

Symbolic Computation for All the Fun

Chad E. Brown¹, Mikoláš Janota¹ and Mirek Olšák²

¹Czech Technical University in Prague, CIIRC, Czechia

²University of Cambridge

Abstract

Motivated by the recent 10 million dollar AIMO challenge, this paper targets the problem of finding all functions conforming to a given specification. This is a popular problem at mathematical competitions and it brings about a number of challenges, primarily, synthesizing the possible solutions and proving that no other solutions exist. Often, there are infinitely many solutions and then the set of solutions has to be captured symbolically. We propose an approach to solving this problem and evaluate it on a set of problems that appeared in mathematical competitions and olympics.

Keywords

AIMO, synthesis, quantifier elimination, SMT

1. Introduction

The *AIMO challenge*¹ promises to award \$10 million to an AI model that will be able to win a gold medal at the *International Mathematical Olympiad (IMO)*². In this paper we draw attention to the fact that symbolic computation tools can give a significant boost to a system aiming at this challenge. We target a popular problem encountered in mathematical competitions, which is finding *all* functions adhering to a certain specification. The general task appears near impossible. Indeed, obtaining a single function is already difficult and typically achieved by synthesis approaches, such as SyGuS [1], or up by specific procedures for dedicated fragments of logic, cf. [2, 3, 4, 5, 6]. However, the problems in the competitions are constructed so that they have “nice” solutions. As a motivational example, consider the problem that asks us to find all functions f from \mathbb{R} to \mathbb{R} satisfying the following identity.

$$\forall xy : \mathbb{R}. f(x + y) = xf(y) + yf(x) \quad (1)$$

Substituting y with 0 gives $\forall x : \mathbb{R}. f(x) = f(0)x$, which tells us that f must be linear and that it is fully determined by the value of $f(0)$. Further, substituting x with 0 shows that f has to be 0 everywhere, i.e., the solution is $f(x) = 0$ (technically, $f = \lambda x. 0$). A solution is expected to give a clear description of how all possible f are calculated, possibly parametrized by constants. To verify the correctness of the solution, one needs to prove the following biconditional.

$$(\forall xy : \mathbb{R}. f(x + y) = xf(y) + yf(x)) \Leftrightarrow (\forall x : \mathbb{R}. f(x) = 0) \quad (2)$$

Note that the problem of finding all solutions is not fully formal in the problem statement, that is, we expect a “reasonable description” of the set of all the solutions. For example, the following would *not* be considered a valid answer: The solution consists of all the functions satisfying $f(x + y) = xf(y) + yf(x)$ for all $x, y \in \mathbb{R}$.

A problem of such a form “find all solutions satisfying X ” sparked some debate within the community about how they should be formally handled. How exactly we handle such questions is explained in Section 3.

9th International Workshop on Satisfiability Checking and Symbolic Computation, July 2, 2024, Nancy, France, Collocated with IJCAR 2024

✉ mikolas.janota@gmail.com (M. Janota)

🌐 <http://people.ciirc.cvut.cz/~janotmik/> (M. Janota)

🆔 0000-0003-3487-784X (M. Janota); 0000-0002-9361-1921 (M. Olšák)

© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

¹<https://aimoprize.com/>

²<https://www.imo-official.org/>

2. Problems Description

As a source for our problems we have used a collection compiled by Vít Musil [7] comprising problems from several sources, including international competition and some (easier) problems from the Prague Seminar in Mathematics³. We have transcribed the problems into SMT2, where each original problem is divided into 3 types of queries.

- The satisfiability of the specification if there is a solution (problems `find`); for the example above:

$$\forall xy : \mathbb{R}. f(x + y) = xf(y) + yf(x)$$

- The unsatisfiability of the specification together with the negation of all proposed solutions (problems `prove`); for the example above:

$$\forall xy : \mathbb{R}. f(x + y) = xf(y) + yf(x) \wedge \exists x : \mathbb{R}. f(x) \neq 0$$

- The check that every proposed solution is indeed a solution to the specifications (problems `check`), i.e., the unsatisfiability of the proposed solution together with the negation of the specification; for the example above:⁴

$$\exists xy : \mathbb{R}. f(x + y) \neq xf(y) + yf(x) \wedge \forall x : \mathbb{R}. f(x) = 0$$

This way, we are asking SMT solvers to solve particular steps, although these steps do not necessarily cover the full solution of a functional equation. We are not providing a formal description of what a valid solution is allowed to consist of (that could be further work using for example SyGuS [1]). Rather, we use our own method for finding the set of all solutions of a particular form, described in Section 3. The SMT2 problems are made available as a GitHub repository⁵ describing 87 problems. In some cases, the check query is split into multiple SMT2 files, if there are multiple distinct solutions.

We remark that during this work we uncovered several bugs in our translation to SMT from the source material. But also, we have found an issue in the original source material. Namely the problem C9 (Cvičení 9) was unsatisfiable in the original [7] and we created two versions of the problem as two different ways of correcting it (C9 and C9a).

2.1. Solutions of Selected Problems

Section 3 describes our approach to solving the problem by trying to prove that it fits a fixed template, in particular, we focus on functions that can be expressed as polynomials. In most cases, this is probably **not** how a human would solve it. It is the subject of future work to make this connection. For comparison, we show here three example problems together with their original (human) solutions:

- Problem U24, which is among the most interesting ones that we were able to fully solve;
- Problem C12, which we were not able to solve but it does not require any advanced technique, so it could be feasible to solve it;
- Problem U2, which showcases a problem requiring induction to be solved, and we assume is currently out of reach.

