

Towards Generalization in QBF Solving via Machine Learning

Mikoláš Janota

INESC-ID / Instituto Superior Técnico,
Universidade de Lisboa, Portugal
mikolas.janota@gmail.com

Abstract

There are well known cases of Quantified Boolean Formulas (QBFs) that have short winning strategies (Skolem/Herbrand functions) but that are hard to solve by nowadays solvers. This paper argues that a solver benefits from generalizing a set of individual wins into a strategy. This idea is realized on top of the competitive RAReQS algorithm by utilizing machine learning, which enables learning shorter strategies. The implemented prototype QFUN has won the first place in the non-CNF track of the most recent QBF competition.

1 Introduction

Against all odds posed by computational complexity, logic-based problem solving had a remarkable success at research but also at industrial level. One of the impressive success stories is the Boolean satisfiability problem (SAT). Quantified Boolean formulas (QBF) go one step further and extend SAT with quantification. This enables targeting a larger class of problems (Benedetti and Mangasarian 2008; Rintanen 2007; Schaefer and Umans 2002; Giunchiglia, Marin, and Narizzano 2009). However, success of QBF solvers comparable to SAT still seems quite far. Nevertheless, we have recently seen a significant progress in the area almost every year, e.g. (Zhang and Malik 2002; Biere 2004; Benedetti 2005; Tu, Hsu, and Jiang 2015; Goultiaeva and Bacchus 2010; Goultiaeva, Seidl, and Biere 2013; Janota and Marques-Silva 2015; Rabe and Tentrup 2015; Janota et al. 2012; Lonsing, Egly, and Seidl 2016; Tentrup 2016; Rabe and Seshia 2016). This paper aims to make a case for the use of machine learning *during* QBF solving.

It has been observed that search is often insufficient. A well-known example is the formula $\forall X \exists Y. \bigwedge x_i \leftrightarrow y_i$ with $X = \{x_1, \dots, x_n\}$ and $Y = \{y_1, \dots, y_n\}$ (Letz 2002). Traditional search will easily find an assignment (valuation) to X and Y satisfying the matrix (the propositional part). However, to prove that there is an assignment for Y given *any* assignment to X is difficult. Traditional search, even with various extensions, will try exponentially many assignments. A human can easily see why the formula is true. Indeed, given an arbitrary assignment to X , setting each y_i to x_i gives a witness for the validity of the formula.

It is useful to see QBFs as two-player games, where the existential player tries to make the formula true and the universal false. A winning strategy for the existential player shows that it is true. The formula above is a good example of a *small winning strategy*—the strategy for y_i is the function $s_{y_i}(x_1, \dots, x_n) \triangleq x_i$. The million dollar question here is, where do we get the strategies? This paper builds on the following idea: *Observe a set of assignments and learn from them strategies using machine learning. In another words, rather than looking at individual assignments, collect a set of them and generalize them into a strategy.*

Learning a strategy is not enough—it must also be incorporated into a solving algorithm. A straightforward approach would be to test for the learned strategy whether it is a winning one (that is possible with a SAT call (Kleine Büning, Subramani, and Zhao 2007)). However, this would put a lot of strain on the learning since we would have to be quite lucky to learn the right strategy and eventually we would have to deal with large training sets.

The algorithm presented in this paper takes inspiration in the existing algorithm RAReQS, which gradually expands the given formula by plugging in the encountered assignments (Janota et al. 2016). Instead of plugging in assignments, we will be plugging in the learned strategies. This forms the second main idea of the paper: *Expand the formula using strategies learned from collected samples and then start collecting a new set of samples.*

2 Preliminaries

A *literal* is a Boolean variable or its negation; complementary literal is denoted as \bar{l} , i.e. $\bar{\bar{x}} = x$, $\bar{\neg x} = x$. For a literal $l = x$ or $l = \neg x$, we write $\text{var}(l)$ for x . Analogously, $\text{vars}(\phi)$ is the set of all variables in formula ϕ . An *assignment* is a mapping from variables to Boolean constants 0, 1. For a formula ϕ and an assignment τ , we write $\phi[\tau]$ for the application of the assignments to ϕ .

2.1 Quantified Boolean Formulas

Quantified Boolean Formulas (QBFs) (Kleine Büning and Bubeck 2009) extend propositional logic by enabling quantification over Boolean variables. Any propositional formula ϕ is also a QBF with all variables *free*. If Φ is a QBF with a free variable x , the formulas $\exists x. \Phi$ and $\forall x. \Phi$ are QBFs

with x bound, i.e. not free. Note that we disallow expressions such as $\exists x.\exists x.x$, i.e., each variable is bound at most once. Whenever possible, we write $\exists x_1 \dots x_k$ instead of $\exists x_1 \dots \exists x_k$; analogously for \forall . For a QBF $\Phi = \forall x.\Psi$ we say that x is *universal* in Φ and is existential in $\exists x.\Psi$. Analogously, a literal l is universal (resp. existential) if $\text{var}(l)$ is universal (resp. existential).

Assignments also can be applied to QBF with $(Qx.\Phi)[\tau]$ defined as $\Phi[\tau]$ if x is in the domain of τ with $Q \in \{\forall, \exists\}$.

A QBF corresponds to a propositional formula: $\forall x.\Psi$ corresponds to $\Psi[x \rightarrow 0] \wedge \Psi[x \rightarrow 1]$ and $\exists x.\Psi$ to $\Psi[x \rightarrow 0] \vee \Psi[x \rightarrow 1]$. Since $\forall x \forall y.\Phi$ and $\forall y \forall x.\Phi$ are semantically equivalent, we allow QX for a set of variables X , $Q \in \{\forall, \exists\}$. A QBF with no free variables is *false* (resp. *true*), iff it is semantically equivalent to the constant 0 (resp. 1).

