

1 Constraint Acquisition Based on Solution Counting

2 Christopher Coulombe ✉

3 Université Laval, Québec, Canada

4 Claude-Guy Quimper ✉ 🏠

5 Université Laval, Québec, Canada

6 — Abstract —

7 We propose CABSC, a system that performs Constraint Acquisition Based on Solution Counting.
8 In order to learn a Constraint Satisfaction Problem (CSP), the user provides positive examples and
9 a Meta-CSP, i.e. a model of a combinatorial problem whose solution is a CSP. This Meta-CSP
10 allows listing the potential constraints that can be part of the CSP the user wants to learn. It also
11 allows stating the parameters of the constraints, such as the coefficients of a linear equation, and
12 imposing constraints over these parameters. The CABSC reads the Meta-CSP using an augmented
13 version of the language MiniZinc and returns the CSP that accepts the fewest solutions among the
14 CSPs accepting all positive examples. This is done using a branch and bound where the bounding
15 mechanism makes use of a model counter. Experiments show that CABSC is successful at learning
16 constraints and their parameters from positive examples.

17 **2012 ACM Subject Classification** Computing methodologies → Modeling methodologies

18 **Keywords and phrases** Constraint acquisition, CSP, Model counting, Solution counting

19 **Digital Object Identifier** 10.4230/LIPIcs.CP.2022.27

20 **1** Introduction

21 Constraint solvers are used to solve complex combinatorial problems. They require an expert
22 to model the problem using the constraints available in the solver. The model creation is a
23 crucial step, but is often time-consuming. One way to save time to the expert is to suggest
24 a model based on sample solutions. For instance, a hospital that wants to automatize the
25 creation of their work schedules for its staff might provide to the experts previous schedules.
26 Assisted with software, the expert wants to discover what constraint generated the examples.
27 While some of these constraints are already known and even written on legal documents,
28 there are as important constraints that are not written but are part of the work culture.
29 These are the constraints for which a constraint acquisition software becomes handy.

30 When two constraints are candidates for a model, the one that was the most likely used
31 to generate the sample solutions is the most restrictive one [14]. Different approaches exist to
32 decide which constraint is the most restrictive. There are mainly statistical approaches [13, 14]
33 and approaches based on a ranking system [6] (that includes many other criteria). Current
34 methods analyze the constraint in isolation. However, adding to a model a constraint that
35 accepts many solutions can reduce more the solution space than adding a constraint that
36 accepts few solutions. It all depends on the interaction between the constraints in the model.
37 We propose the first approach that takes into account this interaction. It uses a model
38 counter to make sure that the constraints suggested to the expert are those that are the
39 most likely to explain the observed sample solutions given the constraints that were already
40 identified to be part of the model.

41 In this paper, we propose CABSC, an algorithm for Constraint Acquisition Based on
42 Solution Counting. CABSC uses examples of solutions to evaluate which constraints to keep
43 from a chosen set of candidates. The selection process is based on solution counting using
44 model counters, an approach which differs from the current methods detailed in Section 2.



© Christopher Coulombe and Claude-Guy Quimper;
licensed under Creative Commons License CC-BY 4.0

28th International Conference on Principles and Practice of Constraint Programming (CP 2022).

Editor: Christine Solnon; Article No. 27; pp. 27:1–27:16



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

45 The definitions for our approach are given in the Section 3, followed by a practical explanation
 46 in Section 4. Experiments are explained in Section 5 and discussed in Section 6.

47 **2** General Background

48 Constraint acquisition is an intricate problem that can be solved in a few ways. A first
 49 idea called passive learning requires examples of solutions and/or non-solutions. A system
 50 chooses which constraint represents best the examples from a preselection of constraints.
 51 The preselected pool of constraints from which the model is built is called a bias. Other
 52 methods use active learning and generate examples of solution and ask an expert to classify
 53 the examples given. From a bias, the system choses the best set of constraints according to
 54 the answer provided.

55 Passive learning systems exploit the idea that the underlying structure of the given
 56 examples gives information about the model to learn. Beldiceanu and Simonis [6] created a
 57 Model Seeker that learns constraints from a catalog given positive and negative examples.
 58 The constraints of the catalog that accepts the positive examples and reject the negative
 59 examples are sorted with the more likely constraints having a higher rank. The sorting system
 60 is based on a ranking value that is a function of multiple parameters, including the number
 61 of solutions satisfying the constraint [5]. A constraint accepting fewer solutions is more
 62 likely to be the constraint that generated the examples as there is a lesser chance that the
 63 examples are a product of a coincidence. To work, this method needs to make the hypothesis
 64 that the constraints learned are independent of each other. That hypothesis is not what
 65 transpires in real applications and may result in errors. Two constraints with a small but
 66 near identical set of solutions would be picked over two constraints accepting more solutions
 67 if picked individually but very few solutions when combined. This is counterproductive as
 68 the idea is often to complete an already existing model or to learn multiple constraints at
 69 the same time.

70 Picard-Cantin et al. [13] approached the problem with a statistical approach with the idea
 71 that the constraint that best explains the examples is the most improbable one. Equation (1)
 72 was therefore used by Picard-Cantin et al. [14] to calculate the probability of the constraints
 73 where $G_C(P)$ is the probability that a random assignment satisfies the constraint C with
 74 the parameters P . The parameters can be, for instance, the coefficients of a linear equation.
 75 $S_C(P)$ is the solution set that satisfies the constraint C with the parameters P . The
 76 probability is calculated for a constraint over n variables. $prob(e)$ is the probability to observe
 77 an assignment e of n variables and $prob(e_i)$ is the probability to observe an assignment of a
 78 single variable.

$$79 \quad G_C(P) = \sum_{e \in S_C(P)} \text{Prob}(e) = \sum_{e \in S_C(P)} \prod_{i=1}^n \text{Prob}(e_i) \quad (1)$$

80 A hypothesis of independence between the variables of the constraints is applied in the
 81 equation. Whenever a variable is in the scopes of multiple constraints, the hypothesis of
 82 independence between the variables becomes an approximation. In all cases, the preferred
 83 constraints are the ones with a small number of solutions but the independence hypothesis
 84 can lead to an erroneous ranking of the constraints. Moreover, this system was not designed
 85 for learning multiple constraints and requires solution counting algorithms specialized for
 86 each constraint.

87 Another approach was suggested by Bessiere et al. [8] which consists of creating a model
 88 from partial queries, a form of active learning, with an algorithm called QuAcq. The system

89 creates an example and asks an expert whether the presented values are valid. The system
90 adapts the learned constraints depending on the provided answer. Recently, QuAcq was
91 improved with a new version called QuAcq2 [7]. In some cases, QuAcq and QuAcq2 can
92 require a number of queries too high to be efficiently answered by a person. The number of
93 queries can go as high as $n^2 \log(n)$ where n is the number of variables of the problem [7].
94 Multiple authors tackled this problem such as Daoudi et al. [10], Addi et al. [2], Addi et al. [1],
95 Arcangiooli and Lazaar [3], Tsouros et al. [20] and Tsouros et al. [19], but up to thousands of
96 queries can still be needed.

