

Learning to Learn in Interactive Constraint Acquisition

Dimos Tsouros, Senne Berden, Tias Guns

KU Leuven, Belgium

dimos.tsouros@kuleuven.be, senne.berden@kuleuven.be, tias.guns@kuleuven.be

Abstract

Constraint Programming (CP) has been successfully used to model and solve complex combinatorial problems. However, modeling is often not trivial and requires expertise, which is a bottleneck to wider adoption. In Constraint Acquisition (CA), the goal is to assist the user by automatically *learning* the model. In (inter)active CA, this is done by interactively posting queries to the user, e.g., asking whether a partial solution satisfies their (unspecified) constraints or not. While interactive CA methods learn the constraints, the learning is related to symbolic concept learning, as the goal is to learn an exact representation. However, a large number of queries is required to learn the model, which is a major limitation. In this paper, we aim to alleviate this limitation by tightening the connection of CA and Machine Learning (ML), by, for the first time in interactive CA, exploiting statistical ML methods. We propose to use probabilistic classification models to guide interactive CA to generate more promising queries. We discuss how to train classifiers to predict whether a candidate expression from the bias is a constraint of the problem or not, using both relation-based and scope-based features. We then show how the predictions can be used in all layers of interactive CA: the query generation, the scope finding, and the lowest-level constraint finding. We experimentally evaluate our proposed methods using different classifiers and show that our methods greatly outperform the state of the art, decreasing the number of queries needed to converge by up to 72%.

Introduction

Constraint Programming (CP) is considered one of the foremost paradigms for solving combinatorial problems in artificial intelligence. In CP, the user declaratively states the constraints over a set of decision variables, defining the feasible solutions to their problem, and then a solver is used to solve it. Although CP has many successful applications on combinatorial problems from various domains, the modeling process is not always trivial and is limiting non-experts from using CP on complex problems. This is considered a major bottleneck for the wider adoption of CP (Freuder and O’Sullivan 2014; Freuder 2018).

Motivated by the need to overcome this obstacle, assisting the user in modeling is regarded as an important re-

search topic (Kolb 2016; De Raedt, Passerini, and Teso 2018; Freuder 2018; Lombardi and Milano 2018). In *Constraint Acquisition (CA)*, which is an area where CP meets Machine Learning (ML), the model of a constraint problem is learned from a set of examples (i.e., assignments to the variables) of solutions, and possibly non-solutions.

In *passive CA*, a set of pre-existing examples is given to the system, and using these examples a set of constraints is returned (Bessiere et al. 2004, 2005; Lallouet et al. 2010; Beldiceanu and Simonis 2012; Bessiere et al. 2017; Kumar, Kolb, and Guns 2022; Berden et al. 2022).

On the other hand, *active* or *interactive* acquisition systems interact with the user to learn a target set of constraints, which represent the problem the user has in mind (Freuder and Wallace 1998; Bessiere et al. 2007, 2017). In the early days, most methods only made use of membership queries (is this a solution or not?) (Angluin 1988; Bessiere et al. 2007), while a more recent family of algorithms also makes use of *partial membership queries* (Arcangoli, Bessiere, and Lazaar 2016; Bessiere et al. 2013; Lazaar 2021; Tsouros and Stergiou 2020, 2021; Tsouros, Stergiou, and Bessiere 2019, 2020; Tsouros, Stergiou, and Sarigianidis 2018). Such (partial) queries ask the user to classify (partial) assignments to the variables as (non-)solution. Recently, a way to guide the top-level query generation was introduced (Tsouros, Berden, and Guns 2023), based on counting-based probabilistic estimates of whether candidate expressions are constraints of the problem or not. Using this method, the number of queries required to converge decreased significantly.

Despite the recent advancements in active CA, there are still significant drawbacks to overcome. One of the most important drawbacks is the large number of queries still required in order to find all constraints. We believe this is due to the search-based learning being mostly *uninformed*. During learning it is not aware of patterns that may appear in the constraints acquired so far, which can guide the rest of the process. An exception is the ANALYZE&LEARN (Tsouros, Stergiou, and Bessiere 2019) function, which tries to detect potential cliques in the constraint network learned.

In this work, we focus on this major limitation and contribute the following elements to alleviate it:

- We show how probabilistic classification can be used to predict whether a candidate expression is a constraint of

the problem or not, based on the constraints learned so far and the ones removed from the candidate set at any point during the acquisition process. We use both relation-based and scope-based features to train ML models that are then exploited to guide interactive CA systems.

- Previously it was shown that top-level query generation can be guided with (counting-based) probabilistic estimates. We show how such guidance can be extended to all layers of interactive CA where queries are asked.
- We make a comprehensive experimental evaluation of our proposed methods, showing the effect of different classifiers, focusing on the number of queries vs. runtime for the ML-guided systems. We also show the effect of guiding all layers where queries are posted to the user.

Background

Let us first give some basic notions regarding constraint satisfaction problems.

A *constraint satisfaction problem (CSP)* is a triple $P = (X, D, C)$, consisting of:

- a set of n variables $X = \{x_1, x_2, \dots, x_n\}$, representing the entities of the problem,
- a set of n domains $D = \{D_1, D_2, \dots, D_n\}$, where $D_i \subset \mathbb{Z}$ is the finite set of values for x_i ,
- a constraint set (also called constraint network) $C = \{c_1, c_2, \dots, c_t\}$.

