

Tensor rank-one decomposition of probability tables

Petr Savicky

Inst. of Comp. Science
Academy of Sciences of the Czech Rep.
Pod vodárenskou věží 2
182 07 Prague, Czech Rep.
<http://www.cs.cas.cz/~savicky/>

Jiří Vomlel

Inst. of Inf. Th. and Automation
Academy of Sciences of the Czech Rep.
Pod vodárenskou věží 4
182 08 Prague, Czech Rep.
<http://staff.utia.cz/vomlel>

Abstract

We propose a new additive decomposition of probability tables - tensor rank-one decomposition. The basic idea is to decompose a probability table into a series of tables, such that the table that is the sum of the series is equal to the original table. Each table in the series has the same domain as the original table but can be expressed as a product of one-dimensional tables. We show that tensor rank-one decomposition can be used to reduce the space and time requirements in probabilistic inference. We provide a closed form solution for minimal tensor rank-one decomposition for some special tables.

Keywords: Graphical Probabilistic Models, Probabilistic Inference.

1 Introduction

A fundamental property of probabilistic graphical models that allows their application in domains with hundreds to thousands variables is the multiplicative factorization of the joint probability distribution. The multiplicative factorization is exploited in inference methods, e.g., in the junction tree propagation [7]. However, in some real applications the models may become intractable using the junction tree propagation and other exact inference methods because after the moralization and triangularization steps the graphical

structure becomes too dense, cliques consist of too many variables, and, consequently, the probability tables corresponding to the cliques are too large to be efficiently manipulated. In such case one usually turns to an approximative inference method.

Following the ideas presented in [3, 9] we propose a new decomposition of probability tables that allows to use exact inference in some models where – without the suggested decomposition – the exact inference using the standard methods is impossible. The basic idea is to decompose a probability table into a series of tables, such that the table that is the sum of the series is equal to the original table. Each table in the series has the same domain as the original table but can be expressed as a product of one-dimensional tables. Entries in tables are allowed to be any real number. Extending the range to negative numbers opens new possibilities for compact representation of probability tables and allows to find a shorter series.

It is convenient to formally specify the task using the tensor terminology. Assume variables X_i , $i \in N \subset \mathbb{N}$ each variable X_i taking values (a value of X_i will be denoted x_i) from a finite set \mathcal{X}_i . Let for any $A \subseteq N$ the symbol x_A denotes a vector of the values $(x_i)_{i \in A}$, where for all $i \in A$: x_i is a value from \mathcal{X}_i .

Definition 1 Tensor

Let $A \subset N$. *Tensor* ψ over A is a mapping

$$\times_{i \in A} \mathcal{X}_i \mapsto \mathbb{R} .$$

The cardinality $|A|$ is called *tensor dimension*.

Note that every probability table can be looked upon as a tensor. Tensor ψ over A is an (unconditional) probability table if for every x_A it holds that $0 \leq \psi(x_A) \leq 1$ and $\sum_{x_A} \psi(x_A) = 1$. Tensor ψ is a conditional probability table (CPT) if for every x_A it holds that $0 \leq \psi(x_A) \leq 1$ and if there exists $B \subset A$ such that for every x_B it holds $\sum_{x_{A \setminus B}} \psi(x_B, x_{A \setminus B}) = 1$.

Next, we will recall the basic tensor notion. If $|A| = 1$ then tensor is a vector. If $|A| = 2$ then tensor is a matrix. The *outer product* $\psi \otimes \varphi$ of two tensors $\psi : \times_{i \in A} \mathcal{X}_i \mapsto \mathbb{R}$ and $\varphi : \times_{i \in B} \mathcal{X}_i \mapsto \mathbb{R}$, $A \cap B = \emptyset$ is a tensor $\xi : \times_{i \in A \cup B} \mathcal{X}_i \mapsto \mathbb{R}$ defined for all $x_{A \cup B}$ as

$$\xi(x_{A \cup B}) = \psi(x_A) \cdot \varphi(x_B) .$$

Now, let ψ and φ are defined on the same domain $\times_{i \in A} \mathcal{X}_i$. The sum $\psi + \varphi$ of two tensors is tensor $\xi : \times_{i \in A} \mathcal{X}_i \mapsto \mathbb{R}$ such that

$$\xi(x_A) = \psi(x_A) + \varphi(x_A) .$$

Definition 2 Tensor rank [6]

Tensor of dimension $|A|$ has *rank* one if it is an outer product of $|A|$ vectors. *Rank* of tensor ψ is the minimal number of tensors of rank one that sum to ψ . *Rank* of tensor ψ will be denoted as $rank(\psi)$.