97 **3 The CABSC approach**

98 The CABSC approach (Constraint Acquisition Based on Solutions Counting) we introduce
99 fulfills three goals:

- 100 1. To lift the hypothesis of independence between variables;
- 101 2. To allow learning multiple constraints;
- 102 3. To work with any set of constraints for which filtering algorithms exist, rather than
103 solution counting algorithms.

104 CABSC models the process of learning constraints as a Meta-CSP. As will be described in
105 Section 3.1, a Meta-CSP is a combinatorial problem whose solution is a CSP. In our case, the
106 solution is the CSP we learn from the examples. When modeling the Meta-CSP, we list the
107 mandatory constraints, i.e. the constraints that we know belong to the model, and also the
108 possible constraints, those that could belong to the model. The variables of the Meta-CSP
109 encode the possible activation of a constraint and also the parameters of the constraints, such
110 as the coefficients of a linear constraint. Solving the Meta-CSP provides the learned model.
111 To do so, we use a branch and bound to decide which constraint to keep and identify the
112 values of the parameters. Our approach uses constraint programming to model a Meta-CSP
113 and to define a family of CSPs from which we can learn. We therefore do not aim to learn
114 any CSP but the optimal CSP among a set programmed through constraint programming.
115 This approach is inspired from regression where one defines a family of functions (e.g. linear
116 functions) and aims at finding the function from this family that best fits the data. Here, we
117 aim at finding the CSP from a family of CSPs defined by the Meta-CSP that best explains
118 the examples.

119 As there are multiple candidate constraints that could belong to the learned model, we
120 follow Beldiceanu and Simonis [6] and Picard-Cantin et al. [13] by selecting the constraints
121 that minimize the number of solutions. However, instead of analyzing the constraints
122 individually like Beldiceanu and Simonis [6] and Picard-Cantin et al. [13], our system reasons
123 globally on all constraints which allows us to consider multiple different constraints at once.

124 In order to lift the hypothesis that variables and constraints in a CSP are independent,
125 we directly count the solutions of a model using a model counter. The solution to our
126 Meta-CSP is therefore a CSP whose constraints are satisfied by all observed examples and is
127 as restrictive as possible, i.e. it minimizes the number of solutions.

128 Our approach has two main differences from existing methods. The first difference is that
129 constraint programming, through the declaration of a Meta-CSP, is used to define a family
130 of CSPs from which we can learn. A second difference from most existing methods is that
131 we use a criterion with a global view on the model to learn by considering the constraints to
132 learn as a whole instead of individually.

133 3.1 Definition of a Meta-CSP

134 Following [16], a CSP \mathcal{P} is a triple $\mathcal{P} = \langle X, \text{dom}, C \rangle$ where X is a n -tuple of variables
 135 $X = \langle X_1, X_2, \dots, X_n \rangle$, dom is a function that maps a variable in $X_i \in X$ to a set of
 136 values, called domain, that can be assigned to the variable X_i , C is a t -tuple of constraints
 137 $C = \langle C_1, C_2, \dots, C_t \rangle$. A constraint C_j is a pair $\langle R_j, S_j \rangle$ where $S_j \subseteq X$ is the scope of the
 138 constraint and R_j is a relation on the variables in S_j . In other words, R_j is a subset of
 139 the Cartesian product of the domains of the variables in S_j . A solution to the CSP \mathcal{P} is an
 140 assignment to the variables $X = v_1, \dots, X_n = v_n$ such that $v_j \in \text{dom}(X_j) \forall 1 \leq j \leq n$ and
 141 each C_j is satisfied in that the tuple $\langle v_1, \dots, v_n \rangle$ projected onto S_j is a tuple in R_j .

142 We extend the definition of a CSP to a Meta-CSP. The solution of a Meta-CSP is a CSP.
 143 In our case, it is the CSP we want to learn. A Meta-CSP is a tuple $M = \langle X, P, \alpha, \text{dom}, E, C \rangle$
 144 where $X = \langle X_1, \dots, X_n \rangle$ are the decision variables, $P = \langle P_1, \dots, P_q \rangle$ are the parameter
 145 variables, $\alpha = \langle \alpha_1, \alpha_2, \dots \rangle$ are the activation variables, dom is a function that maps a variable
 146 in $X \cup P \cup \alpha$ to a set of values that can be assigned to the variable, E is the example matrix
 147 of dimensions $m \times n$, and $C = \{C_1, \dots, C_t\}$ is a set of constraints. A row $e_i = \langle e_{i,1}, \dots, e_{i,n} \rangle$
 148 of matrix E satisfies $e_{i,j} \in \text{dom}(x_j)$ and is a solution to the CSP we want to learn. The
 149 examples of the matrix must satisfy the constraints that we want to learn.

150 A constraint C_j is a quadruple $\langle R_j, S_j, P_j, \alpha_j \rangle$ where $S_j \subseteq X$ is the scope of the constraint,
 151 $P_j \subseteq P \cup \alpha$ its parameters set and $\alpha_j \in \alpha$ its activation variable. For instance, for a linear
 152 constraint, the parameters P_j are the coefficients that need to be learned. To each constraint
 153 C_j is associated the activation variable α_j with domain $\text{dom}(\alpha_j) \subseteq \{\perp, \top\}$. Deciding whether
 154 α_j is true (\top) is equivalent to deciding whether the constraint appears in the learned model.
 155 One can force a constraint to appear in the learned model by setting $\text{dom}(\alpha_j) = \{\top\}$ in the
 156 definition of the Meta-CSP. The relation R_j is a set of the assignments accepted by the
 157 constraint along with the parameters given to the constraint: $R_j \subseteq \times_{x \in S_j} x \times \times_{p \in P_j}$.

158 A solution to the Meta-CSP is an assignment to the parameter variables $P_1 = p_1, \dots, P_q =$
 159 p_q and an assignment to the activation variables $\alpha_1 = r_1, \dots, \alpha_t = r_t$ such that $r_j \in \text{dom}(\alpha_j)$
 160 for all constraints C_j , $p_k \in \text{dom}(P_k)$ for all $1 \leq k \leq q$. Finally, the examples must satisfy the
 161 activated constraint, i.e. $\forall 1 \leq j \leq t, \alpha_j \implies \forall i \langle e_{i,1}, \dots, e_{i,n}, p_1, \dots, p_q \rangle \in R_j$.

162 4 Framework

163 4.1 The Language

164 We augmented the MiniZinc language [12] to model a Meta-CSP. The declaration of constraints
 165 in a Meta-CSP differs from the one in a CSP in two ways. First, the constraints had to
 166 be rewritten in MiniZinc to include the Boolean activation variable. This avoids writing
 167 explicitly, for each constraint, the underlying constraints needed for such variables. Second,
 168 when declaring the scope of a constraint, the indices of the decision variables in X need
 169 to be stored in the constraint. Indeed, the constraint's filtering algorithm needs a map of
 170 the decision variables in its scope to the columns of the matrix of examples E . Therefore,
 171 constraints used for the Meta-CSP have different specifications from what is possible within
 172 MiniZinc, which is why the language had to be augmented. The MiniZinc language was also
 173 modified to better communicate with the solver we developed, i.e. imports and heuristics
 174 were adapted to give a better control. Even though the modifications to MiniZinc do not
 175 change its fundamental structure, the way to write a Meta-CSP is made significantly easier.