A QBF is *closed* if it does not contain any free variables. A QBF is in *prenex form* if it is of the form $Q_1 X_1 \dots Q_k X_k.\phi$, where $Q_i \in \{\exists, \forall\}$, $Q_i \neq Q_{i+1}$, ϕ propositional, and X_i pairwise disjoint sets of variables. The propositional part ϕ is called the *matrix* and the rest *prefix*. For a variable $x \in X_i$ we say that x is at *level* i and write $\text{lv}(x) = i$; we write $\text{lv}(l)$ for $\text{lv}(\text{var}(l))$. Unless specified otherwise, QBFs are assumed to be closed and in prenex form.

2.2 Games and Strategies

For most of the paper, QBFs are seen as two-player games. The existential player tries to make the matrix true and the universal player to make it false. A player assigns value only to variables that belong to the player and may assign a variable only once all variables that precede it in the prefix are assigned. In other words, the two players assign values following the order of the prefix. The game semantic perspective has the advantage that mostly we do not need to distinguish between the player \exists and \forall . Instead, we will be talking about a player and its opponent. **Notation:** We write Q for either of the players, and, \bar{Q} for its opponent.

Given a QBF $Q_1 X_1, \dots, Q_n X_n.\phi$ the *domain* $\text{dom}(x)$ of a variable $x \in X_k$ are all the variables in the preceding blocks, i.e. $\text{dom}(x) = \bigcup_{i \in 1..k-1} X_i$. A *play* is a sequence of assignments τ_1, \dots, τ_n where τ_i is an assignment to X_i .

Definition 1 For a QBF $Q_1 X_1, \dots, Q_n X_n.\phi$ a strategy for a variable $x \in X_k$ is a Boolean function s_x whose arguments are the variable's domain, i.e. $\text{dom}(x)$.

A strategy for a player Q is a set of strategies s_x for each of the variables $x \in Q X_i$. Whenever clear from the context, we simply say strategy for either of the concepts.

Notation. For the sake of succinctness, a strategy for a variable x is conflated with a Boolean formula whose truth value represents the value of the strategy. In another words, a strategy represents both some function $s_x : 2^{\text{dom}(x)} \mapsto \{0, 1\}$ and some formula ψ_x with $\text{vars}(\psi_x) \subseteq \text{dom}(x)$. This convention lets us also treat a set of strategies S for some variables X as a substitution. Hence, $\xi[S]$ represents the formula that results from simultaneously replacing in ξ each variable $x \in X$ with its strategy ψ_x .

Definition 2 (winning strategy) Let Ψ be a closed QBF $(QX \dots \phi)$ with ϕ propositional. A strategy S for \exists is win-

ning in Ψ if $\phi[S]$ is a tautology. A strategy S for \forall is winning in Ψ if $\phi[S]$ is unsatisfiable.

In particular, for a formula $\exists X.\phi$ a winning strategy for \exists corresponds to a satisfying assignment of ϕ .

Observation 1 A closed QBF Φ is true iff \exists has a winning strategy; it is false iff \forall has a winning strategy.

Definition 3 (winning/counter move) For a closed QBF $QX.\Phi$ an assignment τ to X is a winning move if there exists a winning strategy for Q in $\Phi[\tau]$.

For a closed QBF $QX\bar{Q}Y.\Phi$ and an assignment τ to X , an assignment μ to Y is a counter-move to τ if μ is a winning move for $\bar{Q}Y.\Phi[\tau]$.

Observation 2 There exists some winning move for QX in a formula $QX.\Phi$, if and only if there exists a winning strategy for Q in the formula.

Observation 3 For a formula $QX\bar{Q}Y.\Phi$, an assignment to X is a winning move iff there is not a counter-move to it.

3 Algorithm QFUN

As to make it a more pleasant read, this section comes in three installations, each bringing in more detail. The first part quickly overviews the existing algorithm (R)AREQS and sketches the main ideas of the proposed approach, which we will simply call the algorithm QFUN. The second part presents QFUN for the two-level case, i.e., formulas with one quantifier alternation. Finally, the third part details out the algorithm for the general case, i.e., formulas with arbitrary number of quantifier alternations.

3.1 Exposition

Let us quickly review the existing algorithm RAREQS (Janota et al. 2016). For a formula $QX\bar{Q}Y.\Phi$ RAREQS aims to decide whether there exists a winning move for Q . To that end, the algorithm keeps on constructing a sequence of pairs $(\tau_1, \mu_1), \dots, (\tau_k, \mu_k)$. Each τ_i is an assignment to X and μ_i is a counter-move to τ_i (see Def. 3). In each iteration, RAREQS constructs a partial expansion (called *abstraction*) of the original QBF such that no existing μ_i is a counter-move in the original formula to any winning move of the abstraction. In another words, if Q draws the next move so that it wins the abstraction, he is guaranteed not to be beaten by any of the existing counter-moves.

If there is no winning move for the abstraction, there isn't one for the original formula either and therefore there is no winning strategy for Q (we are done). If there is some winning move τ_{k+1} for the abstraction, check whether the opponent still comes up with a counter-move μ_{k+1} . If he does not, τ_{k+1} is a winning move for Q and we are done (see Observation 2). If a counter-move is found, the pair (τ_{k+1}, μ_{k+1}) is added to the sequence and the process repeats.

This setup inspires the use of machine learning. Since each μ_i is a counter-move to τ_i , the constructed sequence of pairs (τ_i, μ_i) can be conceived as a *training set* for the strategies for the variables Y (belonging to the player \bar{Q}). More specifically, for each variable $y \in Y$, the pair (τ_i, μ_i) represents a training sample for the function s_y prescribing that $s_y(\tau_i) = \mu_i(y)$. Observe that there might be other good

Algorithm 1: Playing 1-move multi-game

Function Wins1 ($QX. \{\phi_1, \dots, \phi_n\}$)
input : All ϕ_i propositional.
output: win. move for all $QX. \phi_i$, if exists; \perp otherwise.
1 $\alpha \leftarrow (Q = \exists) ? \bigwedge_{i \in 1..n} \phi_i : \bigwedge_{i \in 1..n} \neg \phi_i$
2 **return** SAT(α)

strategies for the opponent \bar{Q} . However, the pairs (τ_i, μ_i) have already *proven* to be good for Q and therefore we will stick to them.

