



NIL: Learning Nonlinear Interpolants

Mingshuai Chen^{1,2}(✉) , Jian Wang^{1,2} , Jie An³ , Bohua Zhan^{1,2}(✉) ,
Deepak Kapur⁴ , and Najjun Zhan^{1,2}(✉)

¹ State Key Laboratory of Computer Science,
Institute of Software, CAS, Beijing, China
{chenms, bzhan, znj}@ios.ac.cn

² University of Chinese Academy of Sciences, Beijing, China

³ School of Software Engineering, Tongji University, Shanghai, China

⁴ Department of Computer Science, University of New Mexico, Albuquerque, USA

Abstract. Nonlinear interpolants have been shown useful for the verification of programs and hybrid systems in contexts of theorem proving, model checking, abstract interpretation, etc. The underlying synthesis problem, however, is challenging and existing methods have limitations on the form of formulae to be interpolated. We leverage classification techniques with space transformations and kernel tricks as established in the realm of machine learning, and present a counterexample-guided method named NIL for synthesizing polynomial interpolants, thereby yielding a unified framework tackling the interpolation problem for the general quantifier-free theory of nonlinear arithmetic, possibly involving transcendental functions. We prove the soundness of NIL and propose sufficient conditions under which NIL is guaranteed to converge, i.e., the derived sequence of candidate interpolants converges to an actual interpolant, and is complete, namely the algorithm terminates by producing an interpolant if there exists one. The applicability and effectiveness of our technique are demonstrated experimentally on a collection of representative benchmarks from the literature, where in particular, our method suffices to address more interpolation tasks, including those with perturbations in parameters, and in many cases synthesizes simpler interpolants compared with existing approaches.

Keywords: Nonlinear Craig interpolant · Counterexample-guided learning · Program verification · Support vector machines (SVMs)

1 Introduction

Interpolation-based technique provides a powerful mechanism for local and modular reasoning, thereby improving scalability of various verification techniques,

This work has been supported through grants by NSFC under grant No. 61625206 and 61732001, by the CAS Pioneer Hundred Talents Program under grant No. Y9RC585036, and by the National Science Foundation Award DMS-1217054.

e.g., theorem proving, model checking and abstract interpretation, to name just a few. The study of interpolation was pioneered by Krajíček [27] and Pudlák [35] in connection with theorem proving, by McMillan [30] in the context of model checking, by Graf and Saïdi [17], McMillan [31] and Henzinger et al. [20] pertaining to abstraction like CEGAR [8], and by Wang et al. [24] in the context of learning-based invariant generation. Developing efficient algorithms for generating interpolants for various theories and their combination has become an active research area, see e.g., [7, 25, 26, 31, 32, 36, 46].

Though established methods addressing interpolant generation for Presburger arithmetic, decidable fragments of first-order logic, theory of equality over uninterpreted functions (EUFs) as well as their combination have been extensively studied in the literature, there appears to be little work on synthesizing nonlinear interpolants. Dai et al. proposed an algorithm in [11] for generating interpolants for nonlinear polynomial inequalities based on the existence of a witness guaranteed by Stengle’s Positivstellensatz [16] that can be computed using semi-definite programming (SDP). A major limitation of this method is that the two mutually contradictory formulas to be interpolated must share the same set of variables. Okudono et al. extended [11] in [33] to cater for the so-called sharper and simpler interpolants by developing a continuous fraction-based algorithm that rounds off numerical solutions. In [14], Gan et al. considered the interpolation for inequalities combined with EUFs by employing the hierarchical calculus framework proposed in [38] (and its extension [39]), while the inequalities are limited to be of the concave quadratic form. In [15], Gao and Zufferey transformed proof traces from δ -complete decision procedures into interpolants, composed of Boolean combinations of linear constraints, which can deal with certain transcendental functions beyond polynomials. The techniques of encoding interpolants as logical combinations of linear constraints, including [15, 28, 37], however, yield potentially large interpolants (requiring even an infinite length in the worst case) and their usage thus becomes difficult in practical applications (cf. Example 1).

Interpolants can be viewed as classifiers that distinguish, in the context of program verification for instance, positive program states from negative ones (unreachable/error states) and consequently the state-of-the-art classification algorithms can be leveraged for synthesizing interpolants. The universal applicability of classification techniques substantially extends the scope of theories admitting interpolant generation. This idea was first employed by Sharma et al. in [37], which infers linear interpolants through hyperplane-classifiers generated by support vector machines (SVMs) [3, 45] whilst handles superficial nonlinearities by assembling interpolants in the form purely of conjunctions (or dually, disjunctions) of linear half-spaces, which addresses only a limited category of formulae featuring nonlinearities. The learning-based paradigm has also been exploited in the context of nonlinear constraint solving, see e.g., [12].

In this paper, we present a classification-based learning method for the synthesis of polynomial interpolants for the quantifier-free theory of nonlinear arithmetic. Our approach is based on techniques of space transformations and kernel tricks pertinent to SVMs that have been well-developed in the realm of

machine learning. Our method is described by an algorithm called NIL (and its several variants) that adopts the counterexample-guided inductive synthesis framework [22, 40]. We prove the soundness of NIL and propose sufficient conditions under which NIL is guaranteed to converge, that is, the derived sequence of classifiers (candidate interpolants) converges to an actual interpolant, and is complete, i.e., if an interpolant exists, the method terminates with an actual interpolant. In contrast to related work on generation of nonlinear interpolants, which restrict the input formulae, our technique provides a uniform framework, tackling the interpolation problem for the general quantifier-free theory of nonlinear arithmetic, possibly involving transcendental functions. The applicability and effectiveness of NIL are demonstrated experimentally on a collection of representative benchmarks from the literature; as is evident from experimental results, our method is able to address more demands on the nature of interpolants, including those with perturbations in parameters (due to the robustness inherited from SVMs); in many cases, it synthesizes simpler interpolants compared with other approaches, as shown by the following example.