Example Problem U24 (Baltic Way 1998-7): Find all functions $f : \mathbb{R} \rightarrow \mathbb{R}$ such that for any pair of real numbers x, y , the following identity holds

$$f(x) + f(y) = f(f(x)f(y)).$$

³https://prase.cz/info/info_en.php?lang=1

⁴In most cases, the SMT solver can get rid of f completely by macro detection.

⁵<https://github.com/MikolasJanota/FuncProbs>

Solution: First, we notice that the image of f is closed under addition: if a, b are the values of f at points x, y , then also $a + b$ is a value of f , in particular at point $f(x)f(y) = ab$. Fix arbitrary x_1 , and denote $a = f(x_1)$. By the previous observation, we can find also x_2 and x_4 such that $f(x_2) = 2a$, and $f(x_4) = 4a$.

Now, we plug into the functional equation the assignments $(x = x_2, y = x_2)$, and $(x = x_1, y = x_4)$.

$$4a = f(x_2) + f(x_2) = f(f(x_2)f(x_2)) = f(4a^2)$$

$$5a = f(x_1) + f(x_4) = f(f(x_1)f(x_4)) = f(4a^2)$$

We conclude $4a = 5a$, and consequently $a = 0$. Since we started with arbitrary $f(x) = a$, the function must be constant zero, which indeed satisfies the equation. For this problem, the SMT solver cvc5 was able to demonstrate that f must be constantly 0 by the enumerative strategy [8]. The proof found is analogous to the human proof.⁶ It does **not** contain the general observation that the domain is closed under addition, which is possible because this property is only needed for the points appearing in the proof. The proof is also far less legible because it does not denote subexpressions by new constants (such as $a = f(x_1)$).

Example Problem C12 (IMO 1992-2): Find all functions $f : \mathbb{R} \rightarrow \mathbb{R}$ such that for any pair of real numbers x, y , the following identity holds

$$f(x^2 + f(y)) = y + f(x)^2.$$

Solution: Setting $x = 0$ gives $f(f(y)) = y + f(0)^2$. The right-hand side $y + f(0)^2$ is a bijection, so also the left-hand side $f \circ f$ is a bijection $\mathbb{R} \rightarrow \mathbb{R}$, and that can only happen when f itself is a bijection.

Since f is a bijection, we can consider an argument b such that $f(b) = 0$. Let us substitute $(x = b, y = 0)$ and $(x = -b, y = 0)$ to the original equation.

$$0 = 0 + f(b)^2 = f(b^2 + f(0)) = f((-b)^2 + f(0)) = 0 + f(-b)^2 = f(-b)^2.$$

We obtained $f(-b) = 0$. By injectivity, $-b = b$, so $b = 0$, and $f(0) = 0$. This also simplifies our first observation to $f(f(y)) = y$.

Now we prove that the function is increasing. Any pair of numbers $a < b$ can be expressed as $a = f(y), b = x^2 + f(y)$ for some $x, y, x \neq 0$. Combining the original equation with the same equation having substituted $x = 0$, we get

$$f(a) = f(f(y)) = y < y + f(x)^2 = f(x^2 + f(y)) = f(b),$$

concluding that f is increasing.

Consider any y where $y \leq f(y)$. Since f is increasing, also $f(y) \leq f(f(y)) = y$. Analogously, if $f(y) \leq y$, we obtain $y = f(f(y)) \leq f(y)$. Thus the statements $y \leq f(y)$ and $f(y) \leq y$ are equivalent for any y , leaving the only option $f(y) = y$, which is a function satisfying the given equation.

Example Problem U2 (Cauchy equation): Find all **increasing** functions $f : \mathbb{R} \rightarrow \mathbb{R}$ such that for any pair of real numbers x, y , the following identity holds

$$f(x) + f(y) = f(x + y).$$

⁶To obtain instantiations from cvc5, use `-no-e-matching -enum-inst -dump-instantiations -produce-proof`.

Proof sketch: By setting $y = 0, x, 2x, 3x, \dots$, we obtain by induction that $f(nx) = nf(x)$ for every non-negative integer n , in particular $f(0) = 0$, and if we fix $x = 1$ and denote $f(1) = c$, we obtain the following equation for all non-negative integers x .

$$f(x) = cx. \quad (\clubsuit)$$

We gradually extend the domain for which we know this equation. First, substitution ($x = x, y = -x$) into the original equation gives $f(x) = -f(-x)$ which extends (\clubsuit) to all integers x . Let x be an arbitrary number satisfying (\clubsuit) , and n be a positive integer. Then

$$cx = f(x) = f\left(n\frac{x}{n}\right) = nf\left(\frac{x}{n}\right)$$

leading to

$$c\frac{x}{n} = f\left(\frac{x}{n}\right).$$

Therefore, (\clubsuit) holds for any rational number x .

Finally, we show that (\clubsuit) must be satisfied by every real number. Suppose on the contrary that for some x , we have for example $f(x) < cx$, or $f(x) > cx$. We discuss here the first option, the other one is analogous. If $f(x) < cx$, there is a rational number q such that $\frac{f(x)}{c} < q < x$ (note $c > 0$ since f is assumed to be increasing). Then $q < x$ but $f(x) < qc = f(q)$ contradicting that f is increasing.

This proof goes beyond the abstraction level of SMT solvers. Notably, it uses induction, which SMT solvers can do [9, 10, 11] but the induction here is over a subdomain (natural numbers) vs. reals, where it cannot be applied directly.