Remark Note that the standard matrix rank is a special case of the tensor rank for $|A| = 2$. An alternative definition of the matrix rank is that the rank of an $m \times n$ matrix \mathbf{M} is the smallest r for which there exist a $m \times r$ matrix \mathbf{F} and a $r \times n$ matrix \mathbf{G} such that

$$\mathbf{M} = \mathbf{F}\mathbf{G} . \tag{1}$$

We can equivalently write that

$$\mathbf{M} = \sum_{i=1}^r \mathbf{f}^i \otimes \mathbf{g}^i , \tag{2}$$

where \mathbf{f}^i is the i^{th} column of \mathbf{F} and \mathbf{g}^i is the i^{th} row of \mathbf{G} . Since $rank(\mathbf{f}^i \otimes \mathbf{g}^i) = 1$ we can see that the tensor rank of \mathbf{M} is equal to the matrix rank of \mathbf{M} .

Definition 3 Tensor rank-one decomposition

Assume a tensor ψ over A and an integer r . A series of tensors $\{\varrho_b\}_{b=1}^r$ over A such that

- for $b = 1, \dots, r$: $rank(\varrho_b) = 1$, i.e.,

$$\varrho_b = \otimes_{i \in A} \varphi_{i,b} ,$$

where $\varphi_{i,b}, i \in A$ are vectors and

- $\psi = \sum_{b=1}^r \varrho_b$

is called *tensor rank-one decomposition* of ψ of length r .

Note that from the definition of tensor rank it follows that such a series exists iff $r \geq rank(\psi)$. The decomposition is minimal if $r = rank(\psi)$.

Example 1 Let $\psi : \{0, 1\} \times \{0, 1\} \times \{0, 1\} \mapsto \mathbb{R}$ be

$$\left(\begin{array}{c} \left(\begin{array}{c} 1 \\ 2 \end{array} \right) \\ \left(\begin{array}{c} 2 \\ 4 \end{array} \right) \end{array} \right) \cdot \left(\begin{array}{c} \left(\begin{array}{c} 2 \\ 4 \end{array} \right) \\ \left(\begin{array}{c} 4 \\ 9 \end{array} \right) \end{array} \right) .$$

This tensor has rank two since

$$\begin{aligned} \psi &= (1, 2) \otimes (1, 2) \otimes (1, 2) \\ &\quad + (0, 1) \otimes (0, 1) \otimes (0, 1) \end{aligned}$$

and there are no three vectors whose outer product is equal to ψ . \diamond

The rest of the paper is organized as follows. In Section 2 we show using a simple example how tensor rank-one decomposition can be used to reduce the space and time requirements for the probability inference using the junction tree method. We compare sizes of the junction tree for the standard approach, the parent divorcing method, and the junction tree after tensor rank-one decomposition¹. In Section 3 the main theoretical results are presented: the lower bound on the tensor rank for a class of tensors and minimal tensor rank-one decompositions for some special tensors – max, add, xor, and their noisy counterparts.

¹Several other methods were proposed to exploit a special structure of CPTs. For a review of these methods see, for example, [9]. In this paper, due to the lack of space, we do comparisons with the parent divorcing method only.

2 Tensor rank-one decomposition in probabilistic inference

We will use an example of a simple Bayesian network to show computational savings of the proposed decomposition. Assume a Bayesian network having the structure given in Figure 1. Variables X_1, \dots, X_m are binary taking values 0 and 1. For simplicity, assume that $m = 2^d, d \in \mathbb{N}, 2 \leq d$. Further assume a variable $Y \stackrel{df}{=} X_{m+1}$, whose values $y = \sum_{i=1}^m x_i$. This means that Y takes $m + 1$ values.

If we use the standard junction tree construction [7] we connect by edges all parents of Y (this step is often called moralization). We need not perform triangulation since the graph is already triangulated. The resulting junction tree consists of one clique containing all variables, i.e., $C_1 = \{X_1, \dots, X_m, Y\}$.




Figure 1: Bayesian network structure.

Since the CPT $P(Y | X_1, \dots, X_m)$ has a special form we can use the parent divorcing method [8] and introduce a number of auxiliary variables, one auxiliary variable for a pair of parent variables. This is used hierarchically, i.e. we get a tree of auxiliary variables with node Y being the root of the tree. The resulting junction tree consists of $m - 1$ cliques.