Example 1 ([15]). Consider two mutually contradictory inequalities $\phi \hat{=} y \geq x^2$ and $\psi \hat{=} y \leq -\cos(x) + 0.8$. Our NIL algorithm constructs a single polynomial inequality $I \hat{=} 15x^2 < 4 + 20y$ as the interpolant, namely, $\phi \models I$ and $I \wedge \psi$ is unsatisfiable; while the interpolant generated by the approach in [15], only when provided with sufficiently large finite domains, e.g., $x \in [-\pi, \pi]$ and $y \in [-0.2, \pi^2]$, is $y > 1.8 \vee (0.59 \leq y \leq 1.8 \wedge -1.35 \leq x \leq 1.35) \vee (0.09 \leq y < 0.59 \wedge -0.77 \leq x \leq 0.77) \vee (y \geq 0 \wedge -0.3 \leq x \leq 0.3)$. As will be discussed later, we do not need to provide a priori information to our algorithm such as bounds on variables.

The rest of the paper is organized as follows. Section 2 introduces some preliminaries on Craig interpolants and SVMs. In Sect. 3, we present the NIL algorithm dedicated to synthesizing nonlinear interpolants, followed by the analysis of its soundness, conditional completeness and convergence in Sect. 4. Section 5 reports several implementation issues and experimental results on a collection of benchmarks (with the robustness discussed in Sect. 6). The paper is then concluded in Sect. 7.

2 Preliminaries

Let \mathbb{N} , \mathbb{Q} and \mathbb{R} be the set of natural, rational and real numbers, respectively. We denote by $\mathbb{R}[\mathbf{x}]$ the polynomial ring over \mathbb{R} with variables $\mathbf{x} = (x_1, \dots, x_n)$, and $\|\mathbf{x}\|$ denotes the ℓ^2 -norm [4]. For a set $X \subseteq \mathbb{R}^n$, its convex hull is denoted by $\text{conv}(X)$. For $\mathbf{x}, \mathbf{x}' \in X$, $\text{dist}(\mathbf{x}, \mathbf{x}') = \|\mathbf{x} - \mathbf{x}'\|$ denotes the Euclidean distance between two points, which generalizes to $\text{dist}(\mathbf{x}, X') = \min_{\mathbf{x}' \in X'} \text{dist}(\mathbf{x}, \mathbf{x}')$. Given $\delta \geq 0$, define $\mathcal{B}(\mathbf{x}, \delta) = \{\mathbf{x}' \in \mathbb{R}^n \mid \|\mathbf{x}' - \mathbf{x}\| \leq \delta\}$ as the closed ball of radius δ centered at \mathbf{x} . Consider the quantifier-free fragment of a first-order theory of polynomials over the reals, denoted by \mathcal{T}_P , in which a formula φ is of the form

$$\varphi \hat{=} p(\mathbf{x}) \diamond 0 \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \neg \varphi$$

where $p(\mathbf{x}) \in \mathbb{R}[\mathbf{x}]$ and $\diamond \in \{<, >, \leq, \geq, =\}$. A natural extension of our method to cater for more general nonlinearities involving transcendental functions will be demonstrated in subsequent sections. In the sequel, we use \perp to stand for *false* and \top for *true*. Let $\mathbb{R}[\mathbf{x}]_m$ consist of all polynomials $p(\mathbf{x})$ of degree $\leq m \in \mathbb{N}$. We abuse the notation $\varphi \in \mathbb{R}[\mathbf{x}]_m$ to abbreviate $\varphi \hat{=} p(\mathbf{x}) \diamond 0$ and $p(\mathbf{x}) \in \mathbb{R}[\mathbf{x}]_m$ if no ambiguity arises.

Given formulas ϕ and ψ in a theory \mathcal{T} , ϕ is *valid* w.r.t. \mathcal{T} , written as $\models_{\mathcal{T}} \phi$, iff ϕ is true in all models of \mathcal{T} ; ϕ *entails* ψ w.r.t. \mathcal{T} , written as $\phi \models_{\mathcal{T}} \psi$, iff every model of \mathcal{T} that makes ϕ true makes ψ also true; ϕ is *satisfiable* w.r.t. \mathcal{T} , iff there is a model of \mathcal{T} in which ϕ is true; otherwise *unsatisfiable*. It follows that ϕ is unsatisfiable iff $\phi \models_{\mathcal{T}} \perp$. The set of all the models that make ϕ true is denoted by $\llbracket \phi \rrbracket_{\mathcal{T}}$.

2.1 Craig Interpolant

Craig showed in [10] that given two formulas ϕ and ψ in a first-order logic \mathcal{T} s.t. $\phi \models_{\mathcal{T}} \psi$, there always exists an *interpolant* I over the common symbols of ϕ and ψ s.t. $\phi \models_{\mathcal{T}} I$ and $I \models_{\mathcal{T}} \psi$. In the verification literature, this terminology has been abused by [31], which defined an interpolant over the common symbols of ϕ and ψ as

Definition 1 (Interpolant). *Given ϕ and ψ in a theory \mathcal{T} s.t. $\phi \wedge \psi \models_{\mathcal{T}} \perp$, a formula I is a (reverse) interpolant of ϕ and ψ if (i) $\phi \models_{\mathcal{T}} I$; (ii) $I \wedge \psi \models_{\mathcal{T}} \perp$; and (iii) I contains only common symbols shared by ϕ and ψ .*

It is immediately obvious that $\phi \models_{\mathcal{T}} \psi$ iff $\phi \wedge \neg\psi \models_{\mathcal{T}} \perp$, namely, I is an interpolant of ϕ and ψ iff I is a reverse interpolant in McMillan’s sense of ϕ and $\neg\psi$. We follow McMillan in continuing to abuse the terminology.

2.2 Support Vector Machines

In machine learning, support vector machines [3,45] are supervised learning models for effective classification based on convex optimization. In a binary setting, we are given a training dataset $X = X^+ \uplus X^-$ of n sample points $\{(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \dots, (\mathbf{x}_n, y_n)\}$, where $\mathbf{x}_i \in \mathbb{R}^d$, and y_i is either 1, indicating a positive sample $\mathbf{x}_i \in X^+$, or -1 , indicating a negative one in X^- . The goal of classification here is to find a potential hyperplane (a.k.a. *linear classifier*) to separate the positive samples from the negative ones. There however might be various or even infinite number of separating hyperplanes, and an SVM aims to construct a separating hyperplane that yields the largest distance (so-called *functional margin*) to the nearest positive and negative samples. Such a classification hyperplane is called the *optimal-margin classifier* while the samples closest to it are called the *support vectors*.