A *constraint* c is a pair $(rel(c), var(c))$, where $var(c) \subseteq X$ is the *scope* of the constraint, and $rel(c)$ is a relation over the domains of the variables in $var(c)$, that (implicitly) specifies which of their value assignments are allowed. $|var(c)|$ is called the *arity* of the constraint. The constraint set $C[Y]$, where $Y \subseteq X$, denotes the set of constraints from C whose scope is a subset of Y . The set of solutions of a constraint set C is denoted by $sol(C)$.

An *example* e_Y is an assignment on a set of variables $Y \subseteq X$. e_Y is rejected by a constraint c iff $var(c) \subseteq Y$ and the projection $e_{var(c)}$ of e_Y on the variables in the scope $var(c)$ of the constraint is not in $rel(c)$. A complete assignment e that is accepted by all the constraints in C is a *solution* to C , i.e., $e \in sol(C)$. An assignment e_Y is called a *partial solution* iff $e_Y \in sol(C[Y])$. $\kappa_C(e_Y)$ represents the subset of constraints from a constraint set $C[Y]$ that reject e_Y .

In CA, the pair (X, D) is called the *vocabulary* of the problem at hand and is common knowledge shared by the user and the system. Besides the vocabulary, the learner is also given a *language* Γ consisting of *fixed arity* constraints. Using the vocabulary (X, D) and the constraint language Γ , the system generates the *constraint bias* B , which is the set of all expressions that are candidate constraints for the problem. The (unknown) target constraint set C_T is a constraint set such that for every example e it holds that $e \in sol(C_T)$ iff e is a solution to the problem the user has in mind. The goal of CA is to learn a constraint set C_L that is equivalent to the target constraint set C_T .

Algorithm 1: Generic Constraint Acquisition Template

Input: X, D, B, C_{in} (X : the set of variables, D : the set of domains, B : the bias, C_{in} : an optional set of known constraints)

Output: C_L : the learned constraint network

```

1:  $C_L \leftarrow C_{in}$ 
2: while True do
3:    $e \leftarrow \text{QGEN}(C_L, B)$ 
4:   if  $e = \text{nil}$  then return  $C_L$  ▷ converged
5:   if  $\text{ASK}(e) = \text{True}$  then
6:      $B \leftarrow B \setminus \kappa_B(e)$ 
7:   else
8:      $(B, S) \leftarrow \text{FINDSCOPE}(e, B)$ 
9:      $(B, C_L) \leftarrow \text{FINDC}(S, C_L, B)$ 

```

Interactive Constraint Acquisition

In interactive CA, the system interacts with the user while learning the constraints. The classification question $\text{ASK}(e_X)$, asking the user if a complete assignment e_X is a solution to the problem that the user has in mind, is called a *membership query* (Angluin 1988). A *partial query* $\text{ASK}(e_Y)$, with $Y \subset X$, asks the user to determine if e_Y , which is an assignment in D^Y , is a *partial solution* or not, i.e., if $e_Y \in sol(C_T[Y])$. A (complete or partial) query $\text{ASK}(e_Y)$ is called *irredundant* iff the answer is not implied by information already available. That is, it is irredundant iff e_Y is rejected by at least one constraint from the bias B , and not rejected by the network C_L learned thus far.

Algorithm 1 presents the generic process followed in interactive CA through partial queries. The learned set C_L is first initialized either to the empty set or to a set of constraints given by the user that is known to be true (line 1). Then the main loop of the acquisition process begins, where first the system generates an irredundant example (line 3) and posts it as a query to the user (line 5). If the query is classified as positive, then the candidate expressions from B that violate it are removed (line 6). If the example is classified as negative, then the system tries to find one (or more) constraint(s) from C_T that violates it. This is done in two steps. First, the scope of one or more violated constraints is found, by asking queries and possibly shrinking the bias along the way (line 8). Then, the relations of the constraints in this scope(s) are found, again by asking queries and possibly shrinking the bias (line 9).

This process continues until the system converges. The acquisition process has *converged* on the learned network $C_L \subseteq B$ iff C_L agrees with the set of all labeled examples E , and for every other network $C \subseteq B$ that agrees with E , it holds that $sol(C) = sol(C_L)$. This is proved if no example could be generated at line 3, as in this case, all constraints in B are redundant.

Notice that, interactive CA systems consist of three components where (increasingly simpler) queries are asked to the user: (1) Top-level query generation (line 3), (2) Finding the scope(s) of violated constraints (line 8), (3) Finding the relations of constraints in the scopes found (line 9).

State-of-the-art algorithms like QuAcq (Bessiere et al.

2013, 2023), MQuAcq (Tsouros and Stergiou 2020) and MQuAcq-2 (Tsouros, Stergiou, and Bessiere 2019) follow this template. Recently, a *meta*-algorithm named GrowAcq (Tsouros, Berden, and Guns 2023) was introduced, in order to handle large biases and to reduce the number of queries. The key idea is to call a CA algorithm on an increasingly large subset of the variables $Y \subseteq X$, initially using a small number of variables, each time using a (growing) subset of the potentially huge bias.

Guiding Query Generation

When using GROWACQ, only a subset of B needs to be considered at a time, and query generation is often fast, leaving sufficient room for using optimization to find a *good* query in top-level query generation (line 3 of Algorithm 1). Query generation is formulated as a CSP with variables Y and constraints $C_L[Y] \wedge \bigvee_{c_i \in B[Y]} \neg c_i$, in order to find an example e_Y . Hence, when the set of candidates B is reduced, query generation is simplified.

