On Counterfactual Explanations of Cardiovascular Risk in Adolescent and Young Adult Breast Cancer Survivors

Alice Bernasconi[1,2], Antonio Balordi[1], Alessio Zanga[1,3] and Rafael Cabañas[4]

[1]Models and Algorithms for Data and Text Mining Laboratory (MADLab), University of Milano-Bicocca, Milan, Italy
[2]Evaluative Epidemiology Unit, Department of Epidemiology and Data Science, Fondazione IRCCS Istituto Nazionale dei tumori, Milan, Italy
[3]Data Science and Advanced Analytics, F. Hoffmann - La Roche Ltd, Basel, Switzerland
[4]Department of Mathematics and CDTIME, University of Almería, Almería, Spain

Abstract

In the last decades, the growing population of cancer survivors has shifted researchers' focus from primary toward tertiary prevention. Particularly, adolescents and young adults (AYAs) breast cancer (BC) survivors may face long-term outcomes as a result of their treatments, among which cardiovascular diseases (CVDs) are the most life-threatening ones. To plan effective follow-up guidelines for preventing and treating these events, it is essential to disentangle the causal role of cancer treatments in these patients. In this work, we aim to extend the current state of BC treatment guidelines by leveraging on the estimate of the risk of CVDs in AYAs who underwent BC treatments, as provided by a causal Bayesian network. In these regards, we provide counterfactual explanations of a causal query, using real-world data, algorithms and methods from the causal inference domain. We show that while ovarian suppression combined with tamoxifen may be a necessary cause for ischemic heart disease, it is not a sufficient one, i.e., this treatment alone is not enough to cause the disease, other factors must also be present. These findings can contribute to support clinicians in the treatment choice and help in refining treatment strategies and follow-up protocols for AYAs, advancing personalised healthcare in oncology.

Keywords

Causal networks, Counterfactual explanations, Breast cancer survivors, Treatment guidelines, Cardiovascular diseases, Adolescents and young adults

1. Introduction

Chronic diseases are by far the leading causes of mortality world-wide. In the last decades, the prevalence of chronic diseases is rising due to changes in life-style and the population aging, especially in Italy. Most of the times, these disorders present co-concurrence of multiple other chronic diseases (co-morbidities) that require the involvement of several caregivers for proper patient care. However, hospital care, ambulatory specialist care and primary care are subdivided into numerous entities, based mainly on medical specialty. Hence, to provide to these patients the optimal and integrated care is a major challenge for the health care system [1].

Every year in Italy about 400 thousands new cancer diagnosis are registered, with the highest incidence in older adults. On the other hand, cancer survival is continuously improving thanks to the innovations in patient care, thus making the proportion of cured patients growing. Oncologists are in charge of cancer diagnosis and treatment. These tasks require a massive number of visits and examinations especially during the first year since cancer identification. Once cancer treatment is concluded, follow-up visits are scheduled to prevent cancer relapse, with the scheduling decreasing with time and depending on the major cancer prognostic factors. Although oncologists have detailed information about the treatments their patients received, they do not have the ability to monitor all of their effects, especially in the long-term. Nevertheless, there is strong evidence in the scientific literature on the wide variety of long-term effects that cancer therapies can cause, including diseases of the cardiovascular or endocrine system, reproductive disorders, infections and so on [2]. Moreover,

HC@AIxIA 2024: 3rd AIxIA Workshop on Artificial Intelligence For Healthcare
a.bernasconi6@campus.unimib.it,alice.bernasconi@istitutotumori.mi.it (A. Bernasconi)
0000-0001-8522-6882 (A. Bernasconi); 0000-0003-4423-2121 (A. Zanga); 0000-0002-5034-582X (R. Cabañas)

cancer is a complex disease treated with combined treatments: the mixed effect of different treatments makes it even more difficult to predict the possible late outcomes.

General practitioners (GPs) are the only physicians able to monitor the patients throughout their lifetime. However, they might not have complete access to hospital charts nor the knowledge needed in all the medical specialties. Moreover, data about the effect of the most innovative cancer treatments are scarce. Clearly identify who and how should be involved in patients follow-up is essential to help oncologists and GPs to plan effective follow-up guidelines to prevent and treat long-term outcomes in cancer survivors. This aspect is also important from a public health prospective and for policy-makers interested in better organizing resources. Artificial intelligence (AI) and causal inference (CI) can help in enriching the knowledge of the medical stakeholders building effective tools to be used to quantify and causally explain the burden of late outcomes in cancer survivors.

Adolescents and young adults (AYAs, patients aged 15 to 39 at first cancer diagnosis) are an heterogeneous and peculiar group of cancer patients who deserve special attention. This age-group share tumours' case-mix both with the younger and older counterpart. Breast cancer (BC) is the most frequent cancer in AYA females, as in older women. Nevertheless, there is a survival gap attributable to a more aggressive biology of BC in AYA than in older patients.

Most patients with BC receive surgery as main treatment. The major surgical procedure can be preceded or followed by other treatments both systemic (chemotherapy, target therapy or hormones therapy) or local (radiotherapy). Treatment guidelines are the same both for AYAs and older women and depend on several factors related to the cancer (e.g., extent of the disease) and the host (e.g., hormonal status). Oncologists choose the best combination of treatments with two major objective: *i*) to maximise the patient chance of survival and *ii*) to minimise the risk of cancer relapse.

Despite the large knowledge of late effects of cancer treatments in older women, little is known about the magnitude of the impact of cancer treatments in younger patients, like AYAs. The accumulation of stress induced by cancer and its treatment may contribute to accelerate aging in young cancer survivors, inducing premature mortality, frailty and other age-related diseases, like cardiovascular diseases (CVDs) [3].

The main contributions of this manuscript, made by leveraging on the first AI model [4] developed for estimating the risk of CVDs in AYAs that survived after BC treatments, are the following:

- To *enrich* BC treatment guidelines with *knowledge* on CVDs risk in young women;
- To contribute to *disentangling uncertainty in treatment choice* using counterfactual explanations on the most relevant late outcome in these patients;
- To help clinicians in tailoring *personalised follow-up guidelines* for high-risk patients.

The rest of the manuscript is organised as follows. Section 2 introduces the notation and gives the main definitions to make the paper as much as possible self-contained. The main contributions on the case study of adolescent and young adults breast cancer survivors are presented in Section 3. We close the manuscript with the description of the experimental results (Section 3.4) and the discussion of the achievements (Section 4), with some proposals on how to develop further along the same research direction for answering more ambitious and complex counterfactual queries.

