

Value-Based Potentials: Exploiting Quantitative Information Regularity Patterns in Probabilistic Graphical Models

Manuel Gómez-Olmedo^a, Rafael Cabañas^b, Andrés Cano^a, Serafín Moral^a,
Ofelia Paula Retamero^a

^a*Department of Computer Science and Artificial Intelligence, University of Granada, Spain*

^b*Istituto Dalle Molle di Studi sull'Intelligenza Artificiale, Lugano, Switzerland*

Abstract

The efficient representation of quantitative information in probabilistic graphical models (PGMs) is a challenge for complex models (i.e. problems with many variables, high degree of dependence between them or many states per variable). In this work, several alternative structures are introduced for facing this problem. All of them are guided by the values and are named *Value-Based Potentials* (VBPs). VBPs try to take advantage of the regularity patterns founded in the data (repetitions of values), regardless of how they appear. These new structures are compared with respect to standard tables or unidimensional arrays representation (1DA) and probability trees (PTs). This last structure seeks for reducing memory space. But this goal can only be achieved if there are context specific independence patterns, that is, repeated values correspond to consecutive indexes. VBPs structures try to overcome this limitation. The objective of this work is to analyze the properties of these alternative representations. To do so, in addition to their theoretical analysis, they are tested to encode quantitative information, to access their content and to make inference with a set of well-known Bayesian networks.

Alternativo

The efficient representation of quantitative information in probabilistic graphical models (PGMs) is a challenge when dealing with complex models (i.e., models with many variables, high degree of dependence between them, or many states per variable). Several new structures are introduced in this work to face this problem. They are based exclusively on the values and therefore are named *Value-Based Potentials*. VBPs try to take advantage of the presence of repeated values in order to reduce memory requirements. VBPs are compared with respect to some common structures as standard tables or unidimensional arrays (1DA) and probability trees (PT). PT seeks for reducing memory space as well, but this is achieved only if repetitions of values correspond to context specific independence patterns (that is, repeated values are related to consecutive indexes or configurations). VBPs try to overcome this limitation. The objective of this work is to analyze VBPs properties. In addition to VBPs theoretical analysis, they

are used to encode the quantitative information of a set of well-known Bayesian networks, measuring access times to their content and the computational time required for performing some inference tasks.

1. Introduction

Probabilistic graphical models (PGMs) [31, 22, 20] are efficient representations for problems under uncertainty. PGMs encode joint probability or utility distributions and are defined by two parts: first, a qualitative component in the form of a graph that represents a set of dependencies among the variables (i.e., nodes) in the domain being modelled; secondly, a quantitative component consisting of a set of functions quantifying such dependencies. In PGMs over discrete domains, such as *Bayesian networks* (BN) [30, 32] or *influence diagrams* [28, 17], these functions are traditionally represented with tables or unidimensional arrays (marginal or conditional probability tables and utility tables, 1DA in general).

The sizes of 1DAs increase exponentially with the number of variables in their domains. This property may limit the ability to represent certain problems with large 1DAs (memory size requirements may be prohibitive). Moreover, even if the model can be encoded with 1DAs problems may arise when performing inference computations. In order to make inference, these 1DAs are managed for computing marginal or conditional probabilities for a certain variable, most probable explanation [11, 21], decision tables in the case of influence diagrams, etc. Some inference algorithms [10, 12, 18, 19, 24, 25, 26, 33, 39, 40] use basic operations on potentials: combination, restriction and marginalization. Combination computes the product of two potentials $\phi_1(\mathbf{X})$ and $\phi_2(\mathbf{Y})$ producing as result a new potential with higher dimension $\phi(\mathbf{X} \cup \mathbf{Y})$. In this way, 1DAs obtained as intermediate results can be so large that they exceed the memory capacity of the computer.

Therefore, in order to deal with complex problems, it is essential to use efficient representations of the quantitative information of the model. Usually 1DAs encoding probabilities or utilities contain repeated values. For example, some combinations of values are not allowed and are represented with 0's. An efficient representation should take advantage of all these repetitions in order to reduce memory space. Moreover, a useful representation should offer the capability of being approximated with a trade-off between precision and memory space. These features can allow the treatment of more complex models.

The importance of this problem is evidenced by previous attempts to obtain alternative structures to 1DAs. Two examples of alternative approaches are standard and binary probability trees (PTs and BPTs) [3, 7, 34, 14, 6, 4, 5]. These structures can capture context specific independencies [3] and save memory space when repeated values appear under certain circumstances. These representations also have the capability of obtaining approximations through

pruning operation: assuming loss of information, some contiguous values can be substituted by their average value in order to reduce memory space. There are also previous work focused on improving the operations on potentials to alleviate the computational cost when dealing with complex models [2]. Therefore, the existence of efficient structures for storing and managing quantitative information is of interest in all those areas in which PGMs models could be applied, that is, in any problem or system that needs to quantify uncertainty or preferences, [16, 15, 13, 44, 45, 23, 1].

Other data structures exploiting these independencies make possible to compile a full PGM into a more compact representation. This is the case of Algebraic Decision Diagrams (ADDs) [8, 35] or Sequential Decision Diagrams (SDDs) [9, 29]. The former is a graph representation of a function that maps instantiations of Boolean variables to real numbers. A model whose potentials are represented as ADDs can easily be transformed into an arithmetic circuit what minimizes the number of arithmetic operations during the inference. Similarly, a SDD is a full binary tree for representing propositional knowledge bases (a.k.a. Boolean functions). This data structure allows to encode potentials which can be then conjoint to obtain an efficient SDD representation of a full model.

In this work some new alternative representations for potentials are considered. They are based on the properties of the values themselves and not on the contexts in which they appear or the structure of the potential. That is why they have been called *value-based potentials* (VBPs). The paper defines these structures, making a theoretical analysis of their properties and showing concrete examples of how they encode the quantitative information of known and available Bayesian networks in *bnlearn* package repository [38, 37] as well as other networks used in inference *UAI* competitions [42, 43].

The major advantages of VBPs over other related data structures are the following. First, the memory requirements are noticeably reduced for some networks due to a better capacity of exploiting regularity patterns. Secondly, a VBP represents a single potential independently of the rest of parameters of the model. Thanks to that, many inference algorithms could be easily adapted for working with VBPs by simply adapting the basic operations over potentials (e.g., combination and marginalization). This is not the case of ADDs and SDDs, where a complex compilation process is done to obtain a compact representation of the full model.

The structure of the paper is as follows: Section 2 defines some basic concepts and notation and some usual representations for potentials as arrays and trees. Section 3 introduces basic concepts about memory requirement analysis. Section 4 introduces VBPs representations and how to categorize them. Section 5 introduces VBPs alternatives and their properties. Section 6 presents the empirical evaluation performed for testing VBPs capabilities. Finally Section 7 presents conclusions and lines for future research.

2. Basics

2.1. Definitions and notation

Let us first introduce the basic notation. Upper-case *roman* letters will be used to denote random variables and lower-case for their values (or states). Thus, if X_i is a random variable, x_i will denote a generic value of X_i . The finite set of possible values of X_i is called domain and denoted Ω_{X_i} . For simplicity, we will consider variable values as integers starting with 0 and hence possible assignments will be $X_1 = 0$, $X_1 = 1$, $X_1 = 2$, etc. The cardinality of a variable, denoted $|\Omega_{X_i}|$, is the number of values in its domain. Similarly, we use bold-face upper-case *roman* letters to denote sets of variables, e.g. $\mathbf{X} := \{X_1; X_2; \dots; X_N\}$ is a set of N variables ($|\mathbf{X}| = N$). The Cartesian product $\prod_{X_i \in \mathbf{X}} \Omega_{X_i}$ is denoted by $\Omega_{\mathbf{X}}$. The elements of $\Omega_{\mathbf{X}}$ are called configurations of \mathbf{X} and will be represented by $\mathbf{x} := \{X_1 = x_1; X_2 = x_2; \dots; X_N = x_N\}$ or simply $\mathbf{x} := \{x_1; x_2; \dots; x_N\}$ if the variables are obvious from the context.

Formally, a PGM contains three elements $\langle \mathbf{X}; P; G \rangle$ where \mathbf{X} is the set of variables in the problem with a joint probability distribution $P(\mathbf{X})$ and G is a graph that represents the dependence (and independence) relations between the variables. A PGM allows to represent P , which is usually high-dimensional, as a factorisation of lower dimensional local functions. For instance, in case of BNs, these are conditional distributions represented as tables or conditional probability tables (arrays in general, IDAs). However, we will use the term *potential* which is more general: a potential for \mathbf{X} is a function of $\Omega_{\mathbf{X}}$ over \mathbb{R}_0^+ . In other words, each configuration $\mathbf{x} \in \Omega_{\mathbf{X}}$ is associated to a real value. Thus, IDAs or any other function encoding the quantitative information in PGMs can be seen as representations of potentials.

