

Learning Ordinal Preferences on Multiattribute Domains: the Case of CP-nets

Yann Chevaleyre Frédéric Koriche Jérôme Lang
Jérôme Mengin Bruno Zanuttini*

Abstract

A recurrent issue in automated decision making is to extract a preference structure from a set of examples. In this paper, we investigate the problem of learning ordinal preference orderings over discrete multi-attribute, or combinatorial, domains. Specifically, we concentrate on the learnability issue of conditional preference networks, or *CP-nets*, that have recently emerged as a popular graphical language for representing ordinal preferences in a concise and intuitive manner. This paper provides results in both passive and active learning. In passive learning, the learner aims at finding a CP-net compatible with a given set of examples, while in active learning the learner searches for the cheapest interaction policy with the user for acquiring the target CP-net.

1 Introduction

A recurrent issue in both Artificial Intelligence (AI) and Decision Theory (DT) is to devise learning algorithms capable of extracting a preference structure from the user's behavior in presence of different alternatives. In many applications, such as electronic commerce, the space of alternatives has a combinatorial structure, due to the fact that each alternative is defined over a potentially large number of attributes.

For instance, suppose we observe a user expressing her preferences about airplane tickets: (a) she prefers an Aeroflot flight landing at Heathrow to a KLM flight landing at Gatwick and (b) she prefers an Aeroflot flight landing at Heathrow to a KLM flight landing at Heathrow. An intuitively correct hypothesis that explains her behavior is that she prefers Aeroflot to KLM unconditionally, and Heathrow to Gatwick, again unconditionally. Such a hypothesis allows for predicting that she will prefer an Aeroflot flight landing at Heathrow to anything else, and an Aeroflot flight landing at Gatwick to a KLM flight landing at Gatwick. Yet, this hypothesis is not able to predict whether she will prefer an Aeroflot flight landing at Gatwick or a KLM flight landing at Heathrow. Now, if we observe later that she prefers a KLM flight landing at Gatwick to a KLM flight landing at Heathrow, the previous hypothesis must be updated. A new possible hypothesis, among others, could be that she prefers Aeroflot to KLM, and Heathrow to Gatwick when flying on Aeroflot and *vice versa* when flying on KLM.

A key notion underlying this example is *conditional preferential independence*, which stems from multiattribute decision theory [21]. In the first hypothesis above,

*Partially supported by the ANR projects CANAR (ANR-06-BLAN-0383-02) and PHAC (ANR-05-BLAN-0384-01)

the preference over the values of each attribute is independent of the values of other attributes – in such a case, the user’s preference relation is said to be *separable*. The second hypothesis is not separable, because the user’s preference between airports is conditioned by the airline company.

The process of inducing a model of the user’s preferences as in the above toy example is generally called *learning from examples* or *passive learning*. The latter term emphasizes the fact that the learner has no control over the set of examples the user will provide. In another context, we may have the opportunity to learn the local preference relations (on airline companies and airports) by actively querying the user rather than by observing her. In this case we have to find the optimal sequence of queries. This problem, known as *learning by queries*, is related to *preference elicitation*: the system interacts with the user by asking her specific requests, until she has found her target object or left the system [12]. Yet in other contexts, it is relevant to mix both active and passive learning. Consider for instance a system helping a user finding a flat from a large database, such as in [30, 31]. A flat is described by attributes such as price, location, size etc. The system can start to extract a pool of preferences by observing the user’s behavior, and then use queries so as to help her finding her ideal flat while minimizing the number of interactions.

In all the contexts we considered above, the system has to learn preferences of a user over a set of alternatives that possesses a combinatorial structure. Preferences over combinatorial domains have been investigated in detail by researchers in multiattribute Decision Theory and Artificial Intelligence. Researchers in multiattribute DT have focused on *modeling* preferences, that is, giving axiomatic characterizations of classes of preference relations or utility functions, while researchers in AI have concentrated on the development of languages for *representing* preferences that are computationally efficient; such languages have to express preferences as succinctly as possible, and to come with fast algorithms for finding optimal alternatives.

These classes of models and languages can be partitioned by examining the mathematical nature of the preferences they consider. Namely, a distinction is made between *ordinal* preferences that consist in ranking the alternatives, and *numerical* preferences consisting of utility functions mapping each alternative to some number. There have been a number of works for learning or eliciting numerical preferences on multiattribute domains. Elicitation of multi-attribute utility functions is considered in [19, 10, 11, 17, 8]. A related stream of work is preference elicitation in the context of combinatorial auctions [27]: what has to be learnt is the valuation function of every buyer, which associates with every combination of goods the maximum value that she is ready to pay for it. Recently, there has been a growing interest for learning preferences from ordinal comparisons using numerical models. Many standard machine learning methods, such as *neural networks* [9] or *support vector machines* [15], have been adapted to this framework, often called *learning to rank instances* by the machine learning community.

In contrast, the learning of preferences using *ordinal* models has received much less attention. In fact, the most studied model up to now is *lexicographic preferences for binary attributes*. The problem is first addressed in [14], which investigate several problems regarding learning lexicographical preference relations. First they give a polynomial-time algorithm which, from a set of pairwise comparisons between vectors of values, determines the importance order of the criteria if it is possible, or decides whether the sample is insufficient or inconsistent. Then they consider query-based learning and give an oracle based algorithm. Finally, they investigate the worst case situation and deal with adversarial models. Then, [28] show that although it is possible

to determine in polynomial time whether there exists a lexicographical model compatible with a set of examples, described as pairwise comparisons between tuples of values, the corresponding optimization problem is NP-hard, and can even not be approximated in polynomial time to within a constant factor. They also determine the Vapnik-Chervonenkis dimension of lexicographical preference relations: they show that it is equal to the number of attributes. Finally, Yaman et al. [33] do not commit to a single lexicographic preference relation but approximate the target using the votes of a *collection* of consistent lexicographic preference relations. In a nutshell, learning lexicographic preference relations proves not to be so hard, but this comes with a price, namely, the hypothesis made is highly restrictive. Indeed, lexicographic preferences are a very specific form of preferences that is not often met in practice.

In this paper we focus on learning another class of preference relations, where the hypotheses we make bear on the preferential dependence structure. As we said above, a key point, when dealing with ordinal preferences on combinatorial domains, is the dependence structure between attributes. Now, conditional preference networks or *CP-nets* are a graphical language for representing preferences based on conditional preferential independence [6]. A CP-net is composed of a directed graph representing the preferential dependencies between variables, and a set of conditional tables (one for each variable), expressing, for each variable, the local preference on the values of its domain given all possible combinations of values of its parents. The transitive closure of these local preferences is a partial order over the set of alternatives, which can be extended into several total orders. CP-nets and their generalizations are probably the most popular compact representation language for ordinal preferences in multiattribute domains. The local tables in a CP-net express the local preferences for each attribute, given the values of the attributes they depend of. Therefore, learning a preference relation on a multiattribute domain, given some hypothesis on the dependence structure that the preference relation enjoys, comes down, at least to some extent, to *learning a CP-net*. Note that, while many facets of CP-nets have been studied, such as consistency, dominance checking, and optimization (constrained and unconstrained), learning a CP-net from a user has only rarely, and very recently, been addressed. One exception is [5], who proposed an algorithm that, given a set of examples, outputs a CP-net that implies them (see Section 5 for more details). Although not directly concerned with CP-nets, a related work is [26], which proposes to learn a preference theory in the sense of [16], consistent with a set of examples.

The aim of this position paper is to examine the problem of CP-net learning according to several dimensions, and to give a few learnability results. A first dimension that naturally emerges from preference learning is to consider whether the user’s preferences are representable, or not, by a CP-net. If the target preference relation can be described by a CP-net, the goal is to identify this network. Alternatively, if the target preference relation is not representable by a CP-net, we can only hope finding an approximation of it. Among the candidate approximations, some of them are particularly relevant from a reasoning viewpoint. From this perspective, we shall concentrate on finding CP-nets for which the target relation is a “completion” of the hypothesized relation. Orthogonally, a second dimension in CP-net learning is to consider whether the learning process is merely passive, by simply observing examples, or active, by allowing the learner to ask membership queries.

In the different learning models that arise from these dimensions, we will investigate the learnability issues in terms of the worst-case number of resources required to converge toward the desired CP-net, where resources refer both to the running time, the sample complexity in passive learning, and the query complexity in active learning.

Section 2 provides the necessary background about CP-nets. In Section 3, we extend the paradigm of concept learning to preference learning and introduce two frameworks, one for the problem of learning partial orderings that are representable by a CP-net, and the other for the problem of learning linear orderings that are not representable by a CP-net. Our learnability results lie in the next three sections. Namely, section 4 focuses on the VC-dimension and the approximate fingerprint property for classes of CP-nets. Section 5 addresses passive learning of CP-nets. Section 6 considers active learning of CP-nets. Finally, section 7 briefly discusses issues for further work.