176 Listing 1 provides a code snippet written in the augmented MiniZinc language. A set of
 177 two-dimensional points are given as solutions of an unknown CSP problem. We know that

178 the x and y coordinates of these points are nonnegative. We do not know whether these
 179 points are subject to a linear inequality or an elliptic inequality. This Meta-CSP will tell us.

```

180
181 1 set: domain = 1..10;
182 2 array: x = [1]; %Points are (x,y)
183 3 array: y = [2];
184 4 array: x_y = [1..2];
185 5 var domain: a;
186 6 var domain: b;
187 7 var domain: c;
188 8 var 0..1: activation1;
189 9 var 0..1: activation2;
190 10
191 11 constraint Linear(x, [1], ">=", 0, true); % x >= 0
192 12 constraint Linear(y, [1], ">=", 0, true); % y >= 0
193 13 constraint Linear(x_y, [a,b], "<=", c, activation1); % a*x + b*y <= c
194 14 constraint Ellipse(x_y, [a,b], "<=", c, activation2); % a*x^2 + b*y^2 <= c
195 15 constraint Xor(activation1, activation2, true);

```

■ **Listing 1** Code snippet of the augmented MiniZinc

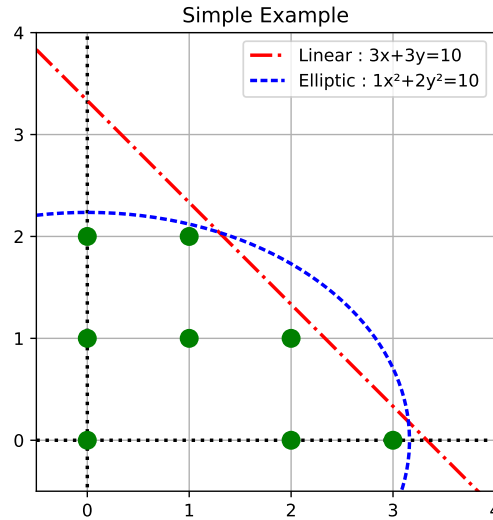
197 The decision variables x and y are declared on lines 2 and 3. As their values are known
 198 for each example, they are not declared as variables using the keyword *var* but rather as
 199 constants corresponding to the column numbers in the example matrix E .

200 Line 11 declares the first constraint of the problem. It is interpreted as follows: It is a
 201 linear constraint whose scope is the decision variable x , whose coefficient vector is $[1]$, whose
 202 comparison operator is \geq , and whose right-hand side is 0. It can be interpreted as $[1]^T x \geq 0$.
 203 The activation variable is set to *true*, which means that this constraint is known to belong to
 204 the CSP. Line 12 imposes $y \geq 0$ with a similar constraint. Line 13 encodes the first constraint
 205 that we want to learn. It is a linear constraint over the variables x and y whose coefficients
 206 and right-hand side are unknown and are represented by the parameter variables a , b , and c .
 207 Finally, it is unknown whether this constraint belongs to the CSP. The activation variable
 208 **activation1** will be set to 1 if it belongs and 0 otherwise. Line 14 encodes the second
 209 constraint that we want to learn. It is an elliptic constraint centered at the origin where
 210 parameter variables a , b , and c are reused. The activation variable **activation2** is used
 211 for this constraint. Line 15 shows an example of a constraint over two activation variables
 212 meaning that exactly one constraint among the linear and the elliptic constraint can be
 213 activated. This is an example of how one can define the bias (i.e. the family of CSPs from
 214 which the CSP is learned) and exploit the full richness of CP to model the learning process.

215 A constraint can be satisfied by all examples even if the solver chooses not to learn it by
 216 setting its activation variable to \perp , unlike a reified constraint which would be set to \perp only
 217 if the examples are not satisfied.

218 Figure 1 is a graphical representation of the problem encoded in Listing 1. The curves
 219 represent both candidate constraints: the linear candidate and the elliptic candidate. The
 220 dots are the sample solutions that are provided.

221 We are looking for the CSP that is the most likely the one that generated the points
 222 provided in the example matrix E , i.e. the CSP that accepts the fewest solutions among
 223 all CSPs that accepts all solutions in E . We see in Figure 1 that the Elliptic constraint
 224 accepts 9 solutions while the linear constraint accepts 10 solutions. Therefore, our approach
 225 learns that an elliptic inequality fits best the examples with parameters $a = 1, b = 2, c =$
 226 $10, \text{activation1} = \perp$ and $\text{activation2} = \top$, which confirms the visual intuition.



■ **Figure 1** Simple example

227 4.2 The Solver

228 We created a custom solver called `CabscSolver` that reads the Meta-CSP written in the
 229 augmented MiniZinc language and the example matrix E . This solver finds the CSP that
 230 accepts the fewest solutions among all CSPs that accept all examples. `CabscSolver` uses
 231 a branch and bound to solve the problem. The branching variables are the activation
 232 and parameter variables $\alpha \cup P$. After branching, constraint propagation is triggered. Let
 233 $C(\vec{x}, \vec{p}, \alpha)$ be a constraint where \vec{x} is the vector of decision variables, \vec{p} is the vector of
 234 parameter variables, and α is the activation variable. Only the domains of \vec{p} and α need to
 235 be filtered as the values of the decision variables are provided by the examples. To filter
 236 the constraint, one needs to filter the expression $\alpha \implies \bigwedge_{i=1}^m C(e_i | \vec{x}, \vec{p}, \top)$ where $e_i | \vec{x}$ is the
 237 projection of the i^{th} example over the decision variables in the scope of the constraint. The
 238 filtering can take place only when the value of the activation variable α is known. Indeed,
 239 if α is false (\perp), the constraint is satisfied and no filtering is required. If α is true (\top), a
 240 conjunction of constraints needs to be filtered. Each component of the conjunction can
 241 be filtered independently, but a more sophisticated algorithm might process the examples
 242 in batch to gain in efficiency. The choice is specific to each constraint. In the example of
 243 Listing 1, if variable `activation1` is set to \top during the search process, the linear constraint
 244 filters values 1 and 2 from the domain of c as the point $(x, y) = (3, 0)$ prevents the linear
 245 constraint to be satisfied when $c \leq 2$.

246 In order to make the branch and bound effective at minimizing the number of solutions
 247 accepted by the CSP we want to learn, one needs to compute a lower bound on this number
 248 of solutions. This computation is carried in two phases. In the first phase, we detect if
 249 a situation occurs where it is possible to deduce which CSP accepts the fewest solutions,
 250 regardless whether this CSP accepts the examples or not. If such a CSP can be deduced, the
 251 second phase launches a model counter to compute the number of solutions for this CSP.

252 Some constraints have monotonic parameters with respect to the number of solutions
 253 they accept [14]. For instance, consider the linear constraint $c^T x \leq b$ where the parameters c

254 and b are a vector of nonnegative coefficients and a nonnegative right-hand side. The vector
255 x contains the decision variables. It is clear that the number of solutions accepted by this
256 constraint decreases as the values in c increases and b decreases. In order to obtain the
257 most restrictive constraint, one needs to fix the parameters c to their greatest values in their
258 domains and b to its smallest value. If all parameter variables with more than one value in
259 their domain are monotonic and all constraints agree to set these variables to the same values
260 (either largest or smallest) in order to minimize the number of solutions, then we can proceed
261 to the second phase and compute a lower bound on the number of solutions. Otherwise, we
262 use the number of examples as the trivial lower bound as this is the minimum number of
263 solutions the CSP can accept. Since parameter variables can be subject to constraints, it is
264 possible that fixing the value of the parameter variables leads to inconsistencies. In such
265 a case, the CSP used to calculate the lower bound has no solution. Even if that CSP has
266 no solution, multiple CSPs can exist further in the search tree. We therefore still use the
267 number of examples as a lower bound on those nodes.