2. Methods

In this section we introduce the notation, together with the main concepts and the mathematical models needed to follow the rest of the paper. In particular, we give the definitions of Bayesian network, causal network and structural causal model, while also describing the three rungs of the *ladder of causation* [5] which are fundamental to understand our contributions.

2.1. Bayesian Networks

Bayesian networks (BNs) [6, 7, 8] are a type of probabilistic graphical model (PGM) used for reasoning under uncertainty. BNs are made of a qualitative component in the form of direct acyclic graph (DAG)

encoding the independence relations between the variables in the problem, while the quantitative component is a set probability distributions measuring such relations. More formally, BNs can be defined as follows.

Definition 1 (Bayesian Network (BN)). *A Bayesian network is a pair $\langle \mathcal{G}, \mathcal{P} \rangle$, where:*

- *$\mathcal{G} = \langle \mathbf{V}, \mathbf{E} \rangle$ is a DAG, with $\mathbf{V}$ a set of vertices and $\mathbf{E} \subset \mathbf{V} \times \mathbf{V}$ a set of directed edges,*
- *$\mathcal{P}$ is a probability distribution over the random vector $\mathbf{X}$.*

Each vertex $V_i \in \mathbf{V}$ is mapped to a variable $X_i \in \mathbf{X}$, so that the global probability distribution $\mathcal{P}$ is factorised over $\mathcal{G}$ into local probability terms $P(X_i|Pa(X_i))$, with $Pa(X_i)$ the parents[1] of X_i.

For each variable $V_i \in \mathbf{V}$, we define the *ancestors* of V_i to be the set of variables $V_j \in \mathbf{V} \setminus \{V_i\}$ such that there exists a directed path[2]. Similarly, we define the *descendants* of V_i to be the set of variables $V_j \in \mathbf{V} \setminus \{V_i\}$ such that there exists a directed path from V_i to V_j. Henceforth, we will refer to a vertex V_i and its corresponding variable X_i interchangeably.

Definition 2 (Causal Network (CN)). *A Causal Network is a BN in which any edge from parents to children represents a cause-effect relationship.*

2.2. Observational and Interventional Rungs

Standard probabilistic inference involves computing the posterior probability distribution for variables of interest given evidence about other variables, commonly referred to as observational queries (e.g., "what if I see this?"). For instance, given states x and y of random variables X and Y, respectively, an observational query might involve computing the conditional probability $P(x|y)$. Here, x and y represent the presence of X and Y, while x' and y' denote their absence. In this context, the *average treatment effect* (ATE) is defined as

$$\mathrm{ATE}(X, Y) = P(y|x) - P(y|x'). \tag{1}$$

Conversely, causal reasoning focuses on hypothetical scenarios where we calculate the probability of a variable given that we intervene on another. For example, the query $P(Y = y|do(X = x))$ represents the probability that Y equals y when X is intervened to take the value x. The notation $\mathrm{do}(X = x)$ explicitly denotes an intervention, distinguishing it from mere observation. The difference between two such interventional queries, known as the *causal effect difference* or *average causal effect* (ACE), is defined as

$$\mathrm{ACE}(X, Y) = P(Y = y \mid \mathrm{do}(X = x)) - P(Y = y \mid \mathrm{do}(X = x')). \tag{2}$$

To calculate an interventional query, a process often referred to as "surgery" is employed. This graphical operation involves removing the incoming arcs to the intervened variable X and setting the node to a specific value $X = x$. The model that results from this surgical intervention is known as the "post-intervention" model. This process is performed to restrict the natural tendency of the variable to change in response to other variables in the environment.
Performing graph surgery is the initial step required to distinguish the associative effect from the purely causal effect. However, a causal estimand cannot be directly estimated using a statistical estimator; it must first be translated into a statistical estimand by removing the intervention. This process is known as the *identification of the causal effect.*

If there exists a set of covariates $\mathbf{Z}$ that satisfies the *back-door criterion* [5] in the model, then there exists a consistent estimator for the causal effect of X on Y:

$$P(Y = y \mid \mathrm{do}(X = x)) = \sum_{\mathbf{z}} P(Y = y \mid X = x, \mathbf{Z} = \mathbf{z}) P(\mathbf{Z} = \mathbf{z}). \tag{3}$$

[1] A vertex V_j is said to be a parent of V_i if there exists a directed edge from V_j to V_i.
[2] A directed path from V_i to V_j is sequence of directed edges starting from V_i and ending in V_j

Under the condition of exogeneity (also known as no-confounding) [9], the way Y would potentially respond to experimental conditions x or x' is independent of the actual value of X. This implies that $P(Y = y|\text{do}(X = x)) = P(Y = y|X = x)$ and $P(Y = y|\text{do}(X = x')) = P(Y = y|X = x')$, thus making $ACE(X, Y) = ATE(X, Y)$. A graphical criterion to identify the condition of exogeneity is the absence of a common ancestor of X and Y connected to Y through a directed path that does not include X.

2.3. Counterfactual Rung

Counterfactual queries [5] explore hypothetical scenarios, such as, "What would the outcome have been if the variable had taken a different value?" For example, $P(Y_x|X = x')$ represents the probability of Y if X had taken the value x instead of x'. Here, Y_x relates to the hypothetical scenario, while X is in the real scenario.

A key concept in this context is the *probability of necessity* (PN), which measures the extent to which one event is a necessary condition for another. The PN is defined as:

$$\text{PN}(X, Y) = P(Y_{x'} = y'|X = x, Y = y). \tag{4}$$

Here, X is considered a necessary cause for Y if y would not have occurred without x, given that both x and y actually occurred. Therefore, PN represents our certainty about X being a necessary cause of Y.

Similarly, we may also be interested in determining whether an event is a sufficient condition. To address this, we define the *probability of sufficiency* (PS) as:

$$\text{PS}(X, Y) = P(Y_x = y|X = x', Y = y'). \tag{5}$$

X is considered a sufficient cause for Y if y occurs whenever x occurs. Thus, PS represents the probability that X is a sufficient cause of Y. In other words, it is the probability that setting x would lead to y in a scenario where both x and y are currently absent.