Given that this is a position paper, the proofs of the results are only sketched, or even omitted. They can be found in the technical reports [13, 22, 24].

2 Conditional Preference Networks

We assume that we have a finite set $\mathcal{V} = \{X_1, \dots, X_n\}$ of n variables, or *attributes*, with associated finite *domains* D_1, \dots, D_n . Then $D = D_1 \times \dots \times D_n$ is the set of all complete assignments, called *outcomes*.

For any nonempty subset \mathcal{X} of \mathcal{V} , we let $D_{\mathcal{X}} = \times_{X_i \in \mathcal{X}} D_i$. Elements of $D_{\mathcal{X}}$ are called \mathcal{X} -assignments and denoted using vectorial notation, e.g., \mathbf{x} . For any disjoint subsets \mathcal{X} and \mathcal{Y} of \mathcal{V} , $\mathbf{x} \in D_{\mathcal{X}}$ and $\mathbf{y} \in D_{\mathcal{Y}}$, the concatenation of \mathbf{x} and \mathbf{y} , \mathbf{xy} , is formally defined as $\mathbf{x} \cup \mathbf{y}$ – i.e., it is the $X \cup Y$ -assignment which assigns to attributes in X (resp. Y) the value assigned by \mathbf{x} (resp. \mathbf{y}). An attribute X_i is *binary* if D_i has two elements, which by convention we note x_i and \bar{x}_i .

A *preference relation* on a multiattribute domain D is a weak order on D , that is, a reflexive and transitive binary relation \succeq . If furthermore \succeq is connected, that is, if for every $\mathbf{x}, \mathbf{y} \in D$ we have either $\mathbf{x} \succeq \mathbf{y}$ or $\mathbf{y} \succeq \mathbf{x}$ then \succeq is a *complete preference relation*. A *strict preference relation* \succ is an order on D , that is, an irreflexive and transitive (thus asymmetric) binary relation. If moreover \succ is connected then \succ is a *linear preference relation*. From a preference relation \succeq we define a strict preference relation in the usual way: $\mathbf{x} \succ \mathbf{y}$ iff $\mathbf{x} \succeq \mathbf{y}$ and not $(\mathbf{y} \succeq \mathbf{x})$.

Preferences between outcomes that differ in the value of one attribute only, all other attributes being equal (or *ceteris paribus*) are often easy to assert, and to understand. CP-nets [6] are a graphical language for representing such preferences. Informally, a CP-net is composed of a directed graph representing the preferential dependencies between attributes, and a set of conditional preference tables (one for each attribute), expressing, for each attribute, the local preference on the values of its domain given all possible combinations of values of its parents.

Let us call a *swap* any pair of outcomes $\{\mathbf{x}, \mathbf{y}\}$ that differ in the value of one attribute only, and let us then call *swapped attribute* the attribute that has different values in \mathbf{x} and \mathbf{y} . A CP-net specifies, for every swap $\{\mathbf{x}, \mathbf{y}\}$, which of $\mathbf{x} \succ \mathbf{y}$ or $\mathbf{y} \succ \mathbf{x}$ is true. This can be achieved in a compact manner when there are many *independencies* among attributes.


Example 1 Consider three attributes A, B and C , with respective domains $\{a, \bar{a}\}$, $\{b, \bar{b}\}$ and $\{c, \bar{c}\}$, and suppose that the four swaps on B are ordered as follows: $abc \succ a\bar{b}c, \bar{a}bc \succ \bar{a}\bar{b}c, ab\bar{c} \succ ab\bar{c}, \bar{a}\bar{b}\bar{c} \succ \bar{a}\bar{b}\bar{c}$. We can see that, irrespective of the value of C , if a is the case, then b is preferred to \bar{b} , whereas if \bar{a} is the case, then \bar{b} is preferred to b . We can represent this ordering on the B -swaps with two conditional preferences: $a : b > \bar{b}$ and $\bar{a} : \bar{b} > b$, and say that, given A, B is (conditionally) preferentially independent of C .

Definition 2 Let $\{\mathcal{X}, \mathcal{Y}, \mathcal{Z}\}$ be a partition of the set \mathcal{V} and \succ a linear preference relation over D . \mathcal{X} is (conditionally) preferentially independent of \mathcal{Y} given \mathcal{Z} (w.r.t. \succ) if and only if for all $\mathbf{x}_1, \mathbf{x}_2 \in D_{\mathcal{X}}, \mathbf{y}_1, \mathbf{y}_2 \in D_{\mathcal{Y}}, \mathbf{z} \in D_{\mathcal{Z}}$,


$$\mathbf{x}_1\mathbf{y}_1\mathbf{z} \succ \mathbf{x}_2\mathbf{y}_1\mathbf{z} \text{ iff } \mathbf{x}_1\mathbf{y}_2\mathbf{z} \succ \mathbf{x}_2\mathbf{y}_2\mathbf{z}$$

Definition 3 A CP-net over attributes $\mathcal{V} = \{X_1, \dots, X_n\}$ with domains D_1, \dots, D_n is a pair $N = \langle G, P \rangle$ where G is a directed graph over X_1, \dots, X_n and P is a set of conditional preference tables, one CPT(X_i) for each $X_i \in \mathcal{V}$. For attribute X_i , we denote by $\text{Pa}(X_i)$ (resp. $\text{NonPa}(X_i)$) the set of parents of X_i in G (resp. $\mathcal{V} - (\{X_i\} \cup \text{Pa}(X_i))$). Each conditional preference table is a list of rows (also called entries or rules) of the form $\mathbf{u} : x_i^{(1)} > \dots > x_i^{(m)}$, where \mathbf{u} is an instantiation of $\text{Pa}(X_i)$ and $x_i^{(1)} > \dots > x_i^{(m)}$ is a linear ordering of the domain D_i (with $m = |D_i|$). It indicates that $\mathbf{u}\mathbf{z}x_i^{(j)} \succ \mathbf{u}\mathbf{z}x_i^{(j+1)}$ for every possible instantiation \mathbf{z} of $\text{NonPa}(X_i)$.

Example 4 A CP-net over attributes A, B and C , with respective domains $\{a, \bar{a}\}, \{b, \bar{b}\}$ and $\{c, \bar{c}\}$ is:



where $\textcircled{X} \longrightarrow \textcircled{Y}$ means “ X is a parent of Y ”. The associated ordering of the swaps is:



where $\mathbf{x} \longrightarrow \mathbf{y}$ means “ \mathbf{x} is preferred to \mathbf{y} ”.

The size of a CP-net N , denoted by $|N|$, is the number of entries in the preference tables of N ; if the dependency graph of N is G then $|N| = \sum_{X_i \in \mathcal{V}} \prod_{X_j \in \text{Pa}(X_i)} |D_j|$, which in the case of binary CP-nets boils down to $|N| = \sum_{X_i \in \mathcal{V}} 2^{|\text{Pa}(X_i)|}$.

Although a CP-net only specifies an ordering of all swaps, we are naturally interested in the transitive closure of this ordering; for a CP-net N , we write \succ_N for this transitive closure. Note that this relation \succ_N may not be total, and it may not be a strict order since it may contain cycles, and thus not be irreflexive. We know from [6] that if G is acyclic, then \succ_N is a strict order (i.e. contains no cycles). In this case we say that N is *consistent*; \succ_N may still not be total, it can then be completed in a number of linear preference relations. If \succ is one of them, we say that \succ is a *completion* of N , or that it is *compatible* with N . When \succ_N is not irreflexive, we say that N is *inconsistent*.

Note that several CP-nets can induce the same preference relation: for example, if a CP-net N over A, B contains the table $b : a > \bar{a}$ and $\bar{b} : a > \bar{a}$, then the CP-net N' in which A has no parent and the table $a > \bar{a}$ is equivalent to N . In general, for every

consistent CP-net N there a unique CP-net N' equivalent to N that is minimal in the number of parents/table entries for each variable.

Importantly, note that every linear preference relation is compatible with exactly one (minimal) CP-net. For instance, $ab \succ \bar{a}\bar{b} \succ a\bar{b} \succ \bar{a}b$ is compatible with the CP-net $N = \{a : b \succ \bar{b}, \bar{a} : \bar{b} \succ b, b : a \succ \bar{a}, \bar{b} : \bar{a} \succ a\}$.

The following property will be frequently used in the remaining sections.