3. Template-and-QE

One could see the problem as quantifier elimination in the theory of uninterpreted functions, accompanied by the theory required for describing the problem itself – in our case, the theory of reals or rationals. Due to the undecidability of such a problem, quantifier elimination algorithms cannot be used directly. A possible approach would be to *synthesize* one function adhering to the specification and then strengthen the specification so that the individual solution is not admitted any more. Then, one would have to repeat this process with the hope that there are finitely many solutions, or, observe the solutions and generalize them into a more compact description. All these steps are highly nontrivial. Instead, we propose an approach that tries a fixed template and then performs quantifier elimination: *template-and-QE*, which comprises the following subtasks.

1. *Identify* a template for the solution, e.g. $f(x) \triangleq ax + b$.
2. *Prove* that all solutions must fall into this template (*template verification*)
3. Perform *quantifier elimination* over input variables of the function(s).

Note that in the case of reals, all tasks are computable, except for task 2. The remainder of this section elaborates on the individual steps.

3.1. Template Verification

We run `cvc5` [12], `Z3` [13], `Vampire` [14] and `Waldmeister` [15] to attempt to prove all solutions to a problem necessarily fit a certain template. Since `Waldmeister` does not support theories directly, it needs special treatment and we elaborate on this below. The other solvers support natively the SMT2 syntax [16] enabling us to use the combination of quantified non-linear reals with uninterpreted functions.⁷

We consider the following set of templates for solutions: *constant*, *(monomial)linear*, *(monomial)quadratic*. All these are subclasses of the quadratic form but we consider these to simplify the task for

⁷The standard does not list UFNRA as one of the logics, therefore the files indicate AUFNIRA as the closest superset.

the solvers—it is conceivable that it is easier to prove that the function must be constant than to prove that it must be arbitrary quadratic. Since the set of templates is small, we always try all of them.

In order to ask an SMT solver if all solutions must be linear, we add the assertion that the function is not linear. An obvious way to state that f is linear is via the formula $\exists ab : \mathbb{R}. \forall x : \mathbb{R}. f(x) = ax + b$. We call the problem resulting from including the properties of the original problem along with the negation of $\exists ab : \mathbb{R}. \forall x : \mathbb{R}. f(x) = ax + b$ the *first variant of the linear template verification*. We had more success using a different test for linearity. A function f is linear if and only if $\forall x : \mathbb{R}. f(x) = (f(1) - f(0))x + f(0)$. Note that this formulation avoids the use of existential quantifiers. We call the problem resulting from including the properties of the original problem along with the negation of $\forall x : \mathbb{R}. f(x) = (f(1) - f(0))x + f(0)$ the *second variant of the linear template verification*. If any SMT solver can determine either variant is unsatisfiable, then we know all solutions must be linear.

We likewise have two variants for each of the other templates. The formulations for the first variants simply use existential quantification for the coefficients. For the second variants, we use the following properties:

- *constant*: $\forall x : \mathbb{R}. f(x) = f(0)$
- *monomial linear*: $\forall x : \mathbb{R}. f(x) = f(1)x$
- *linear*: $\forall x : \mathbb{R}. f(x) = (f(1) - f(0))x + f(0)$
- *monomial quadratic*: $\forall x : \mathbb{R}. f(x) = f(1)x^2$
- *quadratic*: $\forall x : \mathbb{R}. 2f(x) = ((f(1) + f(-1)) - 2f(0))x^2 + (f(1) - f(-1))x + 2f(0)$

3.1.1. Using Waldmeister

In addition to using SMT solvers, we also used Waldmeister [15]. Waldmeister is an automated theorem prover specializing in first-order unit equality. Unlike SMT solvers, Waldmeister has no information about integers or reals. Consequently, in order to ask Waldmeister if all functions f satisfying some properties must be contained within the class given by a template, we first declare a sort R along with constants $0, 1 : R$ and operations $+, -, \cdot$ satisfying properties of a commutative ring with identity. Each of these properties is given by a unit equation. We purposely use rings instead of fields (although \mathbb{R} is a field). This allows us to ignore division since the side condition that the denominator is not zero would result in a clause that is not a unit equation. We only generate such a problem for Waldmeister when all the properties stated for f in the problem are unit equations, only quantify over reals (as opposed to integers), and do not use division. Furthermore, we require that all specific real numbers mentioned in the properties correspond to integers (e.g., 0.0, 1.0, -1.0, etc.). These restrictions allow us to formulate a unit equation to give Waldmeister corresponding to each property of f given in the problem.

In addition to the assumed unit equations, Waldmeister expects a unit equation as a goal to prove. The goal varies based on the template we are targeting and corresponds to the second variant of the problems for the SMT solver. In each case we fix a constant d of type R (not mentioned in the assumptions). For each template, the goal unit equation is as follows:

- *constant*: $f(d) = f(0)$
- *monomial linear*: $f(d) = f(1)d$
- *linear*: $f(d) = (f(1) - f(0))d + f(0)$
- *monomial quadratic*: $f(d) = f(1)d^2$
- *quadratic*: $2f(d) = ((f(1) + f(-1)) - 2f(0))d^2 + (f(1) - f(-1))d + 2f(0)$

Waldmeister uses Knuth-Bendix style completion [17] on the assumed unit equations finding a proof of the goal when both sides of the goal rewrite to a common term.

For 8 of the problems, Waldmeister can prove all solutions are within the expected class. Two of these problems were not covered by SMT solvers: U25 and U87. In both of these problems, the least template class is linear monomials. We briefly discuss these two problems.