Counterfactual queries cannot be directly computed from a CN. Instead, structural causal models (SCMs) [5], which can be viewed as an extension of CNs, are required. SCMs consist of endogenous variables, which represent internal elements of the model, and exogenous variables, which often lack a clear semantic interpretation. SCMs can be formally defined as follows [10].

Definition 3 (Structural Causal Model (SCM)). *A structural causal model is defined as a 4-tuple* $\langle \mathbf{U}, \mathbf{V}, \mathcal{F}, \mathcal{P} \rangle$*, where:*

- $\mathbf{U}$ *is the set of exogenous variables;*
- $\mathbf{V}$ *is the set of endogenous variables;*
- $\mathcal{F} = \{f_i : \mathbf{U}_i \cup Pa(V_i) \rightarrow V_i, \forall V_i \in \mathbf{V}\}$ *is the set of structural equations;*
- $\mathcal{P}$ *is the set containing the exogenous probability distributions* $P(U_i)$ *for each* $U_i \in \mathbf{U}$*.*

Note that the structural equations $\mathcal{F}$ actually define a DAG over the variables in $\mathbf{U} \cup \mathbf{V}$, with an edge from each variable in $\mathbf{U}_i \cup Pa(V_i)$ to V_i.

3. Answering a Challenging Causal Query

In this section we first formulate the causal query representing the subject of this paper, then we show how such a query translates to the language of SCMs. Furthermore, we show how the obtained SCM can be simplified to efficiently answer the causal query. The section closes by answering the causal query.

3.1. The Causal Query

The starting point of this work is the causal network described in [4]. This causal network is the first one developed for estimating the risk of cardiovascular diseases (CVDs) in adolescent and young adults (AYAs) that have been treated and survived breast cancer (BC). It has been developed by combining clinical knowledge with two different patients cohorts, namely a population cohort and a clinical cohort.

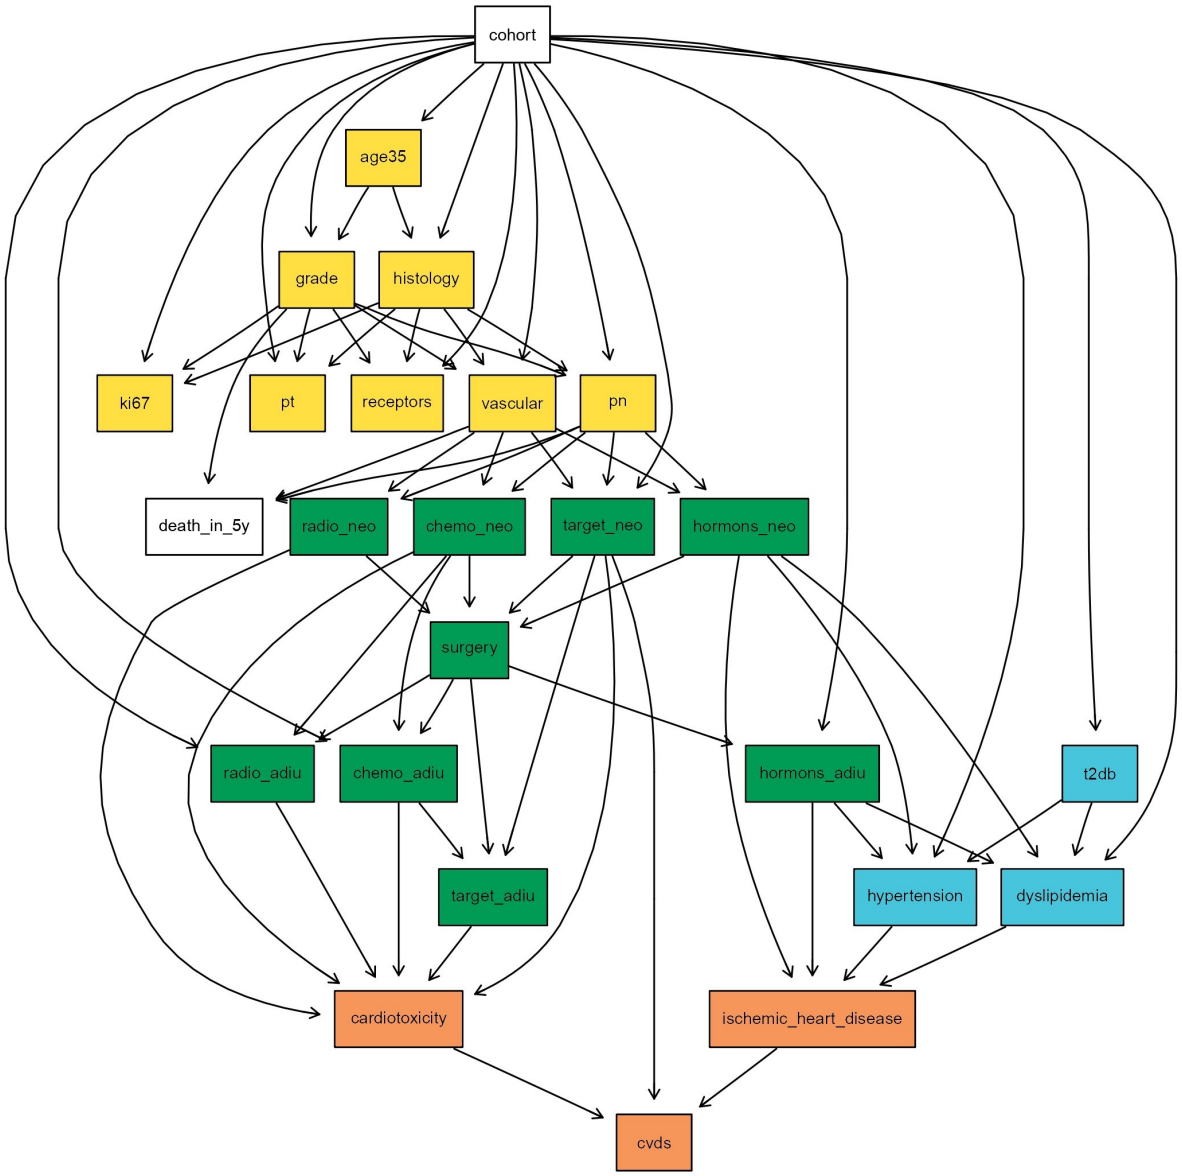

Figure 1: DAG of the CN model describing the interplay of observable factors that contribute to the risk of CVDs in AYA with BC.