As a result of this speed-up, in (Tsouros, Berden, and Guns 2023) a method to guide the top-level query generation was proposed. This method introduces an objective function that uses the prediction of a model $\mathcal{M}(c)$:

$$e = \arg \max_e \sum_{c \in B} \llbracket e \notin \text{sol}(\{c\}) \rrbracket \cdot (1 - |\Gamma| \cdot \llbracket \mathcal{M}(c) \rrbracket) \quad (1)$$

where $\llbracket \cdot \rrbracket$ is the Iverson bracket converting *True/False* to 1/0.

The objective function’s aims are twofold. First, it wants queries that lead to a positive answer to violate many constraints in the bias, shrinking it faster. Second, it wants constraints that lead to a negative answer to violate a small number of constraints from the bias, so that the actual constraint leading to the negative answer can be found more easily. For more exposition on how this objective function achieves these aims, we refer the reader to (Tsouros, Berden, and Guns 2023).

Model \mathcal{M} tries to determine for every constraint c whether violating or satisfying c would lead to the least amount of queries later on in the algorithm, based on a probabilistic estimate $P(c \in C_T)$ of how likely a constraint is to belong to the target set of constraints of the problem

$$\mathcal{M}(c) = \left(\frac{1}{P(c \in C_T)} \leq \log(|Y|) \right) \quad (2)$$

Using Probabilistic Classification To Guide Interactive CA

The model \mathcal{M} leverages a probabilistic estimate of the likelihood of a given candidate constraint belonging to the problem at hand. In (Tsouros, Berden, and Guns 2023), a simple counting-based method was utilized that only uses information about the relation of the constraints. That is, the number of times a constraint with relation $rel(c)$ has been added to C_L is counted, and then divided by the total number of times that such a constraint has been removed from B .

While this technique provides basic guidance, we propose to use more advanced prediction techniques. Specifically, we

propose to use statistical ML techniques, exploiting probabilistic classification in order to calculate $P(c \in C_T)$.

In order to use probabilistic classification in this context, we need to build a dataset to learn from. We formally define a dataset \mathcal{D} as a collection of N instances, each instance corresponding to a constraint. Each instance is a tuple (\mathbf{x}_i, y_i) , $i \in 1, 2, \dots, N$, with \mathbf{x}_i being a vector of features of constraint c_i , and y_i being a (Boolean) label that denotes whether $c_i \in C_T$.

ID	Name	Type	Description
1	Relation	String	Constraint relation
2	Arity	Int	Constraint arity
3	Has_constant	Bool	If a constant value is present
4	Constant	Int	The constant value
5	Var_name_same	Bool	If all variables share the same name
6	Var_Ndims_same	Bool	If the number of dimensions of all variables is the same
7	Var_Ndims_max	Int	The maximum number of dimensions among variables
8	Var_Ndims_min	Int	The minimum number of dimensions among variables
9	Var_dim _{<i>i</i>} _has	Bool	If dimension <i>i</i> is present for all variables
10	Var_dim _{<i>i</i>} _same	Bool	If dimension <i>i</i> of all variables is the same
11	Var_dim _{<i>i</i>} _max	Int	Maximum dimension <i>i</i> value among variables
12	Var_dim _{<i>i</i>} _min	Int	Minimum dimension <i>i</i> value among variables
13	Var_dim _{<i>i</i>} _avg	Float	Average dimension <i>i</i> value among variables
14	Var_dim _{<i>i</i>} _spread	Float	Spread of dimension <i>i</i> values among variables

Table 1: Features for each constraint

To be able to use constraints as instances in our dataset, we need to have a feature representation of constraints. In this paper, for the feature representation, we use both relation-based and scope-based information, exploiting the information we have for the constraint’s relation, the variables of its scope, their indices, name, etc. The features we use are shown in Table 1. Note that variables can be given to the CA system in the form of a matrix or tensor. For example, a natural way to structure the variables representing the cell assignments Sudoku is in a 9x9 2-dimensional matrix. When variables are given in such a form, we represent in the features information about the indices of the variables occurring in the constraints, in each dimension of the tensor they were given in. This allows the system to detect patterns like all variables occurring in the same row or column, not being spread out in some dimension, etc., which are common

patterns in CP problems.

The dataset \mathcal{D} is grown incrementally throughout the CA process, as gradually more information is obtained about constraints from the initial bias. More concretely, whenever a constraint is removed from B (because they were verified to not be part of C_T), it is added to D with a negative label. On the other hand, whenever a constraint gets added to C_L , it also gets added to D with a positive label.

Since dataset \mathcal{D} grows throughout the CA process, the probabilistic classifier should be updated regularly. How often this update should be performed is an important question, as this may affect the waiting time for the user when interacting with the CA system. In this paper, we retrain the classifier on the current dataset \mathcal{D} right before every top-level query generation (line 3 of Algorithm 1), exploiting all the collected information each time. Preliminary experiments showed that this yields the best results.

Guiding All Layers of Interactive CA

As (Tsouros, Berden, and Guns 2023) showed, guiding top-level query generation can reduce the number of queries significantly and improve CA systems. However, CA systems ask queries to the user also in the FINDSCOPE (line 8 of Algorithm 1) and FINDC (line 9 of Algorithm 1) components, which respectively try to find the scope of one or more violated constraints, and then all constraints on that scope. While guiding the generation of top-level queries delivers significant advantages, neglecting guidance within these two layers is a missed opportunity.

In the rest of this section, we show how to use the same logic for guiding query generation in the remaining two layers of CA systems.