It is tempting to learn a strategy for $\bar{Q}Y$ from such samples and then verify that it is a winning one. If it is a winning one, we would be done. If it is not a winning one, we could just learn a better one once we have more samples. However, this approach is unlikely to work. The problem with this approach is twofold. Firstly, it is overly optimistic to hope to hit the right strategy given a set of samples whose number is likely to be much smaller than the full truth table of the strategy. Secondly, it is putting too much strain on machine learning because the set of samples keeps on growing. Instead, this paper proposes the following schema.

1. Collect some suitable set of samples \mathcal{E} .
2. Learn strategies S for the opponent variables.
3. Strengthen the current abstraction using the strategies S .
4. Reset the set of samples \mathcal{E}
5. Repeat.

3.2 QFUN²: 2-level QBF

Let us look at the two-level case, i.e., a QBF of the form $QX\bar{Q}Y.\phi$. This form is particularly amenable to analysis since both the abstraction and candidate-checking is solvable by a SAT solver. Also, 2-level QBF has a number of interesting applications (cf. (Benedetti and Mangassarian 2008; Rintanen 2007; Schaefer and Umans 2002)).

A slight generalization of a game called a *multi-game* (Janota et al. 2016) is useful in the following presentation. A multi-game is a set of sub-games where the top-level player must find a move that is winning for all these sub-games at once. Note that a multi-game can be converted to a standard QBF by prenexing. However, it is useful to maintain this form (see (Janota et al. 2016, Sec. 4.1)).

Definition 4 (multi-game) A *multi-game* is written as $QX. \{\Phi_1, \dots, \Phi_k\}$. An assignment τ to X is a *winning move* for it iff it is a winning move for all $QX. \Phi_i$. Each Φ_i is called a *sub-game* and is either *propositional* or *begins with \bar{Q}* .

When all sub-games are propositional, the multi-game is solvable by a single SAT call. For such we introduce a function Wins1 (Algorithm 1). The function calculates a winning move for the multi-game or returns \perp if it does not exist (the function SAT has the same behavior). Observe that if the set of sub-games is *empty*, the formula α in Wins1 is

Algorithm 2: QFUN²: 2-level QBF Refinement with Learning

Function QFUN²($QX\bar{Q}Y.\phi$)
input : ϕ is propositional.
output: a win. move for QX if exists, \perp otherwise
1 $\mathcal{E} \leftarrow \emptyset$ // start with no samples
2 $\alpha \leftarrow \emptyset$ // empty abstraction
3 **while true do**
4 $\tau \leftarrow \text{Wins1}(QX.\alpha)$ // candidate
5 **if** $\tau = \perp$ **then return** \perp // loss
6 $\mu \leftarrow \text{Wins1}(\bar{Q}Y.\{\phi[\tau]\})$ // countermove
7 **if** $\mu = \perp$ **then return** τ // win
8 $\mathcal{E} \leftarrow \mathcal{E} \cup \{(\tau, \mu)\}$ // record sample
9 **if** ShouldLearn() **then**
10 $S \leftarrow \text{Learn}(\mathcal{E})$ // learn
11 $\alpha \leftarrow \alpha \cup \{\phi[S]\}$
12 $\mathcal{E} \leftarrow \emptyset$ // reset samples
13 **else**
14 $\alpha \leftarrow \alpha \cup \{\phi[\mu]\}$ // refine

the empty conjunction, which is equivalent to true, i.e., the SAT call then returns an arbitrary assignment.

Just as the existing algorithm AREQS, QFUN² (Algorithm 2) maintains an *abstraction* α . The abstraction corresponds to a partial expansion of the inner quantifier. Hence, for a formula $QX\bar{Q}Y.\phi$, the abstraction has the form $QX. \{\Phi[S] \mid S \in \omega\}$, where ω is some set of strategies. Observe that the abstraction is trivially equivalent to the original formula if ω contains all possible constant functions. For instance, $\forall u \exists e. \phi$ is equivalent to $\forall u. \phi[e \rightarrow 0] \vee \phi[e \rightarrow 1]$, which is equivalent to the multi-game $\forall u. \{\phi[e \rightarrow 0], \phi[e \rightarrow 1]\}$.

Example 1 Consider the formula $\forall uw \exists xy. \phi$ with $\phi = (u \Rightarrow (\neg w \Leftrightarrow x \wedge w \Leftrightarrow y)) \wedge (\neg u \Rightarrow (w \Leftrightarrow x \wedge \neg w \Leftrightarrow y))$. The following abstractions of this formula are both losing for \forall . With two sub-games: $\forall uw. \{\phi[x \rightarrow \neg w, y \rightarrow w], \phi[x \rightarrow w, y \rightarrow \neg w]\}$; with single sub-game: $\forall uw. \{\phi[x \rightarrow (u ? \neg w : w), y \rightarrow (u ? w : \neg w)]\}$.

The abstraction α is *refined* with every play losing for Q , which effectively means adding a subgame to the current abstraction. Additionally, QFUN² maintains a set of samples \mathcal{E} . The samples are pairs (τ_i, μ_i) such that τ_i, μ_i is a losing play for Q , i.e., a winning play for \bar{Q} . So for instance, if $Q = \forall$ then $\bar{Q} = \exists$ and $\tau_i \cup \mu_i \models \phi$.

Both the abstraction α and samples \mathcal{E} are initialized as empty. In each iteration, QFUN² calls Wins1 to calculate a candidate for a winning move τ . Subsequently, another call to Wins1 is issued to calculate a counter-move μ . If either candidate or counter-move does not exist, one of the player has lost without recovery.

