n]} \mathbf{b}_{Y[n]}^T &= \begin{bmatrix} \theta_{11}^{(0)} & \cdots & \theta_{1|\mathcal{Y}|}^{(0)} \\ \vdots & \ddots & \vdots \\ \theta_{|\mathcal{V}|1}^{(0)} & \cdots & \theta_{|\mathcal{V}||\mathcal{Y}|}^{(0)} \end{bmatrix} \odot \begin{bmatrix} \phi_1 \\ \vdots \\ \phi_{|\mathcal{V}|} \end{bmatrix} [\beta_1 \cdots \beta_{|\mathcal{Y}|}] \\ &= \begin{bmatrix} \theta_{11}^{(0)} & \cdots & \theta_{1|\mathcal{Y}|}^{(0)} \\ \vdots & \ddots & \vdots \\ \theta_{|\mathcal{V}|1}^{(0)} & \cdots & \theta_{|\mathcal{V}||\mathcal{Y}|}^{(0)} \end{bmatrix} \odot \begin{bmatrix} \phi_1 \beta_1 & \cdots & \phi_1 \beta_{|\mathcal{Y}|} \\ \vdots & \ddots & \vdots \\ \phi_{|\mathcal{V}|} \beta_1 & \cdots & \phi_{|\mathcal{V}|} \beta_{|\mathcal{Y}|} \end{bmatrix} \\ &= \begin{bmatrix} \theta_{11}^{(0)} \phi_1 \beta_1 & \cdots & \theta_{1|\mathcal{Y}|}^{(0)} \phi_1 \beta_{|\mathcal{Y}|} \\ \vdots & \ddots & \vdots \\ \theta_{|\mathcal{V}|1}^{(0)} \phi_{|\mathcal{V}|} \beta_1 & \cdots & \theta_{|\mathcal{V}||\mathcal{Y}|}^{(0)} \phi_{|\mathcal{V}|} \beta_{|\mathcal{Y}|} \end{bmatrix} \end{aligned}$$

whose sum of all the elements can than be written in the form

$$\sum_{l=1}^{|\mathcal{V}|} \sum_{m=1}^{|\mathcal{Y}|} \boldsymbol{\theta}^{(0)} \odot \mathbf{f}_{V[n]} \mathbf{b}_{Y[n]}^T = \sum_{l=1}^{|\mathcal{V}|} \phi_l \sum_{m=1}^{|\mathcal{Y}|} \theta_{lm}^{(0)} \beta_m,$$

which precisely is equal to

$$\begin{aligned} \mathbf{f}_{V[n]}^T \boldsymbol{\theta}^{(0)} \mathbf{b}_{Y[n]} &= [\phi_1 \cdots \phi_{|\mathcal{V}|}] \begin{bmatrix} \theta_{11}^{(0)} & \cdots & \theta_{1|\mathcal{Y}|}^{(0)} \\ \vdots & \ddots & \vdots \\ \theta_{|\mathcal{V}|1}^{(0)} & \cdots & \theta_{|\mathcal{V}||\mathcal{Y}|}^{(0)} \end{bmatrix} \begin{bmatrix} \beta_1 \\ \vdots \\ \beta_{|\mathcal{Y}|} \end{bmatrix} \\ &= [\phi_1 \cdots \phi_{|\mathcal{V}|}] \begin{bmatrix} \sum_{m=1}^{|\mathcal{Y}|} \theta_{1m}^{(0)} \beta_m \\ \vdots \\ \sum_{m=1}^{|\mathcal{Y}|} \theta_{|\mathcal{V}|m}^{(0)} \beta_m \end{bmatrix} = \sum_{l=1}^{|\mathcal{V}|} \phi_l \sum_{m=1}^{|\mathcal{Y}|} \theta_{lm}^{(0)} \beta_m. \end{aligned}$$