Guiding FindScope

The functions used in the literature (Bessiere et al. 2013; Tsouros and Stergiou 2020; Bessiere et al. 2023) to find the scope of violated constraints after a negative answer from the user (line 8 of Algorithm 1) work in a similar way. We will use FINDSCOPE from (Bessiere et al. 2023) (shown in Algorithm 2) to demonstrate our method, but the same logic applies to all existing in the literature. FINDSCOPE methods recursively map the problem of finding a constraint to a simpler problem by removing blocks of variable assignments from the original query (the one asked in line 3 of Algorithm 1, to which the user answered “no”) and asking partial queries to the user. The removed block must contain at least one variable while not including all the present variables, in order to lead to an irredundant query. If after the removal of some variables, the answer of the user changes to “yes”, then the removed block contains at least one variable from the scope of a violated constraint. When this happens, FINDSCOPE focuses on refining this block, adding some variable assignments back to the query. When, after repeating this procedure, the size of the considered block becomes 1 (i.e., the block contains a single variable), this variable is found to be in the scope of a violated constraint we seek (line 5 of Algorithm 2).

In practice, the problem that must be solved in each step is to find a set of variables $Y_1 \subset Y$, splitting the previously

Algorithm 2: FINDSCOPE

Input: e, R, Y, B (e : the example, R, Y : sets of variables, B : the bias)
Output: $B, Scope$ (B : the updated bias, $Scopes$: a set of variables, the scope of a constraint in C_T)

- 1: **function** FINDSCOPE(e, R, Y, B)
- 2: **if** $\kappa_B(e_R) \neq \emptyset$ **then**
- 3: **if** $ASK(e_R) = \text{“yes”}$ **then** $B \leftarrow B \setminus \kappa_B(e_R)$;
- 4: **else return** (B, \emptyset);
- 5: **if** $|Y| = 1$ **then return** (B, Y);
- 6: split Y into $\langle Y_1, Y_2 \rangle$ such that $|Y_1| = \lceil |Y|/2 \rceil$;
- 7: **if** $\kappa_B(e_{R \cup Y_1}) = \kappa_B(e_{R \cup Y})$ **then** $S_1 \leftarrow \emptyset$
- 8: **else** (B, S_1) $\leftarrow FindScope(e, R \cup Y_1, Y_2, B)$;
- 9: **if** $\kappa_B(e_{R \cup S_1}) = \kappa_B(e_{R \cup Y})$ **then** $S_2 \leftarrow \emptyset$
- 10: **else** (B, S_2) $\leftarrow FindScope(e, R \cup S_1, Y_1, B)$;
- 11: **return** ($B, S_1 \cup S_2$);

considered set of variables Y into two parts (line 6). Then, Y_1 is used in the next query posted to the user while Y_2 is the removed set of variables that can be taken into account in the next queries. Depending on the answer of the user we can then update B and decide which part of the problem to focus on next. Existing FINDSCOPE functions naively choose $Y_1 \subset Y$, by splitting the set Y in half. The advantage of this approach is that a logarithmic number of steps is achieved. However, no information about the violated constraints from B is used, and no guidance is utilized. The set of variables to remove, or keep, in the assignment is usually chosen randomly (Bessiere et al. 2023).

The problem of finding a $Y_1 \subset Y$, so that the FINDSCOPE procedure is correct and will lead to finding the scope of a violated constraint, can be formulated as:

$$\text{find } Y_1 \text{ s.t. } \emptyset \subsetneq Y_1 \subsetneq Y \quad (3)$$

That is because, in each query, we get information either for Y_1 , if the answer of the user remains negative, or for Y_2 , if the answer of the user changes to positive. Thus we want both $Y_1, Y_2 \supsetneq \emptyset$.

This problem can be formulated as a CSP with boolean variables BV , with $|BV| = |Y|$, deciding whether a variable $x_i \in Y$ is included in Y_1 . The CSP contains the following constraint:

$$0 < \sum_{bv_i \in BV} \llbracket bv_i \rrbracket < |Y| \quad (4)$$

However, just choosing any (arbitrarily sized) subset of Y can result in many unneeded recursive calls and a large number of queries. Now that we have formally formulated this problem, we can modify the constraints and/or add an objective function in order to improve the performance of FINDSCOPE. In order to achieve the logarithmic complexity from (Bessiere et al. 2023), we can impose the following constraint in our CSP formulation:

$$\sum_{bv_i \in BV} \llbracket bv_i \rrbracket = \left\lfloor \frac{|Y|}{2} \right\rfloor \quad (5)$$

We propose to use this CSP formulation of the problem, and to integrate the objective function from Equation (1) to guide FINDSCOPE queries. Notice that, the queries asked in FINDSCOPE take into account the set R , which is the set $R \cup Y_1$ from the previous recursive call, where we split Y . In addition, we now are not generating assignments, but deciding which variable assignments from the existing example e to include in the next query $ASK(e_{R \cup Y_1})$. Thus, we propose to maximize the following objective function:

$$\sum_{c \in \kappa_B(e)} \llbracket \text{var}(c) \in R \cup Y_1 \rrbracket \cdot (1 - |\Gamma| \cdot \llbracket \mathcal{M}(c) \rrbracket) \quad (6)$$

We also slightly modify the constraint from Equation (5), as, when deciding which constraints to violate in the next query, the number of variables these constraints participate in could be lower than half (but still needs to be at least one, as in Equation (4)). As a result, the constraint becomes

$$0 < \sum_{bv_i \in BV} \llbracket bv_i \rrbracket \leq \left\lfloor \frac{|Y|}{2} \right\rfloor \quad (7)$$

Correctness We now prove that FINDSCOPE is still correct when our modification of line 6 is used, as long as the constraint from Equation (4) holds.