The causal network is depicted in Figure 1. The cancer prognostic factors (coloured in Yellow) and the major CVDs risk factors (coloured in Blue) are non modifiable risk factors. To reduce and prevent the risk of developing a CVD or its sub-forms (i.e., ischemic heart diseases or cardiotoxicity, coloured in Orange), clinicians can intervene on the treatments only (those coloured in Green). Thus, in our work we will interpret the risk of CVDs according to treatment recommendations included within the treatment guidelines as queries.

Breast cancer treatment is regulated in Italy by Italian national guidelines, discussed every year by a panel of experts. The aims of these guidelines are:

- To improve and standardise the clinical practice;

- To offer all patient throughout the country the possibility of best care;
- To ensure an evidence-based reference for national and regional institutions.

In this paper, the major clinical recommendations are presented in the form of clinical queries accompanied by the quality of their supporting evidence together with the strength of the associated recommendation. In particular, we are interested to answer the following causal query:

> **CAUSAL QUERY:** *In pre-menopausal women, with a surgically treated breast cancer, positive to hormonal receptors, HER2 negative, low risk for recurrences, is it recommendable to add ovarian suppression to tamoxifen treatment?*

The clinical recommendation about this casual query is **STRONG IN FAVOUR** to the addition of the ovarian suppression. This recommendation was voted by the panelists in light of the significant improvement in both overall and progression-free survival. Among the side effects of this combination of treatments the more relevant listed were: mood alterations, sexual dysfunction and osteoporosis. Despite the toxicity profile highlighted, the benefit-to-damage ratio was considered in favour of the addition of ovarian suppression to tamoxifen treatment.

Nevertheless, no evidence is provided about the potential role of this treatment combination to the CVD risk. When it is unethical to conduct a randomised clinical trial, observational data and causal inference are the only way to integrate the clinical recommendation with knowledge on CVD risk. Hence, in this work we are answering to three queries, formulated according to the three rungs of the ladder of causation proposed by Judea Pearl [5]:

In pre-menopausal women, with a surgically treated breast cancer, positive to hormonal receptors, HER2 negative, low risk for recurrences...

ASSOCIATION	*...which is the observed risk of CVDs in those patients that received ovarian suppression in combination with tamoxifen treatment?*
INTERVENTION	*...which is the risk of CVDs if we administer ovarian suppression in combination with tamoxifen treatment?*
COUNTERFACTUALS	*...which would have been the risk of CVDs if we did not administered ovarian suppression?*

3.2. Translating the Causal Query into a Structural Causal Model

The causal network depicted in Figure 1 is used to translate the causal query of interest (**CAUSAL QUERY**) into the corresponding structural causal model. This is achieved by first oberving that the ovarian suppression is administered in young patients as neo-adjuvant (i.e., pre-surgical) treatment to blocks the body's ability to produce hormones, particularly estrogen, to preserve ovarian fertility that could be compromised by the toxicity of cancer treatments (like chemotherapy). Nevertheless, estrogens have a cardioprotective effect in young females. Thus, the loss of estrogen during menopause is associated with increased risk of ischemic heart disease [11, 12]. Hence, the causal query investigates the relationship: `[hormons_neo]`$\rightarrow$`[ischemic_heart_diseases]` with the aim to determine whether the former variable is a necessary, a sufficient or both a necessary and sufficient cause of the latter.

Before starting the experiments, we had to select the group of patients described in the causal query (Section 3.1). All patients included in the dataset were surgically treated, thus, no restriction was done in this regard. To select patients positive to hormonal receptors and HER2 negative only, we set the variable `[Receptors]` to "Luminal" OR "Luminal A" OR "Luminal B". Moreover, to select patients receiving tamoxifen, we set the variable `[Hormons_adiu]` to "Yes"; thus, tamoxifen is the elective hormonal adjuvant (i.e., post-surgery) treatment, administered for a minimum of 5-10 years.

The model depicted in Figure 1 was developed using data coming from two different retrospective cohorts of AYA BC patients: a population-based cohort [13], identified in population-based cancer registries, and a single-institution clinical cohort. Population-based cancer registries have the unique opportunity to collect information on all cancer cases diagnosed in a given area with poor clinical details. The clinical cohort, on the opposite, has more detailed information on cancer prognostic factors but suffers from selection bias. For this work, we decided to focus on the population-based data (setting the node `[Cohort]` to "Population-based") because we wanted our results to be valid for all AYAs with BC. By consequence, a group of patients was selected based on the values of the variables `[receptors]`,`[hormons_adiu]` and `[cohort]`. From a graphical perspective, in CNs this translates to removing the outgoing edges from such variables, which brings us to obtain the DAG shown in Figure 2.

Moreover, while in the original model (Figure 1), the variables `[hypertension]`, `[t2db]` and `[dyslipidemia]` (that represent hypertension, dyslipidemia and type 2 diabestes, respectively), were coded as "pre" (if the diseases was diagnosed before BC diagnosis), "post" (if the diseases was diagnosed after BC diagnosis) or "no"; in this work, they were binarised by grouping together "pre" and "post" into the same label named ”yes”.