Machine learning is invoked only ever so often. To decide when, the pseudo-code queries the function ShouldLearn. Whenever ShouldLearn is true, new strategies are learned for Y -variables based on the sam-

ples \mathcal{E} . These strategies are plugged into the formula ϕ and recorded in the abstraction. Then, and that is crucial, the set of samples is reset back to the empty set.

In terms of soundness, the set of samples \mathcal{E} need not be reset after each learning. However, it is crucial in terms of performance. If the set is always augmented, learning will become overly time consuming. Recall that a strategy needs to be learned for each opponents' variable and further, the number of iterations can go to millions.

So what does the abstraction represent and what is the role of the learned strategies? The original AReQS adds a sub-game $\phi[\mu_i]$ for each existing counter-move μ_i . Intuitively, this means that the player Q never plays before making sure that he can successfully defend himself against all the existing counter-moves. Once strategies are also included, the player Q also defends himself against all the strategies devised so far. Strategy-based refinement is a generalization of the traditional one—the traditional refinement corresponds to a set of strategies comprising constant functions.

What do we require from the strategy learning? The good news is that in fact very little. A strategy must be learned in the form of a formula so that it can be plugged into the original formula. This means that the algorithm does not easily allow for neural networks, for instance. The current implementation uses decision-trees (see Section 4). For the sake of soundness, the learned strategy formula *must follow* the definition of a strategy (Definition 1). In practice this means that a strategy formula ξ_y for a variable y must only contain variables from $\text{dom}(y)$.

If the function `ShouldLearn` triggers traditional refinement only finitely many times, the learning method also needs to guarantee termination of the whole algorithm. A natural minimal requirement for this is that the learned strategies will correspond to at least one sample $(\tau_i, \mu_i) \in \mathcal{E}$, i.e. $s_y(\tau_i) = \mu_i(y)$, for each $y \in Y$. This requirement guarantees that τ_i will not appear as a candidate for a winning move in the upcoming iterations. Nevertheless, if `ShouldLearn` alternates between traditional and learning-based refinement, termination is already guaranteed by the traditional refinement and we do not need to worry about what is learned as long as it is sound.

3.3 QFUN: General Case

The general case QFUN generalizes the two-level case QFUN² using recursion (just as RAREQS generalizes AReQS). The basic ideas remain, even though we are faced with a couple of technical complications. The pseudocode is presented as Algorithm 3. Since the abstraction is a multi-game, the recursive call also needs to handle a multi-game. For this purpose, we maintain a set of sequences of samples—each sequence for each given sub-game. Candidates for a winning-move are drawn from the abstraction α by a recursive call. The small technical difficulty here is that the abstraction may return a winning move containing some extra fresh variables coming from refinement. Hence, these need to be filtered out (ln. 8).

If, the candidate move τ is a winning move, it is returned. If, however, there is some counter-move μ_i obtained by playing a sub-game Φ_i , it is used for refinement. This means in-

Algorithm 3: QBF Refinement with Learning

```

Function QFUN( $QX. \{\Phi_1, \dots, \Phi_n\}$ )
input : Each  $\Phi_i$  is propositional or begins with  $\bar{Q}Y$ .
output: a win. move for  $QX$  if exists,  $\perp$  otherwise

1 if all  $\Phi_i$  propositional then
2   return Wins1( $QX. \{\Phi_1, \dots, \Phi_n\}$ )
3  $\mathcal{E}_i \leftarrow \emptyset, i \in 1..n$  // samples
4  $\alpha \leftarrow QX.\emptyset$  // empty abstraction
5 while true do
6    $\tau' \leftarrow$  QFUN( $\alpha$ ) // candidate
7   if  $\tau' = \perp$  then return  $\perp$  // loss
8    $\tau \leftarrow \{l \mid l \in \tau' \wedge \text{var}(l) \in X\}$  // filter
9   if all QFUN( $\Phi_i[\tau]$ ) =  $\perp$  then return  $\tau$  // win
10  let  $l$  be s.t. QFUN( $\Phi_l[\tau]$ ) =  $\mu$  for some
     $l \in \{1..n\}, \mu \neq \perp$ 
11   $\mathcal{E}_l \leftarrow \mathcal{E}_l \cup \{(\tau, \mu)\}$  // record sample
12  if ShouldLearn() then
13     $S \leftarrow$  Learn( $\mathcal{E}_l$ ) // learn
14     $\alpha \leftarrow$  Refine( $\alpha, \Phi_l, S$ )
15     $\mathcal{E}_l \leftarrow \emptyset$  // reset samples
16  else
17     $\alpha \leftarrow$  Refine( $\alpha, \Phi_l, \mu$ ) // refine

```

serting the pair (τ, μ_i) into the sample sequence pertaining to this sub-game, i.e. sequence \mathcal{E}_i , and subsequently, performing refinement. In order to ensure that quantifiers alternate, refinement introduces fresh variables for formulas with more than 2 levels. The refinement function is defined as follows.

$$\text{Refine}(QX.\{\Psi_1, \dots, \Psi_n\}, \bar{Q}YQX_1.\Psi, S) := QXX'_1.\{\Psi_1, \dots, \Psi_n, \Psi'[S]\}$$

$$\text{Refine}(QX.\{\Psi_1, \dots, \Psi_n\}, \bar{Q}Y.\psi, S) := QX.\{\Psi_1, \dots, \Psi_n, \psi[S]\}$$

where X'_1 are fresh duplicates of the variables X_1 and Ψ' is Ψ with X_1 replaced by X'_1 and where ψ is a propositional formula.

4 Implementation

4.1 Formula representation

The algorithm requires nontrivial formula manipulation to achieve refinement. Performing these operations directly on a CNF representation is difficult and further, CNF representation as input has well-known pitfalls (Ansótegui, Gomes, and Selman 2005; Janota and Marques-Silva 2017). Hence, the implementation represents formulas as *And-Inverter graphs* (AIG) (Hellerman 1963), which are simplified by trivial non-invasive simplifications (Brummayer and Biere 2006). All the logical operations (e.g. substitution/conjunction) are performed on AIGs. Only when the time comes to call a SAT solver, the AIG is translated into CNF. This is done in straightforward fashion. Each sub-AIG is mapped to an encoding Boolean variable in the SAT solver. Since the AIGs are hash-consed, each sub-AIG also corresponds to just one variable. All the and-gates are bi-

nary. The input to the solver is the circuit-like format for QBF called *QCIR* (Jordan, Klieber, and Seidl 2016).