Proposition 1. *Given the assumption that C_T is representable by B , FINDSCOPE (with our modification at line 6) is correct.*

Proof. Soundness We will now prove that given an example e_Y , FINDSCOPE will return a set of variables S , such that there exists at least one violated constraint $c \in C_T$ s.t. $\forall x_i \in S \mid x_i \in \text{var}(c)$.

An invariant of *FindScope* is that the example e violates at least one constraint whose scope is a subset of $R \cup Y$ (i.e., $ASK(R \cup Y) = \text{"no"}$).

$$\kappa_{C_T}(e_{R \cup Y}) \supseteq \emptyset \quad (8)$$

That is because it is called only when the example e_Y is classified as non-solution by the user and the recursive calls at lines 8 and 10 are reached only if the conditions at lines 7 and 9 respectively are false.

In addition, *FindScope* reaches line 5 only in the case that e_R does not violate any constraint from C_T (i.e., $ASK(e_R) = \text{"yes"}$ at line 3).

$$\kappa_{C_T}(e_R) = \emptyset \quad (9)$$

In *FindScope* variables are returned (and added in S) only at line 5, in the case Y is a singleton.

$$(8), (9) \implies \exists Y' \subseteq Y \text{ s.t. } Y' \subseteq \text{var}(c) \mid c \in \kappa_{C_T}(e) \\ \xrightarrow[\text{line 5}]{|Y|=1} Y \subseteq \text{var}(c) \mid c \in \kappa_{C_T}(e) \quad (10)$$

Thus, for any $x_i \in S$ we know that $x_i \in \text{var}(c) \mid c \in C_T$.

Completeness We will now prove that given an example e_Y , the set of variables S returned by FINDSCOPE will be

the full scope of a constraint in C_T , i.e. there exists at least one constraint $c \in C_T$ for which $S = \text{var}(c)$.

FINDSCOPE in Algorithm 2 has been proven to be complete in (Bessiere et al. 2023). The key part in that is line 6, splitting Y into 2 parts. The requirement is that in no recursive call we end up with $Y = \emptyset$, so that it continues searching in different subsets of variables in each call. This means that in the recursive call of line 9, $Y_2 \neq \emptyset$ and in the recursive call of line 10, we must have $Y_1 \neq \emptyset$.

Due to the constraint imposed in Equation (4), we know that $Y_1 \supseteq \emptyset$ and also that

$$Y_1 \subsetneq Y \implies Y_1 \subsetneq Y_1 \cup Y_2 \implies Y_2 \neq \emptyset \quad (11)$$

Thus, this constraint guarantees that $Y_1, Y_2 \neq \emptyset$, meaning that FINDSCOPE is still complete. \square

Guiding FindC

After the system has located the scope of a violated constraint, it calls function FINDC (Bessiere et al. 2013, 2023) to find the relation of the violated constraint. To locate this constraint, FINDC asks partial queries to the user in the scope returned by FINDSCOPE. Alternative assignments are used for the variables in the scope given, to discriminate which of the candidate constraints with that scope is part of the target problem. In order to do so, FINDC functions currently use the following query generation step:

$$\text{find } e'_S \in \text{sol}(C_L[S] \wedge \emptyset \subsetneq \kappa_B(e'_S) \subsetneq \Delta), \quad (12)$$

with S being the scope found in the previous step and Δ the set of candidates for this scope, initially being equal to the set of violated constraints in the previous example $\kappa_B(e_S)$.

The objective function typically used in this step, in order to again achieve a logarithmic complexity in terms of the number of queries posted, is to try to half the number of violated candidates, minimizing a slack variable b such that

$$b = \left\lfloor \frac{|\Delta|}{2} \right\rfloor - \kappa_B(e'_S) \quad (13)$$

We propose to replace this objective function with one that guides the query generation in the same way as in Equations (1) and (6):

$$\sum_{c \in \Delta} \llbracket e_S \notin \text{sol}(\{c\}) \rrbracket \cdot (1 - |\Gamma| \cdot \llbracket \mathcal{M}(c) \rrbracket) \quad (14)$$

Experimental Evaluation

In this section, we perform an experimental evaluation of our proposed approaches, aiming to answer the following research questions:

- (Q1) How well can ML classifiers predict whether a candidate constraint is part of the target constraint network?
- (Q2) What is the effect of using probabilistic classification to guide query generation in CA?
- (Q3) What is the added benefit of also guiding the other layers of CA?

Benchmarks

We selected the benchmarks for our experiments to cover different cases, including some puzzle problems that are typically used as benchmarks to evaluate CA systems, some problems closer to real-world applications with a subset of them having a more regular structure, and one randomly generated. The latter was included to evaluate the performance of our system when it cannot learn anything.

More specifically, we used the following benchmarks for the experimental evaluation:

Random. We used a problem with 100 variables and domains of size 5. We generated a random target network with 495 binary constraints from the language $\Gamma = \{\geq, \leq, <, >, \neq, =\}$. The bias was initialized with 19,800 constraints, using the same language.

9x9 Sudoku The Sudoku puzzle is a $n^2 \times n^2$ grid ($n = 3$ for the case we used), which must be completed in such a way that all the rows, columns, and n^2 non-overlapping $n \times n$ squares contain distinct numbers. This gives a *vocabulary* having 81 variables and domains of size 9. The target constraint network consists of $810 \neq$ constraints. The bias was initialized with 12,960 binary constraints, using the language $\Gamma = \{\geq, \leq, <, >, \neq, =\}$.

Jigsaw Sudoku. The Jigsaw Sudoku is a variant of Sudoku in which the 3×3 boxes are replaced by irregular shapes. It consists of 81 variables with domains of size 9. The target network contains 811 binary \neq constraints on rows, columns, and shapes. The bias B was constructed using the language $\Gamma = \{\geq, \leq, <, >, \neq, =\}$ and contains 19,440 binary constraints.