In Section 3.2 we will show that if the CPT corresponds to addition of m binary variables then we can decompose this CPT to a series of $m + 1$ tensors that are products of vectors

$$P(Y | X_1, \dots, X_m) = \sum_{b=1}^{m+1} \xi_b \otimes (\otimes_{i=1}^m \varphi_{i,b}) .$$

As suggested in [3] we can visualize an additive decomposition using one additional variable, which we will denote B . In case of addition of m binary variables variable B will have $m + 1$ states. Instead of moralization

we add variable B into the model and connect it with nodes corresponding to variables Y, X_1, \dots, X_m . We get the structure given in Figure 2. Note that all edges are undirected, which means that we do not perform any moralization. It is not difficult to show (see [9]) that this model can be used to compute marginal probability distributions as in the original model. The resulting junction tree of this model is given in Figure 3.




Figure 2: Bayesian network after the decomposition




Figure 3: Junction tree for the model after the rank-one decomposition

After little algebra we get that the total clique size in the standard case is $(m + 1) \cdot 2^m$, after parent divorcing it is $\frac{1}{3}m^3 + \frac{5}{2}m^2 + 2m \log m - \frac{11}{6}m - 1$, and after the tensor rank-one decomposition (described latter in this paper) it is only $3m^2 + 4m + 1$. In Figure 4 we compare dependence of the total size of junction trees on the number of parent nodes² m of node Y . Note that we use the logarithmic scale for the vertical axis.

It should be noted that the tensor rank-one decomposition can be applied to any probabilistic model. The savings depend on the graphical structure of the probabilistic model.

²It may seem unrealistic to have a node with more than ten parents in an real world application, but it can easily happen, especially, when we need to introduce logical constraints into the model.

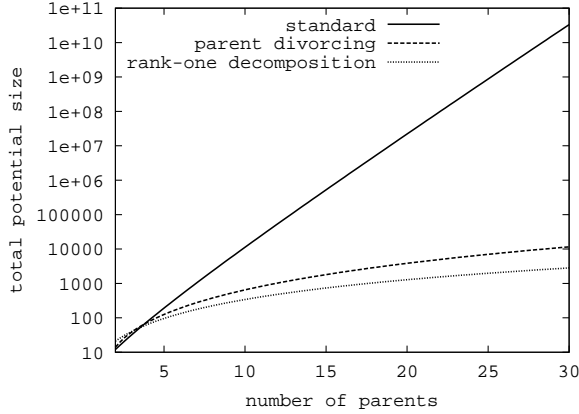


Figure 4: Comparison of the total size of junction tree.

In fact, by avoiding the moralization of the parents, we give the triangulation algorithm more freedom for the construction of a triangulated graph so that the resulting tables corresponding to the cliques of the triangulated graph can be smaller.

3 Minimal tensor rank-one decomposition

In this section we will present the main theoretical results. It was proved in [6] that the computation of tensor rank is an NP-hard problem, therefore determining the minimal rank-one decomposition is also an NP-hard problem. However, we will provide closed-form solution for the problem of finding a minimal rank-one decomposition of some special tensors that play an important role, since they correspond to CPTs that are often used in Bayesian network models in real applications.

The class of tensors of our special interest are tensors ψ_f that represent a functional dependence of one variable Y on variables X_1, \dots, X_m . Let $\mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_m$ and $x = (x_1, \dots, x_m) \in \mathcal{X}$. Further, let

$$I(expr) = \begin{cases} 1 & \text{if } expr \text{ is true} \\ 0 & \text{otherwise.} \end{cases}$$

Then for a function $f : \mathcal{X} \mapsto \mathcal{Y}$ the tensor is defined for all $(x, y) \in \mathcal{X} \times \mathcal{Y}$ as $\psi_f(x, y) = I(y = f(x))$.

Let $r = \text{rank}(\psi_f)$ and

$$\psi_f = \sum_{b=1}^r \xi_b \otimes (\otimes_{i=1}^m \varphi_{i,b}) , \quad (3)$$

where $\xi_b : \mathcal{Y} \mapsto \mathbb{R}$ and $\varphi_{i,b} : \mathcal{X}_i \mapsto \mathbb{R}$ for all $b \in \{1, \dots, r\}$. Formula (3) is called a minimal tensor rank-one decomposition of ψ_f .

First, we will provide a lower bound on the rank of tensors from this class. This bound will be later used to prove the minimality of certain rank-one decompositions.