The problem U25 asks for all real functions f satisfying $f(xf(x) + f(y)) = y + f(x)^2$. Musil [7] gives a proof that the only two solutions are $f(x) = x$ and $f(x) = -x$. Waldmeister's goal is to prove

$f(d) = f(1)d$ from the ring identities and the equation above. It is difficult to directly compare the aforementioned human proof to the proof found by Waldmeister. However, there are some interesting common intermediate results. Both proofs prove f is involutive (i.e., $f(f(x)) = x$), $f(x)^2 = x^2$ and $f(0) = 0$. On the other hand, there are several identities in each proof that do not appear in the other. Musil's proof [7] proves f is surjective and uses this fact, while the Waldmeister proof cannot directly represent the concept of surjectivity.

Waldmeister also proves all solutions to the problem U87 are linear monomials. The problem U87 asks for all real functions f satisfying $f(x + y^2 + z) = f(f(x)) + yf(y) + f(z)$. The (only) two solutions are $f(x) = x$ and $f(x) = 0$. In this case, there is no corresponding proof by Musil [7] as a point of comparison. A number of intermediate identities derived by Waldmeister stand out, e.g., $f(f(0)) = 0$ and $f(f(x)) = f(x)$.

3.2. Quantifier Elimination

Quantifier elimination (QE) is a method to take a general formula in a theory and produce an equivalent quantifier-free formula. So for instance, in $\exists x \in \mathbb{R}. a < x < b$, the quantified variable x is eliminated as $a < b$. In many cases, quantifier elimination serves as a theoretical tool to show that a theory is decidable but also has a long tradition of improvements at the algorithmic level [18, 19, 20, 21, 22, 23, 24]. In our work, QE is used to reveal the particular solutions within a class given by a template.

We assume the original problem only has one uninterpreted function, f , and it is a unary function from reals to reals. If we want to determine all functions f of the form $f(x) = ax + b$, we can simply inline f in the properties, replacing each occurrence $f(t)$ by $at + b$. The resulting problem has no uninterpreted functions and is usually in the theory of the reals. (Exceptional cases are when the properties mention integers as well as reals.) When the problem and template result in such an inlined problem in the theory of reals, we can apply quantifier elimination.

We have made use of two implementations of quantifier elimination for the reals: Z3 [13] and TARSKI [25]. In Z3, QE is invoked by applying the `qe` tactic, with the `qe-nonlinear` parameter turned on. The TARSKI system is highly configurable but since QE is not a bottle-neck for us, we use the `qepcad-qe` [26] function in its default setting. For each system, quantifier elimination should result in a quantifier-free formula involving the uninterpreted constants (e.g., the coefficients a and b of the linear template $ax + b$). One might expect this formula to be in a solved form, e.g., $a = 1 \vee a = 0 \wedge b = 1$, from which one can read off the precise class of linear functions satisfying the original property. However, this is generally not the case. A separate postprocessing step attempts to convert the quantifier-free formula into such a solved form. In practice the postprocessing does not always succeed. In the case of Z3, we also call the tactics `simplify` and `propagate-values` to have Z3 simplify the formula before applying the postprocessing to attempt to find a solved form.

Problem U91 asks for all functions satisfying $f((x - y)^2) = f(x)^2 - 2xf(y) + y^2$. Using the techniques of Section 3.1, we can prove all solutions f must be linear. After replacing $f(x)$ with $ax + b$ and calling Z3 to do quantifier elimination, we obtain the quantifier-free formulas $a = 1$ and $b = 0 \vee b = 1$. In general, Z3 represents the result of QE as a conjunction of several subformulas. The formula $a = 1 \wedge (b = 0 \vee b = 1)$ is almost in solved form, and we can easily obtain the two solutions $f(x) = x$ and $f(x) = x + 1$ via postprocessing. TARSKI's quantifier elimination returns the more complicated formula

$$-1 + a = 0 \wedge b \geq 0 \wedge -1 + b \leq 0 \wedge (b = 0 \vee -1 + b = 0).$$

The postprocessor is also able to obtain the two solutions from this formula.

Sometimes there is a parameterized class of solutions. Problem U3, for example, asks for all functions satisfying $f(x + y) = f(x) + y$. The techniques of Section 3.1 determine all solutions are linear. After inlining $f(x) = ax + b$, we call quantifier elimination. Both Z3 and TARSKI produce the formula $-1 + a = 0$. From this the postprocessor can determine $a = 1$ and b is unconstrained. This gives the class of solutions $f(x) = x + b$ where b is a real number. Indeed, this is the class of all solutions, as desired.