Linear SVMs. Assume that X^+ and X^- are *linearly separable*, meaning that there exists a *linear separating hyperplane* $\mathbf{w}^T \mathbf{x} + b = 0$ such that $y_i(\mathbf{w}^T \mathbf{x}_i + b) > 0$, for all $(\mathbf{x}_i, y_i) \in X$. Then the functional margin can be formulated as

$$\gamma \hat{=} 2 \min_{1 \leq i \leq n} 1 / \|\mathbf{w}\| |\mathbf{w}^T \mathbf{x}_i + b|.$$

Linear SVMs are committed to finding appropriate parameters (\mathbf{w}, b) that maximize the functional margin while adhering to the constraints of separability, which reduces equivalently to the following convex quadratic optimization problem [2] that can be efficiently solved by off-the-shelf packages for quadratic programming:

$$\underset{\mathbf{w}, b}{\text{minimize}} \quad \frac{1}{2} \mathbf{w}^T \mathbf{w} \quad \text{subject to} \quad y_i (\mathbf{w}^T \mathbf{x}_i + b) \geq 1, \quad i = 1, 2, \dots, n. \quad (1)$$

Lemma 1 (Correctness of SVMs [37]). *Given positive samples X^+ which are linearly separable from negative samples X^- , SVMs produce, under computations of infinite precision, a half-space h s.t. $\forall \mathbf{x} \in X^+, h(\mathbf{x}) > 0$ and $\forall \mathbf{x} \in X^-, h(\mathbf{x}) < 0$.*

Corollary 1 (Separation of Convex Hulls [1]). *The half-space h in Lemma 1 satisfies that $\forall \mathbf{x} \in \text{conv}(X^+), h(\mathbf{x}) > 0$ and $\forall \mathbf{x} \in \text{conv}(X^-), h(\mathbf{x}) < 0$.*




Fig. 1. Mapping from a two-dimensional input space into a three-dimensional feature space with linear separation thereof.

Nonlinear SVMs. When ϕ and ψ are formulas over nonlinear arithmetic, often after sampling X , it is not possible to find a linearly separable hyperplane in the common variables. However, a nonlinear surface that can be described as a linear hyperplane in the space of monomials of bounded degree may separate X^+ and X^- . The above construction is generalized by introducing a transformation from \mathbb{R}^d to $\mathbb{R}^{\tilde{d}}$, the vector space of monomials in the common variables up to some bounded degree, with $y_i (\mathbf{w}^T \mathbf{x}_i + b) \geq 1$ in (1) replaced by $y_i (\mathbf{w}^T \Phi(\mathbf{x}_i) + b) \geq 1$, where Φ is a linear expression in monomials in the common variables up to a bounded degree. Here, the vectors $\Phi(\mathbf{x})$ span the *feature space*.

Consider the Lagrangian dual [3] of the modified optimization problem:

$$\begin{aligned} &\underset{\alpha}{\text{minimize}} \quad \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j y_i y_j \Phi(\mathbf{x}_i)^T \Phi(\mathbf{x}_j) - \sum_{i=1}^n \alpha_i \\ &\text{subject to} \quad \sum_{i=1}^n \alpha_i y_i = 0, \quad \text{and} \quad \alpha_i \geq 0 \text{ for } i = 1, 2, \dots, n. \end{aligned}$$

A *kernel function* $\kappa: \mathbb{R}^d \times \mathbb{R}^d \mapsto \mathbb{R}$ is defined as $\kappa(\mathbf{x}, \mathbf{x}') \hat{=} \Phi(\mathbf{x})^T \Phi(\mathbf{x}')$. The introduction of the dual problem and the kernel function [3] reduces the computational complexity essentially from $\mathcal{O}(\tilde{d})$ down to $\mathcal{O}(d)$. For the sake of post-

verifying a candidate interpolant given by SVMs, we adopt an inhomogeneous polynomial kernel function of the form

$$\kappa(\mathbf{x}, \mathbf{x}') \triangleq (\beta \mathbf{x}^T \mathbf{x}' + \theta)^m,$$

where m is the polynomial degree describing complexity of the feature space, $\theta \geq 0$ is a parameter trading off the influence of higher-order versus lower-order terms in the polynomial, and β is a scalar parameter. Henceforth, the optimal-margin classifier (if there exists one) can be derived as $\mathbf{w}^T \Phi(\mathbf{x}) = \sum_{i=1}^n \alpha_i \kappa(\mathbf{x}_i, \mathbf{x}) = 0$, with \mathbf{x}_i being a support vector iff $\alpha_i > 0$. In practice, usually a large amount of α_i s turn out to be zero and this leads to a simple representation of a classifier. Figure 1 illustrates the intuitive idea of the transformation from the original input space to the feature space. We will show in the sequel that the resulting classifier can be viewed as a candidate interpolant, while its optimal-margin feature contributes to a certain “medium” logical strength of the interpolant, which is thus robust to perturbations (in the feature space) in the formulae to be interpolated.

3 Learning Interpolants

In this section, we present the NIL algorithm for synthesizing nontrivial (reverse) Craig interpolants for the quantifier-free theory of nonlinear arithmetic. It takes as input a pair $\langle \phi, \psi \rangle$ of formulas in \mathcal{T}_P as well as a positive integer m , and aims to generate an interpolant I of maximum degree m , i.e., $I \in \mathbb{R}[\mathbf{x}]_m$, if it exists, such that $\phi \models_{\mathcal{T}_P} I$ and $I \wedge \psi \models_{\mathcal{T}_P} \perp$. Here, $\langle \phi, \psi \rangle$ can be decorated as $\langle \phi(\mathbf{x}, \mathbf{y}), \psi(\mathbf{x}, \mathbf{z}) \rangle$ with variables involved in the predicates, and thus \mathbf{x} denotes variables that are common to ϕ and ψ . In the sequel, we drop the subscript \mathcal{T}_P in $\models_{\mathcal{T}_P}$ and $\llbracket \cdot \rrbracket_{\mathcal{T}_P}$ wherever the context is unambiguous.

