



Introducing instance label correlation in multiple instance learning. Application to cancer detection on histopathological images

Pablo Morales-Álvarez^{a,*},¹ Arne Schmidt^a, José Miguel Hernández-Lobato^b, Rafael Molina^a

^a Department of Computer Science and AI, University of Granada, Spain

^b Department of Engineering, University of Cambridge, UK

ARTICLE INFO

Keywords:

Multiple instance learning
Gaussian processes
Ising model
Variational inference
Whole slide images
Histopathology

ABSTRACT

In the last years, the weakly supervised paradigm of multiple instance learning (MIL) has become very popular in many different areas. A paradigmatic example is computational pathology, where the lack of patch-level labels for whole-slide images prevents the application of supervised models. Probabilistic MIL methods based on Gaussian Processes (GPs) have obtained promising results due to their excellent uncertainty estimation capabilities. However, these are general-purpose MIL methods that do not take into account one important fact: in (histopathological) images, the labels of neighboring patches are expected to be correlated. In this work, we extend a state-of-the-art GP-based MIL method, which is called VGPMIL-PR, to exploit such correlation. To do so, we develop a novel coupling term inspired by the statistical physics Ising model. We use variational inference to estimate all the model parameters. Interestingly, the VGPMIL-PR formulation is recovered when the weight that regulates the strength of the Ising term vanishes. The performance of the proposed method is assessed in two real-world problems of prostate cancer detection. We show that our model achieves better results than other state-of-the-art probabilistic MIL methods. We also provide different visualizations and analysis to gain insights into the influence of the novel Ising term. These insights are expected to facilitate the application of the proposed model to other research areas.

1. Introduction

Multiple instance learning (MIL) has caught great attention in fields where there is a challenging lack of labeled data. Although it has been applied in many different areas [1], we will focus on the case of computational pathology. In the last years, thanks to the increasing digitalization of whole-slide images (WSIs), the field of computational pathology is developing computer-aided diagnosis systems based on machine learning for cancer detection [2,3]. The goal of computational pathology is to provide a fast and reliable diagnosis for the most prototypical cases, letting the pathologists focus on the most challenging ones. Ultimately, this will enable a much wider access to early cancer diagnosis [4].

In order to make accurate predictions, machine learning classification methods need to be trained using a labeled set of instances [5]. In the case of computational pathology, these instances are typically patches from the WSIs (and not the complete images themselves) [6–8]. The reason for this is twofold: (i) it is useful to have predictions at patch level in order to know where exactly in the image the cancer is located, and (ii) WSIs are extremely large and cannot be directly fed to a classifier. As a consequence, notice that expert pathologists

must label *every single patch* in the training data as cancerous or not (we will consider the binary problem cancer/no-cancer throughout this work). Given the large number of patches and the limited availability of pathologists, this becomes a daunting task in real practice [9].

To address this problem, different weakly supervised learning paradigms have been proposed in recent years. Here we focus on MIL, which has become very popular in the medical domain [8,10]. The idea in MIL is that instances are grouped in bags, and only bag labels are needed for training. In the case of WSIs, all the patches coming from the same image are considered a bag. Therefore, the labeling workload on pathologists decreases enormously: from labeling every single patch, to only labeling the complete WSI as cancerous or not.

Different machine learning algorithms have been developed to learn under the MIL setting. Notice that dealing with uncertainty is essential in MIL models, since instance-level labels are unknown. To deal with uncertainties, different probabilistic methods have been developed, such as Dirichlet Process Mixture Models [11], Markov chain [12], Monte-Carlo chain [13,14] and Gaussian Processes (GPs) [15,16]. In particular, GPs have attracted plenty of attention in the last years, due to their expressive power and their capacity to handle uncertainty

* Corresponding author.

¹ Work done during the Margarita Salas fellowship outgoing phase at Cambridge University, UK.

in a principled manner. Moreover, we are interested in this type of probabilistic models, since they will allow for introducing correlations in a theoretically sound way.

Among GP-based MIL methods, we will focus on the two most successful ones: VGPMIL and VGPMIL-PR. VGPMIL [17] was proposed in 2017 to overcome the limitations of two earlier formulations [15, 16] (namely, the use of the inefficient Laplace approximation and the impossibility to obtain instance-level predictions, respectively). In short, VGPMIL relies on variational inference and allows for closed-form updates of its parameters. However, the use of the logistic function implies that VGPMIL needs to resort to a theoretical approximation during inference (namely, the Jaakola bound [17, Eq. (10)]). As shown in [18], such approximation hurts predictive performance in practice. As an alternative, the authors of [18] propose the utilization of the probit function, which removes the need for the aforementioned approximation. This method, which will be referred to as VGPMIL-PR, is considered the current state of the art among probabilistic MIL approaches.

Methods such as VGPMIL and VGPMIL-PR are general-purpose MIL models that can be used in any MIL problem (that is, whenever the label is known only at bag level, see different use-cases in [17,19]). However, the underlying MIL assumption that the labels of the instances in a bag are independent of each other is unrealistic in many real problems. For example, in the particular case of WSI images (and in many image-related MIL problems), the labels of neighboring patches are expected to be correlated [20]. We hypothesize that the predictive performance of MIL methods can be enhanced by incorporating this type of prior knowledge into the model.

In this work, we introduce a novel GP-based MIL algorithm that takes into account the correlation between the labels of neighboring patches, and we apply it to the real-world problem of prostate cancer detection on histopathological images. We model the correlation through a coupling term inspired by the Ising model [5, Section 19.4.1], an statistical physics method that has found several applications in computer vision [20,21]. Our GP-MIL modeling builds on VGPMIL-PR, so our method will be referred to as VGPMIL-PR-I (Ising). In VGPMIL-PR-I, a hyperparameter λ regulates the influence of the Ising-inspired terms. Variational inference is used to estimate the model parameters, and the update formulas of VGPMIL-PR are recovered when $\lambda \rightarrow 0$ (that is, when the influence of the coupling term vanishes). In the experimental section, we show that VGPMIL-PR-I outperforms the state-of-the-art GP-based MIL approaches VGPMIL and VGPMIL-PR when predicting at both instance and bag levels, while keeping an analogous computational cost. Moreover, to gain insights into the influence of the new coupling term, we analyze the role of λ , and provide several visualizations for the predictions.

The rest of the paper is organized as follows. Section 2 presents the probabilistic model and inference for the novel VGPMIL-PR-I. Closely related methods such as VGPMIL and VGPMIL-PR are also discussed in this section. Section 3 focuses on the empirical evaluation of the model, including the data description, the experimental framework, and the discussion of results. Section 4 provides the main conclusions and some future outlook.

2. Probabilistic model and inference

In this section we present the theoretical description for VGPMIL-PR-I. Specifically, Section 2.1 explains the problem formulation and the main notation. Section 2.2 explains the closely-related methods VGPMIL and VGPMIL-PR, which are at the base of our formulation. Section 2.3 introduces the novel coupling term that accounts for patch label correlation, which is used to define VGPMIL-PR-I. Section 2.4 shows how to perform variational inference to estimate the parameters in VGPMIL-PR-I. Section 2.5 explains the procedure to make predictions at both instance and bag levels.

2.1. Notation and problem formulation

Our notation follows the state-of-the-art work [22]. The training data is given by a set of bags $\mathbf{X} = \{\mathbf{X}_b\}_{b \in \mathcal{B}}$ and their corresponding labels $\mathbf{y} = \{y_b\}_{b \in \mathcal{B}}$. We deal with a binary problem, i.e. $y_b \in \{0, 1\}$. Each bag $\mathbf{X}_b = \{\mathbf{x}_i\}_{i \in b}$ contains $|b|$ instances, i.e. $b = \{i_1, \dots, i_{|b|}\} \subseteq [N]$ (N is the total amount of instances). Notice that different bags may have different amounts of instances. Each instance \mathbf{x}_i is given by a vector in \mathbb{R}^D . In the MIL setting, one assumes that each instance has its (unknown) label $h_i \in \{0, 1\}$. We write \mathbf{h}_b for the labels of all the instances belonging to bag b . The MIL labeling assumption dictates that a bag is considered positive (class 1) if at least one of its instances is positive. Mathematically, this is

$$p(y_b | \mathbf{h}_b) = 1[y_b = \max_{i \in b} h_i], \quad (1)$$

where $1[\cdot]$ is the indicator function (i.e. it equals one when its argument is true and zero otherwise). Finally, we will collectively denote $\mathbf{h} = \{\mathbf{h}_b\}_{b \in \mathcal{B}}$.

In the case of WSIs, each \mathbf{X}_b is an image, which is composed of its patches $\{\mathbf{x}_i\}_{i \in b}$. Each patch has an unknown label h_i (0 for non-cancerous and 1 for cancerous), and we only have access to the bag label y_b (whether the image is cancerous or not, i.e. whether it contains at least one patch that is cancerous).