Lemma 1 *Let function $f : \mathcal{X} \mapsto \mathcal{Y}$ and $\psi_f : \mathcal{X} \times \mathcal{Y} \mapsto \{0, 1\}$ be a tensor representing the functional dependence given by f . Then $\text{rank}(\psi_f) \geq |\mathcal{Y}|$.*

Proof For a minimal tensor rank-one decomposition of ψ_f it holds for all $(x, y) \in \mathcal{X} \times \mathcal{Y}$ that

$$\psi_f(x, y) = \sum_{b=1}^r \xi_b(y) \cdot \prod_{i=1}^m \varphi_{i,b}(x_i) , \quad (4)$$

where $r = \text{rank}(\psi_f)$. Consider the matrices³

$$\begin{aligned} \mathbf{W} &= \{\psi_f(x, y)\}_{x \in \mathcal{X}}^{y \in \mathcal{Y}} , \\ \mathbf{U} &= \{\xi_b(y)\}_{b \in \{1, \dots, r\}}^{y \in \mathcal{Y}} , \\ \mathbf{V} &= \left\{ \prod_{i=1}^m \varphi_{i,b}(x_i) \right\}_{x \in \mathcal{X}}^{b \in \{1, \dots, r\}} . \end{aligned}$$

Equation (4) can be rewritten as

$$\mathbf{W} = \mathbf{U} \mathbf{V} .$$

Each row of \mathbf{W} contains at least one nonzero entry, since each $y \in \mathcal{Y}$ is in the range of f . Moreover, each column of \mathbf{W} contains exactly one nonzero entry, since f is a function. Hence, no row is a linear combination of other rows. Therefore there are $|\mathcal{Y}|$ independent rows in \mathbf{W} and $\text{rank}(\mathbf{W}) = |\mathcal{Y}|$. Since the rank of a matrix product cannot be higher than the rank of matrices in the product [4] we get that $\text{rank}(\mathbf{W}) \leq \text{rank}(\mathbf{U}) \leq r$. Altogether, $|\mathcal{Y}| \leq r$. \square

³The upper index labels the rows and the lower index labels the columns.

3.1 Maximum and minimum

An additive decomposition of max was originally proposed in [3]. This result is included in this section for completeness and we add a result about its optimality. The proofs are constructive, i.e., they provide a minimal tensor rank-one decomposition for *max* and *min*.

Let us assume that $\mathcal{X}_i = [a_i, b_i]$ is an interval of integers for each $i = 1, \dots, m$. Clearly, the range \mathcal{Y} of max on $\mathcal{X}_1 \times \dots \times \mathcal{X}_m$ is $[\max_{i=1}^m a_i, \max_{i=1}^m b_i]$ and the range \mathcal{Y} of min is $[\min_{i=1}^m a_i, \min_{i=1}^m b_i]$.

Theorem 1 *If $f(x) = \max\{x_1, \dots, x_m\}$ and $x_i \in [a_i, b_i]$ for $i = 1, \dots, m$ then $\text{rank}(\psi_f) = |\mathcal{Y}|$.*

Proof Let for $i \in \{1, \dots, m\}$, $x_i \in \mathcal{X}_i$, and $b \in \mathcal{Y}$

$$\varphi_{i,b}(x_i) = \begin{cases} 1 & x_i \leq b \\ 0 & \text{otherwise} \end{cases}$$

and for $y \in \mathcal{Y}$, $b \in \mathcal{Y}$

$$\xi_b(y) = \begin{cases} +1 & b = y \\ -1 & b = y - 1 \\ 0 & \text{otherwise.} \end{cases}$$

Let $\omega(x, y) = I(\max(x_1, \dots, x_m) \leq y)$. Put

$$\omega(x, y) = \prod_{i=1}^m \varphi_{i,y}(x_i) .$$

Therefore,

$$\psi_f(x, y) = \omega(x, y) - \omega(x, y - 1),$$

where $\omega(x, y_{\min} - 1)$ is considered to be zero. Altogether,

$$\psi_f(x, y) = \prod_{i=1}^m \varphi_{i,y}(x_i) - \prod_{i=1}^m \varphi_{i,y-1}(x_i) .$$