3.3. Postprocessor to Obtain Solved Forms

We say a quantifier-free formula is in *solved form*⁸ if it is a disjunctive normal form where all the literals are of the form $c = v$ where c is a coefficient of the template and v is a real number. We call a formula of the form $c = v$ an *assignment atom*. The postprocessor is given a list of formulas $\bar{\varphi}$ (thought of conjunctively) and attempts to generate a solved form equivalent to $\varphi_1 \wedge \cdots \wedge \varphi_n$. As auxiliary values, a list \bar{t} of equations of the form $s = t$ and a list $\bar{\mu}$ of other literals are generated, so that generally we are working with a triple $(\bar{\varphi}, \bar{t}, \bar{\mu})$ of lists of formulas, with \bar{t} and $\bar{\mu}$ initially empty. Suppose $\bar{\varphi}$ is nonempty, with φ_1 being the first formula on the list. If φ_1 is a conjunction or a disjunction, we can simply make recursive calls to the preprocessor to ask for solved forms of the decomposed formulas. In the case of disjunction, we make two recursive calls (one for each disjunct) and combine the results as a disjunction of the two solved forms (or fail if one call fails). If φ_1 is an equation $s = t$, we add this to the equation list \bar{t} . If φ_1 is of the form $s \leq t, s \geq t, s < t$ or $s > t$, then we add this to the other list $\bar{\mu}$. In every other case for φ_1 , we simply fail to return a solved form. We are finally left with the case where there are no remaining formulas in $\bar{\varphi}$. The hope in this case is that \bar{t} contains sufficient information to obtain assignment atoms, and we only need to verify that the assignments satisfy the other formulas from $\bar{\mu}$. If the assignment atoms satisfy the other formulas, the conjunction of the assignment atoms will be the solved form. If the assignment atoms do not satisfy the other formulas, the empty disjunction \perp will be the solved form. It is also possible (e.g., if $c^2 = 1$ is in the list \bar{t}) a solved form with more than one disjunct. At this point, the algorithm transforms a triple $(\bar{t}, \bar{\mu}, \bar{\alpha})$ where $\bar{\alpha}$ is the list of assignment atoms (initially empty). We will also call $\bar{\alpha}$ an *assignment*.

We say a term t can be evaluated under an assignment $\bar{\alpha}$ if every constant occurring in t is associated with a real number via $\bar{\alpha}$. We say an atom $s = t, s \leq t, s \geq t, s < t$ or $s > t$ can be evaluated under an assignment $\bar{\alpha}$ if both terms s and t can be evaluated under an assignment $\bar{\alpha}$. Clearly, if a term can be evaluated under an assignment $\bar{\alpha}$, the term evaluates to a specific real number. Likewise, if an atom can be evaluated under an assignment $\bar{\alpha}$, the atom evaluates to true or false.

Assume \bar{t} is nonempty and $s = t$ is the first equation on \bar{t} . We consider the following cases (in order), dropping the equation from \bar{t} :

1. Assume both s and t can be evaluated under $\bar{\alpha}$. We drop the equation if they evaluate to the same reals and return \perp otherwise.
2. Assume s is a coefficient c and t can be evaluated to the real v . We add $c = v$ to $\bar{\alpha}$.
3. Assume t can be evaluated under $\bar{\alpha}$ and s is one of a number of special forms which make use of an unassigned coefficient c and all other subterms s_i that can be evaluated under $\bar{\alpha}$. Examples of these special forms are $s_1 + c, s_1 + cs_2$ and c^2 . In each case we solve for values v of c and add $c = v$ to $\bar{\alpha}$ (or fail). If there are no solutions, we return \perp . If there are multiple solutions, we recursively call with both assignments and return the disjunction of the resulting solved forms.
4. Assume the first equation $s = t$ is of the form $c + s_1d = t$ and there is a second equation $s' = t'$ on the list \bar{t} of the form $c + d = t'$ where c and d are unassigned coefficients and s_1, t and t' can all be evaluated under $\bar{\alpha}$. Then we solve these two linear equations for c and d (if possible) and add these to the assignment $\bar{\alpha}$, dropping both equations from \bar{t} .

We finally assume \bar{t} is empty. If some formula in $\bar{\mu}$ cannot be evaluated via the assignment $\bar{\alpha}$, then we fail to return a solved form. Assume every $\bar{\mu}$ can be evaluated via the assignment $\bar{\alpha}$. If every $\bar{\mu}$ evaluates to true, then we return the conjunction of the assignment atoms in $\bar{\alpha}$ as the solved form. Otherwise, we return \perp as the solved form.

The postprocessor included sufficiently many cases to obtain a solved form from either the output of Z3 or TARSKI (or preferably both) in each of the problems we considered. There are many cases in which it will fail to compute a solved form. Easy examples are $c^3 = 1$ and $2c + d = 2 \wedge c + d = 1$.

⁸One may imagine loosening the concept but the general idea is to provide a specific solution as possible.

3.3.1. Lazy Verification

We now briefly consider an alternative to the approach outlined above. When considering a template, we do not need to verify that all solutions are in the template class in advance. Instead, we could use QE to find a class of solutions for a fixed template and then prove there are no other solutions. The advantage of this approach is illustrated by the following example. Problem C1 asks for all functions f such that the following holds

$$f(x + y) + 2f(x - y) - 4f(x) + xf(y) = 3y^2 - x^2 - 2xy + xy^2.$$

The only solution is $f(x) = x^2$.

In this case, the solvers listed in Section 3.1 are unable to prove that all solutions are a monomial quadratic (or that they must be quadratic). However, we can still apply quantifier elimination to find all monomial quadratic solutions. For this example, Z3 timed out even when given 10 minutes. On the other hand, TARSKI returned the formula $-1 + a = 0$ (yielding the solution $a = 1$, i.e., x^2) in less than a second. Since we were unable to prove all solutions must belong to the monomial quadratic template class, it is still possible there are more solutions outside the monomial quadratic template class. However, now that we have the specific solution $f(x) = x^2$ we can ask an SMT solver to prove this is the only solution by assuming the equation and also the negation of $\forall x : \mathbb{R}. f(x) = x^2$. The solver cvc5 is easily able to show this is unsatisfiable, so we know that $f(x) = x^2$ is the only solution.