The goal in MIL is to train a model based only on bag labels $\{y_b\}_{b \in \mathcal{B}}$. And such model must be able to predict at both instance and bag levels. That is, given a previously unseen instance $\mathbf{x}^* \in \mathbb{R}^D$, we are interested in the probability $p(h^* = 1)$. Likewise, given a previously unseen complete bag \mathbf{X}^* , we are interested in $p(y^* = 1)$.

2.2. Background: VGPMIL and VGPMIL-PR

As mentioned in the introduction, our model is inspired by two closely related methods: VGPMIL and VGPMIL-PR. To understand our contribution, it is essential to fully understand their formulations, which we explain next.

2.2.1. VGPMIL formulation

VGPMIL was introduced in [17]. The idea is to consider a sparse GP classification model [23] to describe the relationship between instance features \mathbf{X} and their (unknown) labels \mathbf{h} . Then, an additional bag likelihood must be considered to model the (observed) bag labels \mathbf{y} given the instance labels \mathbf{h} . Both components are described next.

The sparse GP classification model. Instances \mathbf{x}_i are associated latent variables $f_i \in \mathbb{R}$ which are modeled through a GP, $f \sim \mathcal{GP}(0, \kappa)$. We write κ for the GP kernel, which encodes the properties of the considered functions. Then, the instance labels h_i are defined from f_i through a classification likelihood v :

$$p(h_i | f_i) = v(f_i)^{h_i} (1 - v(f_i))^{1-h_i}. \quad (2)$$

Specifically, VGPMIL uses the logistic function $v(x) = (1 + e^{-x})^{-1}$. Intuitively, a large (resp. low) value of f_i implies that the class is likely to be one (resp. zero). Moreover, since standard GPs scale poorly with the number of training instances N , VGPMIL makes use of sparse GPs [23], which summarize the training data through $M \ll N$ inducing points. These inducing points $\mathbf{u} = \{u_1, \dots, u_M\}$ represent the value of the GP at some inducing points locations $\mathbf{Z} = \{\mathbf{z}_1, \dots, \mathbf{z}_M\}$ (just like $\mathbf{f} = \{f_1, \dots, f_N\}$ are the GP values at $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$). Therefore, the distributions of \mathbf{u} and $\mathbf{f} | \mathbf{u}$ are:

$$p(\mathbf{u}) = \mathcal{N}(\mathbf{u} | \mathbf{0}, \mathbf{K}_{ZZ}), \quad (3)$$

$$p(\mathbf{f} | \mathbf{u}) = \mathcal{N}(\mathbf{f} | \mathbf{K}_{XZ} \mathbf{K}_{ZZ}^{-1} \mathbf{u}, \mathcal{K}). \quad (4)$$

The expression of \mathcal{K} is given by the particular sparse GP approach used in VGPMIL, which is FITC [23], so we have $\mathcal{K} = \text{diag}(\mathbf{K}_{XX} - \mathbf{K}_{XZ} \mathbf{K}_{ZZ}^{-1} \mathbf{K}_{ZX})$. As is standard in GP literature, we are writing \mathbf{K}_{XX} for

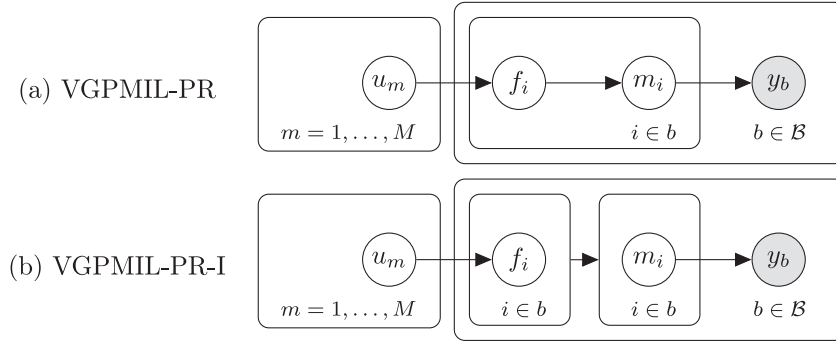


Fig. 1. Probabilistic graphical model for VGPMIL-PR (a) and VGPMIL-PR-I (b). Gray nodes are observed variables, and white ones are latent variables to be estimated. The main difference is that VGPMIL-PR-I introduces correlation between instances in the same bag. Therefore, the distribution of \mathbf{m}_b given \mathbf{f}_b does not factorize across instances. The correlation is introduced through a novel term inspired by the Ising model, see Section 2.3 for details.

the $N \times N$ covariance matrix $\mathbf{K}_{\mathbf{X}\mathbf{X}} = (\kappa(\mathbf{x}_i, \mathbf{x}_j))_{1 \leq i, j \leq N}$. The definitions for $\mathbf{K}_{\mathbf{X}\mathbf{Z}}$ and $\mathbf{K}_{\mathbf{Z}\mathbf{Z}}$ are analogous.

The bag likelihood. VGPMIL introduces the following parameterization to model the bag labels from the instance labels:

$$p(y_b | \mathbf{h}_b) = \frac{H^{G_b}}{H + 1}, \quad (5)$$

where $G_b := 1[y_b = \max_{i \in b} h_i]$ and H is a large and fixed value (in their examples, they use $H = 100$). Eq. (5) approximates the MIL assumption introduced in Eq. (1): if some instance label h_i is one, then the bag label y_b is one with very high probability (namely, with probability $\frac{H}{H+1}$). Otherwise (that is, if all instance labels h_i are zero), the bag label is one with very low probability (namely, $\frac{1}{H+1}$).

In summary, VGPMIL is given by Eqs. (3), (4), (2) and (5). For additional details, the interested reader is referred to the original work [17].

2.2.2. VGPMIL-PR formulation

VGPMIL-PR was recently proposed in [18] as an improvement over VGPMIL. Namely, the logistic function used by VGPMIL in Eq. (2) is not conjugate with the Gaussian distribution coming from the GP, recall Eqs. (3)–(4). This means that, in order to achieve mathematical tractability, VGPMIL needs to resort to the Jaakola bound [17, Eq. (10)]. However, the use of this bound introduces an *approximation* in the training objective. As shown in [18], such approximation damages the predictive performance in practice. Consequently, the authors of [18] introduce an alternative formulation based on the probit function, VGPMIL-PR. They show that, via a variable augmentation approach, VGPMIL-PR allows for directly optimizing the training objective (without approximations).

More specifically, VGPMIL-PR uses the probit function $v(x) = \int_{-\infty}^x \mathcal{N}(t|0, 1) dt$ in Eq. (2). Also, the bag likelihood is given by Eq. (1) (instead of Eq. (5)). Then, to circumvent the need for approximations, VGPMIL-PR leverages a variable augmentation approach [24]. Namely, for each instance we introduce a new variable $m_i \in \mathbb{R}$ between f_i and h_i , which is defined as $m_i \sim \mathcal{N}(f_i, 1)$. Since we are using a probit likelihood, we have that $h_i = 1[m_i > 0]$. Analogously to the rest of variables, we write $\mathbf{m}_b = \{m_i\}_{i \in b}$ for all the m_i 's inside bag b , and we use $\mathbf{m} = \{\mathbf{m}_b\}_{b \in \mathcal{B}}$ to collectively denote all the m_i 's in the model. Then, by marginalizing over \mathbf{h} , we have:

$$p(\mathbf{m} | \mathbf{f}) = \prod_b p(\mathbf{m}_b | \mathbf{f}_b) = \prod_b \mathcal{N}(\mathbf{m}_b | \mathbf{f}_b, I), \quad (6)$$

$$p(\mathbf{y} | \mathbf{m}) = \prod_b p(y_b | \mathbf{m}_b), \quad (7)$$

where I is the identity matrix (of size $|b|$) and $p(y_b = 0 | \mathbf{m}_b) = \prod_{i \in b} 1[m_i < 0]$. Importantly, these augmented variables \mathbf{m} will prove extremely helpful to introduce the Ising correlation in the next section.

In summary, VGPMIL-PR is given by Eqs. (3), (4), (6) and (7). Fig. 1(a) shows the probabilistic graphical model for VGPMIL-PR.




Fig. 2. A simplified illustration of an image composed by five patches: P_1, \dots, P_5 .

2.3. Correlating patch labels: VGPMIL-PR-I

As mentioned in the introduction, VGPMIL and VGPMIL-PR are general MIL approaches that can be used in any MIL problem. Indeed, there are plenty of applications where MIL methods can be used. For instance, think of a recommendation system where a reviewer has not evaluated every single item in the database, but has reviewed “groups” of them (e.g., he/she likes science-fiction movies, although he/she has not rated individual movies). Consider also a task of anomaly detection in which we do not have labels for individual transactions, but we only know whether there was some anomalous behavior in a certain period of time (which contains many different transactions).

Here we focus in the particular use-case of images, where bags are images and their instances are their patches. In this case, there exists very valuable information coming from the structure of the image itself, which can be exploited in the model. For example, it is natural to think that neighboring patches in the same image are likely to have similar labels. The main goal of this work is to introduce such correlation into the VGPMIL-PR formulation. To do so, we are inspired by the Ising model.