4.2 Learning

Recall that learning is invoked with the sequence of pairs of assignments $\mathcal{E} = (\tau_1, \mu_1), \dots, (\tau_k, \mu_k)$, where each τ_i is an assignment to some block of variables X in the prefix and μ_i is an assignment to variables Y , which is the adjacent block in the prefix, belonging to the opposing player.

The objective is to learn a strategy (a function) for each of the variables in Y . A Boolean function can be seen as a *classifier* with two classes: the input assignments where the strategy should return 1 (true) and the input assignments where the strategy should return 0 (false). The implementation uses the popular classifier *Decision trees* (Russell and Norvig 2010). These are constructed by the standard *ID3* algorithm (Quinlan 1986).

For each variable in $y \in Y$, construct the training set \mathcal{E}_y from \mathcal{E} by ignoring all the other Y variables. Subsequently invoke *ID3* on \mathcal{E}_y thus obtaining a decision tree conforming to the sample assignments. Once a decision-tree is constructed, the Boolean formula is constructed as follows.

1. Construct the sets of conjunctions of literals \mathcal{I}_p and \mathcal{I}_n corresponding to the positive and negative branches of the tree, respectively. Hence, if $t \in \mathcal{I}_p$ is true, the tree gives 1.
2. Repeatedly apply subsumption and self-subsumption on each set \mathcal{I}_p and \mathcal{I}_n , until a fixed point is reached.
3. If $|\mathcal{I}_p| < |\mathcal{I}_n|$ return $\bigvee \mathcal{I}_p$, otherwise return $\neg \bigvee \mathcal{I}_n$.

Step 2 would not necessarily be needed but since we are substituting the constructed functions into the input formula, it is desirable to maintain them small. Analogously, either set could be chosen in step 3 but a smaller is preferable. Some heavier methods could be considered here, cf. (Ignatiev, Previti, and Marques-Silva 2015).

When to learn? It is a bad idea to learn too frequently since this would produce poor sample-sets to learn from. However, learning too *infrequently* has two main pitfalls:

1. Learning on large sample-sets will be too costly (recall that a learning algorithm is run for each opponent variable upon refinement).
2. There is a risk of very complicated and therefore large functions to be learned from complicated samples

A straightforward approach was taken to implement the function `ShouldLearn`: learning is triggered every K iterations of the loop, where K is a parameter of the solver. The number of iterations is considered local for each recursive call of `QFUN`. The experimental evaluation examines the solver’s behavior for several values of K (see Section 5).

4.3 Strategy accumulation

Upon each refinement the set of samples is reset. Also, whatever is learned is forgotten in the next rounds—learning starts from scratch on a new set of samples. This might be disadvantageous. The current implementation uses a simple but important improvement. The algorithm records for each

Solver	Quabs	GQ	RAReQS	L-16	L-64	L-128	L-64-f
Solved (320)	103	75	105	110	111	111	104
Wins	63	11	67	55	63	62	60

Table 1: Result summary. A win is counted also for solvers that are not worse than the best time by 1s.

variable y the last learned strategy. This strategy is then evaluated on the next batch of samples when learning is invoked again. If it still fits the data, it is kept. Otherwise it is discarded and a new strategy is learned.

Example 2 Consider the formula from the introduction of the paper: $\forall x_1, \dots, x_n \exists y_1, \dots, y_n. \bigwedge x_i \leftrightarrow y_i$, and, the following sequence of samples.

x_1	x_2	...	x_n	y_1	y_2	...	y_n
0	0	...	0	0	0	...	0
1	0	...	0	1	0	...	0
0	0	...	1	0	0	...	1
0	1	...	1	0	1	...	1

If $K = 2$, the first application of learning gives $y_1 \triangleq x_1$ and the rest of the strategies are constantly 0. In the second refinement, learning gives $y_2 \triangleq x_2$ and the rest constants. If, however, we keep the information from the previous learning, we get both $y_1 \triangleq x_1, y_2 \triangleq x_2$. Hence, accumulating the individual strategies will eventually yield the right strategy.

Some similar functions were in fact learned during the experimental part. For instance, I have observed the solver learn $y \triangleq x$ for $\forall x \dots \exists y \dots (F \wedge (y \Rightarrow (x \wedge G)) \wedge (H \vee \neg x \vee y))$ for some larger formulas F, G, H . Theoretically, learning can be worse than traditional refinement, e.g., $\dots \exists x_1 \dots x_{100}. (x_1 \wedge \dots \wedge x_{100}) \vee G$ it is clearly better refine with $x_i \triangleq 1$ rather than some complicated functions.

4.4 Incrementality

The recursive structure of the algorithm is very elegant but might be too forgetful. If one is to solve $\Phi[\mu_i]$, it could be useful to maintain the abstraction from that solving in order to solve $\Phi[\mu_{i+1}]$. The issue is that then the solvers tend to occupy too much space. Currently, the solver maintains only abstractions that are purely propositional.

5 Experimental Evaluation

The RAReQS algorithm has proven to be highly competitive as it have placed first in several tracks of the recent QBF competitions. So the key question is whether RAReQS benefits, or may benefit, from the proposed learning.

The success of the machine learning techniques can be assessed at various levels. The lowest bar is whether the technique is at all *computationally feasible*. Indeed, it might be that the learning is impractically time-consuming. Second step is whether the *number of iterations decreases when learning is applied*. The third step is whether also *solving time decreases when learning is applied*. Finally, we are interested in variations of the algorithm. Namely, the effect of the learning interval and the effect of the technique of accumulating strategies (see Section 4.3).