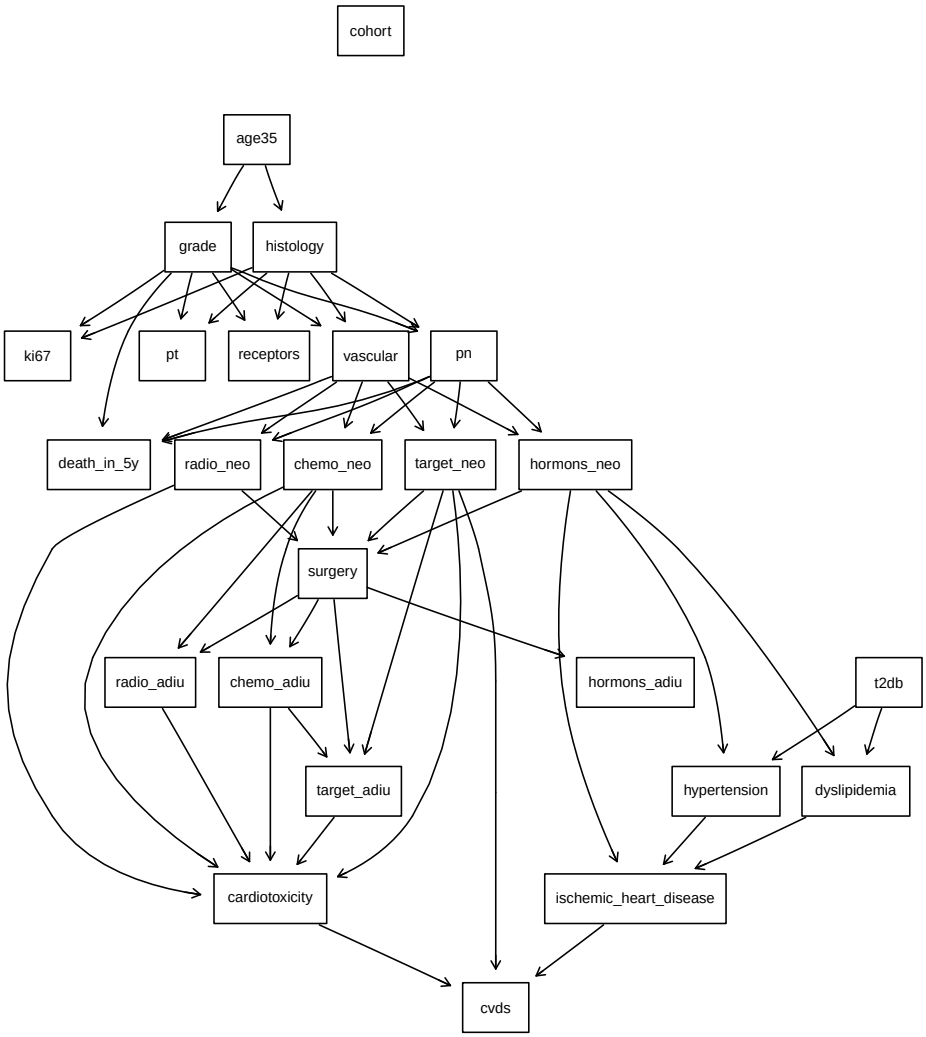

Figure 2: DAG of the CN model after selecting the group of patients.

3.3. Simplifying the Structural Causal Model

Counterfactual reasoning with probabilistic graphical models typically requires significant computational power, making it challenging to manage large problems. To address this issue, we further simplified the model depicted in Figure 2 and the dataset described in Section 1.

In causal and counterfactual queries, a variable is often referred to as either the causes or the effect. In our case study, the node `[hormons_neo]`represents the **cause**, while the node `[ischemic_heart_disease]` is the **effect**.

A *barren variable* (or node) [7] with respect to a causal query is a variable that does not influence the probability distribution of any variables of interest for that causal query. Consequently, such variables can be pruned from the model without affecting the outcome of the inference. In a CN, a variable is considered barren if it is not included in the causal query and either has no descendants or only has descendants that are themselves barren.

In our case study, barren variables are those that are not ancestors of the effect variable `[ischemic_heart_disease]`. The result of pruning barren variables is shown in Figure 3 (left).

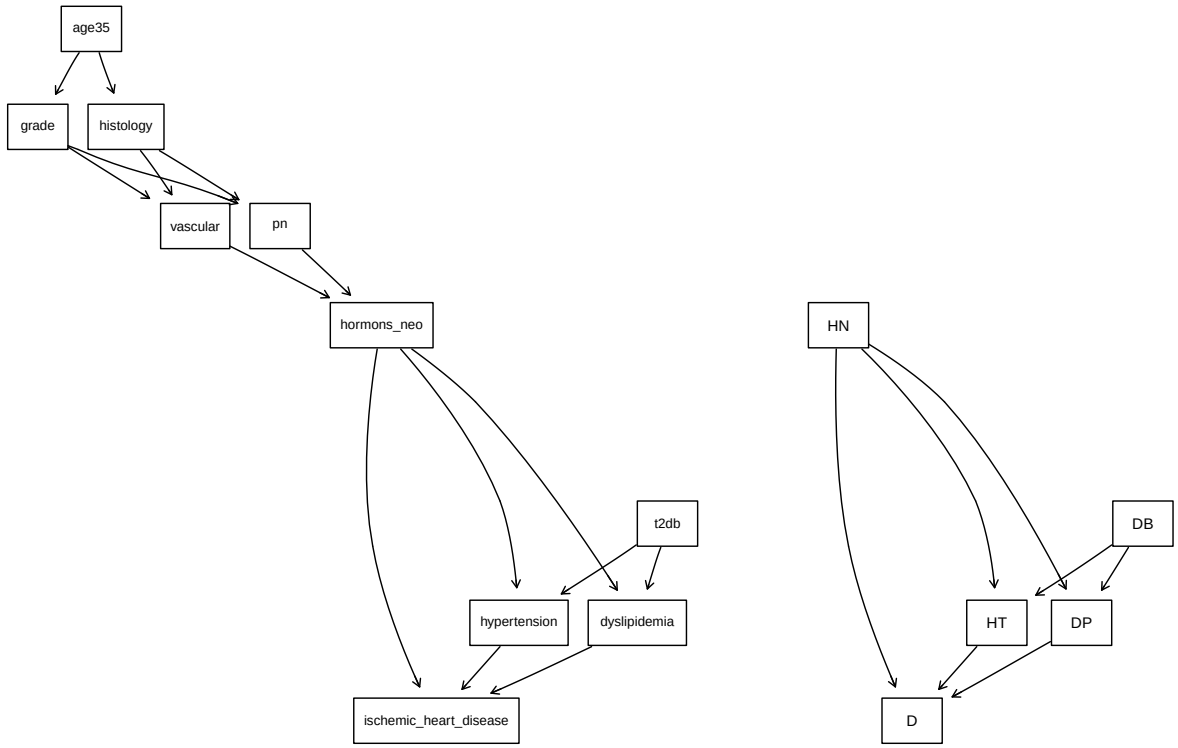

Figure 3: DAGs of the CN model after pruning barren variables (left) and after removing the d-separated variables (right).

However, the model depicted in Figure 3 (left) can be further simplified using d-separation[3]. Specifically, all ancestors of the cause variable `[hormons_neo]`are conditionally independent of `[ischemic_heart_disease]` given `[hormons_neo]`. Consequently, these variables can be removed, resulting in the simplified model shown in Figure 3 (right). For simplicity, variable names were abbreviated as follows: HN represents `[hormons_neo]`, D represents `[ischemic_heart_disease]`, HT represents `[hypertension]`, DB represents `[t2db]` and DP represents `[dyslipidemia]`.