Example 1. *Let us consider the variables X_1 , X_2 and X_3 with 2, 3 and 2 states respectively. Then $\langle X_1; X_2; X_3 \rangle$ is a potential defined on such variables with the values shown in Figure 1. Note that this potential represents the conditional distribution $P(X_3 | X_1; X_2)$.*

The definition of structures for representing potentials requires the introduction of the following concept: *index of a configuration*. It is a unique numeric identifier representing each configuration in a given domain $|\Omega_{\mathbf{X}}|$. We will consider indexes starting with 0 (all the variables take their first value) and ending with $|\Omega_{\mathbf{X}}| - 1$. In the potential given in Figure 1, the index 0 is associated to $\{0; 0; 0\}$, the index 1 to $\{0; 0; 1\}$ and so on until the last one (11) which is associated to $\{1; 2; 1\}$ (indexes are contained in left most column).

It is possible to set a correspondence between indexes and configurations based on the concept of *weight* (a.k.a. *stride* or *step size*). Let us suppose a domain $\mathbf{X} := \{X_1; X_2; \dots; X_N\}$. Each variable X_i has a weight w_i computed as follows:

$$w_i = \begin{cases} 1 & \text{if } i=N \\ |\Omega_{X_{i+1}}| w_{i+1} & \text{otherwise} \end{cases} \quad (1)$$

Example 4. The same potential given in the previous example is presented in Figure 3 as a PT. This PT has a size of 21 nodes (12 leaves and 9 internal nodes).

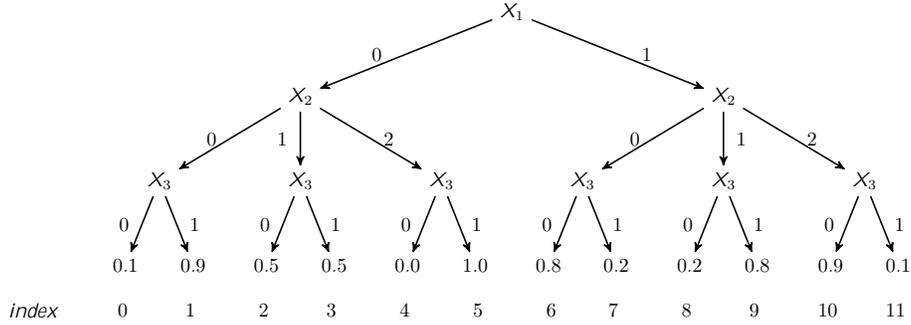


Figure 3: $(X_1; X_2; X_3)$ as PT

With PTs, the most efficient way of access is via configuration: the tree must be traversed from root to leaves selecting the corresponding branch for each variable until reaching a leaf node with its value. In addition, PTs can take advantage of context-specific independencies [3] so that many identical values can be grouped into a single one offering a compact storage. The operation of collapsing identical values is called *pruning*.

Example 5. The potential in previous examples presents a context-specific independence that allows a reduction of its size: the value for $X_1 = 0; X_2 = 1$ is 0:5 regardless of the value of X_3 . If the pruning is done, the result is a PPT (pruned PT) of size 19 shown in Figure 4.

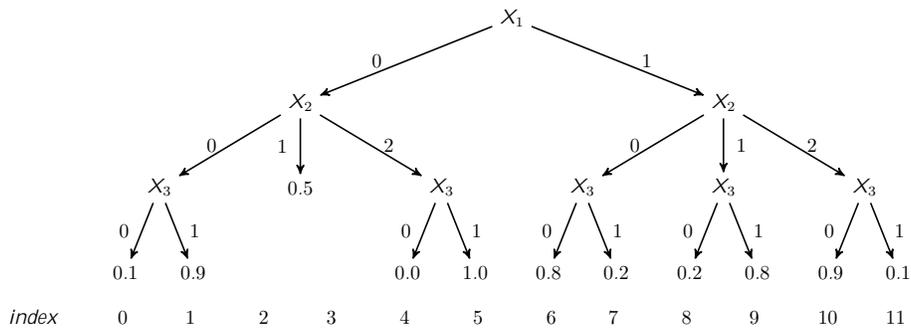


Figure 4: $(X_1; X_2; X_3)$ as PPT

However PTs can not take advantage of another repetition patterns. Let us consider the tree presented in Figure 3. The values for indexes 7 and 8 are 0:2

and consecutive as well, but they can not be pruned because they correspond to configurations varying both in X_2 and X_3 values.

A variant of PTs, called binary trees (BPTs) [6], can divide the domain of each variable into two subsets of states. This would allow finer grain context independencies to be exploited with respect to regular trees. For example, in the PT (left part of Figure 5), the values 0:4 in c configuration (left subtree) can not be pruned. However, with BPTs grouping states 0 and 2 of X_k would allow the value 0:4 to be represented with a single leaf node (as showed in the BPT of the right part of Figure 5). This reduces memory space for c context, although it would require more nodes for the right subtree (c^j context). For this reason, only PTs and PPTs are considered for comparison in this work.

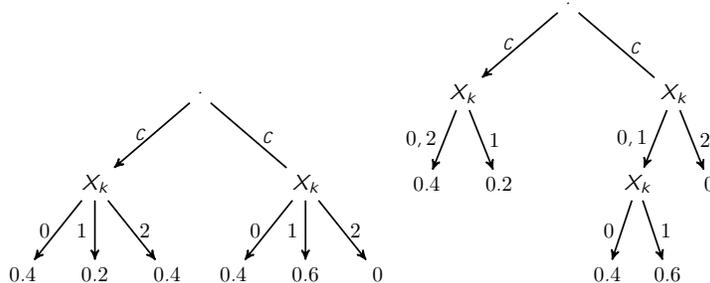


Figure 5: Binary tree representation

3. Memory requirements analysis

Even though the size of a representation gives an idea of its complexity, a more accurate analysis of its memory space requirements is needed: any representation will consume additional elements (e.g., pointers, meta-information, etc.) and each of them use different data types. For this analysis we will consider a potential $\Omega_{\mathbf{X}}$ defined over a set of N variables $\mathbf{X} := \{X_1; X_2; \dots; X_N\}$. Additionally, the following notation related to the different memory sizes is defined.

- S_f will be the memory size required for storing a float value.
- S_i is the memory size for storing an index denoting a concrete configuration of $\Omega_{\mathbf{X}}$.
- S_r denotes the size of a reference to an object.
- S_v represents the memory space required for storing the information about a variable: name, cardinality and state names. As this depends on the names of variables and states, we will assume a fix value for all of them (in fact, this memory space will be negligible respect to the whole memory

used for storing a potential). Moreover we can define a standard way of coding variables with numerical identifiers and employ the same idea for their states.

- S_S denotes the memory size of the data structure used for storing information (array, list, set or dictionary). Specific notation will be used for this term afterwards: S_{arr} , S_{list} , S_{set} and S_{dict} respectively.

As it was mentioned before, the representation by means of 1D-array (1DA) has an important advantage: the values for configurations are stored consecutively. That is, the value in position k corresponds to k -th configuration. In this way it is not necessary to store information about indexes. Therefore, its codification supposes an amount of memory given by the number of values to store, the size of the array data structure and the memory required for its variables:

Proposition 1 (Memory space for an array representing a potential). *Let A be a 1DA representing (\mathbf{X}) , $|\mathbf{X}| = N$. Then the amount of required memory is given by the following expression.*

$$\text{memory}(A) = N \cdot s_v + m \cdot s_f + S_S \quad (5)$$

where $m = \sum_{j \in \mathbf{X}} |\Omega_{\mathbf{X}^j}|$ is the number of entries in the array.

Example 6. *From the previous examples, consider the potential $(X_1; X_2; X_3)$ and its codification as a 1D-array given in Example 3. Then the estimation of the memory size is:*

$$\text{memory}(A) = 3s_v + 12s_f + S_S \quad (6)$$

The tree representation (PT and PPT) is usually less efficient in terms of memory requirements as the full structure of the tree must be stored. Thus, the amount of memory depends on the number of internal nodes, denoted n_I , and the number of leaves n_L . Note that it holds that $n_I + n_L = \text{size}(T)$. Internal nodes store links (or references) to sub-trees (each for a state of the corresponding variable). These links are stored into an array. Then, it is relevant to consider the number of outgoing arcs: $n_I^{(j)}$ denotes the total number of internal nodes for variables with j states.

Proposition 2 (Memory space for a tree representing a potential). *Let T be a tree representing (\mathbf{X}) , $|\mathbf{X}| = N$. Then the amount of required memory is estimated as follows.*

$$\text{memory}(T) = N \cdot s_v + n_L \cdot s_f + \sum_{j=2}^K n_I^{(j)} \cdot (s_v + S_S + j \cdot s_f) \quad (7)$$

where $K = \max_{j \in \mathbf{X}} |\Omega_{X_j}|$ is the maximal cardinality among the variables in the potential.

In case of a non-pruned tree, the number of leaf nodes will be equal to the number of configurations in the potential. As a consequence, the first two terms in Equations(5) and (7) are equal. Thus, the main factors in the size of trees are the data structure employed for storing links to subtrees and the repetition of variables information.

Example 7. Let T be the PT from Example 4 (Figure 3) containing 7 internal nodes for binary variables and 2 internal nodes for the ternary ones and 12 leaves. Similarly, let T^θ be the PPT from Example 5 (Figure 4) with 6 internal nodes for binary variables, 2 internal nodes for ternary variables and 11 leaf nodes. Then, their memory cost can be computed with the following expressions:

$$\text{memory}(T) = 3s_v + 12s_f + 7(s_v + s_s + 2s_r) + 2(s_v + s_s + 3s_r) \quad (8)$$

$$\text{memory}(T^\theta) = 3s_v + 11s_f + 6(s_v + s_s + 2s_r) + 2(s_v + s_s + 3s_r) \quad (9)$$

In the previous example, the pruning operation has reduced the number of internal nodes as well as the number of leaves, but anyway the cost is higher to the size of the table representation. A large number of repeated values are usually required for obtaining memory savings when using PTs.