Proposition 5 ([6]) *Let N be an acyclic CP-net and \mathbf{x}, \mathbf{y} two outcomes. Then $\mathbf{x} \succ_N \mathbf{y}$ iff there is a sequence of swaps $\{\mathbf{x}^0, \mathbf{x}^1\}, \{\mathbf{x}^1, \mathbf{x}^2\}, \dots, \{\mathbf{x}^{k-1}, \mathbf{x}^k\}$ such that $\mathbf{x}^0 = \mathbf{x}$, $\mathbf{x}^k = \mathbf{y}$, and for every $0 \leq i < k$, $\mathbf{x}^i \succ_N \mathbf{x}^{i+1}$, that is, if X_{j_i} is the attribute swapped between \mathbf{x}^i and \mathbf{x}^{i+1} , and if \mathbf{u} is the vector of values commonly assigned by \mathbf{x} and \mathbf{y} to the parents of X_{j_i} , then N contains $\mathbf{u} : x_{j_i}^i > x_{j_i}^{i+1}$;*

Though a CP-net is usually defined as above, most of the results presented here extend to *possibly incomplete CP-nets*. Such a CP-net is one in which each conditional preference table $CPT(X_i)$ contains *at most* one conditional preference rule per instantiation of $\text{Pa}(X_i)$ (instead of *exactly one*). A particular case is when some of the tables are empty. The semantics of an incomplete CP-net is still given by the transitive closure of the dominance relation on swaps induced by the rules. An important difference is that in an incomplete CP-net, not all swaps are comparable.

The rationale for considering incomplete CP-nets can be understood with the following example. A user may well know that she prefers traveling by bus rather than in the subway in Paris, *vice-versa* in London, but be unable to state her preference for a city in which she has never been, say Madrid. This does not mean that she would not have a preference, rather that she does not know it (so far). In this case, a variable encoding the transportation means (with values *subway* and *bus*) would have the variable encoding the city as its parent, with values *Paris*, *London*, and *Madrid*, but would contain only two rules.

3 Learning CP-nets: learning what?

The problem of concept learning is to extrapolate from a collection of examples, each labeled as either positive or negative by some unknown target concept, a representation of this concept that accurately labels future, unlabeled examples. Most concept learning algorithms operate over some *concept class*, which captures the set of concepts that the learner can potentially generate over all possible sets of training examples.

In the setting suggested by our framework, a *concept* is a preference relation \succ over $D_{\mathcal{V}}$. We say that a concept \succ is *representable* by a CP-net N if the induced ordering \succ_N coincides with \succ , that is, $\succ = \succ_N$. For example, the preference relation \succ defined by $\{ab \succ \bar{a}\bar{b}, ab \succ \bar{a}\bar{b}, \bar{a}\bar{b} \succ a\bar{b}, \bar{a}\bar{b} \succ a\bar{b}\}$ is representable by the CP-net N specified by $\{a : b \succ \bar{b}, \bar{a} : \bar{b} \succ b, b : a \succ \bar{a}, \bar{b} : \bar{a} \succ a\}$. A *representation class* is a collection \mathcal{N} of CP-nets, and the *concept class* $\mathcal{C}_{\mathcal{N}}$ defined over \mathcal{N} is the set of all preference relations \succ that are representable by a CP-net N in \mathcal{N} .

The concept class \mathcal{C}_{ACY} of all preference relations that are representable by an acyclic CP-net will play a fundamental role in this study. Since acyclic CP-nets are necessarily consistent, any target concept \succ in this class can be represented by a unique minimal CP-net N . So, with a slight abuse of language, we shall simply say that \mathcal{C}_{ACY} is the class of all acyclic CP-nets. Similar considerations will be applied to subclasses of \mathcal{C}_{ACY} such as, for example, the class $\mathcal{C}_{\text{TREE}}$ of tree-structured CP-nets.

With these notions in hand, we assume that the user has in mind a target preference ordering \succ , and the learner has at its disposal a prefixed and known class of CP-nets \mathcal{C}_N . In this study, we shall consider two different types of target concepts.

- (A) The target concept is a preference relation \succ that belongs to the learner’s concept class \mathcal{C}_N . In other words, there is a CP-net N in \mathcal{C}_N , such that \succ_N coincides with \succ . In this context, the goal of the learner is to find N .
- (B) The target concept is a preference relation \succ that does *not necessarily* belong to the learner’s concept class \mathcal{C}_N . For example, we can easily observe that the linear ordering \succ defined by $\{ab \succ \bar{a}\bar{b} \succ a\bar{b} \succ \bar{a}b\}$ cannot be represented by any CP-net. Still, we shall make the assumption that \succ is a completion of some representation in \mathcal{N} . Specifically, we say that \succ is a *completion* of a CP-net N if $\mathbf{x} \succ_N \mathbf{y}$ implies $\mathbf{x} \succ \mathbf{y}$ for any pair of outcomes (\mathbf{x}, \mathbf{y}) . For example, the above ordering \succ is a completion of $N = \{a : b \succ \bar{b}, \bar{a} : \bar{b} \succ b, b : a \succ \bar{a}, \bar{b} : \bar{a} \succ a\}$. In this “agnostic” setting, the goal of the learner is to find a CP-net N for which \succ is a completion of it.

Note that the distinction between (A) and (B) is not on the set of objects we want to learn, but on their interpretation, which has crucial consequences on how examples are interpreted.

3.1 Learning a preference relation induced by a CP-net

Let us start with context (A) where the target concept is representable by a CP-net N of some given representation class \mathcal{N} . Here, an instance, or *example*, is a pair (\mathbf{x}, \mathbf{y}) of outcomes, and an *instance class* is a set \mathcal{E} of examples. Given a target concept \succ_N , and an example (\mathbf{x}, \mathbf{y}) , we say that (\mathbf{x}, \mathbf{y}) is *positive* for \succ_N if $\mathbf{x} \succ_N \mathbf{y}$, that is, \mathbf{x} dominates \mathbf{y} according to \succ_N . Dually, (\mathbf{x}, \mathbf{y}) is *negative* for \succ_N if $\mathbf{x} \not\succ_N \mathbf{y}$. It is important to keep in mind that, in general, if the pair (\mathbf{x}, \mathbf{y}) is a negative example of \succ_N then the reverse pair (\mathbf{y}, \mathbf{x}) is *not necessarily* a positive example of \succ_N .

In this context, our learning problem can be seen as a standard concept learning problem: given positive and negative examples, we want to find a CP-net that “implies” all positive examples and no negative example.

Definition 6 Let N be a CP-net over \mathcal{V} . An example (\mathbf{x}, \mathbf{y}) is entailed by N if $\mathbf{x} \succ_N \mathbf{y}$. A set of examples E is *implicatively consistent* with N , or *implied by N* , if

- all positive examples in E are entailed by N ;
- no negative example in E is entailed by N .

Finally, we shall say that a training set E is *implicatively compatible* if it is implied by at least one CP-net.

3.2 Learning a CP-net which approximates the user’s preferences

Now, we turn to context (B) and make the assumption that the user’s preference relation \succ is a linear order, in general not exactly representable by any CP-net. In this framework, we start by examining an appropriate notion of consistency between a CP-net and a training set. Consider the following example:

Example 7 We have two binary attributes X_1 and X_2 (with domains $\{x_1, \bar{x}_1\}$ and $\{x_2, \bar{x}_2\}$), and the set of positive examples

$$E = \{x_1x_2 \succ x_1\bar{x}_2, x_1\bar{x}_2 \succ \bar{x}_1x_2, \bar{x}_1x_2 \succ \bar{x}_1\bar{x}_2\}$$

What do we expect to learn from the above set of examples E ? The transitive closure of E is the complete preference relation $x_1x_2 \succ x_1\bar{x}_2 \succ \bar{x}_1x_2 \succ \bar{x}_1\bar{x}_2$. This preference relation is *separable* (the agent unconditionally prefers x_1 to \bar{x}_1 and x_2 to \bar{x}_2). The fact that $x_1\bar{x}_2$ is preferred to \bar{x}_1x_2 simply means that when asked to choose between X_1 and X_2 , the agent prefers to give up X_2 (think of X_1 meaning “getting rich” and X_2 meaning “beautiful weather tomorrow”). Intuitively, since E is separable, we expect to output a structure N that contains $x_1 \succ \bar{x}_1$ and $x_2 \succ \bar{x}_2$. However, no CP-net implies E , whatever the dependencies. The structure N induces a *partial* preference relation in which $x_1\bar{x}_2$ and \bar{x}_1x_2 are incomparable. More generally, no *ceteris paribus* structure can “explain” that $x_1 \succ \bar{x}_1$ is “more important” than $x_2 \succ \bar{x}_2$ (i.e., with no intermediate alternative). Therefore, if we look for a structure *implying* all the examples, we will simply output “failure”. On the other hand, if we look for a separable CP structure that is simply *contingent* with the examples, i.e., that does not imply the contrary of the examples, we will output N .