Since the product $\xi_b(y) \cdot \prod_{i=1}^m \varphi_{i,b}(x_i)$ is nonzero only for $b = y$ and $b = y - 1$, we have

$$\begin{aligned} \psi_f(x, y) &= \xi_y(y) \cdot \prod_{i=1}^m \varphi_{i,y}(x_i) \\ &\quad + \xi_{y-1}(y) \cdot \prod_{i=1}^m \varphi_{i,y-1}(x_i) \\ &= \sum_{b \in \mathcal{Y}} \xi_b(y) \cdot \prod_{i=1}^m \varphi_{i,b}(x_i) . \end{aligned}$$

Taking $b' = b + y_{\min} - 1$ we get the required decomposition

$$\psi_f = \sum_{b'=1}^{|\mathcal{Y}|} \xi_{b'} \otimes \left(\otimes_{i=1}^m \varphi_{i,b'} \right) .$$

By Lemma 1, this is a minimal tensor rank-one decomposition of ψ_f . \square

Theorem 2 *If $f(x) = \min\{x_1, \dots, x_m\}$ and $x_i \in [a_i, b_i]$ for $i = 1, \dots, m$ then $\text{rank}(\psi_f) = |\mathcal{Y}|$.*

Proof Let for $i \in \{1, \dots, m\}$, $x_i \in \mathcal{X}_i$, and $b \in \mathcal{Y}$

$$\varphi_{i,b}(x_i) = \begin{cases} 1 & x_i \geq b \\ 0 & \text{otherwise} \end{cases}$$

and for $y \in \mathcal{Y}$, $b \in \mathcal{Y}$

$$\xi_b(y) = \begin{cases} +1 & b = y \\ -1 & b = y + 1 \\ 0 & \text{otherwise.} \end{cases}$$

and follow an analogous argument as in the proof of Theorem 1 to obtain tensor rank-one decomposition of ψ_f . Again, by Lemma 1, this is a minimal tensor rank-one decomposition. \square

Remark If for $i \in \{1, \dots, m\}$ $\mathcal{X}_i = \{0, 1\}$, then the functions $\max\{x_1, \dots, x_m\}$ and $\min\{x_1, \dots, x_m\}$ correspond to logical disjunction $x_1 \vee \dots \vee x_m$ and logical conjunction $x_1 \wedge \dots \wedge x_m$, respectively. In Example 2 we illustrate how this can be generalized to Boolean expressions consisting of negations and disjunctions.

Example 2 In order to achieve minimal tensor rank-one decomposition of

$$\psi(x_1, x_2, y) = I(y = (x_1 \vee \neg x_2))$$

with variable B having two states 0 and 1, it is sufficient to use functions:

$$\begin{aligned} \varphi_{1,b}(x_1) &= I(x_1 \leq b) \\ \varphi_{2,b}(x_2) &= I(\neg x_2 \leq b) \\ \xi_b(y) &= \begin{cases} +1 & y = b \\ -1 & y = 1, b = 0 \\ 0 & y = 0, b = 1 \end{cases} \end{aligned}$$

\diamond

3.2 Addition

In this section, we assume an integer r_i for each $i = 1, \dots, m$ and assume that \mathcal{X}_i is the interval of integers $[0, r_i]$. This assumption is made for simplicity and without the loss of generality. If \mathcal{X}_i are intervals of integers, which do not start at zero, it is possible to transform the variables by subtracting the lower bounds of the intervals to obtain variables satisfying the assumption. Moreover, let $f : \mathbb{N}^m \rightarrow \mathbb{N}$ be a function, such that $f(x) = f_0(\sum_{i=1}^m x_i)$ where $f_0 : \mathbb{N} \rightarrow \mathbb{N}$. Let \mathcal{A} be the interval of integers $[0, \sum_{i=1}^m r_i]$. Clearly, \mathcal{A} is the range of $\sum_{i=1}^m x_i$.

Theorem 3 *Let f_0, f and \mathcal{A} be as above. Then $\text{rank}(\psi_f) \leq |\mathcal{A}|$. Moreover, if f_0 is the identity function, then $\text{rank}(\psi_f) = |\mathcal{A}|$.*

Proof Consider the Vandermonde matrix

$$\begin{pmatrix} \alpha_1^0 & \alpha_2^0 & \dots & \alpha_{|\mathcal{A}|}^0 \\ \alpha_1^1 & \alpha_2^1 & \dots & \alpha_{|\mathcal{A}|}^1 \\ \dots & \dots & \dots & \dots \\ \alpha_1^{|\mathcal{A}|-1} & \alpha_2^{|\mathcal{A}|-1} & \dots & \alpha_{|\mathcal{A}|}^{|\mathcal{A}|-1} \end{pmatrix}, \quad (5)$$

where $\alpha_1, \dots, \alpha_{|\mathcal{A}|}$ are pairwise distinct real numbers and the upper index is the exponent. This matrix is non-singular, since the corresponding Vandermonde determinant is non-zero [4]. The system of equations

$$\beta_t = \sum_{b=1}^{|\mathcal{A}|} z_b \cdot \alpha_b^t, \quad t \in \mathcal{A} \quad (6)$$

has the Vandermonde matrix as its system matrix and since the Vandermonde determinant is non-zero it has always a solution for variables $z_b, b \in \mathcal{A}$.