4. Experiments

We report the results of the template-and-qe method (Section 3) on the benchmark presented in Section 2. The lazy extension of our approach (Section 3.3.1) was left for future work. The template verification (Section 3.1) has proven to be the most difficult task. Table 1a summarizes the obtained results for template verification. For all these, we were able to solve the QE task. Thirteen instances were solved completely by the automated method. Out of these thirteen, two can be traced to a competition event. Problem U24 comes from Baltic Way competition—see Section 2 for “human” solution. Problem U71 comes from The Prague Seminar (PraSe-18-6-1). In both cases, the only solution is a function that is constantly 0.

The verification tasks results often in satisfiable problems, namely if there exists a solution outside of the proposed template (the \times symbol in Table 1a). These satisfiable problems are nontrivial because they involve creating a counterexample to the template. For this purpose, we also ran the SyGuS approach of cvc5 (option `-sygus-inference`) and its linear model builder [27].

We also report on the verification of the correctness of the handwritten solutions, which comprises two components: *prove* — show that all possible solutions are covered by the suggested solutions; *check* — check that all suggested solutions are indeed solutions to the problem (see Section 2). Table 1b summarizes the obtained results. Sixty-five of the provided handwritten solutions were successfully checked, i.e., the individual solutions satisfy the original specifications. Twenty of the handwritten solutions were shown to cover all the individual solutions. Unsurprisingly, the proving task is harder than the checking task. We remark that checking solutions for U79 is trivial because there are no individual solutions. More detailed results can be found on the authors’ web page [28].

5. Conclusions and Future Work

This paper joins the effort of rising to the challenge of making computers as powerful as the golden medalists at the International Math Olympics (IMO). Efforts such as AlphaGeometry [29], show that machine learning models are useful for the task but at the same time, further research shows that they benefit from existing symbolic methods [30]. Here we show that symbolic methods are indeed already powerful in solving a highly nontrivial task of finding *all* functions fulfilling a certain specification.

Table 1

Result summary. Problems without any result are omitted.

(a) Solution Template. Proven templates are denoted by \checkmark ; disproven template by \times ; dash when no solver provided an answer.

Problem	c	ax	$ax + b$	ax^2	$ax^2 + bx + c$
C2	\times	\times	\checkmark	\times	\checkmark
C6	\times	\times	-	\times	-
C9a	-	-	-	\checkmark	-
C10	\times	\checkmark	\checkmark	\times	\checkmark
C12	\times	-	-	\times	-
C13	\times	\times	\times	\times	\times
U2	\times	-	-	\times	-
U3	\times	\times	\checkmark	\times	\checkmark
U5	\times	\checkmark	\checkmark	\times	\checkmark
U7	\times	\times	-	\times	-
U9	\times	\times	\checkmark	\times	\checkmark
U13	\times	\checkmark	\checkmark	\times	\checkmark
U16	\times	\times	\checkmark	\times	-
U20	\times	\times	-	\times	-
U23	-	\times	-	\times	-
U24	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
U25	\times	\checkmark	\checkmark	\times	-
U26	\times	-	-	\times	-
U27	\times	\times	\times	\times	\times
U41	-	\times	-	\times	-
U42	\times	\times	-	\times	-
U44	\times	-	-	-	-
U45	\times	-	-	\times	-
U48	\times	-	-	\times	-
U49	\times	-	-	\times	-
U50	\times	-	-	\times	-
U51	\times	\times	-	\times	-
U54	\times	\times	\times	\times	\times
U56	\times	\times	-	\times	-
U62	\times	-	-	\times	-
U64	\times	-	-	\times	-
U67	\times	\times	\times	\times	\times
U68	\times	-	-	\times	-
U71	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
U72	-	\times	-	\times	-
U87	\times	\checkmark	\checkmark	\times	\checkmark
U90	\times	\times	\times	\times	-
U91	\times	\times	\checkmark	\times	\checkmark
U92	\times	-	-	\times	-

(b) Prove/check

Problem	Prove	Check
C1	\checkmark	\checkmark
C2	\checkmark	\checkmark
C3	-	\checkmark
C4	-	\checkmark
C5	\checkmark	\checkmark
C6	-	\checkmark
C9	\checkmark	\checkmark
C9a	\checkmark	\checkmark
C10	\checkmark	\checkmark
C12	-	\checkmark
C13	-	\checkmark
U2	-	\checkmark
U3	\checkmark	\checkmark
U4	\checkmark	\checkmark
U5	\checkmark	\checkmark
U6	-	\checkmark
U7	-	\checkmark
U8	-	\checkmark
U9	\checkmark	\checkmark
U10	-	\checkmark
U11	\checkmark	\checkmark
U12	-	\checkmark
U13	\checkmark	\checkmark
U14	-	\checkmark
U16	\checkmark	\checkmark
U17	\checkmark	\checkmark
U19	-	\checkmark
U20	-	\checkmark
U23	-	\checkmark
U24	\checkmark	\checkmark
U25	-	\checkmark
U26	-	\checkmark
U27	-	\checkmark
U39	-	\checkmark
U41	-	\checkmark
U42	-	\checkmark
U44	-	\checkmark
U45	-	\checkmark
U46	-	\checkmark
U48	-	\checkmark
U49	-	\checkmark
U50	-	\checkmark
U51	-	\checkmark
U53	-	\checkmark
U54	-	\checkmark
U56	-	\checkmark
U57	\checkmark	\checkmark
U61	-	\checkmark
U62	-	\checkmark
U64	-	\checkmark
U66	-	\checkmark
U67	-	\checkmark
U68	-	\checkmark
U71	\checkmark	\checkmark
U72	-	\checkmark
U75	-	\checkmark
U76	-	\checkmark
U79	\checkmark	\checkmark
U81	\checkmark	\checkmark
U82	\checkmark	\checkmark
U87	-	\checkmark
U89	-	\checkmark
U90	-	\checkmark
U91	-	\checkmark
U92	-	\checkmark
U93	-	\checkmark

Besides the task leading to undecidable questions, its difficulty also lies in the fact it is *not* a decision problem but the response is a description of a class of mathematical objects.