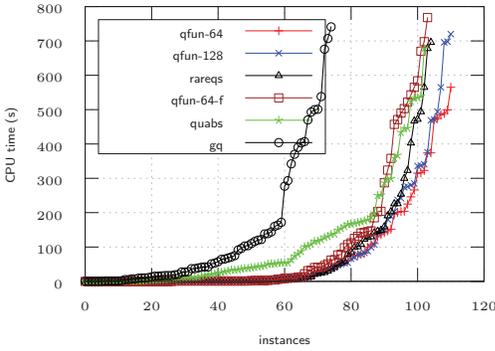


Figure 1: Cactus plot. A point at (x, y) means that the solver solved x instances each within y sec.

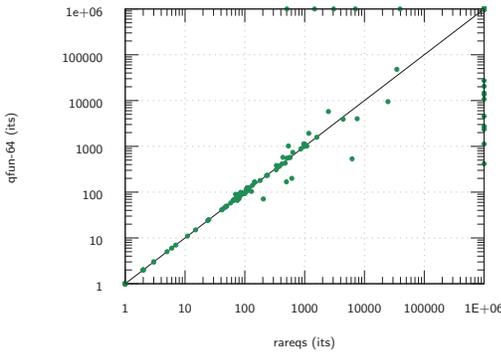


Figure 2: Log-scale scatter plot for #iterations.

The evaluation considers the following configurations of algorithm QFUN (Algorithm 3): QFUN *without* any learning, which is in the fact RAREQS; versions QFUN-16, QFUN-64, and QFUN-128 where learning is triggered every 16/64/128 iterations, respectively; QFUN-64-f forgetful version of QFUN where previously learned strategies are not used in the future. All the other versions accumulate strategies as described in Section 4.3. Additionally we compare to the highly competitive non-CNF solvers GhostQ (Klieber et al. 2010) and QuAbS (Tentrup 2016).

The evaluated prototype is implemented in C^{++} and `minisat 2.2` (Eén and Sörensson 2003) is used as the backend solver. The experiments were carried out on Linux machines with Intel Xeon 5160 3GHz processors and 4GB of memory with the time limit 800 s and memory limit 2GB. For the evaluation we used the non-CNF suite from the *2017 QBF Competition* counting 320 instances.¹

The overall results are summarized in Table 1. The cactus plot in Fig. 1a summarizes the performance. For the sake of readability, the cactus plots omits QFUN-16, whose performance is quite similar to QFUN-64 and QFUN-128, which are already quite close. Fig. 1b is a scatterplot comparing the total number of refinements for QFUN-64 and RAREQS, i.e., machine learning every 64 iterations versus

no learning; all instances are displayed and time/mem-outs are placed on the edges. The learning time observed was the relatively small, for instance on `CM-sat-07-01-07-3` with 425+917 variables, each learning round took 0.8 s with 200 s being the total solving time.

5.1 Results Discussion

Overall, learning gives improvement both in terms of number of solved instances as well as number of iterations. Admittedly, in terms of number of solved instances the gain is modest. However, the difference in performance between RAREQS and QuAbS is even smaller despite each representing a completely different algorithm. Also recall that Fig. 1b is in logarithmic scale so the number of iterations saved are in number of cases in orders of magnitude. Overall this suggests that adding learning in the brings about a new quality in the solver.

The effect of frequency of learning on the performance is relatively small. The best configuration is with learning every 64 refinements (QFUN-64), while QFUN-16 and QFUN-128 perform slightly worse. This is not surprising as too frequent learning will slow down the solving and too infrequent does not give enough opportunity to learn.

The biggest effect has strategy accumulation. Indeed, without it, learning in fact performs worse then without any learning. This suggests that at least for some variables it is important to learn a certain strategy and maintain it. This observation clearly opens opportunities for further investigation as the techniques of accumulating strategies can be further developed.

6 Related Work

The research on QBF solving has been quite active in the last decades and an array of approaches exists. It appears that these different approaches also give us a different classes of instances where they are successful. One of the oldest approaches is conflict/solution learning (Zhang and Malik 2002; Lonsing and Biere 2010; Giunchiglia, Marin, and Narizzano 2010; Lonsing 2012), which essentially generalizes clause learning in SAT. Then there are solvers that perform quantifier expansion into Boolean connectives (Biere 2004; Benedetti 2005; Lonsing and Biere 2008; Pigorsch and Scholl 2010; Tu, Hsu, and Jiang 2015); solvers that target non-CNF inputs (Zhang 2006; Klieber et al. 2010; Goultiaeva and Bacchus 2010; Goultiaeva, Seidl, and Biere 2013; Van Gelder 2013; Balabanov et al. 2016; Tentrup 2016); and solvers that calculate blocking clauses using a SAT solver (Ranjan, Tang, and Malik 2004; Janota and Marques-Silva 2015; Rabe and Tentrup 2015). Recently we have also seen integration of inprocessing with conflict/solution learning (Lonsing, Egly, and Seidl 2016).

This paper builds on the algorithm RAREQS (Janota et al. 2016), which expands quantifiers gradually by substituting them one by one into the formula. This approach is conceptually akin to the *model-based quantifier instantiation* (Wintersteiger, Hamadi, and de Moura 2013).

It is known that QBF solvers *implicitly* trace strategies because a winning strategy can be extracted once the formula is solved (Goultiaeva, Van Gelder, and Bacchus 2011;

¹www.qbflib.org/event_page.php?year=2017

Balabanov and Jiang 2012; Balabanov et al. 2015; Beyersdorff, Chew, and Janota 2014). However, to our best knowledge there are currently only two QBF solvers that *explicitly* target strategy computation. In (Bjørner, Janota, and Klieber 2015) the authors fused clause learning and RAReQS by refining abstractions with strategies calculated from clause learning—with not very promising results. The second solver by Rabe and Seshia works in the context of 2QBF and gradually adds variables to a winning strategy of the inner quantifier (Rabe and Seshia 2016).