For a fixed $y \in \mathcal{Y}$ define $\xi_b(y)$ to be the solution of system (6) with $\beta_t = I(y = f_0(t))$. Therefore it holds for a fixed $y \in \mathcal{Y}$ and for all $t \in \mathcal{A}$ that

$$I(y = f_0(t)) = \sum_{b=1}^{|\mathcal{A}|} \xi_b(y) \cdot \alpha_b^t. \quad (7)$$

By substituting $t = \sum_{i=1}^m x_i$ and taking $\varphi_{i,b}(x_i) = \alpha_b^{x_i}$ for $b = 1, \dots, |\mathcal{A}|$ we obtain

that for all combinations of the values of x and y

$$I(y = f_0(\sum_{i=1}^m x_i)) = \sum_{b=1}^{|\mathcal{A}|} \xi_b(y) \cdot \prod_{i=1}^m \varphi_{i,b}(x_i).$$

This proves the first assertion of the theorem.

If f_0 is the identity, then the range of f is the whole \mathcal{A} . It follows from Lemma 1 that $\text{rank}(\psi_f) \geq |\mathcal{A}|$ and therefore the above decomposition is minimal. \square

Example 3 Let $\mathcal{X}_i = \{0, 1\}$ for $i = 1, 2$, $f(x_1, x_2) = x_1 + x_2$ and $\mathcal{Y} = \{0, 1, 2\}$. We have

$$\begin{aligned} \psi_f(x_1, x_2, y) &= I(y = x_1 + x_2) \\ &= \begin{pmatrix} \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} & \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \end{pmatrix} \\ \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} & \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \end{pmatrix} \end{aligned}$$

As in the proof of Theorem 3, we assume $\varphi_{i,b}(x_i) = \alpha_b^{x_i}$ for $i = 1, 2$ and distinct $\alpha_b, b = 0, 1, 2$. For simplicity of notation, let us assume $\alpha_0 = \alpha, \alpha_1 = \beta$ and $\alpha_2 = \gamma$. Let us substitute these $\varphi_{i,b}(x_i)$ into (3) and rewrite it using tensor product as follows.

$$\begin{aligned} \psi_f(x_1, x_2, y) &= \\ &(\alpha^0, \alpha^1) \otimes (\alpha^0, \alpha^1) \otimes (u_0, u_1, u_2) \\ &+ (\beta^0, \beta^1) \otimes (\beta^0, \beta^1) \otimes (v_0, v_1, v_2) \\ &+ (\gamma^0, \gamma^1) \otimes (\gamma^0, \gamma^1) \otimes (w_0, w_1, w_2) \end{aligned}$$

For each $y = 0, 1, 2$ we require

$$\begin{aligned} \begin{pmatrix} I(y=0) & I(y=1) \\ I(y=1) & I(y=2) \end{pmatrix} &= u_y \cdot \begin{pmatrix} \alpha^0 & \alpha^1 \\ \alpha^1 & \alpha^2 \end{pmatrix} \\ &+ v_y \cdot \begin{pmatrix} \beta^0 & \beta^1 \\ \beta^1 & \beta^2 \end{pmatrix} + w_y \cdot \begin{pmatrix} \gamma^0 & \gamma^1 \\ \gamma^1 & \gamma^2 \end{pmatrix}, \end{aligned}$$

which defines a system of three linear equations with three variables u_y, v_y, w_y

$$\begin{pmatrix} I(y=0) \\ I(y=1) \\ I(y=2) \end{pmatrix} = \begin{pmatrix} \alpha^0 & \beta^0 & \gamma^0 \\ \alpha^1 & \beta^1 & \gamma^1 \\ \alpha^2 & \beta^2 & \gamma^2 \end{pmatrix} \cdot \begin{pmatrix} u_y \\ v_y \\ w_y \end{pmatrix}.$$