Due to the decidability of the first-order theory of real-closed fields established by Tarski [44], \mathcal{T}_P admits *quantifier elimination* (QE). This means that the satisfiability of any formula in \mathcal{T}_P can be decided (in doubly exponential time in the number of variables for the worst case). If the formula is satisfiable, models satisfying the formula can also be constructed algorithmically (following the same time complexity). Though the introduction of general forms of transcendental functions renders the underlying theory undecidable, there does exist certain extension of \mathcal{T}_P with transcendental functions (involving exponential functions, logarithms and trigonometric functions), e.g. that identified by Strzeboński in [43] and references therein, which still admits QE. This allows a straightforward extension of NIL to such a decidable fragment involving transcendental functions. Specifically, the decidability remains when the transcendental functions involved are real univariate exp-log functions [41] or tame elementary functions [42] which admit a real root isolation algorithm.

Algorithm NIL: Learning nonlinear interpolant

```

input :  $\phi$  and  $\psi$  in  $\mathcal{T}_F$  over common variables  $\mathbf{x}$ ;
          $m$ , degree of the polynomial kernel, and hence
         maximum degree of the interpolant.
/* checking unsatisfiability */
1 if  $\phi \wedge \psi \not\models \perp$  then
  | /* no interpolant exists */
2 | abort;
/* generating initial sample points */
3  $\langle X^+, X^- \rangle \leftarrow \text{Sampling}(\phi, \psi)$ ;
/* counterexample-guided learning */
4 while  $\top$  do
  | /* generating a classifier by SVMs */
  |  $C \leftarrow \text{SVM}(X^+, X^-, m)$ ;
  | /* checking classification result */
  | if  $C = \text{Failed}$  then
  | | /* no interpolant exists in  $\mathbb{R}[\mathbf{x}]_m$  */
  | | abort;
  | /* classifier as candidate interpolant */
  | else
  | |  $I \leftarrow C$ ;
  | /* valid interpolant found */
  | if  $\phi \models I$  and  $I \wedge \psi \models \perp$  then
  | | return  $I$ ;
  | /* adding counterexamples */
  | else
  | |  $X^+ \leftarrow X^+ \uplus \text{FindInstance}(\phi \wedge \neg I)$ ;
  | |  $X^- \leftarrow X^- \uplus \text{FindInstance}(I \wedge \psi)$ ;

```




Fig. 2. In NIL, a candidate interpolant (black line as its boundary) is refined to an actual one (red line as its boundary) by adding a counterexample (red dot). (Color figure online)




Fig. 3. In NIL_δ , a counterexample (red dot) stays at least a distance of δ away from the candidate interpolant (black line as its boundary) to be refined, leading to an interpolant (red line as its boundary) with tolerance δ . (Color figure online)

3.1 The Core Algorithm

The basic idea of NIL is to view interpolants as classifiers and use SVMs with the kernel trick to perform effective classification. The algorithm is based on the sampling-guessing-refining technique: in each iteration, it is fed with a classifier (candidate interpolant) for a finite set of sample points from $\llbracket \phi \rrbracket$ and $\llbracket \psi \rrbracket$ (line 5), and verify the candidate (line 10) by checking the entailment problem that defines an interpolant (as in Definition 1). If the verification succeeds, the interpolant is returned as the final result. Otherwise, a set of counterexamples is obtained (line 13 and 14) as new sample points to further refine the classifier. In what follows, we explain the steps of the interpolation procedure in more detail.

Initial Sampling. The algorithm begins by checking the satisfiability of $\phi \wedge \psi$. If the formula is satisfiable, it is then impossible to find an interpolant, and the algorithm stops declaring no interpolant exists.

Next, the algorithm attempts to sample points from both $\llbracket \phi \rrbracket$ and $\llbracket \psi \rrbracket$. This initial sampling stage can usually be done efficiently using the Monte Carlo method, e.g. by (uniformly) scattering a number of random points over certain bounded range and then selecting those fall in $\llbracket \phi \rrbracket$ and $\llbracket \psi \rrbracket$ respectively. However, this method fails when one or both of the predicates is very unlikely to be satisfied. One common example is when the predicate involves equalities. For such situations, solving the satisfiability problem using QE is guaranteed to succeed in producing the sample points.

To meet the condition that the generated interpolant can only involve symbols that are common to ϕ and ψ , we can project the points sampled from $\llbracket\phi\rrbracket$ (resp. $\llbracket\psi\rrbracket$) to the space of \mathbf{x} by simply dropping the components that pertain to \mathbf{y} (resp. \mathbf{z}) and thereby obtain sample points in X^+ (resp. X^-).

Entailment Checking. The correctness of SVM given in Lemma 1 only guarantees that the candidate interpolant separates the finite set of points sampled from $\llbracket\phi\rrbracket$ and $\llbracket\psi\rrbracket$, not necessarily the entirety of the two sets. Hence, post-verification by checking the entailment problem (line 10) is needed for the candidate to be claimed as an interpolant of ϕ and ψ . This can be achieved by solving the equivalent QE problems $\forall \mathbf{x}. \phi(\mathbf{x}, \mathbf{y})|_{\mathbf{x}} \implies I(\mathbf{x})$ and $\forall \mathbf{x}. I(\mathbf{x}) \wedge \psi(\mathbf{x}, \mathbf{z})|_{\mathbf{x}} \implies \perp$, where $\cdot|_{\mathbf{x}}$ is the *projection* to the common space over \mathbf{x} . The candidate will be returned as an actual interpolant if both formulae reduce to \top after eliminating the universal quantifiers. The satisfiability checking at line 1 can be solved analogously. Granted, the entailment checking can also be encoded in SMT techniques by asking the satisfiability of the negation of the universally quantified predicates, however, limitations of current SMT solvers in nonlinear arithmetic hinders them from being practically used in our framework, as demonstrated later in Sect. 5.

Counterexample Generation. If a candidate interpolant cannot be verified as an actual one, then at least one witness can be found as a counterexample to that candidate, which can be added to the set of sample points in the next iteration to refine further candidates (cf. Fig. 2). Multiple counterexamples can be obtained at a time thereby effectively reducing the number of future iterations.