It is hard to do justice to the work that has been done in machine learning, the reader is directed to standard literature (Russell and Norvig 2010). It should be mentioned that strategy learning is a very specific type of learning because we need the result in the form of a formula. This is closely related to function synthesis/learning cf. (Valiant 1984; Kamath et al. 1992; Oliveira and Sangiovanni-Vincentelli 1993; Su et al. 2016). Machine learning has also been used in portfolio solvers e.g. (Xu et al. 2008).

Last but not least, machine learning has been used at a higher level of inference to discover lemmas in the context of first order or higher order reasoning (Urban et al. 2008; Kaliszzyk and Urban 2014).

7 Conclusion and Future Work

This paper presents a QBF solver that periodically *generalizes* a set of observations (plays) into a strategy by machine learning. These strategies are plugged into the original formula in order to gradually strengthen a partial expansion of the formula. The results show that this is feasible and it also helps to reduce the number of refinement iterations but also the solving time. The fact that this results in a competitive QBF solver is already compelling. Indeed, machine learning is invoked many times during solving on a number of variables separately. However, the design of the algorithm enables us to curb the computational burden of machine learning by limiting the size of the training set.

As discussed in Section 4, the current prototype is rather straightforward in its implementation decisions. There is a lot of room for making the solver more intelligent. Besides inprocessing and other implementation issues, number of things are to be investigated for the machine learning part. What kind of machine learning methods are good for this purpose? When to trigger machine learning? Can we improve the training sets (e.g. introduction of don't-cares)?

Another interesting question for future work is whether machine learning can be beneficial in other type of QBF solving. There are opportunities for this. Even if the solver is *not* performing expansion-based refinement (e.g. CAQE (Rabe and Tentrup 2015), QUESTO (Janota and Marques-Silva 2015), CADET (Rabe and Seshia 2016)), it can for instance use a learned strategy to predict the behavior of the opponent.

At the theoretical level, the paper touches a fundamental question: how difficult is it to learn the right strategies? Here, PAC-learnability could give some answers (Valiant 1984).

Acknowledgments. This work was supported by national funds through Fundação para a Ciência e a Tecnologia

(FCT) with reference UID/CEC/50021/2013.

References

- Ansótegui, C.; Gomes, C. P.; and Selman, B. 2005. The Achilles' heel of QBF. In *National Conference on Artificial Intelligence and the Seventeenth Innovative Applications of Artificial Intelligence Conference (AAAI)*, 275–281.
- Balabanov, V., and Jiang, J.-H. R. 2012. Unified QBF certification and its applications. *Formal Methods in System Design* 41(1):45–65.
- Balabanov, V.; Jiang, J. R.; Janota, M.; and Widl, M. 2015. Efficient extraction of QBF (counter)models from long-distance resolution proofs. In *Conference on Artificial Intelligence (AAAI)*, 3694–3701.
- Balabanov, V.; Jiang, J. R.; Mishchenko, A.; and Scholl, C. 2016. Clauses versus gates in CEGAR-Based 2QBF solving. In *Beyond NP, AAAI Workshop*.
- Benedetti, M., and Mangassarian, H. 2008. QBF-based formal verification: Experience and perspectives. *Journal on Satisfiability, Boolean Modeling and Computation (JSAT)* 5(1-4):133–191.
- Benedetti, M. 2005. sKizzo: a suite to evaluate and certify QBFs. In *International Conference on Automated Deduction (CADE)*, 369–376.
- Beyersdorff, O.; Chew, L.; and Janota, M. 2014. On unification of QBF resolution-based calculi. In *International Symposium on Mathematical Foundations of Computer Science (MFCS)*, 81–93. Springer.
- Biere, A.; Heule, M.; van Maaren, H.; and Walsh, T., eds. 2009. *Handbook of Satisfiability*, volume 185 of *Frontiers in Artificial Intelligence and Applications*. IOS Press.
- Biere, A. 2004. Resolve and expand. In *Theory and Applications of Satisfiability Testing (SAT)*, 238–246.
- Bjørner, N.; Janota, M.; and Klieber, W. 2015. On conflicts and strategies in QBF. In *International Conferences on Logic for Programming LPAR-20, Short Presentations*.
- Brummayer, R., and Biere, A. 2006. Local two-level and-inverter graph minimization without blowup. *Proc. MEMICS* 6:32–38.
- Eén, N., and Sörensson, N. 2003. An extensible SAT-solver. In *Theory and Applications of Satisfiability Testing (SAT)*.
- Giunchiglia, E.; Marin, P.; and Narizzano, M. 2009. Reasoning with quantified boolean formulas. In Biere et al. (2009). 761–780.
- Giunchiglia, E.; Marin, P.; and Narizzano, M. 2010. QuBE 7.0 system description. *Journal on Satisfiability, Boolean Modeling and Computation* 7:83–88.
- Goultiaeva, A., and Bacchus, F. 2010. Exploiting QBF duality on a circuit representation. In *Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI)*.
- Goultiaeva, A.; Seidl, M.; and Biere, A. 2013. Bridging the gap between dual propagation and CNF-based QBF solving. In *Design, Automation and Test in Europe (DATE)*, 811–814.
- Goultiaeva, A.; Van Gelder, A.; and Bacchus, F. 2011. A uniform approach for generating proofs and strategies for both