The novel coupling term. The Ising model arose from statistical physics to describe the behavior of magnets. In some magnets, called ferro-magnets, neighboring spins tend to line up in the same direction, whereas in other kinds of magnets, called anti-ferromagnets, the spins are repelled from their neighbors [5]. This type of interactions based on the Ising model have been used previously in machine learning and computer vision to describe relationships between pixels of an image, see e.g. [5, 20, 21]. However, to the best of our knowledge, they have never been used in the context of MIL.

Our first idea was to consider an Ising model over the patch labels of each image, $\{h_i\}_{i \in b}$. However, inference proved very challenging in this case, due to the non-conjugacy of the Ising model and the GP-based MIL formulation. As an alternative, we considered a continuous counterpart of the Ising model over the variables $\mathbf{m}_b = \{m_i\}_{i \in b}$ introduced in VGPMIL-PR, which are directly related to the patch labels (recall from Section 2.2.2 that $h_i = 1[m_i > 0]$). Notice that such continuous version of the Ising model corresponds to the well-known Conditional

Autoregression (CAR) [25]. Importantly, as we will see in the rest of this section, this alternative formulation yields a Gaussian distribution on \mathbf{m}_b , which can be treated analytically together with the GP-based MIL model.

Specifically, we consider the following coupling term for each image, which is defined over the augmented variables \mathbf{m}_b , recall Section 2.2.2:

$$C(\mathbf{m}_b) = \exp \left[-\frac{\lambda}{2} \cdot \sum_{\substack{i,j \in b \\ i < j}} \mathbf{1}[i, j \text{ are contiguous}] \cdot (m_i - m_j)^2 \right] = \exp \left(-\frac{\lambda}{2} \mathbf{m}_b^T \mathbf{C}_b \mathbf{m}_b \right). \quad (8)$$

Notice that $C(\mathbf{m}_b)$ is always in the range $[0, 1]$, and it becomes close to zero when the value of m is very different for neighboring patches. For the second equality in Eq. (8), notice that the sum only produces quadratic terms in m , so it can be written as $\mathbf{m}_b^T \mathbf{C}_b \mathbf{m}_b$ for some positive semidefinite matrix \mathbf{C}_b .

Since m determines the label of each patch (recall from Section 2.2.2 that $h_i = 1[m_i > 0]$), the term $C(\mathbf{m}_b)$ can be used to favor ‘‘smoothness’’ in the labels associated to the different patches. Also, the hyperparameter λ , which can be set to any non-negative value, regulates the strength of the coupling term: the larger λ , the more importance is given to differences in m . For example, when $\lambda = 0$, $C(\mathbf{m}_b)$ becomes constant and it does not account for correlation between patch labels.

An example of $C(\mathbf{m}_b)$. To illustrate the proposed coupling term, consider an image with five patches P_1, \dots, P_5 distributed as in Fig. 2. In this case, the quadratic terms of \mathbf{m}_b are:

$$(m_1 - m_2)^2 + (m_1 - m_3)^2 + (m_2 - m_4)^2 + (m_3 - m_4)^2 + (m_4 - m_5)^2, \quad (9)$$

and therefore we have:

$$\mathbf{C}_b = \begin{pmatrix} 2 & -1 & -1 & 0 & 0 \\ -1 & 2 & 0 & -1 & 0 \\ -1 & 0 & 2 & -1 & 0 \\ 0 & -1 & -1 & 3 & -1 \\ 0 & 0 & 0 & -1 & 1 \end{pmatrix}. \quad (10)$$

In general, it is easy to compute the matrix \mathbf{C}_b for any given image. Notice that it is always a positive semidefinite matrix, and thus it is associated to a (singular) normal distribution.

The VGPMIL-PR-I formulation. To introduce the new coupling term $C(\mathbf{m}_b)$ in the MIL formulation, we modify Eq. (6) and define:

$$p(\mathbf{m}|\mathbf{f}) \propto \prod_b C(\mathbf{m}_b) \cdot \mathcal{N}(\mathbf{m}_b|\mathbf{f}_b, I). \quad (11)$$

Notice that the probability of a configuration \mathbf{m}_b is proportional to the coupling term $C(\mathbf{m}_b)$, which favors smoothness across labels of neighboring patches. Decisively, since both $C(\mathbf{m}_b)$ and $\mathcal{N}(\mathbf{m}_b|\mathbf{f}_b, I)$ only contain (the exponential of) quadratic terms in \mathbf{m}_b , the new distribution can be written as a Gaussian:

$$p(\mathbf{m}|\mathbf{f}) = \prod_b \mathcal{N}(\mathbf{m}_b|\Sigma_b \mathbf{f}_b, \Sigma_b), \quad (12)$$

with $\Sigma_b = (\lambda \mathbf{C}_b + I)^{-1}$. Notice that this new formulation provides a generalization of VGPMIL-PR. Namely, when $\lambda \rightarrow 0$, we have that $\Sigma_b \rightarrow I$ and we recover Eq. (6).

In summary, the proposed model is given by Eqs. (3), (4), (12), and (7). Notice also that, instead of FITC, in Eq. (4) we leverage the more recent sparse GP approach introduced in [26]. Basically, this means that $\mathcal{K} = \mathbf{K}_{XX} - \mathbf{K}_{XZ} \mathbf{K}_{ZZ}^{-1} \mathbf{K}_{ZX}$ in Eq. (4). Our method will be referred to as VGPMIL-PR-I (I denotes Ising). The probabilistic graphical model is depicted in Fig. 1(b).

2.4. Variational inference

In order to make inference in the proposed model, we need to compute the posterior distribution $p(\mathbf{u}, \mathbf{f}, \mathbf{m}|\mathbf{y})$. However, this is not analytically tractable due to the definition of the bag likelihood in Eq. (7), which depends on the sign of the m_i ’s. Following [17,18], we leverage standard mean-field variational inference (VI) theory [27, Section 10.1.1] to calculate an approximate posterior distribution that factorizes as

$$q(\mathbf{u}, \mathbf{f}, \mathbf{m}) = q(\mathbf{u})p(\mathbf{f}|\mathbf{u})q(\mathbf{m}). \quad (13)$$

Applying the well-known mean-field VI update equation [27, Eq. (10.9)], we have that $q(\mathbf{u})$ and $q(\mathbf{m})$ can be iteratively computed as

$$q(\mathbf{u}) = \mathcal{N}(\mathbf{u}|\boldsymbol{\mu}^u, \boldsymbol{\Sigma}^u), \quad (14)$$

$$q(\mathbf{m}) \propto \prod_b p(y_b|\mathbf{m}_b) \mathcal{N}(\mathbf{m}_b|\boldsymbol{\mu}^{m_b}, \boldsymbol{\Sigma}_b), \quad (15)$$

where

$$\boldsymbol{\Sigma}^u = (\mathbf{K}_{ZZ}^{-1} + \mathbf{K}_{ZZ}^{-1} \mathbf{K}_{ZX} \boldsymbol{\Sigma} \mathbf{K}_{XZ} \mathbf{K}_{ZZ}^{-1})^{-1}, \quad (16)$$

$$\boldsymbol{\mu}^u = \boldsymbol{\Sigma}^u \mathbf{K}_{ZZ}^{-1} \mathbf{K}_{ZX} \mathbb{E}_{q(\mathbf{m})}(\mathbf{m}), \quad (17)$$

and

$$\boldsymbol{\mu}^{m_b} = \boldsymbol{\Sigma}_b \mathbf{K}_{bZ} \mathbf{K}_{ZZ}^{-1} \boldsymbol{\mu}^u. \quad (18)$$

Here we are writing $\boldsymbol{\Sigma}$ for the $N \times N$ block-diagonal matrix that contains all the $\boldsymbol{\Sigma}_b$ ’s, $b \in \mathcal{B}$. Also, we are writing \mathbf{K}_{bZ} for the $|b| \times M$ matrix of covariances between \mathbf{X}_b and \mathbf{Z} . Very importantly, notice that these update rules generalize those derived in [18] for VGPMIL-PR. Namely, when $\lambda \rightarrow 0$, we have that $\boldsymbol{\Sigma}_b, \boldsymbol{\Sigma} \rightarrow I$, and then Eqs. (14)–(18) match eqs.(15)–(19) in [18].

All the computations involved in Eqs. (14)–(18) are straightforward, except for $\mathbb{E}_{q(\mathbf{m})}(\mathbf{m})$. Indeed, each $q(\mathbf{m}_b)$ is a multivariate Gaussian truncated to $(-\infty, 0)^{|b|}$ (or $\mathbb{R}^{|b|} \setminus (-\infty, 0)^{|b|}$, depending on whether $y_b = 0$ or $y_b = 1$, respectively). It is well-known that the expectation of a truncated multivariate Gaussian cannot be obtained in closed-form [28]. Notice that this is not an issue for VGPMIL-PR [18], where the absence of Ising terms implies dealing with *univariate* Gaussians, whose expectations can be analytically computed.