For a concrete analysis, the following sizes for data types are assumed (the real sizes can depend on the machine architecture; in fact, real sizes are not relevant as long as the same sizes are used for all the comparisons):

- long: 4 bytes (for indexes).
- float: 8 bytes (for real values).
- pointer or reference: 8 bytes (memory addresses).
- variable: 50 bytes (this includes the space for storing name, state names, etc). That is, $s_v = 50$.
- the concrete value of s_s will depend on the concrete data structure employed:
 - array and list (s_{arr} and s_{list} respectively): 16 bytes.
 - set (s_{set}): 32 bytes.
 - dictionary (s_{dict}): 64 bytes.

Using these sizes, the estimations of memory for the representations A , T and T^θ are 262, 1000 and 910 respectively.

4. Value-based representations

4.1. Motivation

At this point, it is clear the necessity of having efficient mechanisms for handling quantitative information for the tasks of representation, inference and learning with PGMs. We have already considered that PTs allow to capture some patterns of repetition in very specific situations. The underlying idea in VBPs (Value-Based Potentials) is to let the representation process be guided by the values themselves and saving space from repetitions as much as possible. Therefore, the objectives of VBPs are:

- being able to take advantage of all repetition patterns, regardless of the order in which they appear.
- to allow efficient access to values and provide the necessary operations to perform inference tasks. In this work, basic implementations for combination and marginalization operations are provided to obtain a starting estimate of VBPs behavior when using inference algorithms.
- to facilitate the approximation task and the parallel management. These capabilities are not explored in this work but it is important to consider them for reaching a good design including these possibilities (to be explored as future research).

4.2. Alternatives categorization

The proposed alternatives can be classified in two groups depending on how to make the queries:

- *driven by values* approaches, based on the use of dictionaries in which the keys will be values. Two representations belonging to this group are presented: *Value-Driven with Grains* (VDG) and *Value-Driven with Indexes* (VDI).
- *driven by indices* alternatives where keys are indexes. Into this group we present *Index-Driven with Indexes* (IDP) and *Index-Driven with Map* (IDM). Both alternatives use an array for values (V). IDP also uses a second array (L) that stores the indexes and the information needed to link indexes and values. IDM uses a map (M) with indexes as keys and the information to link with V as values.

The particular features for all of them will be presented below. However, a common capability is outlined here. It must be clear that these representations require a search for managing the information. This can be exploited defining a default value to return when the search fails. This default value can be set to 0.0 or fixed after analyzing the values of the potential. In this case it would help to select as default value the most repeated one in order to reduce memory although it requires more computation time. This work considers the first alternative.

5. VBPs description

5.1. VDG: value-driven with grains

Identical values in potentials will often appear in configurations with consecutive indexes. Consider for instance the potential given in Example 3, in which the value 0:5 appears in positions 2 and 3. Similar situation happens with value 0:2 in positions 7 and 8. As a consequence, a compact way of defining sets of configurations associated to the same value could be by means of intervals (i.e., grains). Formally, a grain can be defined as follows.

Definition 1 (Grain). *Let \mathbf{X} be a set of variables and i and j indexes of valid configurations on $\Omega_{\mathbf{X}}$, with $i \leq j$. A grain $g(i;j)$ defines a sequence of consecutive indexes $i; i+1; \dots; j$. Grains will be used for representing sequences of repeated values in VBPs.*

In a VDG encoding a potential, each non-zero value will be associated with one or more grains defining all the indexes for which the potential takes this value. More formally, a VDG can be defined as follows.

Definition 2 (Value-driven with grains). *Let ψ be a potential defined over \mathbf{X} . Then a value-driven with grains for ψ , VDG \mathcal{D} , is a dictionary D in which entries are defined as $\langle v; L_v \rangle$, where $v \geq 0$ (key) is a non-default value and L_v is a list of grains that store its associated indexes. Therefore for each grain $g(i;j) \in L_v$ all its indexes correspond to v (that is, for all $k = i; i+1; \dots; j$ then $\psi(\mathbf{x}_k) = v$).*

Example 8. *The potential $\psi(X_1; X_2; X_3)$ used in previous examples and presented in Figure 1 will be represented with VDG as showed in Figure 6.*

The outermost rectangle with rounded corners represents the dictionary. The circular nodes indicate entry keys. Associated with each key, on the right, it is represented the related list of grains. It can be seen in Figure 6 that each value is stored only once. Some values appear only once in the potential. This is the case of 1:0, with 5 as starting and ending index in the grain. The rest of values appear several times. For example, 0:1 is the value for indexes 0 and 11. As these indexes are not consecutive, they must be stored in two different grains. Values 0:2 and 0:5 appear in consecutive indexes and their corresponding grains capture the sequences of repetitions.

Algorithm 1 (Access to index in VDG). *Given VDG \mathcal{D} , the algorithm for getting the value corresponding to a given index l is described in Algorithm 1.*

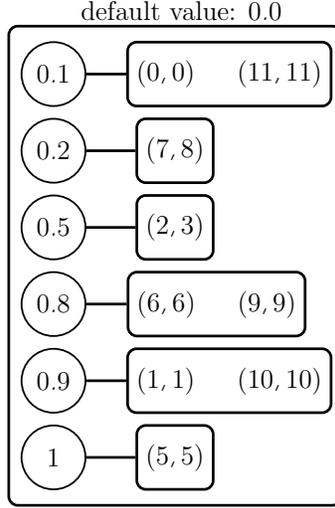


Figure 6: $(X_1; X_2; X_3)$ as VDG

Algorithm 1 Access to l index in VDG

```

1: function access( $VDG ; l$ )
2:    $result = 0.0$  . sets the default value to result
3:   for each  $v$  (key) in  $D$  key set do . loop on dictionary entries
4:      $L_v = D(v)$  . list of grains for  $v$ 
5:     for each  $g$  (grain) in  $L_v$  do
6:       if  $l \geq g$  then .  $l$  is included in  $g$ 
7:          $result = v$  and stop iteration
8:       end if
9:     end for
10:  end for
11:  return  $result$ 
12: end function

```

Assume 6 is the target index. The search progresses through dictionary entries until reaching key 0:8. The internal loop (line 5) iterates on the grains for this value. As the first one contains 6, then 0:8 would be the result of the search. When the search fails and the index is not found, then 0:0 (default value) is returned. This would be the case for index 4.

Proposition 3 (Memory space for a VDG representing a potential). *Let VDG be the representation of (\mathbf{X}) , with $j\mathbf{X}j = N$. Let assume d represents the number of different values in the potential (discarding the default value). The number of grains for each value are denoted by $g_1 :: : g_d$. Then the amount of*

memory required is estimated as follows.

$$\text{memory}(VDG) = N (s_v + s_f + s_{dict} + d \sum_{j=1}^d (s_f + s_{list}) + 2 \sum_{j=1}^d g_j s_i) \quad (10)$$

The terms in Equation (10) consider sizes for: variables; storage for default value; dictionary; values and lists; and grains with 2 indexes per grain. The number of grains for each value will depend of the sequences of repetitions. It will be lower as long as the sequences are longer. Therefore, the critical point in this representation is the required number of grains, given by $\sum_{j=1}^d g_j$.

Example 9. Let VDG be the VDG from Example 8 with 6 different values to store in the dictionary and 0:0 as default value. Therefore, the dictionary stores 6 entries. The sequences of repetitions requires 9 grains. Then, the memory cost can be computed with the following expression:

$$\text{memory}(VDG) = 3s_v + s_f + s_{dict} + 6 (s_f + s_{list}) + 9 \sum_{j=1}^d s_i \quad (11)$$

Using the concrete memory sizes described in Section 3 the complete amount of memory is 438 bytes (sizes of 1DA, PT and PPT are 262, 1000 and 910 bytes respectively).

5.2. VDI: value-driven with indexes

Even though the previous structure with grains is a compact representation for potentials, it could encode unnecessary information when repeated values are not in consecutive positions. This is the case for 0:1 in Figure 6. Its entry includes 2 grains of length 1: (0;0) and (11;11). One way to avoid this repetition would be to associate values with the complete list of indexes in which they appear. In this example, 0:1 value will be related to the list (0 ! 11). This alternative will be advantageous if the sequence of values of a potential does not contain large series of repetitions. Having this idea in mind, the following representation can be defined.

Definition 3 (Value-driven with indexes). Let ψ be a potential defined over \mathbf{X} , then a value-driven with indexes for ψ , VDI , is a dictionary D in which each entry $\langle v; L_v \rangle$ contains a value (as key) and a list of indexes L_v , such that $(\mathbf{x}_i) = v$ for each $i \in L_v$.

Example 10. The potential $(X_1; X_2; X_3)$ used before and described in Figure 1 will be represented as VDI as showed in Figure 7.

The outermost rectangle represents the dictionary: keys of entries are drawn as circles. Keys give access to lists of indexes (rectangles of rounded corners).