The method we employ is anchored in templates, which is a well-known technique in function and program synthesis [31]. The templates we consider are rather simple and they could be made more powerful by adding conditionals (if-then-else). However, the templates need to be kept simple if we wish to be able to apply quantifier elimination.

An alternative to templates would be to attempt extending different synthesis approaches, e.g., such as those that are realized in saturation-based solvers [2, 3] or inside SMT solvers [32, 33, 34] or inductive logic programming [35]. The challenge here would be how to come up with all the solutions.

Acknowledgments

The results were supported by the Ministry of Education, Youth and Sports within the dedicated program ERC CZ under the project *POSTMAN* no. LL1902 and co-funded by the European Union under the project *ROBOPROX* (reg. no. CZ.02.01.01/00/22_008/0004590). C. Brown was supported by *CORESENSE*: the European Union’s Horizon Europe research and innovation programme under grant agreement no. 101070254 CORESENSE. This article is part of the *RICAIP* project that has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 857306.

References

- [1] R. Alur, R. Bodik, G. Juniwal, M. M. K. Martin, M. Raghothaman, S. A. Seshia, R. Singh, A. Solar-Lezama, E. Torlak, A. Udupa, Syntax-guided synthesis, in: 2013 Formal Methods in Computer-Aided Design, 2013, pp. 1–8. doi:10.1109/FMCAD.2013.6679385.
- [2] P. Hozzová, L. Kovács, C. Norman, A. Voronkov, Program synthesis in saturation, in: B. Pientka, C. Tinelli (Eds.), Automated Deduction - CADE 29 - 29th International Conference on Automated Deduction, volume 14132 of *LNCS*, Springer, 2023, pp. 307–324. doi:10.1007/978-3-031-38499-8_18.
- [3] P. Hozzová, D. Amrollahi, M. Hajdu, L. Kovács, A. Voronkov, E. M. Wagner, Synthesis of recursive programs in saturation, in: International Joint Conference on Automated Reasoning IJCAR, 2024.
- [4] S. Ratschan, Deciding predicate logical theories of real-valued functions, in: J. Leroux, S. Lombardy, D. Peleg (Eds.), 48th International Symposium on Mathematical Foundations of Computer Science (MFCS 2023), volume 272 of *LIPICs*, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2023, pp. 76:1–76:15. doi:10.4230/LIPICs.MFCS.2023.76.
- [5] A. R. Bradley, Z. Manna, H. B. Sipma, What’s decidable about arrays?, in: Verification, Model Checking, and Abstract Interpretation Conference, VMCAI, volume 3855, Springer, 2006, pp. 427–442. doi:10.1007/11609773_28.
- [6] Y. Ge, L. M. de Moura, Complete instantiation for quantified formulas in satisfiability modulo theories, in: Computer Aided Verification CAV, volume 5643, Springer, 2009, pp. 306–320. doi:10.1007/978-3-642-02658-4_25.
- [7] V. Musil, Funkcionální rovnice, 2024. URL: <https://prase.cz/library/FunkcionalniRovniceVM/FunkcionalniRovniceVM.pdf>, downloaded 8 March 2024.
- [8] M. Janota, H. Barbosa, P. Fontaine, A. Reynolds, Fair and adventurous enumeration of quantifier instantiations, in: Formal Methods in Computer-Aided Design, IEEE, 2021, pp. 256–260. doi:10.34727/2021/ISBN.978-3-85448-046-4_35.
- [9] P. Hozzová, L. Kovács, A. Voronkov, Integer induction in saturation, in: A. Platzer, G. Sutcliffe (Eds.), Automated Deduction - CADE 28 - 28th International Conference on Automated Deduction, volume 12699 of *LNCS*, Springer, 2021, pp. 361–377. doi:10.1007/978-3-030-79876-5_21.
- [10] M. Hajdú, P. Hozzová, L. Kovács, G. Reger, A. Voronkov, Getting saturated with induction, CoRR abs/2402.18954 (2024). doi:10.48550/ARXIV.2402.18954. arXiv:2402.18954.
- [11] A. Reynolds, V. Kuncak, Induction for SMT solvers, in: D. D’Souza, A. Lal, K. G. Larsen (Eds.), Verification, Model Checking, and Abstract Interpretation - 16th International Conference, VMCAI, volume 8931 of *LNCS*, Springer, 2015, pp. 80–98. doi:10.1007/978-3-662-46081-8_5.
- [12] H. Barbosa, C. W. Barrett, M. Brain, G. Kremer, H. Lachnitt, M. Mann, A. Mohamed, M. Mohamed, A. Niemetz, A. Nötzli, A. Ozdemir, M. Preiner, A. Reynolds, Y. Sheng, C. Tinelli, Y. Zohar, cvc5: A versatile and industrial-strength SMT solver, in: Tools and Algorithms for the Construction and Analysis of Systems, TACAS, volume 13243 of *LNCS*, Springer, 2022, pp. 415–442. doi:10.1007/978-3-030-99524-9_24.
- [13] L. M. de Moura, N. Bjørner, Z3: an efficient SMT solver, in: Tools and Algorithms for the Construction and Analysis of Systems, 14th International Conference, TACAS 2008, volume 4963, Springer, 2008, pp. 337–340. doi:10.1007/978-3-540-78800-3_24.