In general, we have little control over which counterexample will be returned by QE. In the worst case, the counterexample can lie almost exactly on the hyperplane found by SVM. This poses issues for the termination of the algorithm. We will address this theoretical issue by slightly modifying the algorithm, as explained in Sects. 3.3 and 4.

3.2 Comparison with the Naïve QE-Based Method

Simply performing QE on $\exists \mathbf{y}. \phi(\mathbf{x}, \mathbf{y})$ yields already an interpolant for mutually contradictory ϕ and ψ . Such an interpolant is actually the *strongest* in the sense of [13], which presents an ordered family of interpolation systems due to the logical strength of the synthesized interpolants. Dually, the negation of the result when performing QE over $\exists \mathbf{z}. \psi(\mathbf{x}, \mathbf{z})$ is the *weakest* interpolant. However, as argued by D’Silva et al. in [13], a good interpolant (approximation of ϕ or ψ)—when computing invariants of transition systems using interpolation-based model checking—should be coarse enough to enable rapid convergence but strong enough to be contained within the weakest inductive invariant. In contrast, the advantages of NIL are two-fold: first, it produces better interpolants (in the above sense) featuring “medium” strength (due to the way optimal-margin classifier is defined) which are thus more effective in practical use and furthermore resilient to perturbations in ϕ and ψ (i.e., the robustness shown later in Sect. 6); second, NIL always returns a single polynomial inequality as the interpolant which is

often simpler than that derived from the naïve QE-based method, where the direct projection of $\phi(\mathbf{x}, \mathbf{y})$ onto the common space over \mathbf{x} can be as complex as the original ϕ .

These issues can be avoided by combining this method with a template-based approach, which in turn introduces fresh quantifiers over unknown parameters to be eliminated. Note that in NIL the candidate interpolants $I \in \mathbb{R}[\mathbf{x}]_m$ under verification are polynomials without unknown parameters, and therefore, in contrast to performing QE over an assumed template, the learning-based technique can practically generate polynomial interpolants of higher degrees (with acceptable rounds of iterations). For example, NIL is able to synthesize an interpolant of degree 7 over 2 variables (depicted later in Fig. 4(b)), which would require a polynomial template with $\binom{7+2}{2} = 36$ unknown parameters that goes far beyond the capability of QE procedures.

On the other hand, performing QE within every iteration of the learning process, for entailment checking and generating counterexamples, limits the efficiency of the proposed method, thereby confining NIL currently to applications only of small scales. Potential solutions to the efficiency bottleneck will be discussed in Sect. 5.

3.3 Variants of NIL

While the above basic algorithm is already effective in practice (as demonstrated in Sect. 5), it is guaranteed to terminate only when there is an interpolant with positive functional margin between $\llbracket\phi\rrbracket$ and $\llbracket\psi\rrbracket$. In this section, we present two variants of the algorithm that have nicer theoretical properties in cases where the two sets are only separated by an interpolant with zero functional margin, e.g., cases where $\llbracket\phi\rrbracket$ and $\llbracket\psi\rrbracket$ share adjacent or even coincident boundaries.

Entailment Checking with Tolerance δ . When performing entailment checking for a candidate interpolant I , instead of using, e.g., the formula $p(\mathbf{x}) \geq 0$ for I , we can introduce a tolerance of δ . That is, we check the satisfiability of $\phi \wedge (p(\mathbf{x}) < -\delta)$ and $(p(\mathbf{x}) \geq \delta) \wedge \psi$ instead of the original $\phi \wedge (p(\mathbf{x}) < 0)$ and $(p(\mathbf{x}) \geq 0) \wedge \psi$. This means that a candidate that is an interpolant “up to a tolerance of δ ” will be returned as a true interpolant, which may be acceptable in some applications. If the candidate interpolant is still not verified, the counterexample is guaranteed to be at least a distance of δ away from the separating hyperplane. Note the distance δ is taken in the feature space \mathbb{R}^d , not in the original space. We let $\text{NIL}_\delta(\phi, \psi, m)$ denote the version of NIL with this modification (cf. Fig. 3). In the next section, we show $\text{NIL}_\delta(\phi, \psi, m)$ terminates as long as $\llbracket\phi\rrbracket$ and $\llbracket\psi\rrbracket$ are bounded, including the case where they are separated only by interpolants of functional margin zero.

Varying Tolerance During the Execution. A further refinement of the algorithm can be made by varying the tolerance δ during the execution. We also introduce a bounding box B of the varying size to handle unbounded cases. Define algorithm $\text{NIL}_{\delta, B}^*(\phi, \psi, m)$ as follows. Let $\delta_1 = \delta$ and $B_1 = B$. For each iteration i , execute the core algorithm, except that the counterexample must be

a distance of at least δ_i away from the separating boundary, and have absolute value in each dimension at most B (both in $\mathbb{R}^{\bar{d}}$). After the termination of iteration i , begin iteration $i + 1$ with $\delta_{i+1} = \delta_i/2$ and $B_{i+1} = 2B_i$. This continues until an interpolant is found or until a pre-specified cutoff. For any $[\phi]$ and $[\psi]$ (without the boundedness condition), this variant of the algorithm *converges* to an interpolant in the limit, which will be made precise in the next section.

4 Soundness, Completeness and Convergence

In this section, we present theoretical results obtained for the basic NIL algorithm and its variants. Proofs are available in the appendix of [6].

First, the basic algorithm is sound, as captured by Theorem 1.

Theorem 1 (Soundness of NIL). *NIL(ϕ, ψ, m) terminates and returns I if and only if I is an interpolant in $\mathbb{R}[\mathbf{x}]_m$ of ϕ and ψ .*

Under certain conditions, the algorithm is also terminating (and hence complete). We prove two such situations below. In both cases, we require boundedness of the two sets that we want to separate. In the first case, there exists an interpolant with positive functional margin between the two sets.

Theorem 2 (Conditional Completeness of NIL). *If $[\phi]$ and $[\psi]$ are bounded and there exists an interpolant in $\mathbb{R}[\mathbf{x}]_m$ of ϕ and ψ with positive functional margin γ when mapped to $\mathbb{R}^{\bar{d}}$, then NIL(ϕ, ψ, m) terminates and returns an interpolant I of ϕ and ψ .*

The standard algorithm is not guaranteed to terminate when $[\phi]$ and $[\psi]$ are only separated by interpolants of functional margin zero. However, the modified algorithm $\text{NIL}_\delta(\phi, \psi, m)$ does terminate (with the cost that the resulting answer is an interpolant with tolerance δ).