The original dataset was also simplified, thus, only the rows related to the selected patients were considered, all the columns for the irrelevant variables were removed, i.e., the variables except those shown in Figure 3 (right).

According to the the queries of interest, we investigate the causal effects of the variable `[hormons_neo]` on the outcome variable `[ischemic_heart_disease]`in the model shown in Figure 3 (left).

[3]For more details on d-separation, readers can refer to [7].

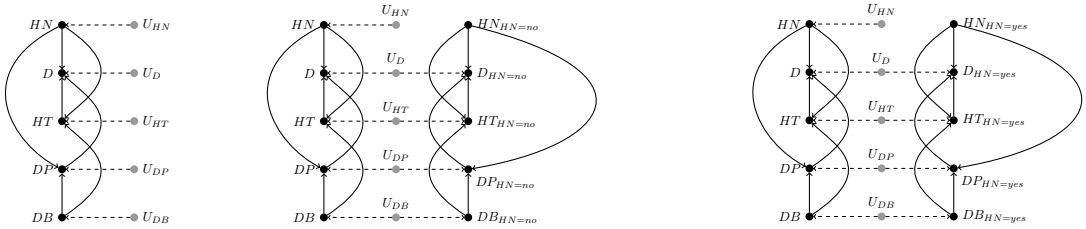

Figure 4: DAG of the SCM of the use case (left) and twin model for calculating $PN(HN, D)$ (center) and for $PS(HN, D)$ (right). Endogenous variables are depicted in black while the exogenous in gray.

These two variables satisfy the exogeneity condition; hence, the $ACE(HN, D)$ was calculated as the variation on the conditional probability $P(D = yes|HN = yes) - P(D = yes|HN = no)$ using inference on junction trees [14].

In our use case, the endogenous variables were those remaining after the removal of barren and d-separated variables. Then, an exogenous variable was added as a parent to each endogenous variable. The resulting DAG is shown in Figure 4 (left). In contrast, exogenous variables are denoted by the letter U followed by the name of the corresponding child variable.

Equations were automatically inferred from the causal graph, without any loss of generality, via a *canonical specification* [15]. Each equation is then a deterministic function, in which the states of an exogenous variable will then represent all possible function mappings between its children domains from their respective endogenous parents domains. In practice, the equations can be represented as a degenerated conditional probability table containing just ones and zeros.

In contrast, the distributions associated with the exogenous variables are initially unknown. To address this, we propose utilizing the innovative technique known as EMCC (*Expectation Maximization for Causal Computation*), as detailed in [16, 17]. This method treats an SCM as a BN with exogenous variables considered latent. The core idea is to repeatedly apply a learning algorithm designed for BNs with latent variables.

While other methods for computing counterfactual queries exist — such as solving structural equations manually or using potential outcomes frameworks — we focus on the EMCC approach due to its efficiency and scalability in handling complex models with latent variables. This method allows for automatic estimation of exogenous distributions directly from data, which is particularly advantageous when dealing with large-scale or high-dimensional problems.

After each run, the specification for the exogenous distributions are available and counterfactual queries can be computed using an extended model known as the *counterfactual model* (or *twin model*).

The twin model is a SCM that includes endogenous variables from both the real and hypothetical scenarios, achieved by duplicating the sub-graph composed of the endogenous nodes for the real scenario and then applying the intervention. In our case study, $PN(HN, D)$ was calculated as the causal query $P(D_{HN=no} = no|HN = yes, D = yes)$ in the model shown in the Figure 4 (center), for which any inference algorithm can be used.

This can be interpreted as the probability that a patient would not have suffered the disease if they had not received the hormonal treatment, given that they actually did receive the treatment and suffered the disease.

Similarly for $PS(HN, D)$ was calculated as the causal query $P(D_{HN=yes} = yes|HN = no, D = no)$ in the model shown in the Figure 4 (right). Analogously, this can be viewed as the probability that a patient would have suffered the disease if they had received the ovarian suppression treatment, given that they actually did not receive the treatment and did not suffered the disease.

3.4. Answering the Causal Query

For the causal analysis, our objective was to calculate the average causal effect (ACE) of the variable `[hormons_neo]`on `[ischemic_heart_disease]`. This can be directly computed as the average

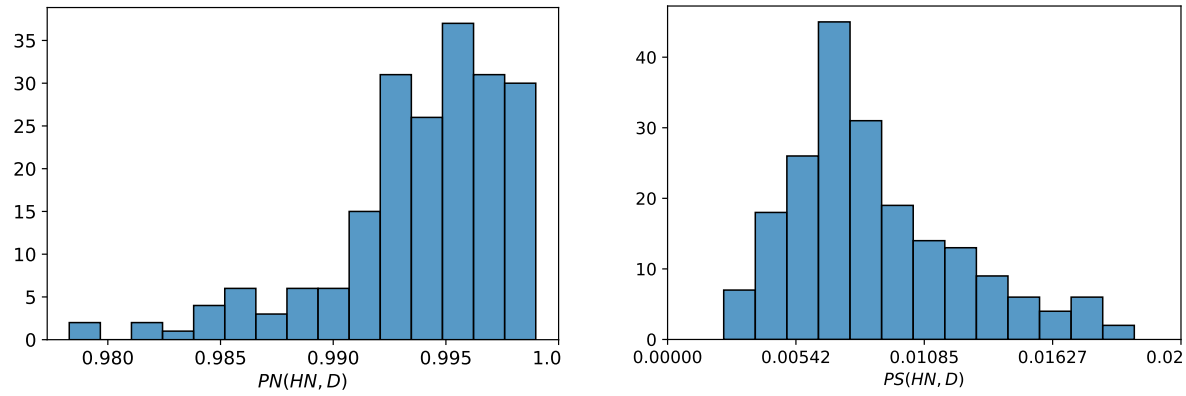

Figure 5: Values of the probability of necessity (left) and probability of sufficiency (right) of variable `[hormons_neo]`for `[ischemic_heart_disease]`.

treatment effect (ATE). We found that $P(D = yes|HN = yes) = 1.45\%$ and $P(D = yes|HN = no) = 0.22\%$. Consequently, we obtain $ACE(HN, D) = 1.23\%$, that means that the probability of developing a CVD increase only about 1% if the patient receives ovarian suppression together with tamoxifen treatment. This result indicate that this treatment does not significantly influence the outcome of the disease.