To overcome the problem, we first tried to leverage numerical methods proposed in [29] to approximate the expectation for the multivariate truncated case. However, these methods proved computationally too expensive to be integrated within our iterative calculation of $q(\mathbf{u})$ and $q(\mathbf{m})$. Therefore, we decided to approximate the multivariate Gaussian $\mathcal{N}(\mathbf{m}_b|\boldsymbol{\mu}^{m_b}, \boldsymbol{\Sigma}_b)$ by the factorized $\mathcal{N}(\mathbf{m}_b|\boldsymbol{\mu}^{m_b}, \text{diag}(\boldsymbol{\Sigma}_b))$ and utilize the expression for one-dimensional truncated Gaussians. Although such approximation reduces the influence of the Ising correlation at this specific computation, notice that the coupling terms, which are included in $\boldsymbol{\Sigma}_b$, affect the update equations in more places across Eqs. (16)–(18).

Specifically, the expression for each $\mathbb{E}_{q(\mathbf{m}_b)}(\mathbf{m}_b)$ is as follows. For bags with $y_b = 0$, we have that each $q(m_i)$, $i \in b$, is a univariate normal distribution $\mathcal{N}((\boldsymbol{\mu}^{m_b})_i, (\boldsymbol{\Sigma}_b)_{ii})$ truncated to $(-\infty, 0)$. The expectation of such a distribution is well-known and can be obtained in closed-form [30]:

$$E_i = \mu_i - \frac{\phi(\mu_i/\sigma_i)}{1 - \Phi(\mu_i/\sigma_i)} \sigma_i, \quad (19)$$

where ϕ and Φ are, respectively, the density and cumulative distribution functions of a standard Gaussian $\mathcal{N}(0, 1)$ (recall that both are efficiently implemented in standard software packages such as Python’s Scipy). We have also abbreviated $\mu_i = (\boldsymbol{\mu}^{m_b})_i$ and $\sigma_i = \sqrt{(\boldsymbol{\Sigma}_b)_{ii}}$. For bags with $y_b = 1$, we proceed analogously to [19] to obtain the normalization constant Z of the distribution of interest (that is, the factorized Gaussian $\mathcal{N}(\mathbf{m}_b|\boldsymbol{\mu}^{m_b}, \text{diag}(\boldsymbol{\Sigma}_b))$ truncated to $\mathbb{R}^{|b|} \setminus (-\infty, 0)^{|b|}$). Then, the expectation of each $q(m_i)$, $i \in b$, is given by:

$$\mathbb{E}_{q(m_i)}(m_i) = \frac{\mu_i - (1 - Z)E_i}{Z}, \quad (20)$$




Fig. 3. Two examples of cancerous (left) and non-cancerous (right) patches in the SICAPv2 test set.

Algorithm 1 Training procedure for VGPMIL-PR-I.

Input : Bags $\mathbf{X} = \{\mathbf{X}_b\}_{b \in \mathcal{B}}$ and bag labels $\mathbf{y} = \{y_b\}_{b \in \mathcal{B}}$.
 Calculate the matrices \mathbf{C}_b that account for the instance correlation inside each bag $b \in \mathcal{B}$, recall Eq. (8) and the example at Eq. (10).
 Initialize GP kernel parameters and inducing points locations, as well as the posterior distributions $q(\mathbf{u})$ and $q(\mathbf{m})$. Details on initializations in the text.
foreach iteration $t = 1, \dots, T$ **do**
 Update $q(\mathbf{u})$, using Eqs. (14), (16) and (17).
 Update $q(\mathbf{m})$ and obtain $\mathbb{E}_{q(\mathbf{m})}(\mathbf{m})$, using Eqs. (15), (18), (19), and (20).
Output: Posterior distributions $q(\mathbf{u})$ and $q(\mathbf{m})$.

where $Z = 1 - \prod_{i \in b} (1 - \Phi(\mu_i/\sigma_i))$, E_i is given by Eq. (19), and we are again abbreviating $\mu_i = (\boldsymbol{\mu}^{\mathbf{m}^b})_i$, and $\sigma_i = \sqrt{(\boldsymbol{\Sigma}_b)_{ii}}$.

The full training algorithm is summarized in Algorithm 1. It is an iterative process that alternates the updates between $q(\mathbf{u})$ and $q(\mathbf{m})$. Details on the GP kernel and initializations used in this work are provided in Section 3.1. The code for the proposed method will be publicly available upon acceptance of the paper.

2.5. Making predictions

Suppose we are given a new bag $\mathbf{X}^* = \{\mathbf{x}_i^*\}_{i \in b^*}$. As explained at the end of Section 2.1, we are interested in both instance-level and bag-level predictions. For this, we first need to compute the predictive distributions over \mathbf{m}^* .

By using the learned posterior $q(\mathbf{u})$ along with $p(\mathbf{f}|\mathbf{u})$, we can obtain the joint distribution over \mathbf{f}^* :

$$p(\mathbf{f}^*) = \int p(\mathbf{f}^*|\mathbf{u})p(\mathbf{u})d\mathbf{u} = \mathcal{N}(\mathbf{f}^*|\boldsymbol{\mu}^*, \mathbf{S}^*), \quad (21)$$

with $\boldsymbol{\mu}^*$ and \mathbf{S}^* given by the standard sparse GP predictions:

$$\boldsymbol{\mu}^* = \mathbf{K}_{*Z}\mathbf{K}_{ZZ}^{-1}\boldsymbol{\mu}^u, \quad \mathbf{S}^* = \mathbf{K}_{**} - \mathbf{K}_{*Z}\mathbf{K}_{ZZ}^{-1}(\mathbf{K}_{ZZ} - \boldsymbol{\Sigma}^u)\mathbf{K}_{ZZ}^{-1}\mathbf{K}_{Z*}. \quad (22)$$

Here, $\boldsymbol{\mu}^u$ and $\boldsymbol{\Sigma}^u$ are the parameters learned during training, recall Eq. (16) and (17). Naturally, the subscript $*$ in the kernel matrices \mathbf{K} indicates that we are using the new bag \mathbf{X}^* . Then, since the distribution $p(\mathbf{m}|\mathbf{f})$ is also Gaussian, recall Eq. (12), we can compute the joint distribution over \mathbf{m}^* in closed-form:

$$p(\mathbf{m}^*) = \int p(\mathbf{m}^*|\mathbf{f}^*)p(\mathbf{f}^*)d\mathbf{f}^* = \mathcal{N}(\mathbf{m}^*|\boldsymbol{\mu}_m^*, \mathbf{S}_m^*), \quad (23)$$

with $\boldsymbol{\mu}_m^*$ and \mathbf{S}_m^* given by

$$\boldsymbol{\mu}_m^* = \boldsymbol{\Sigma}_* \cdot \boldsymbol{\mu}^*, \quad \mathbf{S}_m^* = \boldsymbol{\Sigma}_* + \boldsymbol{\Sigma}_* \cdot \mathbf{S}^* \cdot \boldsymbol{\Sigma}_*^T. \quad (24)$$

Here, $\boldsymbol{\mu}^*$ and \mathbf{S}^* are given by Eq. (22), and $\boldsymbol{\Sigma}_*$ is the matrix that accounts for correlation among instances in the test bag \mathbf{X}^* , which is defined analogously to the training case, recall matrix $\boldsymbol{\Sigma}_b$ in Eq. (12).

Once we have the joint distribution over \mathbf{m}^* , the instance-level and bag-level predictions are given as:

$$p(h_i^* = 1) = p(m_i^* > 0) = \Phi\left(\frac{(\boldsymbol{\mu}_m^*)_i}{\sqrt{(\mathbf{S}_m^*)_{ii}}}\right), \quad (25)$$

$$p(y^* = 1) = 1 - \int_{\mathbf{m}^* \in (-\infty, 0)^{b^*}} p(\mathbf{m}^*)d\mathbf{m}^*. \quad (26)$$

Notice that the integral in Eq. (26) can be computed efficiently with the cumulative distribution function of a multivariate Gaussian, which is also available in most standard statistical packages, such as Python's Scipy.

Interestingly, these predictions generalize those obtained in VGPMIL-PR [18]. Indeed, if we do not consider correlation among instances in the test bag, i.e. $\boldsymbol{\Sigma}_* = \mathbf{I}$, then Eqs. (25) and (26) match those in [18] (last two equations before Section 3.5). Finally, although here we have detailed how to make predictions for a complete previously unseen bag \mathbf{X}^* , the same process can be applied to make predictions on previously unseen individual instances \mathbf{x}^* (patches).

3. Experiments

In this section we thoroughly evaluate VGPMIL-PR-I in a real-world problem of prostate cancer detection. The experimental framework, including data, metrics and baselines, is explained in Section 3.1. The results are discussed in Section 3.2. Finally, in Section 3.3 we evaluate our method in a much larger prostate cancer detection dataset: the well-known PANDA challenge.

3.1. Experimental framework

Data description. In this paper we focus on the problem of prostate cancer detection. However, notice that the algorithm can be applied for any other type of cancer (and more generally, for any other type of image). Prostate cancer is the most commonly occurring cancer in men, and the second most commonly occurring cancer overall, according to the latest 2020 statistics on age-standardized incidence rate from the World Health Organization (WHO) Global Cancer Observatory [31]. We will use the prostate cancer database presented in [32], which is called SICAPv2 and is publicly available. Although this database includes information on the Gleason score, which is used to evaluate the severity of the disease, in this work we will focus on the binary task of presence/absence of cancer.