The explanation is that when an agent expresses a CP-net, the preference relation induced by this CP-net *is not meant to be the whole agent’s preference relation, but a subset (or a lower approximation) of it*. In other terms, when an agent expresses the CP-net N , she simply expresses that she prefers x_1 to \bar{x}_1 *ceteris paribus* (i.e., for a fixed value of X_2) and similarly for the preference $x_2 \succ \bar{x}_2$; the fact that $x_1\bar{x}_2$ and \bar{x}_1x_2 are incomparable in N surely does not mean that the user really sees them incomparable, but, more technically, that CP-nets are not expressive enough for representing the missing preference $x_1\bar{x}_2 \succ \bar{x}_1x_2$ ¹.

Therefore, rather than looking for a CP-net that implies the examples, we should rather look for a CP-net whose preference relation is contingent with the examples. A first way of understanding consistency is to require that the learnt CP-net N be such that the examples are consistent with at least one preference relation extending N . Yet, there are cases where it may even be too strong to require that one of the completions of \succ_N contains all the examples, in particular if they come from multiple users (given that we want to learn the generic preferences of a group of users), or a single user in different contexts:

Example 8 Suppose that we learn that all users in a group unconditionally prefer x_1 to \bar{x}_1 and x_2 to \bar{x}_2 , whereas their preferences between $x_1\bar{x}_2$ and \bar{x}_1x_2 may differ (think as x_1 and x_2 as, respectively, “being invited to a fine dinner” and “receiving a \$50 award”): then $E \supseteq \{\bar{x}_1x_2 \succ x_1\bar{x}_2, x_1\bar{x}_2 \succ \bar{x}_1x_2\}$. E is clearly inconsistent, so there cannot be any preference structure whose ordering can be completed into a linear preference relation that contains E . However, if $\mathcal{N} = \{x_1 \succ \bar{x}_1, x_2 \succ \bar{x}_2\}$, then each example in E is (individually) contained in at least one completion of $\succ_{\mathcal{N}}$.

Such considerations lead us to define two new notions of compatibility of a CP-net with a set of examples. Note that because the target concept is a linear order, (\mathbf{x}, \mathbf{y}) is a negative example if and only if (\mathbf{y}, \mathbf{x}) is a positive one. For this reason, we can make the assumption that all examples in the learner’s training set E are *positive*, with the implicit knowledge that the reverse (\mathbf{y}, \mathbf{x}) of any pair (\mathbf{x}, \mathbf{y}) in E is negative.

¹If we want to do this, we have to resort to a more expressive language such as TCP-nets [7] or conditional preference theories [32].

Definition 9 Let N be a CP-net over \mathcal{V} . An example (\mathbf{x}, \mathbf{y}) is consistent by completion with N if there is a completion \succ of \succ_N such that $\mathbf{x} \succ \mathbf{y}$. Furthermore, we will say that a set of examples E is:

- strongly consistent by completion with N if there is a completion \succ of \succ_N such that for all $(\mathbf{x}, \mathbf{y}) \in E$, $\mathbf{x} \succ \mathbf{y}$;
- weakly consistent by completion with N if every example $(\mathbf{x}, \mathbf{y}) \in E$ is individually consistent by completion with N .

Lastly, we will say that E is *strongly / weakly compatible* if it is strongly / weakly consistent by completion with at least one CP-net. Clearly, strong compatibility implies weak compatibility. Moreover, since an example (\mathbf{x}, \mathbf{y}) is consistent by completion with a CP-net N if and only if (\mathbf{y}, \mathbf{x}) is not implied by N , implicative compatibility implies strong compatibility.

As it stands, it turns out to be significantly more difficult to search for a CP-net strongly or weakly consistent with a set of examples than to search for a CP-net implicatively consistent with it. Therefore, we mainly focus on implicative compatibility; strong and weak compatibility will only be discussed in the context of *separable* CP-nets, that is, CP-nets where all variables are independent.

4 Learnability of CP-nets

In this section we investigate the theoretical limits concerning the learnability issue of CP-nets. In essence, the Vapnik-Chervonenkis dimension of a class gives upper bounds on the difficulty to learn, in terms of numbers of examples, while the approximate fingerprint property gives lower bounds.

4.1 Vapnik-Chervonenkis dimension

The *Vapnik-Chervonenkis (VC) dimension* of a class of concepts is a fundamental complexity measure used in theoretical machine learning. Intuitively, the VC-dimension of \mathcal{C} is the maximum number of informative examples which can be received by the learner, where an “informative” example is an observation that helps the learner reducing the number of consistent hypotheses.

Formally, let \mathcal{C} be a concept class defined over some representation class \mathcal{R} . A set of instances E is said to be *shattered* by \mathcal{C} if, whatever the partition of E into $E^+ \cup E^-$, there is a concept $N \in \mathcal{C}$ which admits all instances in E^+ as positive examples and all instances in E^- as negative examples. The *Vapnik-Chervonenkis dimension* of \mathcal{C} , denoted $VC(\mathcal{C})$, is the maximum size of a set of examples E which is shattered by \mathcal{C} . The intuition is that for such a set E , as long as the learner ignores the label of at least one example, at least two concepts in \mathcal{C} are consistent with the labels it has seen so far.

When the learner has access only to a certain kind of example (e.g., swap examples), it makes sense to adapt the notion of VC-dimension. So if \mathcal{E} is a class of instances, we write $VC_{\mathcal{E}}(\mathcal{C})$ for the VC-dimension of \mathcal{C} with respect to \mathcal{E} , that is, the maximum size of a $E \subseteq \mathcal{E}$ which is shattered by \mathcal{C} . Clearly, if $\mathcal{E} \subseteq \mathcal{E}'$, then $VC_{\mathcal{E}}(\mathcal{C}) \leq VC_{\mathcal{E}'}(\mathcal{C})$.

Observe that, as there are 2^m partitions of a set of m examples into positive and negative examples, each of which must be captured by a different concept, $VC(\mathcal{C}) \leq \log_2 |\mathcal{C}|$ always holds (whatever the class of examples).

We now give the VC-dimension of the class of all CP-nets which have a fixed graph. The intuition here is that since the parents of each variable are known, the quantity of information needed to characterize a CP-net is exactly 1 per possible rule in the CP-net, namely, one pair of outcomes which dictates the conclusion of the rule.

Proposition 10 *Let G be a graph, and let \mathcal{C}_G be the class of all concepts which are representable by a binary complete CP-net whose graph is G . Then the VC-dimension of \mathcal{C}_G with respect to swap examples is exactly the number of conditional preference rules in any such CP-net. If CP-nets are possibly incomplete, then the VC-dimension (w.r.t. swap examples) is still the number of rules in any complete CP-net on G .*

This property can help us finding an upper bound of the VC-dimension of acyclic CP-nets. Intuitively, the number of acyclic graphs with indegree at most k is lower bounded by $|\mathcal{C}_G|$, and upper bounded by $(n-1)^{n(k+1)}$. Since any binary-valued CP-net built over an acyclic graph of degree at most k allows at most $n2^k$ entries, there are at most $(n-1)^{n(k+1)}2^{n2^k}$ binary-valued acyclic CP-nets with degree at most k . We therefore obtain the following result.

Corollary 11 *Let $k \in o(n)$, and let \mathcal{C}_{acy}^k be the class of all concepts which are representable by a possibly incomplete binary CP-net whose graph is acyclic and with indegree at most k . Then, the VC-dimension of \mathcal{C}_{acy}^k with respect to swap or arbitrary examples is $\tilde{\Theta}(n2^k)$.*

Finally, without any restriction over the degree, it can be shown that the VC-dimension grows as $\Theta(2^n)$, which is still much below the VC-dimension of the class of all possible CP-nets.

Corollary 12 *Let \mathcal{C}_{acy} be the class of all concepts which are representable by a possibly incomplete binary CP-net whose graph is acyclic. Then, the VC-dimension of \mathcal{C}_{acy} with respect to swap examples is $\Theta(2^n)$.*

4.2 Approximate fingerprints

Approximate fingerprints are another powerful tool for obtaining non learnability results in active learning. Intuitively, a class of concepts \mathcal{C} has the approximate fingerprint property if there is a subset \mathcal{C}^* of \mathcal{C} such that for any concept $N \in \mathcal{C}$, there is an example with which N is consistent, but with which only a superpolynomially small fraction of the concepts in \mathcal{C}^* are also consistent.

This property can be used to show that in an interactive learning setting, a hypothesis $\hat{N} \in \mathcal{C}$ may fail on an example which only gives clues about superpolynomially few candidate hypotheses. Hence, if the learner only gets information from such failures, in the worst case it *necessarily* makes an exponential number of errors before correctly identifying the target concept. We refer the reader to [2] for formal details.

Obviously, as for the VC-dimension, the definition of approximate fingerprints can be restricted to instance classes \mathcal{E} . Observe that if $\mathcal{E} \subseteq \mathcal{E}'$ and \mathcal{C} has the approximate fingerprint property with respect to \mathcal{E} , then it also has this property with respect to \mathcal{E}' .