Theorem 3 (Completeness of NIL_δ with zero margin). *If $[\phi]$ and $[\psi]$ are bounded, and $\delta > 0$, then $\text{NIL}_\delta(\phi, \psi, m)$ terminates. It returns an interpolant I of ϕ and ψ with tolerance δ whenever such an interpolant exists.*

By iteratively decreasing δ during the execution of the algorithm, as well as introducing an iteratively increasing bounding box, as in $\text{NIL}_{\delta, B}^*(\phi, \psi, m)$, we can obtain more and more accurate candidate interpolants. We now show that this algorithm *converges* to an interpolant without restrictions on ϕ and ψ . We first make this convergence property precise in the following definition.

Definition 2 (Convergence of a sequence of equations to an interpolant). *Given two sets $[\phi]$ and $[\psi]$ that we want to separate, and an infinite sequence of equations I_1, I_2, \dots , we say the sequence I_n converges to an interpolant of ϕ and ψ if, for each point p in the interior of $[\phi]$ or $[\psi]$, there exists some integer K_p such that I_k classifies p correctly for all $k \geq K_p$.*

Theorem 4 (Convergence of $\text{NIL}_{\delta, B}^*$). *Given two regions $[\phi]$ and $[\psi]$. Suppose there exists an interpolant of ϕ and ψ , then the infinite sequence of candidates produced by $\text{NIL}_{\delta, B}^*(\phi, \psi, m)$ converges to an interpolant in the sense of Definition 2.*

5 Implementation and Experiments

5.1 Implementation Issues

We have implemented the core algorithm NIL as a prototype¹ in Wolfram Mathematica with LIBSVM [5] being integrated as an engine to perform SVM classifications. Despite featuring no completeness for adjacent $\llbracket\phi\rrbracket$ and $\llbracket\psi\rrbracket$ nor convergence for unbounded $\llbracket\phi\rrbracket$ or $\llbracket\psi\rrbracket$, the standard NIL algorithm yields already promising results as shown later in the experiments. Key Mathematica functions that are utilized include REDUCE, for entailment checking, e.g., the unsatisfiability checking of $\phi \wedge \psi$ and the post-verification of a candidate interpolant, and FINDINSTANCE, for generating counterexamples and sampling initial points (when the random sampling strategy fails). The REDUCE command implements a decision procedure for \mathcal{T}_P and its appropriate extension to catering for transcendental functions (cf. [43]) based on *cylindrical algebraic decomposition* (CAD), due to Collins [9]. The underlying quantifier-elimination procedure, albeit inducing rather high computation complexity, cannot in practice be replaced by SMT-solving techniques (by checking the negation of a universally quantified predicate) as in the linear arithmetic. For instance, the off-the-shelf SMT solver Z3 fails to accomplish our tasks particularly when the coefficients occurring in the entailment problem to be checked get larger².

Numerical Errors and Rounding. LIBSVM conducts floating-point computations for solving the optimization problems induced by SVMs and consequently yields numerical errors occurring in the candidate interpolants. Such numerical errors may block an otherwise valid interpolant from being verified as an actual one and additionally bring down the simplicity and thereby the effectiveness of the synthesized interpolant, thus not very often proving humans with clear-cut understanding. This is a common issue for approaches that reduce the interpolation problem to numerical solving techniques, e.g. SDP solvers exploited in [11, 14, 33], while an established method to tackle it is known as *rational recovery* [29, 47], which retrieves the nearest rational number from the continued fraction representation of its floating-point approximation at any given accuracy (see e.g. [47] for theoretical guarantees and [33] for applications in interpolation). The algorithm implementing rational recovery has been integrated in our implementation and the consequent benefits are two-fold: (i) NIL can now cope with interpolation tasks where only exact coefficients suffice to constitute an actual interpolant while any numerical error therein will render the interpolant invalid, e.g., cases where $\llbracket\phi\rrbracket$ and $\llbracket\psi\rrbracket$ share parallel, adjacent, or even coincident boundaries, as demonstrated later by examples with ID 10–17 in Table 1; (ii) rationalizing coefficients moreover facilitates simplifications over all of the candidate interpolants and therefore practically accelerating the entailment checking and counterexample generation processes, which in return yields simpler interpolants, as shown in Table 2 in the following section.

¹ Available at <http://lcs.ios.ac.cn/~chenms/tools/NIL.tar.bz2>.

² As can be also observed at <https://github.com/Z3Prover/z3/issues/1765>.

5.2 Benchmark and Experimental Results

Table 1 collects a group of benchmark examples from the literature on synthesizing nonlinear interpolants as well as some geometrically contrived ones. All of the experiments have been evaluated on a 3.6 GHz Intel Core-i7 processor with 8 GB RAM running 64-bit Ubuntu 16.04.

In Table 1, we group the set of examples into four categories comprising 20 cases in total. For each example, **ID** numbers the case, ϕ , ψ and **I** represent the two formulas to be interpolated and the synthesized interpolant by our method respectively, while **Time/s** indicates the total time in seconds for interpolation. The categories are described as follows, and the visualization of a selected set of typical examples thereof is further depicted in Fig. 4.

Cat. I: with/without rounding. This category includes 9 cases, for which our method generates the polynomial interpolants correctly with or without the rounding operation.

Cat. II: with rounding. For cases 10 to 17 in this category, where $\llbracket\phi\rrbracket$ and $\llbracket\psi\rrbracket$ share parallel, adjacent, or even coincident boundaries, our method produces interpolants successfully with the rounding process based on rational recovery.

Cat. III: beyond polynomials. This category encloses two cases beyond the theory \mathcal{T}_P of polynomials: for case 18, a verified polynomial interpolant is obtained in spite of the transcendental term in ψ ; while for case 19, the SVM classification fails since $\llbracket\phi\rrbracket$ and $\llbracket\psi\rrbracket$ are not linearly separable in any finite-dimensional feature space and hence no polynomial interpolant exists for this example. Note that our counterexample-guided learning framework admits a straightforward extension to a decidable fragment of more general nonlinear theories involving transcendental functions, as investigated in [43].