268 In the second phase, the parameter variables are set to their most restrictive value and
269 activation variables that are not set to *false* are forced to be *true* in order to have the
270 maximum number of activated constraints. This results in a CSP \mathcal{A} for which the number of
271 solutions needs to be determined. There exists a few model counters in the literature such
272 as the exact probabilistic model counter GANAK [17] or the approximate model counter
273 ApproxMC4 [9, 18]. Both of these counters can only approximate the number of solutions of a
274 model written as a CNF file. CabscSolver encodes the constraints of \mathcal{A} into a pseudo-Boolean
275 language that is translated to a CNF using the MiniSat+ module NaPS [11]. This CNF is
276 given to the model counter which calculates the number of solutions of the model. This
277 number is used as a lower bound on the number of solutions of the learned CSP for the
278 current node of the branch and bound.

279 Executing the model counter is the most time-consuming operation in the whole search
280 process. Since the parameter variables are often fixed to the same values (due to their
281 monotonicity), it is worth implementing a cache system. Therefore, before calling the model
282 counter, the system checks whether the generated model was previously counted, and if so,
283 returns the number of solutions previously found.

284 The resulting algorithm is summarized in Figure 2. The next branching is defined by the
285 best-first-search heuristics, i.e. the open node with the smallest lower bound is expanded.
286 When the lower bound of a node is greater than the number of solutions of the incumbent
287 CSP, this node is closed.

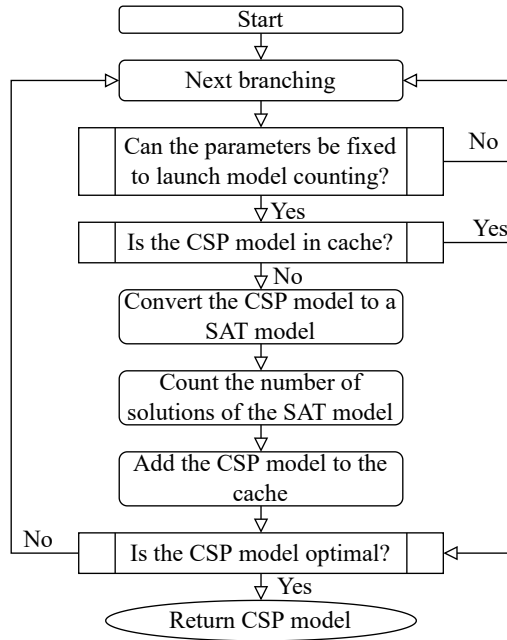
288 5 Experiments

289 5.1 Implementation

290 We implemented CabscSolver in Python¹. While this interpreted language leads to a slow
291 execution, in practice, most of the computation time is spent in the model counters. We use
292 GANAK [17] and ApproxMC4 [9, 18] as model counters that are both efficiently implemented
293 in C/C++.

294 GANAK is a probabilistic exact model counter [17]. Using the parameter δ , GANAK
295 guarantees with a probability of at least $1 - \delta$ that the value provided is an exact count. The
296 approximate model counter ApproxMC4 [9, 18] was also integrated to our solver to count

¹ The code and the benchmarks will be available on the authors' web sites.



■ **Figure 2** Flow chart for the CABSC approach

297 the number of solutions since some calculations are much faster with this counter. Let F be
 298 the real number of solutions of a model. ApproxMC4 gives an approximation of F with a
 299 configurable confidence. More specifically, it returns a count that is guaranteed to be within
 300 $[\frac{F}{(1+\epsilon)}, F \cdot (1 + \epsilon)]$ with a probability of at least $1 - \delta$, where ϵ and δ are the configurable
 301 parameters. The chosen values for the parameters ϵ and δ are discussed in Section 5.3.

302 To read the Meta-CSP models, using the parsing toolkit Lark, we implemented, from
 303 scratch, a parser that interprets a subset of the MiniZinc language [12] to which we add the
 304 necessary augmentations. MiniZinc was not changed in any way other than the required
 305 augmentations. This allows us to efficiently communicate the Meta-CSP models to the solver.

306 5.2 Instances

307 We try to learn the constraints inspired from nurse scheduling problems. The problem
 308 consists of creating a schedule that respects a set of predetermined rules. In these schedules,
 309 the increments used are days, meaning that we are only preoccupied on a daily basis whether
 310 the nurses work or not. Let $\eta \in \{2, 3, 4\}$ and $d \in \{7, 14, 21, 28\}$ be the number of nurses
 311 and days in a schedule (with $\eta d \leq 56$). All instances have a matrix of decision variables
 312 $[[X_{(1,1)}, \dots, X_{(1,d)}], \dots, [X_{(\eta,1)}, \dots, X_{(\eta,d)}]]$. Each variable of the matrix represents a day of
 313 work for a nurse with its domain being $\{0, 1, 2\}$. $X_{i,j}$ takes the value 0 if the nurse i does not
 314 work on day j . If the nurse i does work during day j , $X_{i,j}$ takes the value 1 or 2, depending
 315 on whether the nurse works in room 1 or 2.

316 In the first benchmark, denoted **Sequence**, we want to learn one of these two constraints
 317 on the rows of the matrix.

$$318 \text{SEQUENCE}([X_{i,1}, \dots, X_{i,d}], l, u, k, V) \quad \forall 1 \leq i \leq \eta \quad (2)$$

$$319 \text{AMONG}(t_1, t_2, [X_{i,7w+1}, \dots, X_{i,7(w+1)}], V) \quad \forall 1 \leq i \leq \eta, \forall 0 \leq w < \frac{d}{7} \quad (3)$$

320

321 Constraint (2) is the SEQUENCE constraint [4] that is satisfied when at least l and at most u
 322 variables in a window $X_{i,j}, \dots, X_{i,j+k-1}$ of k consecutive variables are assigned to a value
 323 in the set V . This constraint is used to spread out the workload of the nurses over the
 324 days without underload nor overload. The parameters l , u , and k are unknown and need
 325 to be learned. Their domains are given by $\text{dom}(l) = \text{dom}(u) = \text{dom}(k) = [0, 7]$ and are
 326 subject to $l \leq u < k$. The set V is known and fixed to $\{1, 2\}$ as these are the values that
 327 represent a nurse who is working. Constraint (3) simply constrains the number of work days
 328 to be at least t_1 and at most t_2 every week. The parameter variables t_1 and t_2 have for
 329 domain $\text{dom}(t_1) = \text{dom}(t_2) = [0, 7]$. One, and only one, constraint among (2) or (3) must
 330 be activated. We therefore constrain the activation variables of both constraints with a **Xor**,
 331 just like the line 15 of Listing 1. The benchmark **Sequence** is composed of 368 instances
 332 generated with distinct constraints, parameters, and examples. These instances satisfy the
 333 SEQUENCE constraint and the parameters lie in the intervals $l, u \in [1, 6]$ and $k \in [2, 7]$.