Exam Timetabling There are ns semesters, each containing cps courses, and we want to schedule the exams of the courses in a period of d days. On each day, we have t timeslots and r rooms available for exams. The variables are the courses ($|X| = ns \cdot cps$), having as domains the timeslots they can be assigned on ($D_i = 1, \dots, r \cdot t \cdot d$). There are \neq constraints between each pair of exams. Also, two courses in the same semester cannot be examined on the same day, which is expressed by the constraints $[x_i/spd] \neq [x_j/spd], \forall i, j$ in the same program. We used an instance with $ns = 8$, $cps = 6$, $d = 10$ and $r = 3$. This resulted in a model with 48 variables and domains of size 90. C_T consists of 1,128 constraints. The language given is $\Gamma = \{\geq, \leq, <, >, \neq, =, [x_i/spd] \neq [x_j/spd]\}$, creating a bias of size 7,896.

Nurse rostering There are n nurses, s shifts per day, ns nurses per shift, and d days. The goal is to create a schedule, assigning a nurse to all existing shifts. The variables are the shifts, and there are a total of $d \cdot s \cdot ns$ shifts. The variables are modeled in a 3D matrix. The domains of the variables are the nurses ($D^{x_i} = \{1, \dots, n\}$). Each shift in a day must be assigned to a different nurse and the last shift of a day must be assigned to a different nurse than the first shift of the next day. In the instance used in the experiments, we have $d = 7$, $s = 3$ and $ns = 5$. The available nurses are $n = 18$. This results in $|X| = 105$ with domains $\{1, \dots, 18\}$. C_T consists of 885 \neq constraints. The bias was built using the language $\Gamma = \{\geq, \leq, <, >, \neq, =\}$, resulting in $|B| = 32,760$.

Experimental Settings

All the experiments were conducted on a system carrying an Intel(R) Core(TM) i7-2600 CPU, 3.40GHz clock speed, with 16 GB of RAM. The guiding techniques are integrated within GROWACQ, utilizing MQUACQ-2 as the underlying algorithm. We compare our approach with the counting method from (Tsouros, Berden, and Guns 2023), (“count”), as well as with GROWACQ without guiding (“base”). In the latter, the objective in query generation is simply to maximize the number of violated candidate constraints.

We use the following classifiers: Random Forests (RF), Gaussian Naive Bayes (GNB), Multi-layer Perceptron (MLP), and Support Vector Machines (SVM). We used RF and GNB in their default settings, while we tuned the most important hyperparameters for MLP and SVM. For tuning, we used the final dataset for all benchmarks, having labeled all candidate constraints. A grid search, coupled with 10-fold cross-validation, was conducted, using balanced accuracy as the metric to address class imbalance. Hyperparameter combinations surpassing a 10-second training time were omitted to ensure relevance in interactive scenarios.

All methods and benchmarks were implemented¹ in Python using the CPMpy constraint programming and modeling library (Guns 2019). OR-Tools’ CP-SAT (Perron, Didier, and Gay 2023) solver was used. For query generation, we used PQ-Gen from [Tsouros, Berden, Guns, 2023], with a cutoff of 1 second to return the best query found. Implementation of the ML classifiers was carried out using the Scikit-Learn library (Pedregosa et al. 2011).

The comparison is based on *the number of queries*, which is very important for the applicability of interactive CA systems in real-world scenarios, and *the maximum user waiting time*, which is of paramount importance when human users are involved. The results presented in each benchmark, for each algorithm, are the means of 10 runs.

Results

Q1 In order to answer this question, we performed a 10-fold cross-validation with each classifier on all the datasets, and present the averages. As metrics for the comparison, we use *Accuracy*, *Balanced Accuracy* (the datasets are highly unbalanced, with $< 10\%$ of B typically having a positive label), and *F1-score*. The results are shown in Figure 1.

We can notice that all classifiers considered achieve a decent accuracy and balanced accuracy, with GNB performing slightly worse than the rest, and MLPs performing best. Focusing on F1-score, GNB presents quite bad results, but the rest of the classifiers still achieve a score higher than 70%. The results indicate that based on the way the dataset of constraints is created and the features used, it is possible to successfully train and use ML models to predict whether a constraint is part of the target network or not.

However, in order to use the classifiers to assist during the acquisition process – guiding it to generate promising queries, based on the predictions – it is of high importance to evaluate how they perform not only when the labels for

¹Our code is available online at: <https://github.com/Dimosts/ActiveConLearn>




Figure 1: Classification results with different classifiers


all candidate constraints are available, but when only some parts of the dataset are available (as this is the case during the CA process). Thus, we conducted an experiment to evaluate how the classifiers perform when only a percentage of the dataset is available. We used an increasing portion of the dataset as the training set, to evaluate their performance in different stages of the acquisition process, with the rest of the candidates being the test set. The order of the constraints in the dataset was decided based on the order in which they were added in 5 different runs of CA systems. The averages are presented in Figure 2.

We can observe that RF achieves the best results in all metrics in the beginning when only a small portion of the dataset is labeled, with MLP and SVMs reaching the same performance only when most of the dataset is available. GNB is shown here to have a bad performance throughout the process, having very low accuracy and F1-score.


Q2 Let us now focus on the effect of using probabilistic classification to guide query generation in CA. Figure 3 presents the result when using the different classifiers, compared to guiding using the simple counting method from (Tsouros, Berden, and Guns 2023) (Count) and the GrowAcq without guiding (Base).