If α , β , and γ are pairwise distinct real numbers then the corresponding Vandermonde determinant is non-zero and a solution exists. The solution for $\alpha = 1, \beta = 2, \gamma = 3$ is

$$\begin{aligned} \psi_f(x_1, x_2, y) = & \\ & (1, 1) \otimes (1, 1) \otimes (3, -\frac{5}{2}, \frac{1}{2}) \\ & + (1, 2) \otimes (1, 2) \otimes (-3, 4, -1) \\ & + (1, 3) \otimes (1, 3) \otimes (1, -\frac{3}{2}, \frac{1}{2}) . \end{aligned}$$

◇

3.3 Generalized addition

In this section, we present a tensor rank-one decomposition of ψ_f , where f is defined as $f(x) = f_0(\sum_{i=1}^m f_i(x_i))$. Let \mathcal{A} be the set of all possible values of $\sum_{i=1}^m f_i(x_i)$. The rank of ψ_f depends on the nature of functions f_i , more exactly, on the range of the values of $\sum_{i=1}^m f_i(x_i)$. The decomposition is useful, if this range is substantially smaller than $|\mathcal{X}_1| \cdot \dots \cdot |\mathcal{X}_m|$.

Theorem 4 *If $f(x) = f_0(\sum_{i=1}^m f_i(x_i))$, where f_i are integer valued functions, then $\text{rank}(\psi_f) \leq |\mathcal{A}|$.*

Proof Without a loss of generality, we may assume that $f_i(x_i) \geq 0$ for $i = 1, \dots, m$ and that zero is in the range of f_i . If not, this may be achieved by using $f_i(x_i) - \min_{z_i} f_i(z_i)$ instead of f_i and modifying f_0 so that f does not change. The proof is analogous to the proof of the Theorem 3, where $t = \sum_{i=1}^m f_i(x_i)$ and $\varphi_{i,b}(x_i) = \alpha_b^{f_i(x_i)}$ for $b = 1, \dots, |\mathcal{A}|$. □

3.4 Exclusive-or (parity) function

Let \oplus denote the addition modulo two, which is also known as the exclusive-or operation. By the parity or exclusive-or function, we will understand the function $x_1 \oplus \dots \oplus x_m$.

Theorem 5 *Let $\mathcal{X}_i = \mathcal{Y} = \{0, 1\}$ for $i = 1, \dots, m$ and $f(x) = x_1 \oplus \dots \oplus x_m$. Then $\text{rank}(\psi_f) = 2$.*

Proof The exclusive-or function may easily be expressed as a product, if the values

$\{0, 1\}$ are replaced by $\{1, -1\}$ using substitution $0 \mapsto 1, 1 \mapsto -1$. An odd number of ones in the 0/1 representation is equivalent to a negative product of the corresponding values in the 1/-1 representation. Expressing the required transformations in the form of a linear transformation, we obtain

$$x_1 \oplus \dots \oplus x_m = \frac{1}{2}(1 - (1 - 2x_1) \dots (1 - 2x_m)).$$

Since $\psi_f(x, y) = I(y \oplus x_1 \oplus \dots \oplus x_m = 0) = y \oplus x_1 \oplus \dots \oplus x_m \oplus 1$, we have

$$\begin{aligned} \psi_f(x, y) = & \\ & \frac{1}{2}(1 + (1 - 2y)(1 - 2x_1) \dots (1 - 2x_m)). \end{aligned}$$

Hence, ψ_f may be expressed as a sum of two functions, the first of which is the constant $\frac{1}{2}$ and the second is $(\frac{1}{2} - y)(1 - 2x_1) \dots (1 - 2x_m)$. It is now easy to express ψ_f in the form of (3), if we use tensors defined as follows. Let for $i \in \{1, \dots, m\}$, $x_i \in \{0, 1\}$, and $b \in \{1, 2\}$

$$\varphi_{i,b}(x_i) = \begin{cases} 1 & b = 1 \\ 1 - 2x_i & b = 2 \end{cases}$$

and

$$\xi_b(y) = \begin{cases} \frac{1}{2} & b = 1 \\ \frac{1}{2} - y & b = 2. \end{cases}$$

It follows from Lemma 1 that this defines a minimal tensor rank-one decomposition of exclusive-or. □

3.5 Noisy functional dependence

For every $i = 1, \dots, m$ we define a dummy variable X'_i taking values x'_i from set $\mathcal{X}'_i = \mathcal{X}_i$. The noisy functional dependence of Y on $X = (X_1, \dots, X_m)$ is defined by

$$\psi(x, y) = \sum_{x'} \psi_f(x', y) \cdot \prod_{i=1}^m \varkappa_i(x_i, x'_i) \quad (8)$$

where ψ_f is tensor that represent a functional dependence $y = f(x')$ and for $i = 1, \dots, m$ tensors \varkappa_i represent the noise for variable X_i . Note that models like noisy-or, noisy-and, etc., fall within the scope of the above definition. Actually, the definition covers the whole class of models known as models of independence of causal influence (ICI) [5].