334 The second benchmark, denoted **Complex**, inherits all the characteristics of the **Se-**
 335 **quence** benchmark, including the constraint to learn, to which additional known constraints
 336 are added on the decision variables. These constraints have for goal to encode a more realistic
 337 situation where constraints that we want to learn are mixed with constraints that are known.
 338 For each column $[X_{(1,d)}, \dots, X_{(\eta,d)}]$ of the matrix that represents the schedule for the day
 339 d , we have the constraint **AMONG**($b, 3, [X_{(1,d)}, \dots, X_{(\eta,d)}], V$) where $b = 1$ if d is a Monday,
 340 Tuesday, Wednesday, or Thursday and $b = 2$ otherwise. This constraint and its parameters
 341 are known and added to the Meta-CSP with an activation variable set to \top . This constraint
 342 does not need to be learned. For instances with 3 or more nurses, we also have another
 343 known constraint $X_{(\eta,j)} = 0 \vee X_{(\eta-2,j)} = 0 \quad \forall j \in \{1, \dots, d\}$ in order to prevent nurse η from
 344 working at the same time as $\eta - 2$. When applicable, this constraint is also included in the
 345 Meta-CSP as a known constraint. The **Complex** benchmark has 247 instances that satisfy
 346 the SEQUENCE constraint with the parameters lying in the intervals $l, u \in [1, 6]$ and $k \in [2, 7]$.

347 In the third benchmark denoted **Vacation**, the Meta-CSP is identical to the one of
 348 **Complex**. However, the examples E that are provided to the solver are particular: nurses
 349 can be non-working for 7 consecutive days. This represents a situation where the staff goes on
 350 leave during the vacation period. These leaves violate the SEQUENCE constraint and force the
 351 solver to activate the **AMONG** constraint and learn its parameters t_1 and t_2 . The examples
 352 were created such that nurse η never takes a vacation but other nurses do. For a problem
 353 spanning w weeks, nurses globally take no more than w weeks of vacation. We generated
 354 272 instances for this benchmark such that the instances satisfy the **AMONG** constraint. The
 355 parameters lie in the intervals $t_1 \in [2, 3]$ and $t_2 \in [3, 7]$.

356 The last benchmark **Overtime** uses the same Meta-CSP as **Complex** and **Vacation**,
 357 but the examples E provided to learn the CSP differ from **Vacation** on one point: rather
 358 than leaving for vacations for 7 consecutive days, the nurses in the **Overtime** benchmark
 359 work on a stretch of 7 consecutive days. This represents a situation when the hospital is
 360 understaffed and nurses need to work overtime. This benchmark has 304 instances such that
 361 the instances satisfy the **AMONG** constraint and the parameters lie in the intervals $t_1 \in [2, 7]$
 362 and $t_2 \in [4, 7]$ with the restriction $t_1 \leq t_2$.

363 For all benchmarks, the solver aims to learn exactly one constraint among (2) and (3).
 364 The selection depends on the known constraints added to the Meta-CSPs and the examples.

365 5.3 Experimental Setup

366 For each instance, the CSP we want to learn was written in the MiniZinc language [12] and
 367 used to randomly generate up to a thousand solutions. The Meta-CSP model was written in

our augmented-MiniZinc language in order to learn which constraint, between the SEQUENCE and the AMONG constraints, is activated and what are the parameters that were used to generate the examples.

CabscSolver supports two model counters. We first used the solver with the model counter GANAK [17]. By setting the parameter δ to 0.05, we state that the value returned by the model counter is guaranteed to be exact with a probability of at least 0.95. Tighter guarantees can be used, but the time taken to count the number of solutions of the models increases accordingly. Using this model counter and this configuration, we nevertheless assume the given number of solutions to be exact. GANAK was used with a maximum cache size of 2000 Mb. We ran all benchmarks on the solver using this model counter.

As a second series of tests, we used a mix of ApproxMC4 [9, 18] and GANAK. Some CSP models are faster to evaluate with ApproxMC4, so we tried to make CABSC faster using both model counters. Since ApproxMC4 is not an exact model counter, we did not want to run both model counters at the same time and simply use the result returned by the fastest of the two. When using both model counters, GANAK and ApproxMC4 are simultaneously launched. If GANAK finishes first, ApproxMC4 is terminated. If ApproxMC4 finishes first, GANAK is terminated only if the returned result is conclusive. Indeed, ApproxMC4 returns a solution count that is guaranteed to be within an interval with a parametrized confidence. A solution returned by this model counter could be largely underestimated, which could lead to the wrong CSP model being learned. If F is the exact number of solutions of a CSP, the number of solutions returned by ApproxMC4 lies in $[\frac{F}{(1+\epsilon)}, F \cdot (1 + \epsilon)]$ with probability $1 - \delta$. When ApproxMC4 returns a number of solutions that is $(1 + \epsilon)$ times greater than the number of solutions accepted by the incumbent CSP, the computation of GANAK is halted, and the node is closed, i.e. no children of this node will be explored in the search tree. Otherwise, we draw no conclusion and let GANAK terminate its computation. ApproxMC4 is rather used as a means to close nodes faster than substituting GANAK.

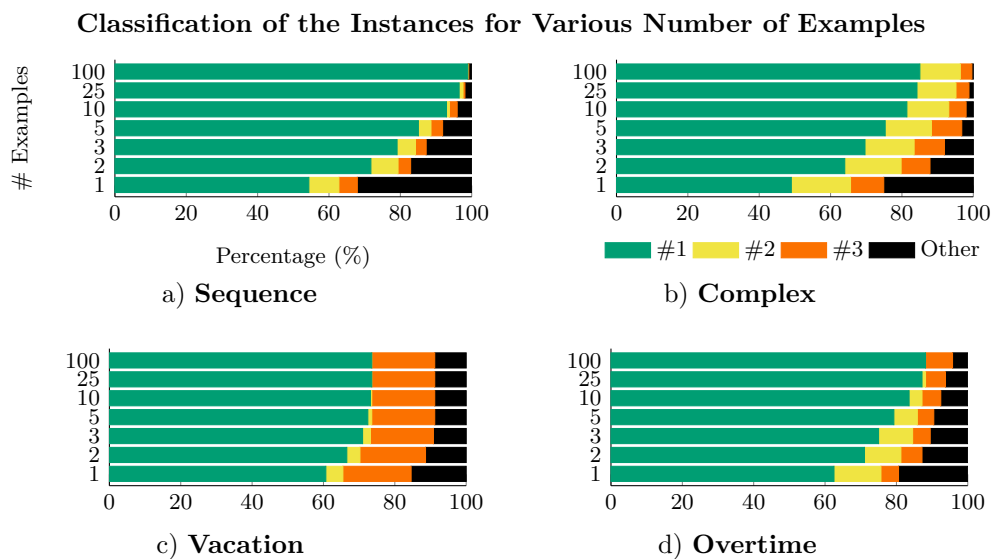
The same way we assumed that GANAK would return exact values, we assume that ApproxMC4 does not give a solution count that is lower than the minimum value of the interval. We used $\delta = 0.10$ and $\epsilon = 0.5$ which means that the count calculated is guaranteed to be in the range $[\frac{F}{1.5}, 1.5F]$ with a probability of at least 0.90. A lower probability is accepted from ApproxMC4 than GANAK since the main focus of using ApproxMC4 is to count CSP models faster than GANAK.