Proposition 13 ([22]) *Let \mathcal{C}_{acy} be the class of all concepts which are representable by a binary-valued complete CP-net whose graph is acyclic. Then \mathcal{C}_{acy} has the approximate fingerprint property with respect to swap examples.*

Proposition 14 ([22]) *Let \mathcal{C}_{tree} be the class of all concepts which are representable by a binary-valued complete CP-net whose graph is a tree. Then \mathcal{C}_{tree} has the approximate fingerprint property with respect to arbitrary examples.*

5 Passive learning of CP-Nets

In this section, we investigate *passive* learning of CP-nets. In this setting, the only information about the target concept available to the learner is a set of examples. We shall concentrate here on the widely studied *Probably Approximately Correct* (PAC) learning model introduced by [29]. The intent of this model is to obtain with high probability a representation that is a good approximation of the target concept. To formalize the notion of good approximation, we need to assume that there is some fixed, but unknown, probability distribution \mathcal{D} defined on the example space \mathcal{E} , from which the available examples were drawn. In our case, \mathcal{D} would define a probability over each instance (\mathbf{x}, \mathbf{y}) . Given a target concept \succ , we then define the *error* of an hypothesized CP-net \hat{N} as the probability that \succ and $\succ_{\hat{N}}$ disagree on an example:

$$error(N) = \Pr_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}} [(\mathbf{x} \succ \mathbf{y} \text{ and } \mathbf{x} \not\succ_{\hat{N}} \mathbf{y}) \text{ or } (\mathbf{x} \not\succ \mathbf{y} \text{ and } \mathbf{x} \succ_{\hat{N}} \mathbf{y})]$$

How does one generate a *good* approximation? In the PAC model, one does this by looking at an example set, in which each example (\mathbf{x}, \mathbf{y}) has been drawn independently from the random distribution \mathcal{D} , and labeled with a “+” (positive) if $\mathbf{x} \succ \mathbf{y}$ and with “−” (negative) if $\mathbf{x} \not\succ \mathbf{y}$.

Thus, in the PAC setting, training and testing use the same distribution, and there is no noise in either phase. A learning algorithm is then a computational procedure that takes a sample of the target concept \succ , consisting of a sequence of independent random examples of \succ , and returns an hypothesis. We can define PAC learnability of CP-nets as follows.

Definition 15 (PAC learning) *A concept class $\mathcal{C}_{\mathcal{N}}$ is PAC learnable by an example class \mathcal{E} if there is a polynomial time learning algorithm A and a polynomial $p(\cdot, \cdot, \cdot)$ such that for any target concept \succ in $\mathcal{C}_{\mathcal{N}}$ over n variables, any probability distribution \mathcal{D} over \mathcal{E} , and any parameters $\delta, \epsilon \in (0, 1)$, if the algorithm A is given at least $p(n, \frac{1}{\epsilon}, \frac{1}{\delta})$ independent random examples of \succ drawn according to \mathcal{D} , then with probability at least $1 - \delta$, A returns a hypothesis $\hat{N} \in \mathcal{N}$ with $error(\hat{N}) \leq \epsilon$. The smallest such polynomial p is called the sample complexity of the learning algorithm A .*

The intent of this definition is that the learning algorithm must process the examples in polynomial time, and must be able to produce a good approximation of the target concept with high probability using only a reasonable number of training examples.

It is important to emphasize that our learnability results are defined over specific instance classes. In particular, if \mathcal{E} is the class of all swap instances, then any distribution \mathcal{D} over \mathcal{E} will assign a zero probability to any “non-swap” instance. This restriction has deep consequences on the predictive power of P-nets. Namely, even if a positive learnability result with swap instances guarantees that the hypothesized CP-net \hat{N} is expected to correctly classify “swap” instances drawn independently at random according to the distribution \mathcal{D} , such a result does *not* ensure that the learner will correctly classify arbitrary outcome pairs. Indeed, even if the probability of making a mistake on swaps is low, the probability of making a mistake on an arbitrary instance (\mathbf{x}, \mathbf{y}) may increase along an improving sequence from \mathbf{y} to \mathbf{x} .

Many positive learnability results in the PAC model are obtained by showing that (1) there is an efficient algorithm capable of finding an hypothesized representation that is *consistent by implication* with a given sample of the target concept (called a *consistent algorithm*), and (2) the sample complexity of any such algorithm is polynomial.

The sample complexity of a consistent learning algorithm is usually measured using the VC-dimension of the concept class $\mathcal{C}_{\mathcal{N}}$. Indeed, it is shown in [4] that the sample complexity of a consistent learning algorithm is at most

$$\frac{1}{\epsilon(1-\sqrt{\epsilon})} \left(2\text{VC}(\mathcal{C}_{\mathcal{N}}) \ln \frac{6}{\epsilon} + \ln \frac{2}{\delta} \right) \quad (1)$$

A preliminary work on the passive learning of CP-nets is [5]. They give an algorithm which, given a set of positive examples, outputs a CP-net that implies them, under some conditions. It is not entirely clear yet which class of CP-nets is learned by this algorithm.

5.1 PAC Learning of Acyclic CP-Nets

We first investigate PAC learnability of various classes of acyclic CP-nets when the examples provided to the learner are swaps. Recall that for such examples, the dominance test with acyclic CP-nets is linear-time solvable (simple lookup in the conditional preference table), contrary to the general case.

We first show that even for the restricted class of concepts which are representable by a CP-net whose graph is a chain, the consistency problem is NP-complete. It follows directly that this class is not PAC-learnable with the very weak restriction that the produced hypothesis classifies correctly the examples received.

Proposition 16 *Deciding whether there is a binary-valued complete CP-net whose graph is a chain and which implies a given set of swaps is NP-complete. The result holds even if all examples are positive.*

A proof based on a reduction from the hamiltonian path problem can be found in [13]. Another interesting class of CP-nets is that of acyclic *singly-connected* CP-nets [6], that is, those acyclic CP-nets in whose graph each pair of vertices is connected by at most one directed path. Unfortunately, again we have a negative result for PAC learnability of such CP-nets.

Proposition 17 *Deciding whether there is a binary-valued complete CP-net whose graph is acyclic singly-connected and which implies a given set of swaps is NP-complete. The result holds even if all examples are positive.*

This result [13] can be proven with a reduction from propositional satisfiability (SAT). We conjecture that a similar negative result holds for more general classes of acyclic CP-nets.

We now turn to positive results with swaps. Observe that if $\mathcal{C}_{\mathcal{N}}$ is PAC learnable with swaps and $\mathcal{C}_{\mathcal{N}} \subseteq \mathcal{C}_{\mathcal{N}'}$, then $\mathcal{C}_{\mathcal{N}'}$ is not necessarily PAC learnable with swaps, as well. So our positive results do not contradict the negative ones. The main result is that tree CP-nets are PAC-learnable from swaps.

Proposition 18 *The class of all concepts which are representable by a (possibly incomplete) tree binary-valued CP-net is PAC-learnable from swap examples.*

In addition, the learning the structure of such a CP-net can be reduced to finding a spanning tree in a directed graph [13]. Since in general there may be several CP-nets which imply a given set of examples, it is interesting to impose some restrictions, e.g., on the degree of the forest (maximum number of children of a node). The next result states that the class of CP-nets whose graph is a forest with degree at most k is *improperly* PAC-learnable (in quasi-polynomial time) [13]. That is, it is “PAC-learnable”, but the hypothesis may be in a larger representation class than the target concept.

Proposition 19 *There is a quasi-polynomial time algorithm which, given a set of swaps E over n variables implied by a (possibly incomplete) binary-valued CP-net whose graph is a forest of degree k , computes a binary-valued CP-net which implies E and whose graph is a forest of degree at most $k + \log n$.*

Finally, we give a more general result about tree CP-nets with a bounded number of tables on *arbitrary* examples. By Cayley’s formula, we know that there are k^{k-1} rooted trees with k vertices. Each root is labeled by an unconditional rule of the form $p \succ \bar{p}$, and all other nodes are labeled by conditional rules of the form $p' : p \succ \bar{p}$, where p and p' are literals. There are $2n$ (possibly incomplete) tables with no condition per rule and $6n(n-1)$ tables with one condition per rule. It follows that the number C_k of CP-trees with at most k tables is bounded by $\sum_{i=0}^k 2n(6kn(n-1) + 1)^{k-1}$, which is indeed polynomial in k . So the VC-dimension of such CP-trees is polynomial in n .

Based on this result, we can use a simple consistent algorithm specified as follows. Start with the hypothesis set \mathcal{N} of all CP-trees with at most k tables. For each example (\mathbf{x}, \mathbf{y}) in E , remove any hypothesis N in \mathcal{N} that is inconsistent with (\mathbf{x}, \mathbf{y}) , that is, any hypothesis N for which the dominance test over (\mathbf{x}, \mathbf{y}) disagrees with its label. If the resulting set \mathcal{N} is empty then E is not consistent with N . Otherwise, pick an arbitrary tree from \mathcal{N} . Because the dominance test is quadratic in the number of variables for binary-valued CP-trees, the running time is polynomial in C_k .