Cat. IV: unbalanced. The case 20, called Unbalanced, instantiates a particular scenario where ϕ and ψ have extraordinary “unbalanced” number of models that make them true respectively. For this example, there are an infinite number of models satisfying ϕ yet one single model (i.e., $x = 0$) satisfying ψ . The training process in SVMs may fail when encountering extremely unbalanced number of positive/negative samples. This is solved by specifying a *weight* factor for the positive set of samples as the number of negative ones, and dually for the other way around, to balance biased number of training samples before triggering the classification. Such a balancing trick is supported in LIBSVM.

Remark that examples named CAV13-1/3/4 are taken from [11] (and the latter two originally from [28] and [18] respectively), where interpolation is applied to discovering *inductive invariants* in the verification of programs and hybrid systems. For instance, CAV13-3 is a program fragment describing an accelerating car and the synthesized interpolant by NIL suffices to prove the safety property of the car concerning its velocity.

Applicability and Comparison with Existing Approaches. As shown in Table 1, our learning-based technique succeeds in all of the benchmark examples

Table 1. Benchmark examples for synthesizing nonlinear interpolants.

Category	ID Name	ϕ	ψ	I	Times
1	Dumny	$x < -1$	$x > 1$	$x < 0$	0.11
		$y - x^2 - 1 = 0$	$y + x^2 + 1 = 0$	$-y < 0$	0.21
2	Necklace			$\frac{x^3}{x^4} - \frac{x^3 y}{x^4} + \frac{2y^2}{45} - \frac{y}{170} - \frac{2}{9} + \frac{223}{316} + \frac{x^3 y^2}{170} + \frac{y^4}{9}$	
		$(x+4)^2 + y^2 - 1 \leq 0 \vee (x-4)^2 + y^2 - 9 \leq 0$	$x^2 + y^2 - 64 \leq 0 \wedge (x+4)^2 + y^2 - 9 \geq 0 \wedge (x-4)^2 + y^2 - 1 \geq 0$	$\pi(\frac{89}{48} + \frac{74}{55}) + \frac{146}{5} + \frac{3^3 y^2}{30} + \frac{y}{37} + 1 < 0$	0.33
3	Face			$x^2 - 2xy^2 + 3xz - y^2$	
		$-xy + x^2 - 1 \geq 0 \wedge \frac{120}{(-x^6 - y^6) + z^2 z^2} - x^2 + \frac{1}{6}(x^4 + 2xz^2 y^2 + y^4) + y^2 z^2 - y^2 - 4 \leq 0$	$xy^2 + x^2 - 1 \geq 0 \wedge \frac{120}{(-x^6 - y^6) + z^2 z^2} - x^2 + \frac{1}{6}(x^4 + 2xz^2 y^2 + y^4) + y^2 z^2 - y^2 - 4 \leq 0$	$-\frac{x^4}{160} + \frac{x^3}{170} - \frac{1}{113} + x^2 \left(-\frac{y^2}{225} + \frac{y}{76} + \frac{3}{57} \right) + x \left(\frac{y^3}{259} + \frac{y^2}{43} + \frac{5y}{51} - \frac{1}{316} \right) - \frac{y^4}{183} - \frac{y^3}{94} + \frac{y^2}{14} + \frac{2y}{205} - 1 < 0$	140.62
4	Twisted			$\frac{x^7}{27} + x^6 \left(-\frac{y}{5} - \frac{1}{96} \right) + x^5 \left(\frac{2y^2}{9} - \frac{y}{32} - \frac{1}{2} \right) + x^4 \left(-\frac{2y^3}{9} + \frac{y^2}{3} + \frac{1}{31} \right) + \frac{1}{3} \frac{y^3}{11} + \frac{10y^2}{13} - \frac{10y}{18} + \frac{15}{16}$	
		$(x^2 + y^2 - 3.8025 \leq 0 \wedge y \geq 0 \vee (x-1)^2 + y^2 - 0.9025 \leq 0) \wedge (x-1)^2 + y^2 - 0.09 > 0 \wedge (x+1)^2 + y^2 - 1.1025 \geq 0 \vee (x+1)^2 + y^2 - \frac{1}{25} \leq 0$	$(-3.8025 + x^2 + y^2 \leq 0 \wedge -y \geq 0 \vee -0.9025 + (-1-x)^2 + y^2 \leq 0) \wedge -0.09 + (-1-x)^2 + y^2 > 0 \wedge -1.1025 + (1-x)^2 + y^2 \geq 0 \vee -\frac{1}{25} + (1-x)^2 + y^2 \leq 0$	$x^2 \left(-\frac{y^6}{25} - \frac{18}{32} + \frac{y^2}{10} - \frac{1}{32} \right) + x \left(\frac{y^6}{71} + \frac{2y^4}{11} - \frac{y^3}{25} - \frac{y}{45} - \frac{3}{8} \right) + \frac{y^6}{6} - \frac{y^5}{4} + \frac{y^4}{3} - \frac{y^3}{2} - \frac{y}{1} < 0$	48.82
5	Ultimate without rounding			$\frac{y^6}{48} - \frac{y^5}{7} + \frac{y^4}{6} - \frac{y^3}{2} - \frac{y}{1} < 0$	
		$-x^2 + 4x - y^4 + (x+y)^2 - 8y \leq 0 \wedge -xy^2(x-y)^2 + 100(x+y)^2 - 3000 \geq 0$	$xy^2 + x^2 - 1 \geq 0 \wedge \frac{120}{(-x^6 - y^6) + z^2 z^2} - x^2 + \frac{1}{6}(x^4 + 2xz^2 y^2 + y^4) + y^2 z^2 - y^2 - 4 \leq 0$	$1 - \frac{321}{221} - \frac{y^2}{52} < 0$	0.16