We use the original partition of the dataset, which contains 95 training and 31 test WSIs, respectively. These very large images are split in 512×512 patches, resulting in a total amount of 15132 patches for training and 5246 for testing. Following the MIL paradigm, for the training set we only use binary labels benign/malign at the level of images (bags), but we do not have information at the level of patches (instances). In order to evaluate the predictive performance at instance-level, we do have labels for the patches in the test set. The amount of cancerous (resp. non-cancerous) images for the train set is 70 (resp. 25). For the test set, it is 25 (resp. 6). For illustration purposes, a couple of cancerous and non-cancerous patches are shown in Fig. 3. In order to train our model, each patch is represented through a 128-dimensional feature vector extracted in previous work [7].

Baselines and metrics. Since our model is framed in the field of probabilistic GP-based MIL methods, we compare with the two most popular approaches VGPMIL [17] and VGPMIL-PR [18], which

Table 1

Predictive performance at the level of patches (instances). In bold, we highlight the values of λ for which VGPMIL-PR-I gets better (or equal) performance than both baselines in all the metrics. The results are the mean and standard deviation over five independent runs.

	λ	Accuracy	Precision	Recall	F1-score
VGPMIL	–	92.22 ± 0.00	96.40 ± 0.00	92.29 ± 0.00	94.30 ± 0.00
VGPMIL-PR	–	92.38 ± 0.03	96.44 ± 0.06	92.48 ± 0.06	94.42 ± 0.02
VGPMIL-PR-I	0.1	92.94 ± 0.05	96.24 ± 0.15	93.52 ± 0.12	94.86 ± 0.04
	0.5	93.85 ± 0.04	97.17 ± 0.19	93.90 ± 0.14	95.51 ± 0.02
	1.0	94.58 ± 0.03	97.32 ± 0.06	94.83 ± 0.04	96.06 ± 0.02
	5.0	95.11 ± 0.06	97.74 ± 0.14	95.18 ± 0.09	96.44 ± 0.04
	10.0	95.03 ± 0.14	97.72 ± 0.16	95.09 ± 0.06	96.39 ± 0.10

Table 2

Predictive performance at the level of images (bags). In bold, we highlight the values of λ for which VGPMIL-PR-I gets better (or equal) performance than both baselines in all the metrics. The results are the mean and standard deviation over five independent runs.

	λ	Accuracy	Precision	Recall	F1-score
VGPMIL	–	83.87 ± 0.00	83.33 ± 0.00	100.00 ± 0.00	90.91 ± 0.00
VGPMIL-PR	–	90.32 ± 0.00	89.29 ± 0.00	100.00 ± 0.00	94.34 ± 0.00
VGPMIL-PR-I	0.1	93.55 ± 0.00	92.59 ± 0.00	100.00 ± 0.00	96.15 ± 0.00
	0.5	93.55 ± 0.00	92.59 ± 0.00	100.00 ± 0.00	96.15 ± 0.00
	1.0	90.32 ± 0.00	92.31 ± 0.00	96.00 ± 0.00	94.12 ± 0.00
	5.0	87.10 ± 0.00	95.65 ± 0.00	88.00 ± 0.00	91.67 ± 0.00
	10.0	83.87 ± 0.00	95.45 ± 0.00	84.00 ± 0.00	89.36 ± 0.00

Table 3

Confusion matrices obtained at the level of images for the compared methods.

		VGPMIL		VGPMIL-PR	
		Neg.	Pos.	Neg.	Pos.
Actual	Neg.	1	5	3	3
	Pos.	0	25	0	25

		VGPMIL-PR-I									
		$\lambda = 0.1$		$\lambda = 0.5$		$\lambda = 1$		$\lambda = 5$		$\lambda = 10$	
		Neg.	Pos.	Neg.	Pos.	Neg.	Pos.	Neg.	Pos.	Neg.	Pos.
Actual	Neg.	4	2	4	2	4	2	5	1	5	1
	Pos.	0	25	0	25	1	24	3	22	4	21

were reviewed in Section 2.2. For a fair comparison, the parameters used for the baselines are the same as those used for our method (described in next paragraph). For those parameters that do not have an analogous in our method (e.g. the initialization of $q(\mathbf{y})$ in VGPMIL), we use the default values proposed in the original papers. To evaluate the performance of the compared methods we use four metrics: accuracy, precision, recall and F1-score (which provides a trade-off between precision and recall). For all the metrics, we use the standard implementations in the popular Python scikit-learn library [33].

Experimental details. For the underlying GPs, in this work we use the well-known squared exponential kernel [34], i.e. $\kappa(\mathbf{x}, \mathbf{y}) = \gamma \cdot \exp(-\|\mathbf{x} - \mathbf{y}\|^2 / (2\ell^2))$. Following [17,18], we use standard values for the kernel hyperparameters, i.e. $\gamma = 1$ and ℓ equals the square root of the number of features of \mathbf{x}, \mathbf{y} (in this work we set $\ell = 11 \approx \sqrt{128}$). The number of inducing points is set to $M = 200$, and their locations are initialized through K-means clustering as in previous work [17,18] (namely, 100 of them are obtained by doing clustering on the patches that belong to the positive images, and the other 100 on the patches that belong to the negative ones). The number of iterations is set to $T = 200$, which was enough to achieve convergence in practice. The expectation of the posterior distribution $\mathbb{E}_{q(\mathbf{m})}(\mathbf{m})$ is initialized with a standard Gaussian for each instance independently. Notice that the initialization of $q(\mathbf{u})$ is irrelevant since it gets updated first in Algorithm 1. As for the value of λ , which regulates the strength of the Ising correlation (recall Eq. (8)), we will analyze five different values in the experiments, $\lambda \in \{0.1, 0.5, 1.0, 5.0, 10.0\}$. This will allow us to empirically illustrate the effect of λ .

Table 4

Computational cost for training and testing the compared methods (in seconds). We are using 200 iterations in all cases, recall the experimental details in Section 3.1. The results are the mean and standard deviation over five independent runs.

	λ	Training time	Testing time
VGPMIL	–	15.45 ± 0.53	3.14 ± 0.16
VGPMIL-PR	–	11.78 ± 0.75	2.63 ± 0.12
VGPMIL-PR-I	0.1	12.21 ± 0.26	2.23 ± 0.09
	0.5	12.18 ± 0.57	2.21 ± 0.07
	1.0	11.79 ± 0.44	2.24 ± 0.15
	5.0	11.80 ± 0.74	2.22 ± 0.08
	10.0	11.64 ± 0.91	2.26 ± 0.09

3.2. Experimental results

In this section we evaluate the performance of the compared methods on the aforementioned prostate cancer problem. We analyze eight different research questions, which are discussed in the following paragraphs.

Predictions at the level of instances (patches). Although they only use bag labels for training, the compared methods can make predictions at the level of instances, recall Section 2.5. This is important to determine more precisely in which region (patch) the cancer is present.

Table 1 shows the results when making predictions at patch level. We observe that VGPMIL-PR-I outperforms both baselines in all the metrics for four out of five values of λ (and for $\lambda = 0.1$, the baselines are only better in terms of precision). We also appreciate that VGPMIL-PR-I results are robust across different runs, obtaining low values of standard deviation. This stability is important for real-world applications,




Fig. 4. Patch-level predictions inside each one of the six negative (non-cancerous) WSIs in the test set. Each column is an image (the header is the image identifier in the SICAPv2 dataset). The rows refers to the three compared methods. Each subplot has red/green axis depending on whether the image is correctly classified or not by that method. The blue bars inside the subplots represent the probability of cancer for the different patches inside the image (for ease of visualization, they are sorted in increasing order). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

where one wants to avoid high sensitivity to random initializations. Finally, notice that, as argued in [19], we also observe that VGPMIL-PR (slightly) outperforms VGPMIL.

Predictions at the level of bags (images). Table 2 shows the results when making predictions at the level of images. We observe that VGPMIL-PR-I outperforms both baselines in all the metrics when $\lambda \in \{0.1, 0.5\}$. However, when λ becomes larger, the results of VGPMIL-PR-I get worse. This fact can be explained theoretically because, whereas having low-to-moderate correlation among patches can be helpful, having strong ones tends to make the predictions too homogeneous, damaging the bag-level prediction (which takes into account the correlation among patches). Indeed, in the next research question we analyze with greater detail how λ is affecting the predictions on cancerous and non-cancerous images separately, which will provide additional insights. Finally, similar to the patch-level results, we observe that VGPMIL-PR obtains better results than VGPMIL, as expected. Also, the predictions at the level of images are very stable across runs (notice the zero standard deviation).