Proposition 20 *The class $\mathcal{C}_{\text{TREE}}^k$ of all concepts representable by a (possibly incomplete) binary-valued CP-tree with at most k tables is PAC learnable from arbitrary examples.*

5.2 PAC Learning of Separable CP-nets

We now consider the task of learning a CP-net of the simplest form: the variables are independent. With binary variables, this means that if the possible values for variable X are x and \bar{x} , then the preference table for X contains either $x \succ \bar{x}$ or $\bar{x} \succ x$. In this case, checking if a given CP-net N entails $\mathbf{x} \succ_N \mathbf{y}$ for an arbitrary example is easy: let $\text{Diff}(\mathbf{x}, \mathbf{y}) = \{x_i \mid (\mathbf{x})_i = x_i \text{ and } (\mathbf{y})_i = \bar{x}_i\} \cup \{\bar{x}_i \mid (\mathbf{x})_i = \bar{x}_i \text{ and } (\mathbf{y})_i = x_i\}$; we can then use the following characterization of \succ_N (a corollary of Theorems 7 and 8 by [6]):

Lemma 21 *Let N be a separable CP-net over binary variables, and $\mathbf{y} \neq \mathbf{x}$. Then $\mathbf{x} \succ_N \mathbf{y}$ if and only if N contains $x_i \succ \bar{x}_i$ for every $x_i \in \text{Diff}(\mathbf{x}, \mathbf{y})$ and $\bar{x}_i \succ x_i$ for every $\bar{x}_i \in \text{Diff}(\mathbf{x}, \mathbf{y})$.*

Now, with each example (\mathbf{x}, \mathbf{y}) we associate the clause $C_{\mathbf{x}, \mathbf{y}}^-$ that contains $\neg x_i$ iff $x_i \in \text{Diff}(\mathbf{x}, \mathbf{y})$ and x_i iff $\bar{x}_i \in \text{Diff}(\mathbf{x}, \mathbf{y})$. The intended meaning of the literal $\neg x_i$ is that \bar{x}_i is preferred to x_i , whereas the meaning of the literal x_i is that x_i is preferred

to \bar{x}_i ; hence the meaning of the clause $C_{\mathbf{x},\mathbf{y}}^-$ is that $\mathbf{x} \not\prec_N \mathbf{y}$ for every separable CP-net N in which at least one of these local preferences is true, by virtue of the lemma above. For instance, if $\mathbf{x}=\bar{x}_1x_2x_3x_4$ and $\mathbf{y}=x_1\bar{x}_2x_3x_4$ then $\text{Diff}(\mathbf{x}, \mathbf{y}) = \{x_1, \bar{x}_2, x_4\}$ and $C_{\mathbf{x},\mathbf{y}}^- = x_1 \vee \neg x_2 \vee x_4$. This clause expresses that x_1 is preferred to \bar{x}_1 , or \bar{x}_2 is preferred to x_2 , or x_4 is preferred to \bar{x}_4 .

With each positive example (\mathbf{x}, \mathbf{y}) we can also associate the cube (conjunction of literals) $C_{\mathbf{x},\mathbf{y}}^+ \equiv \neg C_{\mathbf{y},\mathbf{x}}^-$. Given a set of examples E , let $\Gamma_E = \bigwedge \{C_e^* \mid e \in E\}$, where $C_{\mathbf{x},\mathbf{y}}^* = C_{\mathbf{x},\mathbf{y}}^+$ if (\mathbf{x}, \mathbf{y}) is a positive example, and $C_{\mathbf{x},\mathbf{y}}^* = C_{\mathbf{x},\mathbf{y}}^-$ if the example is negative. Clearly, Γ_E is equivalent to a set of clauses.

We now define the following one-to-one correspondence between truth assignments over $\{x_1, \dots, x_n\}$ and separable CP-nets over \mathcal{V} . If M is such a truth assignment, then N_M contains the preference $x_i \succ \bar{x}_i$ for every i such that $M \models x_i$ and the preference $\bar{x}_i \succ x_i$ for every i such that $M \models \neg x_i$. For instance, if $M(x_1) = \top$, $M(x_2) = \perp$, $M(x_3) = \perp$ and $M(x_4) = \top$ then N_M contains the preference tables $\{x_1 \succ \bar{x}_1, \bar{x}_2 \succ x_2, \bar{x}_3 \succ x_3, x_4 \succ \bar{x}_4\}$.

Given an interpretation M , it is not difficult to see that $M \models C_{\mathbf{x},\mathbf{y}}^-$ if and only if (\mathbf{x}, \mathbf{y}) is not implied by N_M or, equivalently, that $M \models C_{\mathbf{x},\mathbf{y}}^+$ if and only if (\mathbf{x}, \mathbf{y}) is implied by N_M (This is a straightforward consequence of Lemma 21). So, a set of examples E is implicatively consistent with N_M for a given model M if and only if $M \models \Gamma_E$. Therefore, searching for a CP-net that is implicatively consistent with a given set of examples amounts to searching for a model of the corresponding set of clauses, the size of which grows polynomially with the size of the set of examples.

This technique easily extends to nonbinary variables: we can use one propositional variable x_i^{kl} for every pair of distinct values $\{x_i^k, x_i^l\}$ for every variable X_i , where the intended meaning of x_i^{kl} is $x_i^k \succ_N x_i^l$, and add clauses to represent the transitivity of the relation \succ_N ; there is a polynomial number of them (details can be found in [23]).

The one-to-one correspondence given above is a reduction from our learning problem to satisfiability. It is actually possible to find a reduction on the opposite direction (see [23]), from which we get the following result:

Proposition 22 [23] *Deciding whether there is a (binary or non-binary) complete separable CP-net which is implicatively consistent with a given set of arbitrary examples is NP-complete. The result holds even if all examples are negative.*

It follows directly that this class is not PAC-learnable with the very weak restriction that the produced hypothesis classifies correctly the examples received.

However:

Proposition 23 [23] *Deciding whether there is a binary-valued complete separable CP-net which is implicatively consistent with a given set of positive examples can be done in polynomial time.*

Observe that this result can be extended to PAC-learnability of (possibly incomplete) separable CP-nets from positive examples, in the setting of one-sided errors [29]. This is because there is always a unique minimal (in terms of rules) incomplete separable CP-net which implies a given set of positive examples.

Now, as soon as E becomes large with respect to the number of attributes n , the chances that E is implicatively consistent with a separable CP-net become low. In this case, we may want to determine a separable CP-net that is implicatively consistent with as many examples of E as possible. This problem amounts to solving a MAXSAT problem, when each example is translated into exactly one clause of Γ_E , that is when we

have no positive example: the separable CP-net that best fits a set of positive examples corresponds to the interpretation maximizing the number of clauses from Γ_E satisfied. In this case, we can reuse algorithms for MAXSAT for computing a separable CP-net that best fits a set of positive examples, as well as polynomial approximation schemes. This extends to nonbinary variables, with the difference that the clauses representing transitivity of the local preference tables are protected.

Lastly, again using the same kind of translation, we easily get the following results:

Proposition 24 *If all variables are binary and all examples in E differ at most on two variables, then deciding whether there exists a separable CP-net consistent with E can be done in polynomial time; however, the corresponding optimization problem remains NP-hard.*

5.3 Learning a complete preference relation

We close this section by providing results about the learning context (B) specified in section 3: the target concept is a linear order, not necessarily representable by a CP-net. Our goal is to find a CP-net that would be a good representation for this relation. Recall that since the target is a linear order, we only need to consider *positive* examples.

In the rest of this section, we investigate in turn the problems of finding a CP-net that is weakly consistent with a given set of examples, then strongly consistent with it. We focus on the problem of finding *separable* CP-nets. A set of examples is said to be *weakly separable* (resp. *strongly separable*) if there exists a separable CP-net with which it is weakly (resp. strongly) consistent.

We start by showing how the search for a separable CP-net that is weakly consistent with a set of examples can be rewritten as an instance of propositional satisfiability (SAT). Recall from Section 5.2 that an example (x, y) can be translated into a clause $C_{x,y}^-$, the models of which correspond to separable CP-nets that are consistent with (x, y) . Given a set of examples E , let $\Phi_E = \{C_{x,y}^- \mid (x, y) \in E\}$. Then, given an interpretation M , a set of examples E is weakly consistent with N_M if and only if $M \models \Phi_E$. As a consequence, E is weakly separable if and only if Φ_E is satisfiable.