- true and false QBF formulas. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 546–553.
- Hellerman, L. 1963. A catalog of three-variable or-invert and and-invert logical circuits. *IEEE Trans. Electronic Computers* 12(3):198–223.
- Ignatiev, A.; Previti, A.; and Marques-Silva, J. 2015. SAT-based formula simplification. In *Theory and Applications of Satisfiability Testing (SAT)*, 287–298.
- Janota, M., and Marques-Silva, J. 2015. Solving QBF by clause selection. In *International Conference on Artificial Intelligence (IJCAI)*, 325–331. AAAI Press.
- Janota, M., and Marques-Silva, J. 2017. An Achilles’ heel of term-resolution. In *Conference on Artificial Intelligence (EPIA)*, 670–680.
- Janota, M.; Klieber, W.; Marques-Silva, J.; and Clarke, E. M. 2012. Solving QBF with counterexample guided refinement. In *Theory and Applications of Satisfiability Testing (SAT)*, 114–128.
- Janota, M.; Klieber, W.; Marques-Silva, J.; and Clarke, E. 2016. Solving QBF with counterexample guided refinement. *Artificial Intelligence* 234:1–25.
- Jordan, C.; Klieber, W.; and Seidl, M. 2016. Non-CNF QBF solving with QCIR. In *Proceedings of BNP (Workshop)*.
- Kaliszyk, C., and Urban, J. 2014. Learning-assisted automated reasoning with Flyspeck. *Journal of Automated Reasoning* 53(2):173–213.
- Kamath, A. P.; Karmarkar, N.; Ramakrishnan, K. G.; and Rensde, M. G. C. 1992. A continuous approach to inductive inference. *Math. Program.* 57:215–238.
- Kleine Büning, H., and Bubeck, U. 2009. Theory of quantified boolean formulas. In Biere et al. (2009). 735–760.
- Kleine Büning, H.; Subramani, K.; and Zhao, X. 2007. Boolean functions as models for quantified boolean formulas. *J. Autom. Reasoning* 39(1):49–75.
- Klieber, W.; Sapra, S.; Gao, S.; and Clarke, E. M. 2010. A non-prenex, non-clausal QBF solver with game-state learning. In *SAT*, 128–142.
- Letz, R. 2002. Lemma and model caching in decision procedures for quantified boolean formulas. In *Automated Reasoning with Analytic Tableaux and Related Methods (TABLEAUX)*, 160–175.
- Lonsing, F., and Biere, A. 2008. Nenofex: Expanding NNF for QBF solving. In *Theory and Applications of Satisfiability Testing (SAT)*, 196–210.
- Lonsing, F., and Biere, A. 2010. DepQBF: A dependency-aware QBF solver. *Journal on Satisfiability, Boolean Modeling and Computation (JSAT)* 7(2-3):71–76.
- Lonsing, F.; Egly, U.; and Seidl, M. 2016. Q-resolution with generalized axioms. In *Theory and Applications of Satisfiability Testing (SAT)*, 435–452.
- Lonsing, F. 2012. *Dependency Schemes and Search-Based QBF Solving: Theory and Practice*. Ph.D. Dissertation, Johannes Kepler Universität.
- Oliveira, A. L., and Sangiovanni-Vincentelli, A. L. 1993. Learning complex boolean functions: Algorithms and applications. In *Advances in Neural Information Processing Systems NIPS*, 911–918.
- Pigorsch, F., and Scholl, C. 2010. An AIG-based QBF-solver using SAT for preprocessing. In *Design Automation Conference, DAC*, 170–175.
- Quinlan, J. R. 1986. Induction of decision trees. *Machine learning* 1(1):81–106.
- Rabe, M. N., and Seshia, S. A. 2016. Incremental determination. In *Theory and Applications of Satisfiability Testing (SAT)*, 375–392.
- Rabe, M. N., and Tentrup, L. 2015. CAQE: A certifying QBF solver. In *Formal Methods in Computer-Aided Design, FMCAD*, 136–143.
- Ranjan, D. P.; Tang, D.; and Malik, S. 2004. A comparative study of 2QBF algorithms. In *Theory and Applications of Satisfiability Testing (SAT)*, 292–305.
- Rintanen, J. 2007. Asymptotically optimal encodings of conformant planning in QBF. In *AAAI Conference on Artificial Intelligence*, 1045–1050. AAAI Press.
- Russell, S. J., and Norvig, P. 2010. *Artificial intelligence: a modern approach*. Prentice Hall.
- Schaefer, M., and Umans, C. 2002. Completeness in the polynomial-time hierarchy: A compendium. *SIGACT news* 33(3):32–49.
- Su, G.; Wei, D.; Varshney, K. R.; and Malioutov, D. M. 2016. Learning sparse two-level boolean rules. In *IEEE 26th International Workshop on Machine Learning for Signal Processing (MLSP)*, 1–6.
- Tentrup, L. 2016. Non-prenex QBF solving using abstraction. In *Theory and Applications of Satisfiability Testing (SAT)*, 393–401.
- Tu, K.; Hsu, T.; and Jiang, J. R. 2015. QELL: QBF reasoning with extended clause learning and leveled SAT solving. In *Theory and Applications of Satisfiability Testing (SAT)*, 343–359.
- Urban, J.; Sutcliffe, G.; Pudlák, P.; and Vyskočil, J. 2008. MaLAREa SG1- machine learner for automated reasoning with semantic guidance. In *International Joint Conference on Automated Reasoning (IJCAR)*, 441–456.
- Valiant, L. G. 1984. A theory of the learnable. *Commun. ACM* 27(11):1134–1142.
- Van Gelder, A. 2013. Primal and dual encoding from applications into quantified boolean formulas. In *CP*, 694–707.
- Wintersteiger, C. M.; Hamadi, Y.; and de Moura, L. M. 2013. Efficiently solving quantified bit-vector formulas. *Formal Methods in System Design* 42(1):3–23.
- Xu, L.; Hutter, F.; Hoos, H. H.; and Leyton-Brown, K. 2008. SATzilla: Portfolio-based algorithm selection for SAT. *J. Artif. Intell. Res.* 32:565–606.
- Zhang, L., and Malik, S. 2002. Conflict driven learning in a quantified Boolean satisfiability solver. In *International Conference On Computer Aided Design (ICCAD)*, 442–449.
- Zhang, L. 2006. Solving QBF by combining conjunctive and disjunctive normal forms. In *National Conference on Artificial Intelligence and the Eighteenth Innovative Applications of Artificial Intelligence Conference (AAAI)*.