Analyzing the confusion matrices at bag level. Here we analyze more in detail the results presented in the previous paragraph (i.e. at the level of images). Table 3 shows the confusion matrices for the compared methods. Notice that both baselines classify all the 25 positive (cancerous) images correctly. However, the difficulties arise at the non-cancerous images. This happens because of the nature of the MIL problem: as soon as a few patches obtain a non-negligible probability of cancer, the image will be likely predicted as cancerous (recall that the MIL formulation establishes that a bag has positive class as soon as one instance inside the bag has positive class).

Interestingly, the Ising model can help to avoid isolated false positive predictions on the patch-level (which lead to false positive bag predictions) by using the patch-level correlation. This is reflected in Table 3, where we observe that increasingly more negative images are predicted correctly as λ gets higher. In contrast, notice that strong correlation damage the performance in the positive class, as they penalize the appearance of positive patches (which would break the homogeneity of the bag, where most patches do not contain cancer). Therefore, we conclude that instance correlation is beneficial when

used with a low-to-medium intensity. For instance, in this application $\lambda = 0.5$ is the best performing value, and it will be the one used by default in the sequel.

Analyzing the instance-level predictions for negative bags. In the previous paragraph, we have explained that negative images are incorrectly classified because a few patches inside them get classified as positive. Here we provide a visualization to support this. Fig. 4 shows how the patch-level predictions are distributed inside the six negative images available in the test set. We observe that the amount of patches with a non-negligible probability of cancer gets reduced as we move from VGPMIL to VGPMIL-PR, and then to VGPMIL-PR-I. This translates into better performance at bag-level (observe that the amount of green-axis subplots increases in the same sequence VGPMIL \rightarrow VGPMIL-PR \rightarrow VGPMIL-PR-I). Notice that the improvement from VGPMIL to VGPMIL-PR is larger than from VGPMIL-PR to VGPMIL-PR-I. This may be due to the simplification that was introduced when computing the expectation of the truncated multivariate Gaussian in VGPMIL-PR-I, recall the third-to-last paragraph in Section 2.4.

Visualizing the predictions. In the last paragraph, we have analyzed how the patch-level predictions are distributed inside non-cancerous images *quantitatively*. Indeed, Fig. 4 represents each patch through a bar. However, this hampers the *qualitative* visualization from a medical viewpoint. Here we focus on such qualitative assessment by visualizing the predictions obtained for image 16B0028138, see Fig. 5. We have chosen this image because it illustrates best the effect of the coupling term. Notice that, thanks to these terms, the proposed VGPMIL-PR-I manages to keep all patches with a probability closer to zero than VGPMIL and VGPMIL-PR. As a consequence, VGPMIL-PR-I is the only method that correctly classifies this image as non-cancerous, recall Fig. 4. For the other two methods, there are some patches that trigger the image prediction to be cancerous.

An explicit analysis on the role of λ . The hyperparameter λ is at the core of the novel VGPMIL-PR-I. It was introduced in the probabilistic model to regulate the strength of the coupling term, recall Eq. (8). This role has been confirmed *indirectly* in Table 3: when λ gets higher, the predictive performance improves for negative bags and degrades for positive ones. This can be explained because a higher λ homogenizes

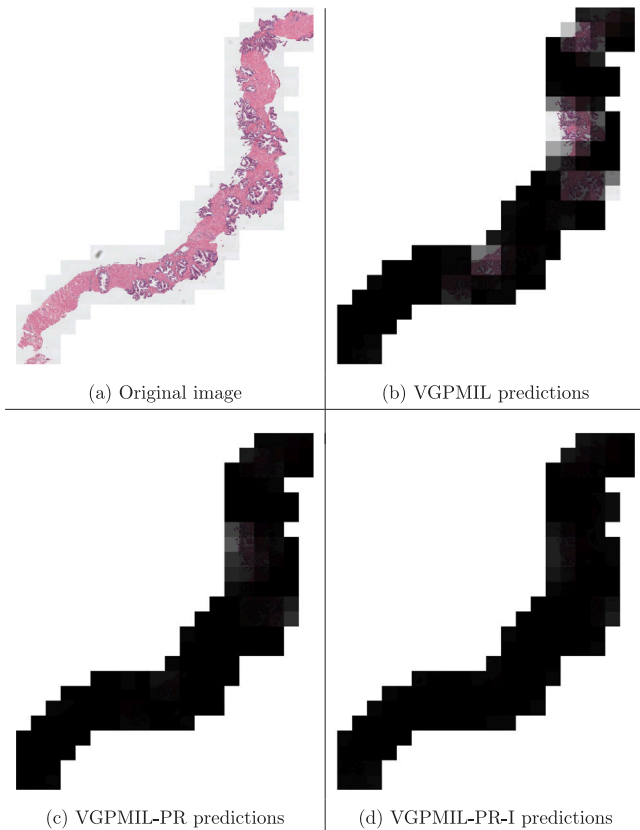


Fig. 5. Patch level predictions obtained by the compared methods for image 16B0028138, which is non-cancerous. The original image is shown in (a). For predictions (b)–(d), the brightness of the patch is proportional to the probability of cancer (the brighter, the more probability).




Fig. 6. Variability in the patch-level predictions inside images. In VGPMIL-PR-I, the variability decreases as the strength of the Ising terms, given by λ , increases. VGPMIL-PR, which does not include Ising terms (i.e. $\lambda = 0$), gets larger variability. As explained in the text, the variability inside an image is measured as the standard deviation of the probability of cancer for all the patches inside that image.

the patch-level predictions and difficulties the appearance of positive patches. Here we perform a more *direct* measure to gain insights into the role of λ . Specifically, we define the “variability inside a bag” as the standard deviation of the probability of cancer for all the patches inside that bag. Therefore, this metric measures the dispersion in the patch-level predictions obtained inside a bag. Fig. 6 shows the evolution of this metric for VGPMIL-PR-I as λ increases. As theoretically expected, the metric decreases as λ gets higher. Also, notice that the metric

value for VGPMIL-PR is higher. This is explained because VGPMIL-PR does not incorporate Ising correlation, i.e. $\lambda = 0$. The value for VGPMIL is even higher, 0.30, and it is not included in Fig. 6 for ease of visualization. This greater value is probably due to the additional approximations that VGPMIL involves, which deepens the independence among patches.

Computational cost. Finally, we report the computational training and testing time for the compared methods, see Table 4. The results are in the same order of magnitude in all cases, which justifies the practical utility of the novel VGPMIL-PR-I, which obtained better predictive performance, recall Tables 1 and 2. In fact, VGPMIL-PR-I is slightly faster than VGPMIL, since the Jaakola bound approximation leveraged in the latter introduces additional parameters ξ to be estimated. As theoretically expected, the computational cost of VGPMIL-PR-I and VGPMIL-PR is analogous, since the update equations for the former are just a generalization of those for the latter, recall Sections 2.4 and 2.5. Finally, notice that the value of λ does not affect the computational cost of VGPMIL-PR-I, as λ only regulates the intensity of the Ising terms (but it does not introduce any additional computation).

Comparison to other related MIL approaches. So far we have focused on the comparison of VGPMIL-PR-I with VGPMIL and VGPMIL-PR. Since VGPMIL-PR-I builds on the same type of GP-based modeling, this is the most meaningful comparison in order to evaluate our main contribution (the Ising term to account for correlations among patches). However, to provide a wider perspective, it is interesting to compare the novel VGPMIL-PR-I to other state-of-the-art and popular families of MIL methods. We consider three families: attention-based methods, where the two algorithms proposed in [35] are the most popular approaches; MIL methods based on pseudo-labels such as the recent [7]; and classical pooling/aggregation methods such as the mean aggregation [36]. These will be referred to as Att-MIL, Gated-Att-MIL, PS-MIL and Mean-Agg, respectively.

Let us discuss the results both at instance (patch) and bag (image) levels. For the former, notice that the formulation of attention-based methods (Att-MIL and Gated-Att-MIL) and classical aggregation methods (Mean-Agg) do not allow for making predictions at instance level in a natural way. Namely, the instance-level labels are not modeled explicitly in this type of methods, and this is precisely one of their main limitations [37]. Compared to PS-MIL, which does model instance labels explicitly, the novel VGPMIL-PR-I achieves higher predictive performance (95.11 vs 85.01 in accuracy and 96.44 vs 88.05 in F1-Score). Regarding bag-level performance, the results are shown in Table 5. We observe that the best results are obtained by attention-based methods Att-MIL and Gated-Att-MIL, followed by the novel VGPMIL-PR-I.

In conclusion, we observe that the results are quite different depending on the nature of the model and the user requirements. If one is interested in predictions at patch level, then the novel VGPMIL-PR-I is the best choice. However, if one is only interested in image level performance, then attention-based approaches are the best option for this data. Indeed, we hypothesize that the performance of attention-based approaches could be even enhanced by leveraging correlation among patches in a similar way to VGPMIL-PR-I. This is an interesting line of future research, see Section 4.

3.3. Evaluation on a larger dataset: PANDA

The SICAPv2 dataset used so far is of medium-size (total amount of 126 WSIs, leading to 20,378 patches; recall Section 3.1). This has allowed us to carry out a very detailed analysis of the results. In this section we show that the novel VGPMIL-PR-I also performs well on larger datasets, such as the well-known PANDA set. Although scalability is not an issue from a theoretical perspective, since the model is based on sparse GPs, it is important to verify it in practice.