In the context of counterfactual analysis[4], we investigated the likelihood that the variable `[hormons_neo]` is a necessary and sufficient cause for `[ischemic_heart_disease]`. To achieve this, we conducted 100 runs of the previously described EMCC algorithm. For each run, we obtained a value for the causal query $PN(HN, D)$ and another one for $PS(HN, D)$. The distribution of these values is depicted in Figure 5. In the case of PN, all values exceeded 97.83%. Conversely, the values for PS remained below 1.97%.

These results suggest that, with a high probability, the ovarian suppression treatment is a necessary cause but not a sufficient cause for the ischemic heart disease. That means that the latter factor is essential for the disease to occur, i.e., the disease cannot happen without the presence of this hormonal treatment. However, this treatment alone is not enough to cause the disease; other factors must also be present.

4. Conclusions and Future work

In this work we showed how causal networks can be effectively used to disentangle uncertainty in treatment choices, while helping clinicians in better tailoring personalised follow-up guidelines for chronic patients, like cancer survivors. Moreover, the evidence derived from the experimental results can integrate the actual state of the treatment guidelines, enriching them with knowledge on the CVD risk of AYA BC patients. Starting from a clinical question on the recommendability of the ovarian suppression addition to tamoxifen treatment, the counterfactual explanations made it evident that while `[hormons_neo]`is necessary to induce `[ischemic_heart_disease]`it is not a sufficient cause. Thus, while it is true that no patient not receiving ovarian suppression will develop a CVD, receiving ovarian suppression is not sufficient to explain an increase in the risk of CVDs.

This work is an important use case which concretely shows how effective observational real-world data can be to answer clinical questions. In these regards, the overlapping of association and intervention results support the idea that, despite confounding is present, clinicians are able to deal with it in the everyday clinical practice, even without ad hoc strict guidelines. This result makes even more important and relevant the evidence driven by the third ladder of causation (counterfactuals) that mimics the

[4]Code available at: https://github.com/AlessioZanga/cardiovascular-counterfactuals.git

results of a randomised controlled trial, the gold standard in medicine. This approach, applied to observational data, can be particularly important especially when evidence from trials is not available nor ethic to be obtained [18].

However, despite its relevance this work has some limitations too. First of all, while ATE and ACE are calculated using a statistical estimand, the computation of PN and PS is more complex and requires a SCM. As described in Section 3.3, in our proposed approach an exogenous variable was added as a parent to each endogenous variable. Given the unknown distributions of the exogenous variable we learn it by repeatedly applying a learning algorithm. This approach makes all the process extremely computationally expensive which allows to handle only queries in which the effect node can have maximum of three parents. Methodological work is needed to extend this approach to be able to answer to more complex queries. Moreover, attention should be paid when interpreting the results of ACE, PN and PS with regards to the rarity of the events. Thus, CVDs are really rare in AYA surviving BC [19], so even though ACE and PS are very low they may be relevant for this specific population. Furthermore, a discussion with clinicians (oncologists and cardio-oncologists) will be needed to validate the clinical plausibility of the presented results.

Finally, as illustrated in the experimental results, the ovarian suppression is not a sufficient cause of the ischemic events, under the assumption of causal sufficiency (all the causes needed to explain the causal mechanisms are included in the model). Nevertheless, this assumption is difficult to be valid especially considering the vast literature that describes the role of lifestyle factors (like smoking, physical inactivity, obesity and poor diet) on the development of cardiovascular diseases both directly and indirectly through type 2 diabetes, hypertension and dyslipidemia [20, 21, 22]. Planning interventions on lifestyle modifiable factors would be more effective than treating their effects only, hence, the addition of these variables to the model would be essential to better stratify the CVD risk and develop more personalised follow-up strategies. To conclude, to achieve transportability of results, the model needs external validation. The external validation is already ongoing in similar cohorts of AYA BC survivors in 6 different areas of Italy (Veneto, Friuli-Venezia-Giulia, Tuscany, Apulia regions and two Sicilian provinces) and in 4 European Countries (Estonia, Norway, Denmark and Belgium) thanks to pilot studies nested in international Joint Actions, namely Innovative Partnership for Action Against Cancer (iPAAC) and Prevent Non-Communicable Diseases and Cancer (Prevent NCD).

Acknowledgments

Alice Bernasconi is funded by an *AIRC* 2020 project, grant number 24864, titled *"pRedicting cardiOvascular diSeAses iN adolescent and young breast caNcer pAtients (ROSANNA)*".

Alessio Zanga is funded by F. Hoffmann-La Roche Ltd.

This work was partially supported by the MUR under the grant "Dipartimenti di Eccellenza 2023-2027" of the Department of Informatics, Systems and Communication of the University of Milano-Bicocca, Italy.

Grant PID2022-139293NB-C31 funded by MCIN/AEI/10.13039/501100011033 and by ERDF A way of making Europe partially supported this work. Rafael Cabñas was also supported by "Plan Propio de Investigación y Transferencia 2024-2025" from University of Almería under the project P_LANZ_2024003.

This work was possible thanks to the data collected in the *Ada project ("Adolescents and young adults with cancer in Italy. How to ensure access to the best care and quality of survival").* We thank the project working group: Alessandra Andreotti, Carlotta Buzzoni, Sabrina Fabiano, Adele Caldarella, Stefania Giorgetti, Salvatore Sciacca, Rosa Angela Filiberti, Cinzia Gasparotti, Giuliano Carrozzi, Rosalba Amodio, Lucia Mangone, Silvia Iacovacci, Mario Fusco, Fabrizio Stracci, William Mantovani, Giuseppe Cascone, Sante Minerba, Giuseppe Sampietro, Anna Melcarne, Paolo Ricci, Federica Manzoni, Maria Letizia Gambino, Elisabetta Merlo, Rossella Bruni, Alessandra Sessa, Giancarlo D'Orsi, Anna Clara Fanetti, Lucia De Lorentis, Tiziana Scuderi, Fabio Savoia, Iolanda Grappasonni, Salvatore Bongiorno, Antonio Romanelli.