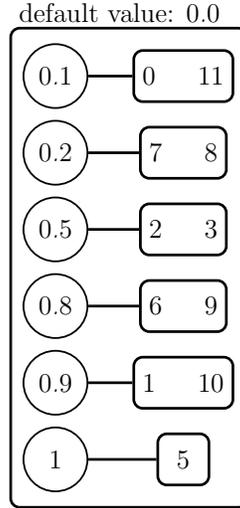


Figure 7: $(X_1; X_2; X_3)$ as VDI

Algorithm 2 (Access to index in VDI). *Given VDI , the algorithm for getting the value corresponding to a given index l is described in Algorithm 2.*

Algorithm 2 Access to l index in VDI

```

1: function access(VDI ;  $l$ )
2:    $result = 0.0$  . sets the default value to result
3:   for each  $v$  (key) in  $D$  key set do . loop on dictionary entries
4:      $L_v = D(v)$  . list of indexes for  $v$ 
5:     for each  $p$  in  $L_v$  do . loop on list of indexes
6:       if  $p == l$  then .  $l$  is included in  $L_v$ 
7:          $result = v$  and stop iteration
8:       end if
9:     end for
10:  end for
11:  return  $result$ 
12: end function

```

Assume 6 is the target index. The search progress through dictionary entries until reaching key 0:8. The internal loop (line 5) iterates on the list. As it contains 6, then 0:8 would be the result. A failed search produces 0:0 as a result. This is the case for index 4.

Proposition 4 (Memory space for a VDI representing a potential). *Let VDI be the representation of (\mathbf{X}) with $|\mathbf{X}| = N$. Let assume d represents the number of different values in the potential (discarding the default value). The number*

of indexes for each value are denoted by $i_1 :: i_d$. Then the required amount of memory is estimated as follows.

$$\text{memory}(VDI) = N (s_v + s_f + s_{dict} + d \sum_{j=1}^d (s_f + s_{list}) + s_j) \quad (12)$$

The terms of Equation 12 consider sizes for: variables, default value, dictionary, values, lists and indexes. In VDI, the main source of savings comes from avoiding the storage of repeated values, since the indexes in which the significant values (non zero) appear must be explicitly stored.

Example 11. Let VDI be the VDI from Example 10 with 6 different values to store in the dictionary and 0.0 as default value. Therefore, the dictionary stores 6 pairs and 6 lists with 11 indexes. Then, the memory cost can be computed with the following expression:

$$\text{memory}(VDI) = 3s_v + s_f + s_{dict} + 6 (s_f + s_{list}) + 11 s_j \quad (13)$$

Using the memory sizes described in Section 3 the complete amount of memory is 410 bytes (a bit lower than VDG).

5.3. IDP: index-driven with pairs

The problem with the access to value-driven structures is the need to perform a double iteration. The search for a target index, l , requires iterating over the list of entries and over the associated lists (of grains or indexes). To avoid this double iteration and make the search more efficient, new structures are introduced in which the search is based on the indexes themselves. This is the case of IDP and IDM.

Definition 4 (Index-driven with pairs). Let f be a potential defined over \mathbf{X} . Then a structure index-driven with pairs (IDP) representing f , IDP_f , is a pair of arrays: V and L . Non-repeated values in f (excluding 0.0 as default value) are stored in $V := [v_0; v_1; \dots; v_{d-1}]$. Let nd represents the set of indexes storing non-default values. The array L is defined as follows.

$$L := f(i;j) : (x_i) = v_j; i \in nd \quad (14)$$

That is, IDP is based on two components. First, an array storing the values (without repetitions, as before, and excluding 0.0 as default value). Secondly, an array of pairs (index in potential - index in array of values). The second index of the pair keeps the relation between indexes and values.

default value: 0.0										
0	1	2	3	4	5					
0.1	0.2	0.5	0.8	0.9	1.0					
(0,0)	(1,4)	(2,2)	(3,2)	(5,5)	(6,3)	(7,1)	(8,1)	(9,3)	(10,4)	(11,0)

Figure 8: $(X_1; X_2; X_3)$ as IDP

Example 12. The representation as IDP of the potential $(X_1; X_2; X_3)$ presented in Figure 1 is shown in Figure 8.

As explained before, IDP uses two coherent arrays. V stores non repeated values (except the default value). L contains pairs of indexes. Let us consider value 0.5 in $(X_1; X_2; X_3)$, presented in indexes 2 and 3. Then, the array of pairs, L (bottom part of Figure 8), requires two pairs for this relation: (2;2) and (3;2). Both indicate that potential indexes 2 and 3 stores the value in $V(2)$.

Algorithm 3 (Access to index in IDP). Given a IDP, the algorithm for getting the value corresponding to a given index l is described in Algorithm 3.

Algorithm 3 Access to l index in IDP

```

1: function access(IDP ; l)
2:   result = 0.0                                     . sets the default value to result
3:   for each pair  $t = (i ; i_V)$  in  $L$  do           . loop on  $L$  array pairs
4:     if  $i == l$  then                               .  $l$  is found; value stored in  $V(i_V)$ 
5:       result =  $V(i_V)$  and stop iteration
6:     end if
7:   end for
8:   return result
9: end function

```

Assume 6 is the target index. The search progress through L until reaching 6-th position (pair (6;3)). The value to return is stored in $V(3) = 0.8$. If index 4 is searched, then 0.0 will be returned (this index is not present in L).

Proposition 5 (Memory space for an IDP representing a potential). Let IDP be the structure representing $(\mathbf{X}); j \in \mathbf{X} = N$. Let assume d represents the number of different values in the potential (discarding the default value). The number of indexes corresponding to non-default value is p . Then the amount of memory is estimated as follows.

$$memory(IDP) = N \cdot s_V + s_f + 2 \cdot s_{arr} + d \cdot s_f + 2 \cdot s_i \cdot p \quad (15)$$

The terms in Equation 15 considers the sizes for: variables, default value, both arrays, values and pairs of indexes. This representation tries to use simple structures and favours the direct search on indices rather than on values.

Example 13. Let IDP be the IDP from Example 12 with 6 different values to store and 0.0 as default value. Therefore, V stores 6 values and L 11 pairs. Then, the memory cost can be computed with the following expression:

$$\text{memory}(IDP) = 3s_v + s_f + 2 \cdot s_{arr} + 6 \cdot s_f + 11 \cdot 2 \cdot s_i \quad (16)$$

Using the concrete memory sizes described in Section 3 the complete amount of memory is 326 bytes.

5.4. IDM: index-driven with map

This structure aims to take a further step in the idea of facilitating access to the structure, using a dictionary in which the keys are the indexes. Below is the definition of this new index-driven alternative.

Definition 5 (Index-driven with map). Let ϕ be a potential defined over \mathbf{X} , then a structure index-driven with map (IDM) representing ϕ , IDM consists of a dictionary D and an array V . Non-repeated values in ϕ are stored in $V := [v_0; v_1; \dots; v_{d-1}]$. D entries $\langle i; j \rangle$ are formed by i indexes (keys) and V indexes. Let nd represents the set of indexes storing non-default values. Given an entry $\langle i; j \rangle$, then: $(\mathbf{x}_i) = v_j$ and $i \in nd$.

Example 14. The representation as IDM of the potential $(X_1; X_2; X_3)$ presented in Figure 1 is shown in Figure 9.

The dictionary D is represented in the left part of Figure 9 and the array of values V in the right one. Keys in D are drawn as circles and give access to V indexes (boxes linked to keys).

Algorithm 4 (Access to index in IDM). Given IDM , the algorithm for getting the value corresponding to a given index l is described in Algorithm 4.

Algorithm 4 Access to l index in IDM

```

1: function access( $IDM ; l$ )
2:    $result = 0.0$  . sets the default value to result
3:    $entry(\langle l; j \rangle) = D(l)$  . search dictionary for  $l$ 
4:   if  $entry \neq \text{null}$  then . dictionary contains  $l$  as key
5:      $result = V(j)$ 
6:   end if
7:   return  $result$ 
8: end function

```

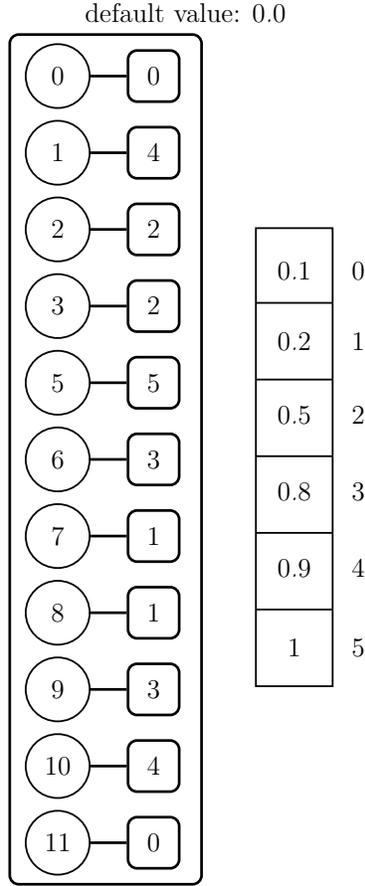


Figure 9: $(X_1; X_2; X_3)$ as IDM

Assume 6 is the target index. As this is a valid key, entry $\langle 6;3 \rangle$ is retrieved. The value to return is stored in $V(3) = 0.8$. If index 4 is searched, then 0.0 will be returned (this index is not present in D).