In all benchmarks, except Random and JSS, the decrease in the number of queries is significant compared to both the baseline and the simple counting method, for most classifiers. When SVMs are used, the performance is similar to the baseline because it has a lower accuracy in earlier stages of the acquisition process and thus does not offer any meaningful guidance early enough. GNB presents decent results in some benchmarks, but its overfitting is shown in the Random benchmark, where guiding should not detect any patterns and thus have a similar performance as the baseline, which is true for the rest of the classifiers. Using RF and MLPs is the most promising, giving the best results in all benchmarks, with RF being superior in some cases. We attribute RF’s superior performance to the fact that it already achieves good prediction performance when only a small portion of the constraints is labeled, i.e., at the beginning of the acquisition process (Figure 2).


Regarding the waiting time for the user, it includes 1 second for query generation (based on the imposed cutoff), and the rest of the waiting time involves mainly the training and prediction time. As a result, we can see higher waiting times



(a) Accuracy



(b) Balanced Accuracy

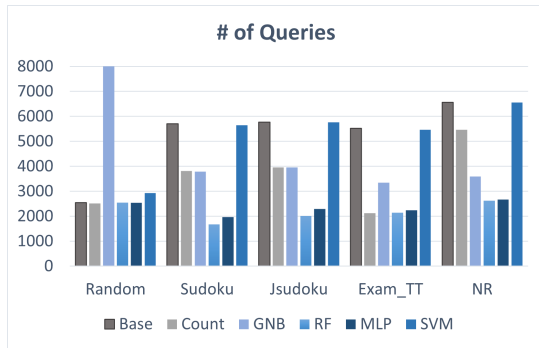


(c) F1 Score


Figure 2: Classification results when only part of the dataset is available for training

when SVM or MLPs are used, which need a larger training time, while GrowAcq with no guiding (Base) and the simple counting method (Count) have similar waiting times because they do not need any training. We can also observe that the training time for GNB and RF is small and very reasonable for interactive settings, as the maximum observed waiting time is less than 2s.

Overall, we can see that using RF to predict probabilities for the candidate constraints, and then guiding query generation based on these predictions, seems to be the best choice, both in terms of the number of queries and the user waiting time. It can decrease the number of queries required by up to 70% compared to the baseline (and up to 56% compared to the counting method), with the average decrease in the benchmarks that have structure (i.e., all except Random) being 52% (and 32% compared to counting). At the same time, the increase in the user waiting time is minor and acceptable for interactive scenarios.



(a)



(b)

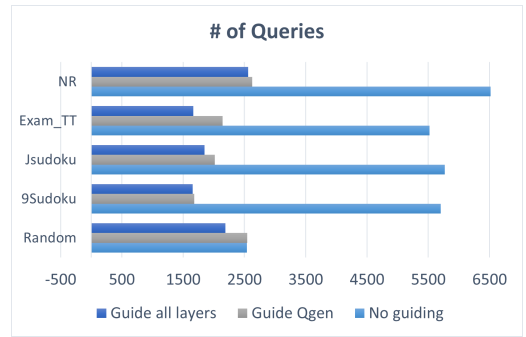
Figure 3: Results when guiding query generation using probabilistic classification

Q3 We now evaluate the effect of also guiding the other layers of interactive CA where queries are asked to the user. We only use RF, as it presented the best performance in the previous experiment, and we compare it against the baseline (without guiding) and against only guiding query generation. Figure 4 presents the results.

We can see a comparatively small but additional improvement when guiding all layers compared to guiding only top-level query generation. The improvement is relatively small because guiding query generation already led to significantly fewer queries needed in FINDSCOPE and FINDC. However, the additional decrease still goes up to 22% (in Exam TT), with the average decrease in the number of queries reaching 10%. In addition, we can observe a slight reduction in the maximum waiting time for the user when we use the predictions to guide all layers of CA. We believe that this is because by prioritizing the removal of candidate constraints in all layers, B is shrinking during FINDSCOPE and FINDC. Thus, fewer top-level queries will be needed leading to a smaller amount of retraining steps.

Conclusions

One bottleneck of major importance in interactive CA is the number of queries needed to converge. The search-based learning used in CA is often not able to detect patterns existing in the problem while learning the constraints, and thus, does not use such patterns to better guide its search. In this



(a)



(b)

Figure 4: Evaluating the effect of guiding all layers of CA

work, we tighten the connection between ML and CA, by using for the first time statistical ML methods that can learn during the acquisition process and predict whether a candidate constraint is part of the problem or not. We propose to use probabilistic classification, using the predictions from the ML models in order to guide the search process of CA. In doing so, we extend recent work that guided query generation using probabilities derived via a simple counting-based method. We also extend guidance to the other components of CA that post queries to the user, further reducing the number of queries. Our experimental evaluation showed that the number of queries was decreased by up to 72% compared, greatly outperforming the state of the art. These findings confirm that statistical ML methods can indeed detect patterns in constraint models, while they are being learned. This can be a stepping stone to further reducing the number of queries in interactive CA.

Future work should investigate the use of online learning in this setting, as data becomes available gradually. Other opportunities include learning a prior distribution over constraints and transfer learning across different problems. We also think that our closer integration with statistical ML techniques can be a stepping stone towards handling wrong answers from the user, which is an important part of future work, in order to make interactive CA more realistic. Finally, extending interactive CA systems to also be able to learn global constraints and linear inequalities with constants is important for expanding the reach of learnable problems.

Acknowledgments

This research received funding from the European Research Council (ERC) under the EU Horizon 2020 research and innovation programme (Grant No 101002802, CHAT-Opt) and from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 101070149, project Tuples.