We ran the experiments on a computer with the following configuration: CentOS 7.6.1810, 32 GB ram, Processor Intel(R) Xeon(R) Silver 4110 CPU @ 2.10GHz, 32 Cores. We simultaneously launch 7 instances of the solver.

From each instance, random subsets of 1, 2, 3, 5, 10, 25, and 100 examples were used. Each time, the top 3 solutions are returned by the solver, and we verify that one of these solutions is the one used to generate the examples. For the **Sequence** and **Complex** benchmarks, the expected constraint to be learned is the SEQUENCE constraint with parameters l , u , and k . For the **Vacation** and **Overtime** benchmarks, the examples violate the SEQUENCE constraint, and the AMONG constraint is expected to be learned with parameters t_1 and t_2 .

6 Results and Discussion

Figure 3 presents the results obtained when running CabscSolver using only GANAK for the four benchmarks presented at Section 5.2. On the y -axis is the number of examples that are given to the solver. On the x -axis is the proportion of instances for which the solution is the best one returned by the solver, the second best, the third best, or whether the CSP that



■ **Figure 3** Classification of the instances in percentage for each number of examples. CabscSolver uses GANAK as the only model counter.

414 was used to produce the examples does not appear at all in the top-3 learned models. We
 415 recall that the solver returns the CSP that minimizes the number of solutions.

416 6.1 Accuracy

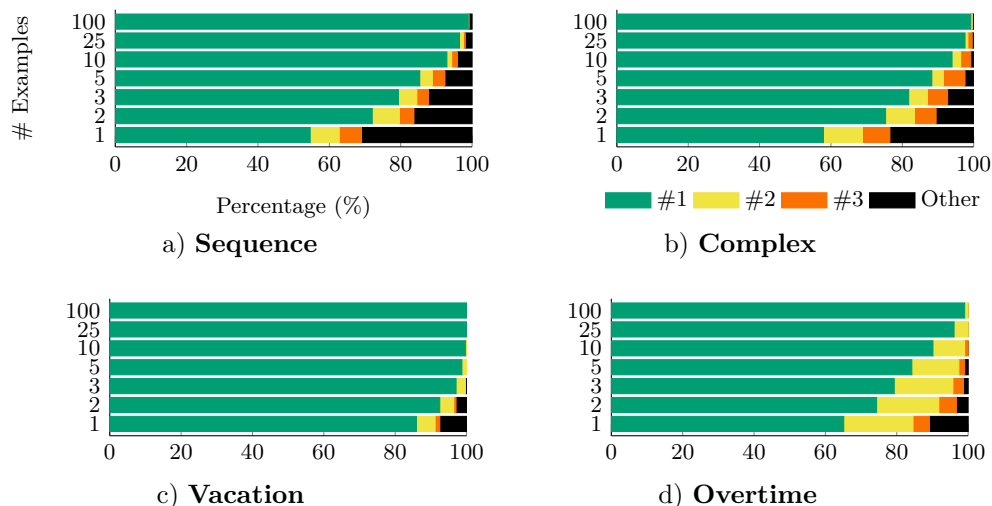
417 CABSC performs generally well as seen in Figure 3. Each benchmark presents a distinct
 418 behavior regarding the quality of the results. The first observable behavior is that CABSC
 419 succeeds in learning the CSP that was used to generate the data as seen with the benchmark
 420 **Sequence**. In this simpler case, the solver has to count the solutions of a conjunction of
 421 **SEQUENCE** constraints, i.e. the constraints to learn. With few examples, our approach stays
 422 coherent with the results of Picard-Cantin et al. [13] where they reach above 70% accuracy
 423 with a single example of solution, and around 85% accuracy with 5 examples. Extending the
 424 number of examples drastically reduces the margin of incorrectly learned instances while
 425 the number of examples needed is still relatively low. With only 25 examples, 96.47% of the
 426 instances resulted in a correctly learned **SEQUENCE** constraint at the first try. A few instances
 427 could not be resolved even with 100 examples. The unsolved instances occur when the solver
 428 finds a more restrictive constraint than the one that was used to generate the examples. This
 429 can happen if all the examples given are not enough to filter out parameters that would make
 430 the constraints more restrictive. This is why we see that with more examples given, fewer
 431 instances remain unsolved. The same phenomenon happens with the **Complex** benchmark
 432 where we see an efficient progression as the number of given examples increases.

433 Finding a more restrictive constraint is not the only way to get an incorrect model.
 434 As Figure 3 c) shows, the results for the **Vacation** benchmark converge toward a point
 435 where increasing the number of examples does not affect the results while still having a
 436 non-negligible proportion (8.82%) of unsolved instances. This is caused by multiple CSPs
 437 that are tied. A tie occurs when two distinct CSPs have the same number of solutions. In
 438 an instance from **Vacation**, the constraint we want to learn restricts 2 nurses to work a

27:12 Constraint Acquisition Based on Solution Counting

439 minimum of 3 days and a maximum of 4 days from Monday to Sunday. Since at least one
 440 nurse is required to work each day and that a nurse can work a maximum of 4 days within
 441 the week, the only way to satisfy the requirements is by having a first nurse working 4 days
 442 and the second nurse working 3 or 4 days. It is impossible for one of the nurses to work
 443 fewer than 3 days without violating the constraints. The problem comes when setting the
 444 value for the minimum number of days a nurse can work during the week. Consider a second
 445 selection of parameters where a minimum of 2 working days is required instead of 3. The
 446 same solutions are available since this change in parameters does not add solutions. The
 447 same goes with a minimum of 1 or 0 working day. This situation leads to four distinct CSPs
 448 with the same solution space. Since the objective is to find the CSP accepting the fewest
 449 solutions, these four CSPs are equivalent and the solver returns them in an arbitrary order.
 450 Most of the unsolved instances in the benchmark **Vacation** have the correct CSP in fourth
 451 position, which would have been first if the branching heuristics broke ties differently. We did
 452 not observe in our benchmarks situations where the solution spaces differ which let us believe
 453 that these models are equivalent. If we pretend for a moment that the **Vacation** benchmark
 454 was completed using heuristics that break ties without errors, we obtain the Figure 4.

Classification of the Instances for Various Number of Examples



■ **Figure 4** Hypothetical best results for each benchmark

455 Figure 4 shows that this hypothetical heuristic allows solving perfectly the **Vacation**
 456 benchmark using as few as 10 examples. Improvements are also present with the other
 457 benchmarks. This confirms that finding equivalent CSPs is the main reason why the solver
 458 does not succeed to correctly learn some CSPs.

459 The unsolved instances from the **Complex** benchmark are mainly caused by constraints
 460 found more restrictive than the correct one while the unsolved instances from the **Vacation**
 461 benchmark are mostly caused by equivalent CSPs. The unsolved instances of the **Overtime**
 462 benchmark are caused by a mix of these two reasons.