Proposition 6 (Memory space for an IDM representing a potential). *Let IDM be the structure representing $(\mathbf{X}); j \in \mathbf{X} = N$. Let assume d represents the number of different values in the potential (discarding the default value) and p the number of indexes storing non-default values. Then the amount of memory is estimated as follows.*

$$memory(IDM) = N \cdot s_v + s_f + s_{dict} + s_{arr} + d \cdot s_f + 2 \cdot p \cdot s_i \quad (17)$$

The terms of the Equation (17) represents sizes for: variables, default value, dictionary, arrays of values, different values and indexes in dictionary entries.

Example 15. Let IDM be the VBP from Example 14 with 6 different values to store and 0:0 as default value. Then, the array of values has 6 elements. The dictionary contains 11 entries. Memory cost can be computed with the following expression:

$$memory(IDM) = 3s_v + s_f + s_{dict} + s_{arr} + 6 \cdot s_f + 2 \cdot 11 \cdot s_j \quad (18)$$

Considering the concrete values as specified in Section 3 the complete amount of memory is 374 bytes.

5.5. Example of a extreme case

Below it is included an example of the use of the representation structures under consideration for a potential with extreme features: only three different values (0, 0.5 and 1) and many repetitions of the default value (around 70% of the indexes are assigned to 0). The potential has 1024 possible values, with 5 variables in its domain, each with 4 states. The values are randomly generate. We have considered 10 random potentials with these features in order to get a reliable idea of the behavior of the representations. Results are presented in Table 1. $1DA$ and PT representations do not depend on the concrete values and requires memory sizes of 8574 and 53816 respectively.

iteration	1	2	3	4	5	6	7	8	9	10
PPT	39090	38332	37658	38396	40946	38910	42610	38120	42706	38544
VDG	1822	1814	1670	1774	1934	1854	1990	1734	2030	1694
VDI	1206	1210	1110	1190	1310	1218	1346	1166	1374	1194
IPD	1862	1870	1670	1830	2070	1886	2142	1782	2198	1838
IDM	1910	1918	1718	1878	2118	1934	2190	1830	2246	1886
values	180	166	156	176	206	183	215	170	222	177

Table 1: Memory sizes for random potential representations

In all the random potentials the number of non-zero values ranges from 156 to 222 (last row in Table 1) and the longest sequence of repeated values is 2. Although better results could be obtained with longest sequences of repeated values (at least with VDG representation), the results show that the savings in memory consumption with respect $1DA$, PT and PPT are noticeable.

6. Empirical evaluation

Two sets of Bayesian networks are used for evaluating VBPs capabilities with respect to previous representations for potentials in PGMs: conditional probability tables ($1DA$) and trees (PT and PPT). The first set is taken from the *bnlearn* repository ([38, 37]) and the second one from *UAI* competitions ([42, 43]). The quantitative information of these models is represented with the structures previously defined (VDG , VDI , IDP and IDM). Experiments are organized in three different blocks: comparison of memory sizes, access times

and computation time of posterior distributions using the Variable Elimination (VE) algorithm [41, 46, 12].

The representations compared in experiments are:

- 1DA, PT, PPT, VDG, VDI, IDP and IDM for memory sizes and access times comparisons on *bnlearn* and *UAI* networks.
- 1DA, PT, VDI and IDM for posterior computations with *UAI* networks.

6.1. Bayesian networks features

Some basic information about the Bayesian networks employed is gathered in Table 2 and Table 3: name of network; number of nodes, number of arcs; minimum, average and maximum number of states of variables; and complete number of parameters for quantifying networks uncertainty. Networks are ordered according the number of parameters.

network	nodes	arcs	min. st.	avg. st.	max. st.	parameters
cancer	5	4	2	2	2	20
asia	8	8	2	2	2	36
survey	6	6	2	2.33	3	37
sachs	11	17	3	3	3	267
child	20	25	2	3	6	344
alarm	37	46	2	2.83	4	752
win95pts	76	112	2	2	2	1148
insurance	27	52	2	3.29	5	1419
hepar2	70	123	2	2.31	4	2139
andes	223	338	2	2	2	2314
hail nder	56	66	2	3.98	11	3741
pigs	441	592	3	3	3	8427
water	32	66	3	3.625	4	13484
munin1	186	273	2	5.33	21	19226
link	724	1125	2	2.53	4	20502
munin2	1003	1244	2	5.36	21	83920
munin3	1041	1306	2	5.38	21	85615
path nder	109	195	2	4.11	63	97851
munin4	1038	1388	2	5.44	21	97943
munin	1041	1397	2	5.43	21	98423
barley	48	84	2	8.77	67	130180
diabetes	413	602	3	11.34	21	461069
mildew	35	46	3	17.6	100	547158

Table 2: *bnlearn* Bayesian networks features

Observe that networks in the *UAI* set require more parameters than those from *bnlearn*.

network	nodes	arcs	min. st.	avg. st.	max. st.	parameters
BN_76	2155	3686	2	7.01	36	627298
BN_87	422	867	2	2	2	933776
BN_29	24	30	10	10	10	1132080
BN_125	50	375	2	2	2	2117680
BN_115	50	375	2	2	2	2285616
BN_119	50	375	2	2	2	2410544
BN_121	50	375	2	2	2	2564144
BN_123	50	375	2	2	2	3249200
BN_27	3025	7040	3	6	10	3698565
BN_117	50	375	2	2	2	4003888
BN_22	2425	4239	2	18.743	91	4073904
BN_111	50	375	2	2	2	4238512
BN_109	50	375	2	2	2	4581936
BN_113	50	375	2	2	2	4669488
BN_20	2483	5272	2	18.92	91	5009364
BN_107	50	375	2	2	2	5154864
BN_105	50	375	2	2	2	6431792

Table 3: *UAI competition* Bayesian networks features

6.2. Memory sizes comparisons

In order to compare the memory sizes required for each representation with respect to 1DA, PT and PPT, we proceed to convert all the potentials to VDG, VDI, IDP and IDM. The memory size used by 1DA is taken as a reference and does not appear in the table. Given a certain network, let m_{1DA} be the memory space for 1DA representation and m_{rep} the memory size for another one. Then the value s included in table cells and representing the gain (or loose) for rep is computed as:

$$s = \frac{m_{rep}}{m_{1DA}} - 100 \quad (19)$$

Thus, a negative value for s indicates that rep requires less memory space than 1DA. On the contrary, positive values indicate higher memory consumption.

6.2.1. Memory sizes for bnlearn networks

The results for this set of networks are presented in Figure 10. Some comments about these results are included below.

- PTs and PPTs always require more memory space than 1DA. Both of them are quite similar except for **win95pts**. This network contains several potentials with repeated values where the prune operation substantially reduces the number of leaf nodes and consequently memory size.

149.35	149.35	98.70	88.31	36.36	67.53	cancer
161.31	154.61	76.79	67.86	27.38	55.95	asia
184.01	184.01	112.50	98.90	40.44	66.91	survey
291.77	283.59	129.22	108.44	38.07	51.65	sachs
242.25	239.06	83.37	63.31	30.77	47.78	child
272.07	266.47	40.74	29.85	20.23	35.88	alarm
412.49	245.47	21.32	21.89	9.40	26.76	win95pts
329.88	327.65	51.07	27.94	9.45	17.47	insurance
343.07	342.14	171.08	143.23	58.86	70.45	hepar2
312.66	303.61	49.52	33.63	18.27	38.20	andes
340.31	337.86	45.20	14.13	9.86	16.96	hailfinder
264.54	264.54	11.32	3.79	-6.31	9.58	pigs
390.78	390.58	27.39	4.57	-23.36	-22.01	water
263.39	261.70	10.99	-5.66	-27.12	-22.23	munin1
322.50	317.04	-19.18	-26.98	-30.53	-18.11	link
241.63	240.63	25.95	8.53	-20.87	-14.97	munin2
238.27	237.85	28.34	10.80	-20.04	-14.06	munin3
429.55	370.87	-48.60	-63.05	-40.49	-39.84	pathfinder
252.88	251.91	18.81	1.55	-24.04	-18.72	munin4
254.53	250.29	17.31	0.34	-24.25	-18.95	munin
230.45	228.87	59.95	35.08	28.25	28.47	barley
167.13	167.13	-63.48	-74.85	-71.08	-70.56	diabetes
112.44	112.43	-84.86	-88.29	-90.25	-90.21	mildew
PT	PPT	VDG	VDI	IDP	IDM	

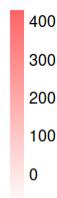


Figure 10: Memory sizes for *bnlearn* networks

- In almost all networks, the most competitive representation is IDP. In networks with a number of parameters under 8427 (from **cancer** to **hailfinder**), IDP requires more memory than 1DA. But for the rest of networks IDP offers memory savings ranging from 6.31% until 90.25%; **barley** network is an exception. The potentials in this network have short sequences of repeated values including few indices. For example, the potential for **jordn** variable has 4752 possible values, but only 4 different ones. However, the sequences are arranged in such a way that they cannot be exploited by PPT. This is the reason why VDG uses many grains. Moreover, potentials contain few zeros, which means that there are many indexes to store. In these cases, it would be appropriate to make the default value be the most repeated one, although this would complicate the way of performing combination and marginalization operations.
- More important savings correspond to **diabetes** and **mildew**. In these two networks there are several potentials with a high number of repeated values. For **diabetes** there are 25 potentials with 7056 parameters but only 44 different values. In **mildew** the same circumstance arises (a potential with 39040 possible values but only 1756 different ones; another one with 201300 parameters and 4508 different values are two examples). In these potentials repeated values can not be collapsed when using PPTs.
- VDI is the best representation for **pathfinder** and **diabetes**. This is

explained by the reduced number of different values respect to the number of parameters. In **pathfinder** a potential with 8064 parameters needs only 29 different values.