Example 25 *Consider three binary attributes A, B, C , and the set of examples*

$$E = \{abc \succ \bar{a}bc, \bar{a}\bar{b}c \succ ab\bar{c}, \bar{a}b\bar{c} \succ a\bar{b}\bar{c}, \bar{a}bc \succ \bar{a}\bar{b}\bar{c}\}$$

Φ_E has a unique model, corresponding to the separable CP-net $N = \{a \succ \bar{a}, b \succ \bar{b}, c \succ \bar{c}\}$. Therefore, N is the unique separable CP-net weakly consistent with N , and E is weakly separable.

As in section 5.2, a similar translation can be used with non-binary variables. And algorithms for solving MAXSAT can be used to search for a CP-net that is weakly consistent with as many examples as possible.

Proposition 26 *Deciding whether a set of examples over (binary or non-binary) attributes is weakly separable is NP-complete.*

Now, let us turn to the notion of strong compatibility. Characterizing such a property is less easy. Indeed, the difference between weak and strong compatibility is that while in weak compatibility we look for a separable CP-net which is consistent with each individual example in E , in strong compatibility we look for a separable CP-net which is consistent with the whole set of examples E .

Example 25, continued E is not strongly consistent with N , because $E \cup \succ_N$ has the following cycle:

$$\bar{a}\bar{b}c \succ_N \bar{a}\bar{b}c \succ_E ab\bar{c} \succ_N \bar{a}\bar{b}c \succ_E \bar{a}\bar{b}c$$

Since E is not strongly consistent with any other separable CP-net than N (because N is the unique separable CP-net with which E is weakly compatible), E is not strongly separable².

Note that all alternatives of the cycle on the example above appear in E . More generally, if we denote by $\mathcal{O}(E)$ the set of outcomes that appear in E , it can be proved that E is strongly consistent with N if and only if the restriction of $\succ_N \cup E$ to $\mathcal{O}(E)$ is acyclic. Since this restriction has at most $2|E|$ vertices, checking if it possesses a cycle can be done in polynomial time. Thus, checking whether E is strongly consistent with a given CP-net N is in \mathbf{P} , and we have the following result:

Proposition 27 [24] *Checking whether E is strongly separable is NP-complete.*

Note that although weak and strong separability have the same complexity, weak separability enjoys the nice property that there is a simple solution-preserving translation into SAT (the models of Φ_E correspond bijectively to the CP-nets that are weakly consistent with E), which allows weak separability to be computed in practice using algorithms for SAT³. In order to compute a separable CP-net strongly consistent with E , we can generate structures N weakly consistent with E , and test for acyclicity of $\succ_N \cup E$ using graph algorithms.

6 Active learning of CP-nets

In this section, we investigate the learnability issues of CP-nets in the paradigm of active learning. Recall that in the standard PAC learning model, examples are drawn at random according to an unknown but fixed distribution. This model of learning is merely passive in the sense that the learner has no control over the selection of examples. One can increase the flexibility of this model by allowing the learner to ask about particular examples, that is, the learner makes *membership queries* [1]. This capability appears to increase the power of polynomial-time learning algorithms. For instance, it is known that propositional Horn formulas are PAC-learnable with membership queries [3], but the results of [20] show that without membership queries, Horn formulas are no easier to learn than general CNF or DNF formulas.

In the setting of active preference learning, we assume that the user has in mind a target preference structure \succ , but doesn't know how to represent this structure into a CP-net. However, the user is disposed to help the learner by answering membership queries of the form “does \mathbf{x} dominates \mathbf{y} ”, where \mathbf{x} and \mathbf{y} are outcomes chosen by the learner. A membership query for a target concept \succ is a map MQ that takes as input a pair of outcomes (\mathbf{x}, \mathbf{y}) and returns as output `yes` if $\mathbf{x} \succ \mathbf{y}$, and `no` if $\mathbf{x} \not\succ \mathbf{y}$.

From a practical perspective, one must take into account the fact that outcomes are typically *not* comparable with an equivalent cost. As observed in [18], users can meaningfully compare outcomes if they differ only on very few attributes. To this end,

²Note that E is both weakly separable and does not contain any cycles as it was the case for Example 8, yet is not strongly separable.

³Such a translation exists for strong separability (which we do not give here), but unfortunately, the set of clauses generated uses $\mathcal{O}(n^2)$ variables (where n is the set of examples), which limits its practical applicability.

we define the *width* of $\text{MQ}(\mathbf{x}, \mathbf{y})$ is defined to be the number of variables on which \mathbf{x} and \mathbf{y} differ. A membership query of width 1 is called a *swap membership query*.

Based on these considerations, a minimal requirement behind active learning is to ask as few membership queries as possible. An additional desiderata for minimizing the cognitive effort spent by the user in answering preference queries is restrict on swap membership queries.

Definition 28 (PAC learning with membership queries) *A concept class $\mathcal{C}_{\mathcal{N}}$ is PAC learnable with swap membership queries by an instance class \mathcal{E} if there is a polynomial time learning algorithm A and two polynomials $p(\cdot, \cdot, \cdot)$ and $q(\cdot)$ such that for any target concept \succ in $\mathcal{C}_{\mathcal{N}}$, any probability distribution \mathcal{D} over \mathcal{I} , and any parameters $\delta, \epsilon \in (0, 1)$, after receiving $p(n, \frac{1}{\epsilon}, \frac{1}{\delta})$ random examples of \succ drawn independently according to \mathcal{D} , and asking $q(n)$ swap membership queries, then with probability at least $1 - \delta$, A returns a hypothesis $\hat{N} \in \mathcal{N}$ with $\text{error}(\hat{N}) \leq \epsilon$. The smallest such polynomial q is called the query complexity of the learning algorithm A .*

6.1 Active Learning with swap examples

We now investigate the problem of active preference learning, where the target concept can be represented by an *acyclic* CP-net, and the questions are restricted to swap membership queries.

In this setting, we can build an online algorithm for learning actively acyclic CP-nets. Recall that online learning proceeds into trials. Initially the learner chooses an hypothesis \hat{N} . During each trial, the learner first receives an example (\mathbf{x}, \mathbf{y}) , next predicts the label “+” or “-” of this instance according to its current hypothesis, and then receives the correct label from the user. If the prediction was incorrect then the learner is charged one mistake.

The basic idea underlying our online learning algorithm is to start from the empty CP-net $\hat{N} = \emptyset$ and, during each trial, iteratively revise \hat{N} by maintaining two invariants. The first invariant is that each rule (or entry) in the learner’s hypothesis \hat{N} is subsumed by a rule in the minimal representation of the target CP-net N . In other words, for each rule $\hat{u} : x \succ \bar{x}$ in \hat{N} , there is a rule $u : x \succ \bar{x}$ in the minimal representation of N , with $\hat{u} \subseteq u$. The second invariant is that each such rule $r = \hat{u} : x \succ \bar{x}$ in \hat{N} is *supported* by an instance $(\mathbf{x}_r, \mathbf{y}_r)$ of N , that is, $\mathbf{x}_r \succ_N \mathbf{y}_r$, \mathbf{x}_r and \mathbf{y}_r satisfy \hat{u} , \mathbf{x}_r satisfies x , and \mathbf{y}_r satisfies \bar{x} .

Technically, the algorithm proceeds as follows. On seeing an example (\mathbf{x}, \mathbf{y}) , if the learner predicts this instance as negative, while it is positive, then it expands its CP-net with a new rule $u : x \succ \bar{x}$, where the support (\mathbf{x}, \mathbf{y}) is stored. Here, x is the literal in \mathbf{x} whose value differs from \mathbf{y} , and u is the projection of the outcome \mathbf{x} onto the current parent set of the variable X in \hat{N} . Dually, if the learner predicts (\mathbf{x}, \mathbf{y}) as positive, while it is negative, then it expands the condition u of the misclassifying rule $u : x \succ \bar{x}$ with a new parent. Using the support $(\mathbf{x}_r, \mathbf{y}_r)$ of this rule, a new parent can be found by asking at most $n - 1$ membership queries. To this end, we simply need to incrementally transform \mathbf{x} into \mathbf{x}_r by iteratively flipping those literals that differ from \mathbf{x} and \mathbf{x}_r , until we find the first literal x_j for which the label of $(\mathbf{x}_j[x], \mathbf{x}_j[\bar{x}])$ is positive; this literal is kept as a new parent of X . In fact, only $\log_2 n$ membership queries are needed to find this parent, by performing a binary search over this sequence of transformations. A detailed implementation of this algorithm is given in [22].

By using a well-known conversion from online learning to PAC-learning [25], we derive the following result.

Proposition 29 *Binary-valued acyclic (and possibly incomplete) CP-nets are PAC-learnable with swap membership queries, over the instance class of swaps. There is an online learning algorithm A for this class, such that for any target concept \succ of description size s , the algorithm makes at most s mistakes and uses at most $s \log_2 n$ membership queries.*

6.2 Active Learning with Unrestricted Examples

When the instances supplied to the learner are unrestricted, even predicting their label is a difficult task, because dominance testing is NP-hard for acyclic CP-nets. As observed in the previous section, an important class of concepts for which dominance testing can be accomplished in polynomial time is the class of *tree CP-nets*. In this section, we briefly discuss an online algorithm for learning tree CP-nets.