463 The final results show that our model can accurately learn the constraints even when the
 464 schedules contain vacations, overtime, or constraints that interfere with the constraints one
 465 wants to learn. Few examples are needed to obtain good results. These figures demonstrate
 466 that CABSC can learn constraints with the right parameters in diverse situations.

467 6.2 Execution Time


468 6.2.1 Using GANAK alone

469 For the **Complex** benchmark, model counting represents on average 93.6% of the time
 470 spent in the solver. Solution counting is a #P-difficult problem with few effective algorithms.
 471 Even with state-of-the-art tools, computing a lower bound on the number of solutions
 472 can take several minutes. The bound that took the longest time to compute by GANAK
 473 took 648 seconds. Figure 5 represents the time taken to solve all instances, i.e. the
 474 $(368 + 247 + 272 + 304) \times 7$ instances that come from the four benchmarks that were solved
 475 with 1, 2, 3, 5, 10, 25, and 100 examples using only GANAK as a model counter. Most of
 476 the instances are solved within a minute, but the solving time quickly and abruptly rises.
 477 This time limitation comes from a few main elements.

478 First, the size of the Meta-CSP greatly impacts the time needed for CABSC to find a
 479 solution. This size is measured in the number of parameter variables and activation variables
 480 since their number affects the depth of the search tree, thus the number of nodes explored
 481 in the branch and bound. For our instances, a few hundreds nodes could be observed on
 482 average resulting in around 30 to 60 unique calls to a model counter.

483 Second, the examples also impact the total runtime in two ways. With a higher number
 484 of examples, the solver is able to filter out more values from the domain of the parameter
 485 variables which directly decreases the number of potential calls to a model counter. Using a
 486 single example, the instances in the **Complex** benchmark takes on average 385.3 seconds to
 487 solve. With a hundred examples, the average time drops to 306.9 seconds, an improvement of
 488 20.35%. The second way the solving time is impacted by the examples is with their length, i.e.
 489 the number of decision variables. The more decision variables, the more Boolean variables
 490 in the SAT model to count. For this reason, we were not able to learn the constraints of
 491 schedules with a horizon of 56 days or more.

492 Lastly, all bounds do not take the same computation time. Indeed, we obtain SAT
 493 instances with various numbers of Boolean variables and clauses. The internal structure of
 494 these SAT instances can also vary. The bound that is the slowest to calculate uses a SAT
 495 instance with 672 Boolean variables and 1172 clauses and takes 648 seconds to count. The



■ **Figure 5** Measures of time for all instances using GANAK alone

496 Boolean model with the greatest number of variables has 804 variables and 2052 clauses and
 497 is counted in 0.11 seconds. This demonstrates that the counting time does not only depend
 498 on the number of decision variables, but also the structure of the problem.

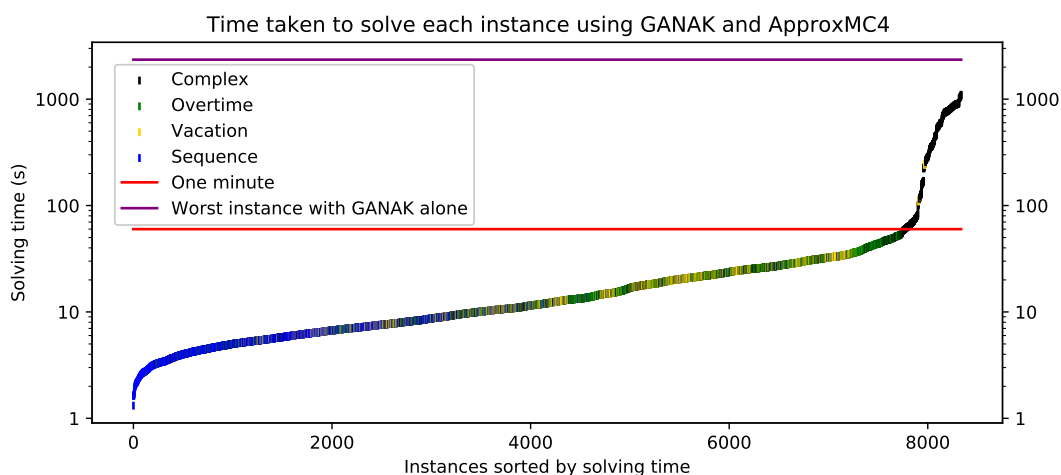
499 6.2.2 Using both GANAK and ApproxMC4

500 One method used to improve the time needed to solve a Meta-CSP is by combining a
 501 probabilistic exact model counter with an approximate model counter. This allows some
 502 CSP models to have their solutions counted quicker. The way ApproxMC4 was added to
 503 CabscSolver was to use it to prune CSP models from the search tree when the number of
 504 solutions was reasonably far from the number of solutions of the best CSP model found so
 505 far, as explained in Section 5.3.

506 This method is a lot faster than using GANAK as the only model counter as demonstrated
 507 by the Figure 6. The worst instance with GANAK alone lasted 2350 seconds while the same
 508 instance lasted 1099 seconds using ApproxMC4. The arithmetic average solving time of
 509 the **Complex** drops from 333.0 seconds to 158.7 seconds. This represents an improvement
 510 of 52.3% in average. The geometric average drops from 54.2 seconds to 41.0 seconds, an
 511 improvement of 24.4%.

512 The results obtained using both GANAK and ApproxMC4 have a lower accuracy by a
 513 small margin. While the accuracy of the results for the **Sequence**, **Vacation** and **Overtime**
 514 benchmarks remain unchanged, **Complex** suffers slight changes when few examples of
 515 solutions are given. Since the results have no significant differences to be seen on a graph,
 516 the changes are textually reported. With a single example of solution, the percentage of
 517 correctly learned CSP models drops from 48.99% to 48.48%. When using two examples of
 518 solutions, the percentage of correctly learned CSP models drops from 63.97% to 63.56% and
 519 with three examples, it drops from 69.64% to 68.83%. When using five examples of solutions
 520 or more, adding ApproxMC4 do not change the results anymore. All the other accuracy
 521 results are exactly the same, whether ApproxMC4 was used or not.

522 The lack of changes in the accuracy of **Sequence**, **Vacation** and **Overtime** benchmarks
 523 is mainly caused by the fact that ApproxMC4 returns approximations that are often too



■ **Figure 6** Measures of time for all instances using GANAK with ApproxMC4

524 close to take into account. The solver then has to ask GANAK to finish calculating the
 525 number of solutions of the CSP model regardless of the time needed by ApproxMC4. For the
 526 **Complex** benchmark, many CSP models were approximated by ApproxMC4 a lot faster
 527 than GANAK could and with values that allow pruning many nodes. ApproxMC4 sometimes
 528 overestimates the count of solutions outside the wanted interval of values. Since we used
 529 $\delta = 0.10$ for the model counters, ApproxMC4 therefore has a probability of at most 0.10
 530 to return values outside the wanted interval. This can cause many of the evaluations to
 531 accidentally prune correct CSP models, which can cause the Meta-CSP not to be properly
 532 solved. On the opposite side, it is possible to see improvements in the CSP learned due
 533 to overestimations that prune CSP models that would be learned if counted exactly. This
 534 happened on few instances from the **Complex** benchmark where the correct CSP went from
 535 being the third suggestion to the second. Since the correct CSP was not suggested as a first
 536 choice, the accuracy of correctly learned CSP models did not improve from these.