6.2.2. Memory sizes for UAI networks

The results for these networks are presented in Figure 11.



Figure 11: Memory sizes for UAI networks

The following conclusions can be outlined from these results:

- Percentages for PTs and PPTs are always over 200% except for **BN_27**. In this network there is an important difference between PT (488%) and PPT (-95:55%). Furthermore, PPT is the best representation, but VDG offers a similar saving. This network presents 1005 potentials with 3645 parameters but a single value (therefore PPTs contain a single leaf node).
- IDP is the best representation for most of the networks, with substantial savings for **BN_76**, **BN_22** and **BN_20**; moderated savings for **BN_111**, **BN_109**, **BN_113**, **BN_107** and **BN_106**. For the rest of network, the percentages of increase respect to 1DA is the lowest one.
- All the proposed representations offer a competitive alternative respect to PTs and PPTs. In general, the significant savings are produced in networks with a very high number of parameters, where it is really necessary to have efficient representations to be able to apply inference algorithms.

6.3. Access time

Although the treatment of complex models implies the need of alternative models that imply worse calculation times (more calculation time is assumed in exchange for saving memory space), it is important to take into account the

speed of access to potential values and the efficiency of the operations necessary for the inference tasks. It would be completely useless to be able to represent complex models but having excessively long computation times.

Therefore, the efficiency of access to potential values is an indispensable requirement. For this reason, part of the experimentation is focused on testing this operation. This experiment is based on: a) random selection of 10000 pairs (potential, index). This set will be used for all the representations of each network. Tables including the results show times in milliseconds.

For the reliable estimation access times, the **Scalometer** library [36] has been used. This tool allows to configure the time-taking experiments ensuring that the machine has reached a steady state (after a warm-up phase); and after that, repeats several times the procedure of interest and finally offers the average time.

6.3.1. Access times for *bnlearn* networks

Access times for *bnlearn* networks are presented in Figure 12. These times show that IDM representation is very competitive, with similar times that 1DA. Times for the rest of alternatives are in most cases shorter than those required for PT and PPT, except for **barley** and **mildew**. In these networks, savings in memory space entail a more complex structure that requires longer access times.

6.3.2. Access times for *UAI* networks

Times for *UAI* networks are shown in Figure 13. In this set, *IDM* representation is the most advantageous, with times similar to those required for *1DA*. In these complex networks, where VBPs offer significant reductions in memory space, the resulting structures for *VDG*, *VDI* and *IDP* lead to worse access times. This is especially relevant for structures where the search is value-driven (*VDG* and *VDI*).

6.4. Posterior computation

Since the objective of this work is to investigate the possibilities of using VBPs in inference algorithms with Bayesian networks it is required to define marginalization and combination operations on these structures.

In general if ψ is a potential defined on \mathbf{X} and $\mathbf{Z} \subseteq \mathbf{X}$, then the marginalization of ψ in \mathbf{Z} is computed by:

$$\psi^{\#Z}(\mathbf{z}) = \sum_{\mathbf{x} \downarrow \mathbf{Z} = \mathbf{z}} \psi(\mathbf{x}); \quad \forall \mathbf{z} \in \Omega_{\mathbf{Z}} \quad (20)$$

where $\mathbf{x}^{\#Z}$ denotes the projection of \mathbf{x} configuration on \mathbf{Z} . This operation can be done by iteratively marginalizing out each variable $Y \in \mathbf{X} \setminus \mathbf{Z}$.

80.23	185.30	191.69	103.27	81.59	81.37	80.33	cancer
80.14	191.56	193.11	120.11	81.57	81.40	80.51	asia
80.11	191.88	192.26	139.74	82.32	81.58	80.34	survey
84.31	201.43	200.16	88.12	152.50	86.64	84.85	sachs
84.37	203.67	202.61	158.12	86.86	87.05	85.05	child
82.32	201.85	207.13	135.98	87.07	87.56	85.49	alarm
81.43	202.33	203.20	125.04	87.89	88.22	86.46	win95pts
84.67	210.76	205.57	90.00	143.92	88.90	85.29	insurance
81.26	205.67	211.07	93.00	186.39	149.86	86.32	hepar2
86.01	208.47	206.46	128.51	92.35	92.79	91.23	andes
80.98	200.34	208.40	144.50	138.83	195.80	85.97	hailfinder
93.17	200.01	201.96	100.76	99.86	100.47	98.79	pigs
80.64	196.65	203.51	218.65	159.43	166.19	85.47	water
86.19	194.50	192.43	102.61	177.59	127.41	90.29	munin1
102.17	217.49	208.57	110.02	108.04	109.95	107.76	link
111.52	198.04	184.45	172.79	153.92	169.81	117.50	munin2
112.97	201.51	196.53	177.70	169.53	169.66	118.90	munin3
83.46	163.16	162.41	128.73	145.64	213.23	88.15	pathfinder
114.93	212.08	216.38	166.33	144.81	174.60	118.59	munin4
112.82	221.60	214.37	173.61	152.52	173.87	119.05	munin
80.60	169.13	152.91	400.59	341.94	366.37	86.18	barley
92.42	145.48	149.97	184.06	180.03	168.23	98.21	diabetes
82.15	132.21	122.57	376.68	318.27	298.26	85.47	mildew
1DA	PT	PPT	VDG	VDI	IDP	IDM	

Figure 12: Access times for *bnlearn* networks

154.14	191.86	192.61	209.16	175.00	188.44	174.74	BN_76
95.77	144.29	133.81	243.29	206.96	485.75	89.49	BN_87
84.88	106.60	116.61	4981.93	4068.71	5569.48	85.42	BN_29
85.41	223.65	227.75	9843.88	9354.15	4515.37	115.40	BN_125
85.45	183.66	168.31	10921.45	10550.74	5028.24	86.39	BN_115
85.33	144.45	138.25	10440.35	9222.65	4640.05	86.14	BN_119
85.46	144.02	141.79	12378.30	10439.09	5782.90	86.45	BN_121
85.32	148.88	161.15	13518.58	12811.29	6856.87	86.31	BN_123
198.96	213.17	249.33	257.38	187.09	362.90	137.23	BN_27
85.28	180.07	187.33	18970.51	15463.65	8602.86	86.08	BN_117
187.97	141.40	144.54	194.30	246.33	189.03	122.83	BN_22
85.39	222.55	225.89	9683.49	8955.63	5577.46	85.79	BN_111
85.33	142.11	149.75	9743.39	9087.85	5757.92	132.04	BN_109
85.39	144.49	152.26	12387.95	8895.18	6372.12	105.94	BN_113
211.16	154.88	156.08	167.42	185.23	239.19	128.23	BN_20
85.41	185.18	173.14	12255.72	10530.21	6452.88	85.91	BN_107
85.40	142.99	160.58	13998.57	13033.97	8379.92	86.35	BN_105
1DA	PT	PPT	VDG	VDI	IDP	IDM	

Figure 13: Access times for *UAI* networks

In the case of combination operation, given two potentials $\phi_1(\mathbf{X})$ and $\phi_2(\mathbf{Y})$ then the combination of ϕ_1 and ϕ_2 is the potential denoted $\phi_1 \otimes \phi_2$ defined on $\mathbf{Z} = \mathbf{X} \cup \mathbf{Y}$ by means of pointwise multiplication:

$$VBP_1 \otimes VBP_2(\mathbf{z}) = VBP_1(\mathbf{z}^{\#X}) \otimes VBP_2(\mathbf{z}^{\#Y}); \quad \mathcal{Z} = \mathbf{X} \sqcup \mathbf{Y} \quad (21)$$

We have developed simple and direct algorithms for marginalization and combination operations (see Algorithms 5 and 6). These algorithms are based on accessing the values of the potentials (operations defined in Algorithms 1, 2, 3 and 4). This makes the access operation so important. As there is a one to one correspondence between indexes and configurations we refer to them interchangeably (as an example, l index on \mathbf{Z} corresponds to \mathbf{z}_l configuration).