This suggests that moving the Hadamard product inside the summation

$$\boldsymbol{\theta}^{(1)} \leftarrow \sum_{n=1}^N \frac{\boldsymbol{\theta}^{(0)} \odot \mathbf{f}_{V[n]} \mathbf{b}_{Y[n]}^T}{\mathbf{f}_{V[n]}^T \boldsymbol{\theta}^{(0)} \mathbf{b}_{Y[n]}}$$

and calculating the sum of all the elements of the matrix $\boldsymbol{\theta}^{(0)} \odot \mathbf{f}_{V[n]} \mathbf{b}_{Y[n]}^T$ at the same time of their generation, computational complexity of the algorithm can overall be reduced from $N(3|\mathcal{Y}| + 1)|\mathcal{V}|$ to $2N|\mathcal{V}||\mathcal{Y}|$ multiplications, which is rather significant if considering that this reduction is relative to each algorithm recall.

6 Performance on LVM

To evaluate the computational advantages obtained by the proposed improvements let us consider the more general situation of Fig. 3, in which the LVM (single-layer) model

Table 1 Computational cost differences for the inference phase between the classical and the optimized algorithm

	Computational cost
Direct	$\mathcal{O}((N(\mathcal{Y} + H + N) + H^2) \mathcal{P})$
Optimized	$\mathcal{O}((N \mathcal{Y} + H) \mathcal{P})$

foresees H sources (hidden) and N output (observed) variables $Y_1, \dots, Y_N$. Since the variables $Y_1, \dots, Y_N$ can have different dimensions (never less than 2), before starting it is important to underline that we will assume that all the observed variables have the same dimension $|\mathcal{Y}|$. This statement does not compromise in any way the goodness of the results obtained, for we can certainly decide to choose $|\mathcal{Y}|$ as the highest value possible being interested only in an upper-bound description of computational complexity (in terms of big-O notation). For the same reason, we will avoid considering the case (more advantageous) in which some variables are not known. Note that, the previous problem does not arise in the case of input variables to the Diverter because they already have the same dimension $|\mathcal{P}|$ by construction.

6.1 Cost of the inference phase

In the inference mode, at the beginning of the process, the backward messages of the N output variables will be post-multiplied by the probabilistic matrix inside the SISO blocks; thus producing $\mathcal{O}(N|\mathcal{P}||\mathcal{Y}|)$ operations before being sent to the Diverter. As seen previously, $\mathcal{O}((H + N)|\mathcal{P}|)$ operations will be performed inside the Diverter, related to the multiplication of all the $H + N$ incoming messages to it. Finally, the messages must be propagated again to the SISO blocks, which will be pre-multiplied by the matrices still producing $\mathcal{O}(N|\mathcal{P}||\mathcal{Y}|)$ operations, and thus making the total computational cost equal to $\mathcal{O}((N|\mathcal{Y}| + H)|\mathcal{P}|)$. At this point the forward messages of the $Y_1, \dots, Y_N$ output variables will be available for analysis. Table 1 summarizes the differences in computational terms determined by the introduced optimizations.

To provide a practical example, suppose to perform an inferential process on an LVM network with 10 binary output variables and a single hidden variable, with an embedding space $|\mathcal{S}|$ of 10. It will be noted that while the direct algorithm will require 1500 multiplications, the optimized one will require only 670.

Moreover, the additional memory required by the optimized algorithm is still linear with the size of the inputs, since it depends only on the temporary vectors inside the Diverter. In fact, in the previous example, the significant computational advantage obtained will correspond to an increase in memory equal to only 10 vectors of size 10.

Table 2 Computational cost differences for the batch learning phase between the classical and the optimized algorithm

	Computational cost
Direct	$\mathcal{O}(TL(N(K \mathcal{Y} + H + N) + H^2) \mathcal{P})$
Optimized	$\mathcal{O}(TL(KN \mathcal{Y} + H) \mathcal{P})$

6.2 Cost of the batch learning phase

For what concerns the batch learning session, assume to have L different examples given in input through the backward messages of the output variables. As already mentioned, when the process begins backward messages entering the network will be saved inside the SISO blocks, requesting an amount of memory of $\mathcal{O}(L|\mathcal{Y}|)$ individually. After storing the vector, the process will continue by multiplying it and the probability matrix; sending the result to the Diverter.