PANDA also tackles the problem of prostate cancer detection, and was presented at the MICCAI 2020 conference as a challenge.² Since the

² <https://panda.grand-challenge.org/>

Table 5

Comparison with other related MIL methods which are not based on the GP modeling. The predictive performance at the level of images (bags) is shown. The results are the mean and standard deviation over five independent runs. The algorithm PS-MIL was run only once because of its high computational training cost.

	Accuracy	Precision	Recall	F1-score
VGPMIL-PR-I	93.55 ± 0.00	92.59 ± 0.00	100.00 ± 0.00	96.15 ± 0.00
Att-MIL	96.80 ± 0.00	96.20 ± 0.00	100.00 ± 0.00	98.00 ± 0.00
Gated-Att-MIL	96.80 ± 0.00	96.20 ± 0.00	100.00 ± 0.00	98.00 ± 0.00
Mean-Agg	87.10 ± 0.00	95.70 ± 0.00	88.00 ± 0.00	91.70 ± 0.00
PS-MIL	90.32 ± NA	89.28 ± NA	100.00 ± NA	94.33 ± NA

Table 6

Predictive performance at the level of images (bags) in the PANDA dataset. The results are the mean and standard deviation over five independent runs. The algorithm PS-MIL was run only once because of its high computational training cost.

	Accuracy	Precision	Recall	F1-score
VGPMIL	74.87 ± 0.08	74.25 ± 0.06	99.77 ± 0.00	85.13 ± 0.04
VGPMIL-PR	90.64 ± 0.06	90.97 ± 0.05	96.60 ± 0.04	93.70 ± 0.04
VGPMIL-PR-I	92.57 ± 0.14	95.42 ± 0.24	94.22 ± 0.11	94.82 ± 0.09
Att-MIL	90.92 ± 0.01	94.17 ± 0.01	93.43 ± 0.01	93.79 ± 0.01
Gated-Att-MIL	91.70 ± 0.01	94.31 ± 0.01	94.40 ± 0.01	94.34 ± 0.01
Mean-Agg	88.09 ± 0.00	91.57 ± 0.01	92.27 ± 0.01	91.91 ± 0.00
PS-MIL	88.36 ± NA	87.99 ± NA	97.11 ± NA	92.33 ± NA

test set of PANDA is not publicly available, we use the train/test split proposed in [38], where each split follows the overall class proportions. Namely, the dataset used here features a total amount of 10,503 WSIs, which leads to 1,107,931 patches. Notice that this is much larger than SICAPv2 (83 times larger in terms of WSIs).

Table 6 shows the predictive performance at image level, an aspect where attention-based methods stood out in the previous dataset. In this case, we observe that VGPMIL-PR-I obtains consistently better results. Additionally, as outlined in the last research question in Section 3.2, VGPMIL-PR-I is able to provide instance-level predictions, which is not the case for attention-based models. We conclude that, for the PANDA dataset, the proposed method is the best choice in comparison to the other tested approaches.

4. Conclusions, limitations and future work

In this work we have introduced VGPMIL-PR-I, a novel MIL methodology that incorporates instance label correlation through a coupling term inspired by the Ising model. VGPMIL-PR-I is a generalization of another probabilistic MIL method, whose formulation is theoretically recovered when the influence of the Ising term converges to zero. In the experimental section, we have shown that VGPMIL-PR-I outperforms other related state-of-the-art probabilistic MIL approaches in two real-world problems of prostate cancer detection, effectively reducing false positive bag predictions and providing instance-level predictions. We have also provided different visualizations to better understand the behavior of the proposed model, specially the influence of the new coupling term.

As discussed along the paper, our model presents several limitations which we summarize next. Firstly, we needed to introduce a diagonal approximation to compute the expectation of the truncated multivariate Gaussian in VGPMIL-PR-I, recall Section 2.4. This is probably reflected in the empirical performance, as the improvement when moving from VGPMIL to VGPMIL-PR is generally larger than when moving from VGPMIL-PR to VGPMIL-PR-I. Secondly, we have observed that the behavior of VGPMIL-PR-I depends on the value of λ , which regulates the strength of the coupling term. Although we have discussed the role of λ and tested different values, it remains a hyperparameter that has to be found empirically using the validation set. We believe that its value (or even distribution over it) could be estimated from the data by introducing λ in the probabilistic modeling. Even more, λ could be estimated per image, since the level of correlation could be image-dependent. Thirdly, we have observed that the image-level performance of VGPMIL-PR-I is not generally better than that of

attention-based methods. This is probably due to the different nature of the models. Indeed, the explicit modeling of instance label in GP-based models, which allows them to provide instance-level predictions, may come at the cost of less accurate bag-level predictions.

In addition to the aforementioned ideas, this work opens other future research lines. First, seeing the performance boost obtained in GP-based methods through the novel coupling term, and taking into account the good results of attention-based methods in bag-level prediction, it is very interesting to explore the modeling of instance label correlations in the context of attention-based methods. Second, notice that we are using mean-field variational inference to estimate the model parameters in VGPMIL-PR-I. A promising alternative is to estimate them by directly optimizing the evidence lower bound (ELBO). Finally, although we have focused on modeling correlation between neighboring patches in histopathological images, we expect that the ideas behind our proposal can boost further research in MIL, by exploiting the particular structure of the data used in different applications.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The data is publicly available, as explained in the paper.

Acknowledgments

This work has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska Curie grant agreement No 860627 (CLARIFY Project), from the Spanish Ministry of Science and Innovation under project PID2019-105142RB-C22, and by the University of Granada and FEDER/Junta de Andalucía under project B-TIC-324-UGR20 (Proyectos de I+D+i en el marco del Programa Operativo FEDER Andalucía). PMA has been supported by the Margarita Salas postdoctoral fellowship (Spanish Ministry of Universities with Next Generation EU funds).