The algorithm can be specified as follows. Initially, the learner starts from the hypothesis $\hat{N} = \emptyset$ and iteratively expands \hat{N} until it finds the target representation N . An invariant of the algorithm is that \hat{N} is always included in N so the learner can only make mistakes on positive examples (\mathbf{x}, \mathbf{y}) . In such cases, the algorithm considers in turn each variable on which \mathbf{x} and \mathbf{y} differ, and builds its CP-table and that of all its ascendants in the tree. It stops whenever such a variable already has a CP-table in its current hypothesis. Again, to find a parent for each candidate variable, the algorithm can use a binary search strategy.

Proposition 30 *Binary-valued Tree (possibly incomplete) CP-nets are PAC-learnable with swap membership queries, over arbitrary examples: there is an online learning algorithm A for this class, such that for any target concept \succ with k nodes, the algorithm makes at most k mistakes and uses $O(k \log_2 n)$ membership queries.*

We conclude this section by emphasizing that some classes of CP-nets are *teachable*, that is, learnable by asking a polynomial number of membership queries, without the need of observing any sample. Thus after making those queries, the learner is guaranteed to correctly predict any instance.

We focus on the class of acyclic CP-nets whose graph has indegree at most k . The learner proceeds by using a levelwise generate-and-test procedure and membership queries for uncovering the set of parents of each variable. Clearly, this approach is acceptable only for small bounded degrees, such as tree CP-nets.

Proposition 31 *The class of concepts representable by a binary-valued acyclic CP-net whose graph has degree at most k is teachable: there is an algorithm which outputs such a CP-net using $O(kn^{k+1}2^{k-1})$ membership queries and no other queries or examples, where n is the number of variables and k is the degree of the target concept.*

7 Conclusion and open problems

In this paper we addressed many issues related to learning CP-nets. We argued that a first important problem is whether the CP-net that we aim at learning is such that the user's preference relation coincides with its induced preference relation, or is an approximation by below of the user's complete preference relation. Then we gave a few theoretical results on the learnability of CP-nets, and then we considered two different learning frameworks: passive learning (from a set of examples), and active learning (by queries).

We hope that our preliminary results in the learnability issue of CP-net open the door to new theoretical results and practical learning algorithms. First of all, we do not have a general method (other than brute-force search) for computing a CP-net which is weakly or strongly consistent with a set of examples in the nonseparable case, nor do we have algorithms for outputting a CP-net realizing an optimal trade-off between simplicity and accuracy.

We already emphasized the lack of expressivity of CP-nets. Although CP-nets are a representation language well-suited to expressing preferential (in)dependencies, they do not allow, for instance, to express relative important statements between variables, as lexicographic orders do. We may desire to learn preferences that combine both aspects (preferential dependencies and relative importance); for this, it is attack preference learning with more expressive languages such as TCP-nets [7] or (even more general) conditional preference theories [32].

References

- [1] D. Angluin. Queries and concept learning. *Machine Learning*, 2(4):319–342, 1988.
- [2] D. Angluin. Negative results for equivalence queries. *Machine Learning*, 5:121–150, 1990.
- [3] D. Angluin, M. Frazier, and L. Pitt. Learning conjunctions of Horn clauses. *Machine Learning*, 9:147–164, 1992.
- [4] M. Anthony and N. Biggs. *Computational Learning Theory*. Cambridge Univ. Press, 1992.
- [5] F. Athienitou and Y. Dimopoulos. Learning CP-networks: a preliminary investigation. In *3rd Multidisciplinary Workshop on Advances in Preference Handling (PREF)*, 2007.
- [6] C. Boutilier, R. Brafman, C. Domshlak, H. Hoos, and D. Poole. CP-nets: a tool for representing and reasoning with conditional ceteris paribus preference statements. *Journal of Artificial Intelligence Research*, 21:135–191, 2004.
- [7] R. Brafman, C. Domshlak, and S. Shimony. On graphical modeling of preference and importance. *Journal of Artificial Intelligence Research*, 25:389–424, 2006.
- [8] D. Braziunas and C. Boutilier. Local utility elicitation in GAI models. In *Proceedings of the 21st Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 42–49, 2005.
- [9] C.J.C. Burges, T. Shaked, E. Renshaw, A. Lazier, M. Deeds, N. Hamilton, and G. Hullender. Learning to rank using gradient descent. In *Proceedings of the 22nd International Conference on Machine Learning (ICML)*, 2005.
- [10] U. Chajewska, L. Getoor, J. Norman, and Y. Shahar. Utility elicitation as a classification problem. In *Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 79–88, 1998.
- [11] U. Chajewska, D. Koller, and R. Parr. Making rational decisions using adaptive utility elicitation. In *Proceedings of the 17th National Conference on Artificial Intelligence (AAAI)*, pages 363–369, 2000.
- [12] L. Chen and P. Pu. Survey of preference elicitation methods. Technical Report 200467, EPFL, 2004.
- [13] Y. Chevaleyre. A short note on passive learning of cp-nets. Rapport de recherche, Lamsade, mars 2009.
- [14] J. Dombi, C. Imreh, and N. Vincze. Learning lexicographic orders. *European Journal of Operational Research*, 183:748–756, 2007.
- [15] C. Domshlak and T. Joachims. Efficient and non-parametric reasoning over user preferences. *User Modeling and User-Adapted Interaction (UMUAI)*, 17(1-2):41–69, 2007.

- [16] J. Doyle, Y. Shoham, and M. Wellman. A logic of relative desire (preliminary report). In *Proceedings of the 6th International Symposium on Methodologies for Intelligent Systems (ISMIS)*, pages 16–31. Springer, 1991.
- [17] C. Gonzales and P. Perny. GAI networks for utility elicitation. In *Principles of Knowledge Representation and Reasoning: Proceedings of the 9th International Conference (KR)*, pages 224–234, 2004.
- [18] P. Green and V. Srinivasan. Conjoint analysis in consumer research: Issues and outlook. *Journal of Consumer Research*, 5(2):103–123, 1978.
- [19] V. Ha and P. Haddawy. Problem-focused incremental elicitation of multi-attribute utility models. In *Proceedings of the 13th Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 215–222, 1997.
- [20] M. J. Kearns, M. Li, L. Pitt, and L. G. Valiant. On the learnability of boolean formulae. In *Proceedings of the Nineteenth Annual ACM Symposium on Theory of Computing*, pages 285–295, 1987.
- [21] R. Keeney and H. Raiffa. *Decision with Multiple Objectives: Preferences and Value Trade-offs*. Wiley, 1976.
- [22] F. Koriche and B. Zanuttini. Learning conditional preference networks with queries, 2009. Submitted for publication.
- [23] J. Lang and J. Mengin. Learning preference relations over combinatorial domains. In *Workshop on Preference Learning at ECML’08*, 2008.
- [24] J. Lang and J. Mengin. The complexity of learning separable ceteris paribus preferences. Rapport de recherche RR-2009-3-FR, IRIT, Université Paul Sabatier, Toulouse, mars 2009.
- [25] N. Littlestone. From on-line to batch learning. In *Proceedings of the Second Annual Workshop on Computational Learning Theory*, pages 269–284. Morgan Kaufmann, 1989.
- [26] M. Sachdev. On learning of ceteris paribus preference theories. Master’s thesis, Graduate Faculty of North Carolina State University, 2007.
- [27] T. Sandholm and C. Boutilier. *Preference Elicitation in Combinatorial Auctions*, chapter 10. In *Combinatorial Auctions*, Cramton, Shoham, and Steinberg Ed., MIT Press, 2006.
- [28] M. Schmitt and L. Martignon. On the complexity of learning lexicographic strategies. *Journal of Machine Learning Research*, 7:55–83, 2006.
- [29] Leslie G. Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.
- [30] P. Viappiani, B. Faltings, and P. Pu. Evaluating preference-based search tools: a tale for two approaches. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI)*, pages 205–210, 2006.
- [31] P. Viappiani, B. Faltings, and P. Pu. Preference-based search using example-critiquing with suggestions. *Journal of Artificial Intelligence Research*, 27:465–503, 2006.
- [32] N. Wilson. Extending CP-nets with stronger conditional preference statements. In *Proceedings of the 19th National Conference on Artificial Intelligence (AAAI)*, pages 735–741, 2004.
- [33] F. Yaman, Th. Walsh, M. Littman, and M. desJardins. Democratic approximation of lexicographic preference models. In *Proceedings of the 35th International Conference in Machine Learning (ICML)*, pages 1200–1207, 2008.