Theorem 6 Let tensor ψ represent the noisy functional dependence f defined by formula (8). Then $\text{rank}(\psi) \leq \text{rank}(\psi_f)$.

Proof Let $r = \text{rank}(\psi_f)$. Then

$$\psi_f(x', y) = \sum_{b=1}^r \xi_b(y) \cdot \prod_{i=1}^m \varphi_{i,b}(x'_i) .$$

Substituting this to formula (8) we get

$$\begin{aligned} \psi(x, y) &= \sum_{x'} \sum_{b=1}^r \xi_b(y) \cdot \prod_{i=1}^m (\varphi_{i,b}(x'_i) \cdot \varkappa_i(x_i, x'_i)) \\ &= \sum_{b=1}^r \xi_b(y) \prod_{i=1}^m \sum_{x'_i} (\varphi_{i,b}(x'_i) \cdot \varkappa_i(x_i, x'_i)) \\ &= \sum_{b=1}^r \xi_b(y) \cdot \prod_{i=1}^m \varphi'_{b,i}(x_i) , \end{aligned}$$

where $\varphi'_{b,i}(x_i) = \sum_{x'_i} (\varphi_{i,b}(x'_i) \cdot \varkappa_i(x_i, x'_i))$. The last equation proves that $\text{rank}(\psi) \leq \text{rank}(\psi_f)$. \square

4 Related work

Higher-dimensional tensors are studied in multilinear algebra [1]. The problem of tensor rank-one decomposition is also known as *canonical decomposition* (CANDECOMP) or *parallel factors* (PARAFAC). A typical task is to find a tensor of rank one that is a best approximation of a tensor ψ . This task is usually solved using an alternating least square algorithm (ALS) that is a higher-order generalization of the power method for matrices [2].

Acknowledgments

The authors were supported by the Ministry of Education of the Czech Republic under the project nr. 1M0545 (P. Savicky) and nr. 1M0572 (J. Vomlel). J. Vomlel was also supported by the Grant Agency of the Czech Republic under the project nr. 201/04/0393.

References

[1] L. De Lathauwer and B. De Moor. From matrix to tensor: multilinear algebra and

signal processing. In *4th Int. Conf. on Mathematics in Signal Processing, Part I*, IMA Conf. Series, pages 1–11, Warwick, 1996. Keynote paper.

- [2] L. De Lathauwer, B. De Moor, and J. Vandewalle. On the best rank-1 and rank- (R_1, R_2, \dots, R_N) approximation of higher-order tensors. *SIAM Journal on Matrix Analysis and Applications*, 21(4):1324–1342, 2000.
- [3] F. J. Díez and S. F. Galán. An efficient factorization for the noisy MAX. *International Journal of Intelligent Systems*, 18:165–177, 2003.
- [4] G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, 3rd edition edition, 1996.
- [5] D. Heckerman. Causal independence for knowledge acquisition and inference. In D. Heckerman and A. Mamdani, editors, *Proc. of the Ninth Conf. on Uncertainty in AI*, pages 122–127, 1993.
- [6] J. Håstad. Tensor rank is NP-complete. *Journal of Algorithms*, 11:644–654, 1990.
- [7] F. V. Jensen, S. L. Lauritzen, and K. G. Olesen. Bayesian updating in recursive graphical models by local computation. *Computational Statistics Quarterly*, 4:269–282, 1990.
- [8] K. G. Olesen, U. Kjærulff, F. Jensen, F. V. Jensen, B. Falck, S. Andreassen, and S. K. Andersen. A MUNIN network for the median nerve — a case study on loops. *Applied Artificial Intelligence*, 3:384–403, 1989. Special issue: Towards Causal AI Models in Practice.
- [9] J. Vomlel. Exploiting functional dependence in Bayesian network inference. In *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 528–535. Morgan Kaufmann Publishers, 2002.