References

- [1] M.-A. Carbonneau, V. Cheplygina, E. Granger, G. Gagnon, Multiple instance learning: A survey of problem characteristics and applications, *Pattern Recognit.* 77 (2018) 329–353.
- [2] M. Cui, D.Y. Zhang, Artificial intelligence and computational pathology, *Lab. Invest.* 101 (4) (2021) 412–422.
- [3] S. Huang, Z. Liu, W. Jin, Y. Mu, Bag dissimilarity regularized multi-instance learning, *Pattern Recognit.* 126 (2022) 108583.
- [4] M.K.K. Niazi, A.V. Parwani, M.N. Gurcan, Digital pathology and artificial intelligence, *Lancet Oncol.* 20 (5) (2019) e253–e261.
- [5] K.P. Murphy, *Machine Learning: A Probabilistic Perspective*, MIT Press, 2012.
- [6] M. López-Pérez, M. Amgad, P. Morales-Álvarez, P. Ruiz, L.A. Cooper, R. Molina, A.K. Katsaggelos, Learning from crowds in digital pathology using scalable variational Gaussian processes, *Sci. Rep.* 11 (1) (2021) 1–9.
- [7] A. Schmidt, J. Silva-Rodríguez, R. Molina, V. Naranjo, Efficient cancer classification by coupling semi supervised and multiple instance learning, *IEEE Access* 10 (2022) 9763–9773.
- [8] G. Campanella, M.G. Hanna, L. Geneslaw, A. Mirafior, V. Werneck Krauss Silva, K.J. Busam, E. Brogi, V.E. Reuter, D.S. Klimstra, T.J. Fuchs, Clinical-grade computational pathology using weakly supervised deep learning on whole slide images, *Nature Med.* 25 (8) (2019) 1301–1309.
- [9] A. Schmidt, J. Silva-Rodríguez, R. Molina, V. Naranjo, Coupling semi-supervised and multiple instance learning for histopathological image classification, *IEEE Access* (2022).
- [10] Y. Wu, A. Schmidt, E. Hernández-Sánchez, R. Molina, A.K. Katsaggelos, Combining attention-based multiple instance learning and Gaussian processes for CT hemorrhage detection, in: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, Springer, 2021, pp. 582–591.
- [11] M. Kandemir, F.A. Hamprecht, et al., Instance label prediction by Dirichlet process multiple instance learning, in: *UAI*, 2014, pp. 380–389.
- [12] J. Read, L. Martino, J. Hollmén, Multi-label methods for prediction with sequential data, *Pattern Recognit.* 63 (2017) 45–55.
- [13] J. Read, L. Martino, D. Luengo, Efficient monte carlo methods for multi-dimensional learning with classifier chains, *Pattern Recognit.* 47 (3) (2014) 1535–1546.
- [14] J. Read, L. Martino, Probabilistic regressor chains with Monte Carlo methods, *Neurocomputing* 413 (2019) 471–486.
- [15] M. Kim, F. De la Torre, Gaussian processes multiple instance learning, in: *ICML*, 2010.
- [16] M. Kandemir, M. Hausmann, F. Diego, K.T. Rajamani, J. Van Der Laak, F.A. Hamprecht, Variational weakly supervised Gaussian processes, in: E.R.H. Richard C. Wilson, W.A.P. Smith (Eds.), *Proceedings of the British Machine Vision Conference*, BMVC, BMVA Press, 2016, pp. 71.1–71.12.
- [17] M. Hausmann, F.A. Hamprecht, M. Kandemir, Variational bayesian multiple instance learning with gaussian processes, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 6570–6579.
- [18] F. Wang, A. Pinar, The multiple instance learning Gaussian process probit model, in: *International Conference on Artificial Intelligence and Statistics*, PMLR, 2021, pp. 3034–3042.
- [19] Z. Wang, A. Lamb, E. Saveliev, P. Cameron, J. Zaykov, J.M. Hernandez-Lobato, R.E. Turner, R.G. Baraniuk, C. Barton, S.P. Jones, et al., Results and insights from diagnostic questions: The neurips 2020 education challenge, in: *NeurIPS 2020 Competition and Demonstration Track*, PMLR, 2021, pp. 191–205.
- [20] N. Ding, J. Deng, K.P. Murphy, H. Neven, Probabilistic label relation graphs with ising models, in: *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 1161–1169.
- [21] P. Qin, J. Zhao, A polynomial-time algorithm for image segmentation using ising models, in: *2011 Seventh International Conference on Natural Computation*, Vol. 2, 2011, pp. 932–935.
- [22] F. Wang, A. Pinar, The multiple instance learning Gaussian process probit model, in: A. Banerjee, K. Fukumizu (Eds.), *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics*, in: *Proceedings of Machine Learning Research*, vol. 130, PMLR, 2021, pp. 3034–3042.
- [23] E. Snelson, Z. Ghahramani, Sparse Gaussian processes using pseudo-inputs, in: *Advances in Neural Information Processing Systems*, Vol. 18, 2006.
- [24] M. Girolami, S. Rogers, Variational Bayesian multinomial probit regression with Gaussian process priors, *Neural Comput.* 18 (8) (2006) 1790–1817.
- [25] B.D. Ripley, *Spatial Statistics*, John Wiley & Sons, 2005.
- [26] J. Hensman, A. De G. Matthews, Z. Ghahramani, Scalable variational Gaussian process classification, in: *International Conference on Artificial Intelligence and Statistics*, 2015, pp. 351–360.
- [27] C.M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*, Springer-Verlag, Berlin, Heidelberg, 2006.
- [28] S. Wilhelm, B. Manjunath, Tmvtnorm: A package for the truncated multivariate normal distribution, *R J.* 2 (2) (2010) 1–25.
- [29] Y. Li, S.K. Ghosh, Efficient sampling methods for truncated multivariate normal and student-t distributions subject to linear inequality constraints, *J. Stat. Theory Pract.* 9 (4) (2015) 712–732.
- [30] N.L. Johnson, S. Kotz, N. Balakrishnan, *Continuous Univariate Distributions*, Volume 2, 289, John Wiley & sons, 1995.
- [31] Global cancer observatory, World Health Organisation, 2010, <https://gco.iarc.fr/>. (Accessed 15 March 2010).
- [32] J. Silva-Rodríguez, A. Colomer, M.A. Sales, R. Molina, V. Naranjo, Going deeper through the gleason scoring scale: An automatic end-to-end system for histology prostate grading and cribriform pattern detection, *Comput. Methods Programs Biomed.* 195 (2020) 105637.
- [33] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, E. Duchesnay, Scikit-learn: Machine learning in Python, *J. Mach. Learn. Res.* 12 (2011) 2825–2830.
- [34] C. Williams, C. Rasmussen, *Gaussian Processes for Machine Learning*, Vol. 2, MIT press Cambridge, MA, 2006.
- [35] M. Ilse, J. Tomczak, M. Welling, Attention-based deep multiple instance learning, in: *International Conference on Machine Learning*, PMLR, 2018, pp. 2127–2136.
- [36] W. Zhu, Q. Lou, Y.S. Vang, X. Xie, Deep multi-instance networks with sparse label assignment for whole mammogram classification, in: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, Springer, 2017, pp. 603–611.
- [37] A. Schmidt, P. Morales-Álvarez, R. Molina, Probabilistic attention based on Gaussian processes for deep multiple instance learning, *IEEE Trans. Neural Netw. Learn. Syst.* (2023) 1–14, <http://dx.doi.org/10.1109/TNNLS.2023.3245329>.
- [38] J. Silva-Rodríguez, A. Colomer, J. Dolz, V. Naranjo, Self-learning for weakly supervised gleason grading of local patterns, *IEEE J. Biomed. Health Inform.* 25 (8) (2021) 3094–3104.

Pablo Morales-Álvarez received the B.Sc. degree in Mathematics from the University of Granada (UGR), Spain, in 2014. He obtained the First End of Studies Award by the Spanish Ministry of Science to the most outstanding undergraduate in mathematics in Spain. Then, he received the M.Sc. degrees in Mathematical Physics (2015) and Data Science (2016), both from UGR. Funded by the highly competitive Ph.D. fellowship from La Caixa Banking Foundation, he got his Ph.D. degree in 2020 at UGR under the supervision of Prof. Rafael Molina (UGR) and Prof. Aggelos K. Katsaggelos (Northwestern University, USA). During his Ph.D. he visited several research groups, including the Machine Learning Group at the University of Cambridge (UK), where he worked for five months with Prof. Hernandez-Lobato. From October 2020 to May 2021, he did postdoctoral research at Microsoft Research Cambridge (UK) with Cheng Zhang. He has published in top machine learning conferences and journals, such as ICLR 2021, NeurIPS 2022 and IEEE Trans. PAMI. He has also served as a reviewer for top-tier conferences (ICML 2020, NeurIPS 2020, ICLR 2021, NeurIPS 2021), and journals (Science). He has obtained the SCIE-FBBVA Award 2022 to the Most Outstanding Young Spanish Researchers in Computer Science. His main research interests are probabilistic machine learning methods, specially (Deep) Gaussian Processes and Bayesian Neural Networks.

Arne Schmidt studied Mathematics and received his Bachelor degree at the Freie Universität Berlin in 2015 and his Master degree at the Technische Universität Berlin in 2018. He worked as a software engineer specialized in deep learning at Astrofein GmbH, Fraunhofer Heinrich Hertz Institut, and TomTom. In 2020, he started his Ph.D. under the supervision of Prof. Rafael Molina at the University of Granada. His Ph.D. is part of the EU-funded H2020-MSCA-ITN CLARIFY, which focuses on digital pathology for cancer classification. During his first year of Ph.D., he published a paper on MIL for CT hemorrhage detection at MICCAI 2021. His research interests are probabilistic deep learning, Gaussian processes and crowdsourcing with applications to medical images.

José Miguel Hernández-Lobato is a Professor in Machine Learning at the Department of Engineering, University of Cambridge, UK. He is also the director of the ELLIS unit at Cambridge, the premier European Machine Learning network. He serves as a senior scientific advisor for Boltzbit Limited (London, UK), and is a faculty member at the Cambridge Center for AI in Medicine (Cambridge, UK). From 2017, he has raised more than 8.5M€ in research projects funded by different institutions such as Google Deep Mind, Microsoft, Samsung, the British EPSRC, the EU, and Cambridge Innovation Capital. His research revolves around model based machine learning with a focus on probabilistic learning techniques and with a particular interest on Bayesian optimization, matrix factorization methods, copulas, Gaussian processes and sparse linear models. He publishes very regularly at top machine learning conferences (NeurIPS, ICLR, ICML), where he has also served as committee senior member. He is a frequent evaluator for prestigious international funding programs such as ERC Starting (2020), Swiss Data Science Centre (EPFL and ETH Zurich 2019), and AI Singapore Research Program (2020).

Rafael Molina received the M.Sc. degree in mathematics (statistics) and the Ph.D. degree in optimal design in linear models from the University of Granada, Spain, in 1979 and 1983, respectively. He was the Dean of the Computer Engineering School, University of Granada, from 1992 to 2002, where he became a Professor of computer

science and artificial intelligence in 2000. He was the Head of the Computer Science and Artificial Intelligence Department, University of Granada, from 2005 to 2007. He has co-authored an article that received the runner-up prize at reception for early stage researchers at the House of Commons in 2007. He has co-authored an awarded Best Student Paper at the IEEE International Conference on Image Processing in 2007, the ISPA Best Paper in 2009, and the EUSIPCO 2013 Best Student Paper. Dr. Molina has

served as an Associate Editor for Applied Signal Processing from 2005 to 2007, and the IEEE TRANSACTIONS ON IMAGE PROCESSING from 2010 to 2014. He has been serving as an Area Editor for Digital Signal Processing since 2011. He publishes regularly in top AI journals (IEEE TPAMI, IEEE TIP, Pattern Recognition) and conferences (e.g. ICLR, MICCAI). In the last years, he has obtained funding from different bodies (European Union, Spanish Ministry, Andalusia regional government, and University of Granada) to study probabilistic machine learning for medical images.