In the propagation phase the computational costs will not change with respect to the inference phase, but it must be remembered that the messages that will return to the SISO blocks will be stored again; individually producing an additional memory cost equal to $\mathcal{O}(L|\mathcal{P}|)$. The total cost of additional memory required is thus $\mathcal{O}(LN(|\mathcal{Y}| + |\mathcal{P}|))$, while the computational cost is still $\mathcal{O}(L(N|\mathcal{Y}| + H)|\mathcal{P}|)$.

Once this first phase has been completed, the ML algorithm will be executed at most K times (with K fixed a priori), trying to make conditional probability matrices converge, and the whole process is then repeated T times. Thus, the total computational cost of the batch learning session is equal to

$$\mathcal{O}(TL(KN|\mathcal{Y}| + H)|\mathcal{P}|).$$

The comparison between the computational costs of the direct and the optimized case is shown more clearly in Table 2.

Furthermore, it is easy to state that the repetition of the process for T epochs, as well as the various calls of the ML algorithm, do not imply the need for any additional memory units with respect to the non-optimized case.

Table 3 Comparison of the classification accuracy on three different datasets between the incremental and the batch algorithm

	Training set		Test set	
	Batch (%)	Incremental (%)	Batch (%)	Incremental (%)
Breast cancer	96.6	96.2	95.48	97.99
Mammographic mass	86.62	85.88	79.5	78.88
Contraceptive method	58.1	58.9	49.89	50.52

Datasets composition: Breast Cancer = 10 useful variables, 699 instances, 16 missing values; Mammographic Mass = 6 variables, 961 instances, 162 missing values; Contraceptive Method = 10 variables, 1473 instances, 0 missing values

7 Incremental algorithm

The ML algorithm has many undoubted advantages, being very stable and generally converging in a few steps, but it obviously has the disadvantage of being batch (i.e. able to perform learning using only the entire training set). In order to obtain a lighter implementation we have changed the previous structure by requiring that at each epoch of the learning phase only one ML cycle (i.e. $K = 1$) is included. In other words, the algorithm has been made incremental, making it unnecessary to store backward messages within SISO blocks; and thus eliminating the need for the previous storage space equal to $L(|\mathcal{P}| + |\mathcal{Y}|)$ for each SISO block.

Despite the great advantage in terms of both memory and computational costs, this type of approach has surprisingly proved to be as robust as the previous one, being less likely to provide overfitting of the data. Referring to a One-To-Many LVM structure (Fig. 4b), in which the only hidden variable has an embedding space $|\mathcal{S}|$ of 20, various tests were performed on different datasets.

Table 3 present the classification success rates, both on the training and on the test set, of three databases from the UCI repository: Wisconsin Breast Cancer Dataset (Dheeru and Taniskidou 2017), Mammographic Mass Data Set (Elter et al. 2007) and Contraceptive Method Choice Data Set (Dheeru and Taniskidou 2017). The values presented are obtained by making sure that the latent variable is not learned, and therefore represent the learning ability of a single layer; being a good index to represent layer-by-layer learning of more complex networks. Note that, in this particular situation, the incremental algorithm does not give results that deviate much from the values obtained with the batch learning, providing in some cases even better results. It should also be noted that the results do not improve according to the adaptability of the paradigm to the specific case, since they are better even when the realized one-layer LVM network is probably not suitable for capturing the underlying implication scheme (as seen in the case of Contraceptive Method database).

Finally, the results prove even more interesting if we consider that in both cases (batch and incremental) they are obtained using the same number of epochs (in particular equal to 20). In fact, an incremental algorithm should typi-

cally employ many more steps to achieve the performance of a batch algorithm, whereas in the cases under examination the change seems to bring only advantages.