Algorithm 5 (Marginalization in VBPs). *Given $VBP(\mathbf{X})$ a potential defined over \mathbf{X} and $Y \subseteq \mathbf{X}$. The method for marginalizing out Y from VBP is described in Algorithm 5. Observe that this algorithm can be employed for all VBP alternatives.*

Algorithm 5 Marginalization of Y from $VBP(\mathbf{X})$

```

1: function marginalize( $VBP ; Y$ )
2:    $\mathbf{Z} = \mathbf{X} \cap Y$  . make result domain
3:   creates  $VBP_r(\mathbf{Z})$  . empty result potential
4:   for each  $l = 1 \dots k; k = |\Omega_{\mathbf{Z}}|$  do . loop on  $VBP(\mathbf{Z})$  indexes
5:     for each  $y \in \Omega_Y$  do
6:        $\mathbf{x}_{ly} = \mathbf{z}_l^{\#Y=y}$  . get  $\mathbf{x}_{ly}$  configuration compatible with  $\mathbf{z}_l$ 
7:        $\mathbf{v}_{ly} = VBP(\mathbf{x}_{ly})$  . value in  $VBP(\mathbf{X})$  for  $\mathbf{x}_{ly}$ 
8:     end for
9:      $VBP_r(\mathbf{z}_l) = \sum_{y \in \Omega_Y} \mathbf{v}_{ly}$ 
10:  end for
11:  return  $VBP_r(\mathbf{Z})$ 
12: end function

```

Algorithm 5 removes a variable Y from $VBP(\mathbf{X})$. In lines 2 and 3 the final potential domain $\mathbf{Z} = \mathbf{X} \cap Y$ is used to create the resulting potential, which will initially be empty. Line 4 iterates over $VBP_r(\mathbf{Z})$ indexes. Let assume l corresponds to a specific configuration of variables in \mathbf{Z} (denoted by \mathbf{z}_l). Compatible indices \mathbf{x}_l refer to configurations produced completing \mathbf{z}_l with the possible values of Y . This operation is noted as $\mathbf{z}_l^{\#Y}$. Internal loop (lines 5 to 8) iterates on Y values. The sum of \mathbf{v}_{ly} values is assigned to the resulting potential (line 9).

Algorithm 6 (Combination in VBPs). *Given $VBP_1(\mathbf{X})$ and $VBP_2(\mathbf{Y})$ two potentials defined over \mathbf{X} and \mathbf{Y} . The method for combining both potentials is presented in Algorithm 6. As it happens with Algorithm 5, this is a general method applicable to all the alternatives previously described: VDG, VDI, IDP and IDM.*

Algorithm 6 combines two potentials $VBP_1(\mathbf{X})$ and $VBP_2(\mathbf{Y})$. Line 2 produces the domain \mathbf{Z} as $\mathbf{Z} = \mathbf{X} \sqcup \mathbf{Y}$. This is the domain of the potential

Algorithm 6 Combination of $VBP_1(\mathbf{X})$ and $VBP_2(\mathbf{Y})$

```

1: function combine( $VBP_1(\mathbf{X}); VBP_2(\mathbf{Y})$ )
2:    $\mathbf{Z} = \mathbf{X} \sqcup \mathbf{Y}$  . make result domain
3:   creates  $VBP_r(\mathbf{Z})$  . empty result potential
4:   for each  $l = \overline{f}0 :: kg; k = j\Omega_{\mathbf{Z}}j - 1$  do . loop on  $VBP_r(\mathbf{Z})$  indexes
5:      $v_l = 0.0$ 
6:      $\mathbf{x}_l = \mathbf{z}_l^{\#X}$  . project  $\mathbf{z}_l$  index on  $\mathbf{X}$ 
7:      $v_1 = VBP_1(\mathbf{x}_l)$  . value in  $VBP_1(\mathbf{X})$  for  $\mathbf{x}_l$ 
8:     if  $v_1 \neq 0.0$  (default value) then
9:        $\mathbf{y}_l = \mathbf{z}_l^{\#Y}$  . project  $\mathbf{z}_l$  index on  $\mathbf{Y}$ 
10:       $v_2 = VBP_2(\mathbf{y}_l)$  . value in  $VBP_2(\mathbf{Y})$  for  $\mathbf{y}_l$ 
11:      if  $v_2 \neq 0.0$  then
12:         $v_l = v_1 \vee v_2$ 
13:      end if
14:    end if
15:     $VBP_r(\mathbf{z}_l) = v_l$ 
16:  end for
17:  return  $VBP_r(\mathbf{Z})$ 
18: end function

```

$VBP_r(\mathbf{Z})$ to return. Line 4 iterates over $VBP_r(\mathbf{Z})$ indexes. Let be l the index under consideration (it corresponds to a given configuration \mathbf{z}_l). The value to assign to l is initialized to 0.0 (line 5). The configuration \mathbf{z}_l must be projected into $VBP_1(\mathbf{X})$ (line 6) and $VBP_2(\mathbf{Y})$ (line 9). These operation are termed $\mathbf{z}_l^{\#X}$ and $\mathbf{z}_l^{\#Y}$ and consists of removing from \mathbf{z}_l the values of the variables that do not belong to \mathbf{X} and \mathbf{Y} respectively. The index of configuration \mathbf{x}_l is employed for getting v_1 (line 7). If v_1 is 0.0 then for sure $v_l = 0.0$ and no more operations are required. Otherwise, it is also necessary to access $VBP_2(\mathbf{y}_l)$ (see line 10). Finally v_l is assigned to $VBP_r(\mathbf{Z})$ in line 15.

This section presents the computation times required for obtaining the posterior on 10 randomly selected variables on each *bnlearn* network with Variable Elimination algorithm and using the algorithms for marginalization and combination previously described. Experiments have been limited *bnlearn* networks because for most of *UAI* networks computations with 1DA, PT and PPT overcomes the memory capacity of the computer used for the experimental work.

In any case, it is important to highlight that these experiments aim to obtain a first idea of the behavior of VBPs structures. A future line of work will be to carry out specific implementations for marginalization and combination operations, which take into account the special properties of each representation.

We have selected one alternative for each category: VDI for value-driven approximation and IDM for index-driven approach. These two representations are the ones that show the best compromise between memory use and access times within their category. They are compared respect to 1DA and PT (when

using trees there is not much difference between PT and PPT in general). We have also employed **Scalometer** library for measuring computation times. The results for this section are presented in Figure 14.

1.60	0.47	0.97	1.29	cancer	12000
1.84	2.19	1.60	1.51	asia	10000
1.57	0.52	1.42	1.27	survey	8000
5.74	1.22	5.34	4.17	sachs	6000
3.58	1.58	4.36	4.07	child	4000
10.91	4.87	15.45	13.29	alarm	2000
4.60	3.29	5.09	4.85	win95pts	
21.68	5.01	30.78	17.83	insurance	
14.54	8.31	12.53	16.39	hepar2	
46.87	24.87	36.04	34.25	andes	
21.87	7.55	50.47	20.77	hailfinder	
49.85	42.20	41.73	46.55	pigs	
3965.73	249.52	5932.28	1120.53	water	
54.75	17.79	104.08	60.12	munin1	
139.74	121.00	141.60	144.76	link	
183.35	192.84	208.69	202.05	munin2	
214.49	194.18	195.84	216.74	munin3	
18.24	5.31	36.65	13.92	pathfinder	
342.58	206.08	646.60	380.25	munin4	
231.36	205.52	257.04	225.19	munin	
1237.16	58.86	9599.25	2435.81	barley	
2967.32	615.03	12116.15	3647.13	diabetes	
66.68	7.35	363.57	54.98	mildew	
1DA	PT	VDI	IDM		

Figure 14: VE times for *bnlearn* networks

Regarding these results, it is observed that the inference with PT is the most efficient one. This is because PT has specific implementations of marginalization and combination operations (see [34]). The implementation of these methods is recursive. It should be noted that in some of *UAI* networks, the execution of the algorithm produces a stack overflow error when PTs with many variables are produced. In these cases the evaluation with 1DA also fails producing out of memory errors.

It is also observed that, in most cases, times for IDM are similar to those for 1DA. This is remarkable and it is possible to think that more refined implementations of marginalize and combine operations, which take into account their structure, will improve current computation times. More efficient implementations of these operations should try to avoid iterating over all indexes of the resulting potential, limiting it to those stored. This, in some of the networks, can suppose significant reductions of time.

7. Conclusion

Regarding the use of memory space, it is observed that all the alternatives proposed offer competitive results with respect to 1DA, PT and PPT. In most

of the networks, VDI (value-driven) and IDP (index-driven) alternatives stand out. In terms of access times, the best alternative is IDM. For this reason, this representation was selected as an alternative for VE tests.

The basic versions of marginalization and combination allows to observe that VDI and IDM also offer reasonable execution times, similar to those necessary for 1DA. We think that these results are promising and that more efficient implementations will produce better results. This task will be the subject for a line of future research.

Although it has not been used in this work, it should be noted that another important feature of VBPs lies in its ability to be approximated. The approximation operation assumes loss of information. Intuitively, the idea is to group nearby values and replace them with their average (or some other measure), so that repetition patterns are expanded and, therefore, the number of values to be stored is reduced. With this operation any algorithm that involves the use of approximated potentials will become approximate as well and will ultimately offer non-exact solutions. In very complex problems it is always better to have at least one approximate solution (see [14, 6, 4, 5] as examples of approximation with PTs).

The software used in this paper is implemented with **Scala** programming language and it is available at <https://github.com/mgomez-olmedo/bnetSbtV2> with explanations about how to reproduce the experiments. The functional programming paradigm (combined with object orientation) offered by **Scala** can be exploited for getting well defined operations easily converted into parallel ones on multi-core CPUs when possible. Some of these benefits are investigated in [27].

Acknowledgements

This research was jointly supported by the Spanish Ministry of Education and Science under projects PID2019-106758GB-C31 and TIN2016-77902-C3-2-P and the European Regional Development Fund (FEDER).