537 6.3 Potential Improvements

538 There exist several open source model counters that are efficient at counting SAT models, but
 539 fewer available programs to count the solutions of a CSP. Translating SEQUENCE constraints
 540 into pseudo-Boolean constraints and then to CNF offers no guarantee in the efficiency of the
 541 model. Directly counting the solution of a CSP could be faster and would certainly prevent
 542 from translating the model.

543 Parallelization could also speed up the exploration of the search tree. An approach like
 544 Embarassingly Parallel Search [15] could be appropriate, but also parallelization within the
 545 model counters would be suited as it is offered by ApproxMC3 [9, 18].

546 7 Conclusion

547 We introduced CABSC, a technique for Constraint Acquisition Based on Solution Counting.
 548 Our approach learns the CSP that accepts all provided examples but that minimizes the size
 549 of its solution space. This criterion has proven to return good solutions. The branch and
 550 bound uses model counters to compute a bound on the number of solutions for a given CSP.
 551 Experimental results show that CABSC successfully learns models and require few examples
 552 for our benchmarks.

553 ——— References ———

- 554 1 Hajar Ait Addi and Redouane Ezzahir. P_a -QUACQ: Algorithm for constraint acquisition system.
 555 In *Smart Data and Computational Intelligence*, pages 249–256, 2019.
- 556 2 Hajar Ait Addi, Christian Bessiere, Redouane Ezzahir, and Nadjib Lazaar. Time-bounded
 557 query generator for constraint acquisition. In *Integration of Constraint Programming, Artificial
 558 Intelligence, and Operations Research (CPAIOR 2018)*, pages 1–17, 2018.
- 559 3 Robin Arcangioli and Nadjib Lazaar. Multiple constraint acquisition. In *Proceedings of
 560 the 2015 International Conference on Constraints and Preferences for Configuration and
 561 Recommendation and Intelligent Techniques for Web Personalization*, pages 16—20, 2015.
- 562 4 Nicolas Beldiceanu and Évelyne Contejean. Introducing global constraints in chip. *Mathematical
 563 and Computer Modelling*, 20(12):97–123, 1994.
- 564 5 Nicolas Beldiceanu and Helmut Simonis. A constraint seeker: Finding and ranking global
 565 constraints from examples. In *Proceedings of the 17th International Conference on Principles
 566 and Practice of Constraint Programming (CP 2011)*, pages 12–26, 2011.

27:16 Constraint Acquisition Based on Solution Counting

- 567 6 Nicolas Beldiceanu and Helmut Simonis. A model seeker: Extracting global constraint models
568 from positive examples. In *Proceedings of the 18th International Conference on Principles and*
569 *Practice of Constraint Programming (CP 2012)*, pages 141–157, 2012.
- 570 7 Christian Bessiere, Clément Carbonnel, Anton Dries, Emmanuel Hebrard, George Katsirelos,
571 Nadjib Lazaar, Nina Narodytska, Claude-Guy Quimper, Kostas Stergiou, Dimosthenis C.
572 Tsouros, and Toby Walsh. Partial queries for constraint acquisition. Technical Report
573 abs/2003.06649, CoRR, 2020.
- 574 8 Christian Bessiere, Remi Coletta, Emmanuel Hebrard, George Katsirelos, Nadjib Lazaar, Nina
575 Narodytska, Claude-Guy Quimper, and Toby Walsh. Constraint acquisition via partial queries.
576 In Francesca Rossi, editor, *Proceedings of the 23rd International Joint Conference on Artificial*
577 *Intelligence (IJCAI-13)*, pages 475–481, 2013.
- 578 9 Supratik Chakraborty, Kuldeep S. Meel, and Moshe Y. Vardi. Algorithmic improvements
579 in approximate counting for probabilistic inference: From linear to logarithmic sat calls. In
580 *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI-16)*,
581 pages 3569–3576, 2016.
- 582 10 Abderrazak Daoudi, Younes Mechqrane, Christian Bessiere, Nadjib Lazaar, and El-Houssine
583 Bouyakhf. Constraint acquisition with recommendation queries. In *Proceedings of the 25th*
584 *International Joint Conference on Artificial Intelligence (IJCAI-16)*, pages 720–726, 2016.
- 585 11 Niklas Eén and Niklas Sörensson. Translating pseudo-boolean constraints into sat. *Journal on*
586 *Satisfiability, Boolean Modeling and Computation*, 2(1-4):1–26, 2006.
- 587 12 Nicholas Nethercote, Peter J. Stuckey, Ralph Becket, Sebastian Brand, Gregory J. Duck, and
588 Guido Tack. Minizinc: Towards a standard cp modelling language. In *Proceedings of the 13th*
589 *International Conference on Principles and Practice of Constraint Programming (CP 2007)*,
590 pages 529–543, 2007.
- 591 13 Émilie Picard-Cantin, Mathieu Bouchard, Claude-Guy Quimper, and Jason Sweeney. Learning
592 parameters for the sequence constraint from solutions. In *Proceedings of the 22nd International*
593 *Conference on Principles and Practice of Constraint Programming (CP 2016)*, pages 405–420,
594 2016.
- 595 14 Émilie Picard-Cantin, Mathieu Bouchard, Claude-Guy Quimper, and Jason Sweeney. Learning
596 the parameters of global constraints using branch-and-bound. In *Proceedings of the 23rd*
597 *International Conference on Principles and Practice of Constraint Programming (CP 2017)*,
598 pages 512–528, 2017.
- 599 15 Jean-Charles Régin, Mohamed Rezgui, and Arnaud Malapert. Embarrassingly parallel search.
600 In *Proceedings of the 19th International Conference on Principles and Practice of Constraint*
601 *Programming (CP 2013)*, pages 596–610, 2013.
- 602 16 Francesca Rossi, Peter van Beek, and Toby Walsh, editors. *Handbook of Constraint Program-*
603 *ming*, volume 2 of *Foundations of Artificial Intelligence*. Elsevier, 2006.
- 604 17 Shubham Sharma, Subhajit Roy, Mate Soos, and Kuldeep S. Meel. Ganak: A scalable
605 probabilistic exact model counter. In *Proceedings of the 28th International Joint Conference*
606 *on Artificial Intelligence (IJCAI-19)*, pages 1169–1176, 2019.
- 607 18 Mate Soos and Kuldeep S. Meel. Bird: Engineering an efficient cnf-xor sat solver and its
608 applications to approximate model counting. In *Proceedings of the 33rd AAAI Conference on*
609 *Artificial Intelligence (AAAI-19)*, pages 1592–1599, 2019.
- 610 19 Dimosthenis C. Tsouros, Kostas Stergiou, and Christian Bessiere. Structure-driven multiple
611 constraint acquisition. In *Proceedings of the 25th International Conference on Principles and*
612 *Practice of Constraint Programming (CP 2019)*, pages 709–725, 2019.
- 613 20 Dimosthenis C. Tsouros, Kostas Stergiou, and Panagiotis G. Sarigiannidis. Efficient methods
614 for constraint acquisition. In *Proceedings of the 24th International Conference on Principles*
615 *and Practice of Constraint Programming (CP 2018)*, pages 373–388, 2018.