References

- Angluin, D. 1988. Queries and concept learning. *Machine learning*, 2(4): 319–342.
- Arcangioli, R.; Bessiere, C.; and Lazaar, N. 2016. Multiple Constraint Acquisition. In *IJCAI: International Joint Conference on Artificial Intelligence*, 698–704.
- Beldiceanu, N.; and Simonis, H. 2012. A model seeker: Extracting global constraint models from positive examples. In *Principles and practice of constraint programming*, 141–157. Springer.
- Berden, S.; Kumar, M.; Kolb, S.; and Guns, T. 2022. Learning MAX-SAT Models from Examples using Genetic Algorithms and Knowledge Compilation. In *28th International Conference on Principles and Practice of Constraint Programming (CP 2022)*.
- Bessiere, C.; Carbonnel, C.; Dries, A.; Hebrard, E.; Katsirelos, G.; Narodytska, N.; Quimper, C.-G.; Stergiou, K.; Tsouros, D. C.; and Walsh, T. 2023. Learning constraints through partial queries. *Artificial Intelligence*, 319: 103896.
- Bessiere, C.; Coletta, R.; Freuder, E. C.; and O’Sullivan, B. 2004. Leveraging the learning power of examples in automated constraint acquisition. In *International Conference on Principles and Practice of Constraint Programming*, 123–137. Springer.
- Bessiere, C.; Coletta, R.; Hebrard, E.; Katsirelos, G.; Lazaar, N.; Narodytska, N.; Quimper, C.-G.; Walsh, T.; et al. 2013. Constraint Acquisition via Partial Queries. In *IJCAI*, volume 13, 475–481.
- Bessiere, C.; Coletta, R.; Koriche, F.; and O’Sullivan, B. 2005. A SAT-based version space algorithm for acquiring constraint satisfaction problems. In *European Conference on Machine Learning*, 23–34. Springer.
- Bessiere, C.; Coletta, R.; O’Sullivan, B.; Paulin, M.; et al. 2007. Query-Driven Constraint Acquisition. In *IJCAI*, volume 7, 50–55.
- Bessiere, C.; Koriche, F.; Lazaar, N.; and O’Sullivan, B. 2017. Constraint acquisition. *Artificial Intelligence*, 244: 315–342.
- De Raedt, L.; Passerini, A.; and Teso, S. 2018. Learning constraints from examples. In *Proceedings in Thirty-Second AAAI Conference on Artificial Intelligence*.
- Freuder, E. C. 2018. Progress towards the Holy Grail. *Constraints*, 23(2): 158–171.
- Freuder, E. C.; and O’Sullivan, B. 2014. Grand challenges for constraint programming. *Constraints*, 19(2): 150–162.
- Freuder, E. C.; and Wallace, R. J. 1998. Suggestion strategies for constraint-based matchmaker agents. In *International Conference on Principles and Practice of Constraint Programming*, 192–204. Springer.
- Guns, T. 2019. Increasing modeling language convenience with a universal n-dimensional array, CPpy as python-embedded example. In *Proceedings of the 18th workshop on Constraint Modelling and Reformulation at CP (Modref 2019)*, volume 19.
- Kolb, S. M. 2016. Learning constraints and optimization criteria. In *Workshops at the Thirtieth AAAI Conference on Artificial Intelligence*.
- Kumar, M.; Kolb, S.; and Guns, T. 2022. Learning Constraint Programming Models from Data Using Generate-And-Aggregate. In *28th International Conference on Principles and Practice of Constraint Programming (CP 2022)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik.
- Lallouet, A.; Lopez, M.; Martin, L.; and Vrain, C. 2010. On learning constraint problems. In *Tools with Artificial Intelligence (ICTAI), 2010 22nd IEEE International Conference on*, volume 1, 45–52. IEEE.
- Lazaar, N. 2021. Parallel Constraint Acquisition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 3860–3867.
- Lombardi, M.; and Milano, M. 2018. Boosting Combinatorial Problem Modeling with Machine Learning. *arXiv preprint arXiv:1807.05517*.
- Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; Vanderplas, J.; Passos, A.; Cournapeau, D.; Brucher, M.; Perrot, M.; and Duchesnay, E. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12: 2825–2830.
- Perron, L.; Didier, F.; and Gay, S. 2023. The CP-SAT-LP Solver (Invited Talk). In *29th International Conference on Principles and Practice of Constraint Programming (CP 2023)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik.
- Tsouros, D.; Berden, S.; and Guns, T. 2023. Guided Bottom-Up Interactive Constraint Acquisition. In *International Conference on Principles and Practice of Constraint Programming*.
- Tsouros, D. C.; and Stergiou, K. 2020. Efficient multiple constraint acquisition. *Constraints*, 25(3): 180–225.
- Tsouros, D. C.; and Stergiou, K. 2021. Learning Max-CSPs via Active Constraint Acquisition. In *27th International Conference on Principles and Practice of Constraint Programming (CP 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik.
- Tsouros, D. C.; Stergiou, K.; and Bessiere, C. 2019. Structure-Driven Multiple Constraint Acquisition. In *International Conference on Principles and Practice of Constraint Programming*, 709–725. Springer.
- Tsouros, D. C.; Stergiou, K.; and Bessiere, C. 2020. Omissions in Constraint Acquisition. In *International Conference on Principles and Practice of Constraint Programming*, 935–951. Springer.
- Tsouros, D. C.; Stergiou, K.; and Sarigiannidis, P. G. 2018. Efficient Methods for Constraint Acquisition. In *24th International Conference on Principles and Practice of Constraint Programming*.