- [14] L. Kovács, A. Voronkov, First-order theorem proving and Vampire, in: N. Sharygina, H. Veith (Eds.), *Computer Aided Verification - 25th International Conference, CAV*, volume 8044 of *LNCs*, Springer, 2013, pp. 1–35. doi:10.1007/978-3-642-39799-8_1.
- [15] T. Hillenbrand, B. Löchner, The next Waldmeister loop, in: A. Voronkov (Ed.), *Automated Deduction - CADE-18, 18th International Conference on Automated Deduction*, volume 2392 of *Lecture Notes in Computer Science*, 2002, pp. 486–500. doi:10.1007/978-3-540-45085-6_27.
- [16] C. Barrett, P. Fontaine, C. Tinelli, The SMT-LIB Standard: Version 2.6, Technical Report, Department of Computer Science, The University of Iowa, 2017. Available at www.SMT-LIB.org.
- [17] L. Bachmair, N. Dershowitz, D. A. Plaisted, Completion without failure, in: *Rewriting Techniques*, Elsevier, 1989, pp. 1–30.
- [18] G. E. Collins, Quantifier elimination for real closed fields by cylindrical algebraic decomposition, in: H. Brakhage (Ed.), *Automata Theory and Formal Languages*, Springer, Berlin, Heidelberg, 1975, pp. 134–183.
- [19] D. C. Cooper, Theorem proving in arithmetic without multiplication, *Machine intelligence* 7 (1972) 300.
- [20] J. Ferrante, C. Rackoff, A decision procedure for the first order theory of real addition with order, *SIAM Journal on Computing* 4 (1975) 69–76. doi:10.1137/0204006.
- [21] R. Loos, V. Weispfenning, Applying linear quantifier elimination, *Comput. J.* 36 (1993) 450–462. doi:10.1093/COMJNL/36.5.450.
- [22] J. H. Davenport, Y. Siret, E. Tournier, *Computer algebra - systems and algorithms for algebraic computation* (2. ed.), Academic Press, 1993.
- [23] A. R. Bradley, Z. Manna, *The calculus of computation - decision procedures with applications to verification*, Springer, 2007. doi:10.1007/978-3-540-74113-8.
- [24] N. S. Bjørner, M. Janota, Playing with quantified satisfaction, in: *LPAR, EasyChair*, 2015, pp. 15–27. doi:10.29007/vv21.
- [25] F. Vale-Enriquez, C. W. Brown, Polynomial constraints and unsat cores in Tarski, in: J. H. Davenport, M. Kauers, G. Labahn, J. Urban (Eds.), *Mathematical Software - ICMS - 6th International Conference*, volume 10931 of *LNCs*, Springer, 2018, pp. 466–474. doi:10.1007/978-3-319-96418-8_55.
- [26] G. E. Collins, H. Hong, Partial cylindrical algebraic decomposition for quantifier elimination, *J. Symb. Comput.* 12 (1991) 299–328. doi:10.1016/S0747-7171(08)80152-6.
- [27] M. Janota, B. Piotrowski, K. Chvalovský, Towards learning infinite SMT models, 2023. URL: <http://people.ciirc.cvut.cz/~janotmik/>, 25th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing, SYNASC.
- [28] M. Janota, 2024. URL: https://sat.inesc-id.pt/~mikolas/sc2_2024/.
- [29] T. H. Trinh, Y. Wu, Q. V. Le, H. He, T. Luong, Solving olympiad geometry without human demonstrations, *Nature* 625 (2024) 476–482. doi:10.1038/s41586-023-06747-5.
- [30] S. Sinha, A. Prabhu, P. Kumaraguru, S. Bhat, M. Bethge, Wu’s method can boost symbolic AI to rival silver medalists and AlphaGeometry to outperform gold medalists at IMO geometry, 2024. arXiv:2404.06405.
- [31] S. Srivastava, S. Gulwani, J. S. Foster, Template-based program verification and program synthesis, *Int. J. Softw. Tools Technol. Transf.* 15 (2013) 497–518. doi:10.1007/S10009-012-0223-4.
- [32] A. Reynolds, V. Kuncak, C. Tinelli, C. W. Barrett, M. Deters, Refutation-based synthesis in SMT, *Formal Methods Syst. Des.* 55 (2019) 73–102. doi:10.1007/S10703-017-0270-2.
- [33] A. Abate, H. Barbosa, C. W. Barrett, C. David, P. Kesseli, D. Kroening, E. Polgreen, A. Reynolds, C. Tinelli, Synthesising programs with non-trivial constants, *J. Autom. Reason.* 67 (2023) 19. doi:10.1007/S10817-023-09664-4.
- [34] H. Barbosa, A. Reynolds, D. Larraz, C. Tinelli, Extending enumerative function synthesis via SMT-driven classification, in: C. W. Barrett, J. Yang (Eds.), *Formal Methods in Computer Aided Design, FMCAD, IEEE*, 2019, pp. 212–220. doi:10.23919/FMCAD.2019.8894267.
- [35] D. M. Cerna, A. Cropper, Generalisation through negation and predicate invention, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 2024, pp. 10467–10475. doi:10.1609/aaai.v38i9.28915.