8 Discussion and conclusions

In this work, an in-depth analysis of the individual elements necessary to create a Bayesian network through the FG_{rn} paradigm was conducted, showing how it is possible to reduce memory and computational costs during implementation. The analysis led to the creation of a C++ library able to provide excellent results from a computational point of view, transforming polynomial costs into linear (respectively to the number of variables involved). The incremental use of the ML algorithm has finally demonstrated how it is further possible to reduce both the computational and memory costs of the learning phase in an unsupervised context. All these algorithmic choices are the basis for extending the FG_{rn} paradigm to higher-scale problems.

Funding Open access funding provided by Università degli Studi della Campania Luigi Vanvitelli within the CRUI-CARE Agreement. This work has been partially sponsored by PON03PE-00185-1-2 MAR.TE., with Consorzio Nazionale Interuniversitario per le Telecomunicazioni (CNIT) - Italy.

Availability of data and material The only data used in the paper refer to the three public databases from the UCI repository: Wisconsin Breast Cancer Dataset (Dheeru and Taniskidou 2017), Mammographic Mass Data Set (Elter et al. 2007) and Contraceptive Method Choice Data Set (Dheeru and Taniskidou 2017).

Compliance with ethical standards

Conflicts of interest Author Giovanni Di Gennaro declares that he has no conflict of interest. Author Amedeo Buonanno declares that he has no conflict of interest. Author Francesco A.N. Palmieri declares that he has no conflict of interest.

Code availability As stated in the paper, the C ++ library created is made available at <https://github.com/mlunicampania/FGrnLib>.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Barber D (2012) Bayesian reasoning and machine learning. Cambridge University Press, Cambridge
- Bishop CM (1999) Latent variable models. Learning in graphical models. MIT Press, Cambridge, pp 371–403
- Buonanno A, Palmieri FAN (2015) Simulink implementation of belief propagation in normal factor graphs. In: Advances in neural networks: computational and theoretical issues. Springer, New York, pp 11–20
- Buonanno A, Palmieri FAN (2015) Towards building deep networks with bayesian factor graphs. CoRR abs/1502.04492, [arXiv:1502.04492](https://arxiv.org/abs/1502.04492)
- Buonanno A, Palmieri FAN (2015) Two-dimensional multi-layer factor graphs in reduced normal form. In: Proceedings of the international joint conference on neural networks, IJCNN2015, July 12–17, 2015, Killarney, Ireland
- Cheeseman P, Stutz J (1996) Bayesian classification (AutoClass): theory and results. <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.44.7048>
- Dheeru D, Taniskidou EK (2017) Uci machine learning repository. <http://archive.ics.uci.edu/ml>
- Elter M, Schulz-Wendtland R, Wittenberg T (2007) The prediction of breast cancer biopsy outcomes using two cad approaches that both emphasize an intelligible decision process. Med Phys 34(11):4164–4172
- Forney GD (2001) Codes on graphs: normal realizations. IEEE Trans Inf Theory 47:520–548
- Koller D, Friedman N (2009) Probabilistic graphical models: principles and techniques. MIT Press, Cambridge
- Loeliger HA (2004) An introduction to factor graphs. IEEE Signal Process Mag 21(1):28–41
- Murphy K (2014) Software packages for graphical models. <http://www.cs.ubc.ca/~murphyk/Software/bnsoft.html>
- Murphy KP (2012) Machine learning: a probabilistic perspective (adaptive computation and machine learning series). The MIT Press, Cambridge
- Palmieri FAN (2016) A comparison of algorithms for learning hidden variables in Bayesian factor graphs in reduced normal form. IEEE Trans Neural Netw Learn Syst 27(11):2242–2255
- Palmieri FAN, Buonanno A (2014) Belief propagation and learning in convolution multi-layer factor graph. In: Proceedings of the 4th international workshop on cognitive information processing, Copenhagen, Denmark
- Palmieri FAN, Buonanno A (2015) Discrete independent component analysis (dica) with belief propagation. In: Proceedings of IEEE machine learning for signal processing conference, MLSP2015, Sept. 17–20, Boston, MA, US
- Pearl J (1988) Probabilistic reasoning in intelligent systems: networks of plausible inference. Morgan Kaufmann Publishers Inc., San Francisco

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.