References

- [1] A. M. Alaa and M. van der Schaar. Bayesian Nonparametric Causal Inference: Information Rates and Learning Algorithms. *IEEE Journal of Selected Topics in Signal Processing*, **2018**, doi: 10.1109/JSTSP.2018.2848230.
- [2] M. Arias, F. Díez. Operating with potentials of discrete variables. *International Journal of Approximate Reasoning*, **2007**, *46* (1), pp. 166–187.
- [3] C. Boutilier, N. Friedman, M. Goldszmidt, D. Koller. Context-specific independence in Bayesian networks. *In: Proceedings of the 12th International Conference on Uncertainty in Artificial Intelligence, Morgan Kaufmann Publishers Inc.*, **1996**, pp. 115–123.

- [4] R. Cabañas, M. Gómez, A. Cano. Approximate inference in influence diagrams using binary trees. *In: Proceedings of the Sixth European Workshop on Probabilistic Graphical Models (PGM-12)*, **2012**.
- [5] R. Cabañas, M. Gómez-Olmedo, A. Cano. Using binary trees for the evaluation of influence diagrams. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, **2016**, *24 (01)*, pp. 59–89.
- [6] A. Cano, M. Gómez-Olmedo, S. Moral. Approximate inference in Bayesian networks using binary probability trees. *International Journal of Approximate Reasoning*, **2011**, *52 (1)*, pp. 49–62.
- [7] A. Cano, S. Moral, A. Salmerón, Penniless propagation in join trees. *International Journal of Intelligent Systems*, **2000**, *15 (11)*, pp. 1027–1059.
- [8] M. Chavira, A. Darwiche. Compiling Bayesian networks using variable elimination. *In: Proceedings of the 20th International Joint Conference on Artificial Intelligence*, **2007**, pp. 2443–2449.
- [9] A. Choi, D. Kisa, A. Darwiche. Compiling probabilistic graphical models using sentential decision diagrams. *In: European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, Springer, **2013**, pp. 121–132.
- [10] P. Dagum, M. Luby, An optimal approximation algorithm for Bayesian inference. *Artificial Intelligence*, **1997**, *93 (1)*, pp. 1–27.
- [11] A. P. David, Applications of a general propagation algorithm for probabilistic expert systems. *Statistics and Computing*, **1992**, doi:10.1007/BF01890546.
- [12] R. Dechter, Bucket elimination: A unifying framework for probabilistic inference. *In: Learning in graphical models*, Springer, **1998**, pp. 75–104.
- [13] F.J. Díez, M. Luque, M. Arias, J. Pérez-Martín. Cost-Effectiveness Analysis with Unordered Decisions. *Artificial Intelligence in Medicine*, **2021**, <https://doi.org/10.1016/j.artmed.2021.102064>.
- [14] M. Gómez-Olmedo, A. Cano. Applying numerical trees to evaluate asymmetric decision problems. *In: T. Nielsen, N. Zhang (Eds.), Symbolic and Quantitative Approaches to Reasoning with Uncertainty, Vol. 2711 of Lecture Notes in Computer Science*, Springer Berlin Heidelberg, **2003**, doi:10.1007/978-3-540-45062-7_16.
- [15] M. Gómez-Olmedo. Real-World Applications of Influence Diagrams. *In: J.A. Gámez, S. Moral and A. Salmerón (eds) Advances in Bayesian Networks. Studies in Fuzziness and Soft Computing*, Springer, Berlin, Heidelberg, **2004**, doi:10.1007/978-3-540-39879-0_9.

- [16] D. Heckerman, A. Mamdani, M.P. Wellman. Real-World Applications of Bayesian Networks. *In: Association for Computing Machinery*, **1995**, doi: 10.1145/203330.203334.
- [17] R. A. Howard, J. E. Matheson. Influence diagram retrospective. *Decision Analysis*, **2005**, 2 (3), 144–147.
- [18] C. S. Jensen, U. Kjærulff, A. Kong, Blocking Gibbs sampling in very large probabilistic expert systems. *International Journal of Human-Computer Studies*, **1995**, 42 (6), pp. 647–666.
- [19] F. V. Jensen, T. D. Nielsen, Bayesian networks and decision graphs. *Springer Verlag*, **2007**.
- [20] D. Koller, N. Friedman. Probabilistic graphical models: principles and techniques. *MIT Press*, **2009**.
- [21] J. Kwisthout. Most probable explanations in Bayesian networks: Complexity and tractability. *Int. J. Approx. Reasoning*, **2011** doi:10.1016/j.ijar.2011.08.003.
- [22] S. L. Lauritzen. Graphical Models. *Oxford University Press*, **1996**.
- [23] W. Lee and N. Zabaraz. Parallel probabilistic graphical model approach for nonparametric Bayesian inference. *Journal of Computational Physics*, **2018**, doi: 10.1016/j.jcp.2018.06.057.
- [24] Z. Li, B. D’Ambrosio, Efficient inference in Bayes networks as a combinatorial optimization problem. *International Journal of Approximate Reasoning*, **1994**, 11 (1), pp. 55–81.
- [25] A. L. Madsen, F. V. Jensen, Lazy evaluation of symmetric Bayesian decision problems. *In: Proceedings of the 15th Conference on Uncertainty in AI, Morgan Kaufmann Publishers Inc.*, **1999**, pp. 382–390.
- [26] A. L. Madsen, F. V. Jensen, Lazy propagation: a junction tree inference algorithm based on lazy evaluation. *Artificial Intelligence*, **2004**, 113 (1-2), pp. 203–245.
- [27] A. R. Masegosa, A. M. Martinez, H. Borchani. Probabilistic graphical models on multi-core cpus using Java 8. *IEEE Computational Intelligence Magazine*, **2016**, 11 (2), pp. 41–54.
- [28] S. M. Olmsted. On Representing an Solving Decision Problems. *PhD Thesis, Department of Engineering-Economic Systems, Stanford University*, **1983**.
- [29] U. Oztok, A. Darwiche. A top-down compiler for sentential decision diagrams. *In: Proceedings of the 24th International Conference on Artificial Intelligence*, **2015**, pp. 3141–3148.

- [30] J. Pearl. Bayesian networks: A model of self-activated memory for evidential reasoning. University of California (Los Angeles). Computer Science Department, **1985**.
- [31] J. Pearl. Probabilistic reasoning in Intelligent Systems: Networks of Plausible Inference. *Morgan Kaufmann*, **1988**.
- [32] J. Pearl, S. Russell. Bayesian networks. *Computer Science Department, University of California*, **1998**.
- [33] J. Pearl, Evidential reasoning using stochastic simulation of causal models. *Artificial Intelligence*, **1987**, *32 (2)*, pp. 245–257.
- [34] A. Salmerón, A. Cano, S. Moral, Importance sampling in Bayesian networks using probability trees, *Computational Statistics & Data Analysis* *34 (4)* (2000) 387–413.
- [35] S. Sanner, D. McAllester. Affine algebraic decision diagrams (AADDs) and their application to structured probabilistic inference. *In: Proceedings of the 19th International Joint Conference on Artificial Intelligence*, **2005**, pp. 1384–1390.
- [36] Scalometer: automate your performance testing today, **2021**, <http://https://scalometer.github.io/>.
- [37] M. Scutari. Learning Bayesian networks with the bnlearn R package. *Journal of Statistical Software*, **2010**, doi:10.18637/jss.v035.i03.
- [38] M. Scutari. Bayesian network constraint-based structure learning algorithms: Parallel and optimized implementations in the bnlearn R package. *Journal of Statistical Software*, **2017**, doi:10.18637/jss.v077.i02.
- [39] R. D. Shachter, Evaluating influence diagrams. *Operations Research*, **1986**, *34 (6)*, pp. 871–882.
- [40] R. D. Shachter, B. D’Ambrosio, B. Del Favero, Symbolic probabilistic inference in belief networks. *In: AAAI Proceedings*, **1990**, pp. 126–131.
- [41] P. Shenoy, G. R. Shafer. Axioms for probability and belief-function propagation. *In: R. D. Shachter, T. S. Levitt, L. N. Kanal, J. F. Lemmer (Eds.), Uncertainty in Artificial Intelligence, Vol. 9 of Machine Intelligence and Pattern Recognition*, **1990**, doi.org/10.1016/B978-0-444-88650-7.50019-6.
- [42] UAI 2016 Inference Competition, **2016**, <http://www.hlt.utdallas.edu/~vgogate/uai16-evaluation/>
- [43] UAI 2014 Inference Competition, **2014**, <http://www.hlt.utdallas.edu/~vgogate/uai14-competition/index.html>

- [44] T. L. Wong, H. Xie, W. Lam, and F. L. Wang. A learning framework for information block search based on probabilistic graphical models and Fisher Kernel. *International Journal on Machine Learning and Cybernetics*, **2018**, doi: 10.1007/s13042-017-0657-9.
- [45] L. Yang and Y. Guo. Combining pre- and post-model information in the uncertainty quantification of non-deterministic models using an extended Bayesian melding approach. *Information Sciences*, **2019**, doi: 10.1016/j.ins.2019.06.029.
- [46] N. L. Zhang, D. Poole. Exploiting causal independence in Bayesian network inference. *Journal of Artificial Intelligence Research*, **1996**, *5*, pp. 301–328.