References

[1] H.-P. B.-L. Rocca, L. Fleischhacker, O. Golubnitschaja, F. Heemskerk, T. Helms, T. Hoedemakers, S. H. Allianses, T. Jaarsma, J. Kinkorova, J. Ramaekers, P. Ruff, I. Schnur, E. Vanoli, J. Verdu, B. Zippel-Schultz, Challenges in personalised management of chronic diseases—heart failure as prominent example to advance the care process, EPMA Journal 7 (2015) 2. doi:10.1186/s13167-016-0051-9.

[2] K. D. Miller, L. Nogueira, A. B. Mariotto, J. H. Rowland, K. R. Yabroff, C. M. Alfano, A. Jemal, J. L. Kramer, R. L. Siegel, Cancer treatment and survivorship statistics, 2019, CA: A Cancer Journal for Clinicians 69 (2019) 363–385. doi:10.3322/caac.21565.

[3] S. Wang, N. E. Jurdi, B. Thyagarajan, A. Prizment, A. H. Blaes, Accelerated aging in cancer survivors: Cellular senescence, frailty, and possible opportunities for interventions, International Journal of Molecular Sciences 25 (2024) 3319. doi:10.3390/ijms25063319.

[4] A. Bernasconi, A. Zanga, P. J. Lucas, M. Scutari, F. Stella, Towards a transportable causal network model based on observational healthcare data, volume 3578, 2023, p. 122 – 129.

[5] J. Pearl, Causality: Models, Reasoning and Inference, Cambridge University Press, New York, 2000.

[6] J. Pearl, Probabilistic Reasoning in Intelligent Systems, Elsevier, 1988. URL: https://linkinghub.elsevier.com/retrieve/pii/C20090276094. doi:10.1016/C2009-0-27609-4.

[7] F. V. Jensen, T. D. Nielsen, Bayesian networks and decision graphs, volume 2, Springer, New York, NY, 2007.

[8] D. Koller, N. Friedman, Probabilistic Graphical Models: Principles and Techniques, The MIT Press, 2009.

[9] J. Tian, J. Pearl, Probabilities of causation: Bounds and identification, Annals of Mathematics and Artificial Intelligence 28 (2000) 287–313.

[10] E. Bareinboim, J. D. Correa, D. Ibeling, T. Icard, On Pearl's hierarchy and the foundations of causal inference, in: Probabilistic and causal inference: the works of Judea Pearl, ACM, 2022, pp. 507–556.

[11] A. Christinat, S. D. Lascio, O. Pagani, Hormonal therapies in young breast cancer patients: when, what and for how long?, Journal of thoracic disease 5 Suppl 1 (2013) S36–46. doi:10.3978/j.issn.2072-1439.2013.05.25.

[12] I. Alfaris, J. Asselah, H. Aziz, N. Bouganim, N. Mousavi, The cardiovascular risks associated with aromatase inhibitors, tamoxifen, and gnrh agonists in women with breast cancer, Current Atherosclerosis Reports 25 (2023) 145–154. doi:10.1007/s11883-023-01085-3.

[13] A. Bernasconi, G. Barigelletti, A. Tittarelli, L. Botta, G. Gatta, G. Tagliabue, P. Contiero, S. Guzzinati, A. Andreano, G. Manneschi, F. Falcini, M. Castaing, R. A. Filiberti, C. Gasparotti, C. Cirilli, W. Mazzucco, L. Mangone, S. Iacovacci, M. F. Vitale, F. Stracci, S. Piffer, R. Tumino, S. Carone, G. Sampietro, A. Melcarne, P. Ballotari, L. Boschetti, S. Pisani, L. C. D'Oro, F. Cuccaro, A. D'Argenzio, G. D'Orsi, A. C. Fanetti, A. Ardizzone, G. Candela, F. Savoia, C. Pascucci, M. Castelli, C. Storchi, A. Trama, Adolescent and young adult cancer survivors: Design and characteristics of the first nationwide population-based cohort in italy, Journal of Adolescent and Young Adult Oncology 9 (2020) 586–593. doi:10.1089/jayao.2019.0170.

[14] R. Nagarajan, M. Scutari, S. Lèbre, Bayesian Network Inference Algorithms, Springer New York, New York, NY, 2013, pp. 85–101. URL: https://doi.org/10.1007/978-1-4614-6446-4_4. doi:10.1007/978-1-4614-6446-4_4.

[15] J. Zhang, J. Tian, E. Bareinboim, Partial counterfactual identification from observational and experimental data, in: International Conference on Machine Learning, PMLR, 2022, pp. 26548–26558.

[16] M. Zaffalon, A. Antonucci, R. Cabañas, D. Huber, D. Azzimonti, Efficient computation of counterfactual bounds, International Journal of Approximate Reasoning (2024) 109111.

[17] M. Zaffalon, A. Antonucci, R. Cabañas, D. Huber, Approximating counterfactual bounds while fusing observational, biased and randomised data sources, International Journal of Approximate Reasoning 162 (2023) 109023.

[18] G. Hammerton, M. R. Munafò, Causal inference with observational data: the need for triangulation

of evidence, Psychological Medicine 51 (2021) 563–578. doi:10.1017/S0033291720005127.
[19] A. M. Berkman, C. R. Andersen, M. E. Roth, S. C. Gilchrist, Cardiovascular disease in adolescent and young adult cancer survivors: Impact of sociodemographic and modifiable risk factors, Cancer 129 (2023) 450–460. doi:10.1002/cncr.34505.
[20] D. O. Ondimu, G. M. Kikuvi, W. N. Otieno, Risk factors for hypertension among young adults (18-35) years attending in tenwek mission hospital, bomet county, kenya in 2018, Pan African Medical Journal 33 (2019). doi:10.11604/pamj.2019.33.210.18407.
[21] A. Saavedra, E. Rodrigues, D. Carvalho, Dislipidemia secundária a hipotiroidismo e colestase, Acta Médica Portuguesa 33 (2020) 204–207. doi:10.20344/amp.9944.
[22] J. Muhandiramge, J. R. Zalcberg, G. J. van Londen, E. T. Warner, P. R. Carr, A. Haydon, S. G. Orchard, Cardiovascular disease in adult cancer survivors: a review of current evidence, strategies for prevention and management, and future directions for cardio-oncology, Current Oncology Reports 24 (2022) 1579–1592. doi:10.1007/s11912-022-01309-w.