6	UCAR16-[14]			$103x^4 + x^2(140y^2 + 24y(5z+7) + 35z(3z+8)) + 2(70y^2x + 5y^2(12z^2 + 21z + 28) - 14y(6z^3 + 5z^2 + 10) - 35(3z^4 + 8z^2 + 4z - 9)) < 14z(20x^2(z+1) + 10y^2(z+2) - 3y(4z^2 - 3z + 4) - 20z(z^2 + 2))$	
		$-x^2 + 4x - y^4 + (x+y)^2 - 8y \leq 0$	$-3x_1^2 - x_2^2 + 1 \geq 0 \wedge x_2 - x^2 \geq 0$	$-1 + \frac{x^2}{2} - \frac{y}{3} + \frac{2y}{4} < 0$	3.25
7	CAV13-[11]			$103x^4 + x^2(140y^2 + 24y(5z+7) + 35z(3z+8)) + 2(70y^2x + 5y^2(12z^2 + 21z + 28) - 14y(6z^3 + 5z^2 + 10) - 35(3z^4 + 8z^2 + 4z - 9)) < 14z(20x^2(z+1) + 10y^2(z+2) - 3y(4z^2 - 3z + 4) - 20z(z^2 + 2))$	
		$6 + bx + 1 - y = 0$	$x^2 - 2y^2 - 4 > 0$	$-1 + \frac{x^2}{2} - \frac{y}{3} + \frac{2y}{4} < 0$	365.89
8	CAV12-[11]			$103x^4 + x^2(140y^2 + 24y(5z+7) + 35z(3z+8)) + 2(70y^2x + 5y^2(12z^2 + 21z + 28) - 14y(6z^3 + 5z^2 + 10) - 35(3z^4 + 8z^2 + 4z - 9)) < 14z(20x^2(z+1) + 10y^2(z+2) - 3y(4z^2 - 3z + 4) - 20z(z^2 + 2))$	
		$1.2x^2 + y^2 + xz = 0$	$20 - 3x^2 - 4y^3 - 10z^2 \geq 0 \wedge x^2 + y^2 - x - 1 = 0$	$10y^2(z+2) - 3y(4z^2 - 3z + 4) - 20z(z^2 + 2)$	46.63
9	CAV13-[11]			$-1 + \frac{x^2}{2} - \frac{y}{3} + \frac{2y}{4} < 0$	
		$vc < 40.61 \wedge fa = 0.5418vc \wedge vc_1 = vc + ac$	$vc_1 \geq 49.61$	$-1 + \frac{x^2}{2} - \frac{y}{3} + \frac{2y}{4} < 0$	4.90
10	Parallel parabola			$\frac{1}{3} + x^2 < y$	
		$y - x^2 - 1 \geq 0$	$y - x^2 < 0$	$\frac{1}{3} + x^2 < y$	2.46
11	Shaper-[13]			$x < y$	
		$y - 1 < 0$	$x^2 + y^2 - 1 \leq 0$	$2 + y < y^2$	2.19
12	Shaper2-[13]			$y > 0$	
		$y - x > 0 \wedge x + y > 0$	$y + x^2 < 0$	$2 + y < y^2$	2.38
13	Concider			$y > 0$	
		$x + y^2 > 0 \wedge x + y > 0$	$y + x^2 < 0$	$2 + y < y^2$	0.38
14	Adjacent			$(\frac{5}{2} + y)^2 > 0$	
		$x + y^2 > 0 \wedge x + y > 0$	$y + x^2 < 0$	$2 + y < y^2$	0.25
with rounding	16	UCAR16-[14]		$10y^2(z+2) - 3y(4z^2 - 3z + 4) - 20z(z^2 + 2)$	
			$-y_1 + x_1 - 2 \geq 0 \wedge 2x_2 - x_1 - 1 > 0 \wedge -y_1^2 - x_1^2 + 2x_1y_1 - 2y_1 + 2x_1 \geq 0 \wedge x_1 < x_2$	$-\frac{1}{2} - 4x_2^2 + 4x_2x_1 + 3x_1 - 6x_2 - 2 \geq 0 \wedge x_1 < x_2$	$10y^2(z+2) - 3y(4z^2 - 3z + 4) - 20z(z^2 + 2)$
17	CAV13-[11]			$2x_1 + 2y_1 - 1 \geq 0 \wedge x_1 - x_2 - 1 = 0 \wedge x_1 + 2y_1 - x_1 = 0$	
		$2x_1 + 2y_1 - 1 \geq 0 \wedge x_1 - x_2 - 1 = 0 \wedge x_1 + 2y_1 - x_1 = 0$	$-x_1^2 - x_2^2 + 2x_1y_1 - 2y_1 + 2x_1 \geq 0$	$2x_1 + 4y_1 > 5$	3.10
beyond polynomials	18	TACAS16-[15]		$2x_1 + 2y_1 - 1 \geq 0 \wedge x_1 - x_2 - 1 = 0 \wedge x_1 + 2y_1 - x_1 = 0$	
			$y - x^2 \geq 0$	$y + \cos x - 0.8 \leq 0$	$15x^2 < 4 + 20y$
unbalanced	19	Transcendental		$\sin x \leq 0.4$	
			$x = 0$	$x = 0$	$x^2 > 0$
20	Unbalanced			$x^2 > 0$	
		$x \geq 0 \vee x < 0$	$x = 0$	$x^2 > 0$	0.11




Fig. 4. Visualization in NIL on a selected set of examples. Legends: gray region: $[\phi]$, blue region: $[\psi]$, pink region: $[I]$ with a valid interpolant I , red dots: X^+ , blue dots: X^- , circled dots: support vectors. Sample points are hidden in 3D-graphics for a clear presentation.

that admit polynomial interpolants. Due to theoretical limitations of existing approaches as elaborated in Sect. 1, none of the aforementioned methods can cope with as many cases in Table 1 as NIL can. For instances, the Twisted example as depicted in Fig. 4(c) falls beyond the scope of concave quadratic formulas and thus cannot be addressed by the approach in [14], while the Parallel parabola example as shown in Fig. 4(d) needs an infinite combination of linear constraints as an interpolant when performing the technique in [15] and hence not of practical use, to name just a few. Moreover, we list in Table 2 a comparison of the synthesized interpolants against works where the benchmark examples are collected from. As being immediately obvious from Table 2, our technique

Table 2. Comparison of the synthesized interpolants.

Name	Interpolants by NIL	Interpolants from the sources
ICAR16-1 [14]	$1 - \frac{3x_1}{4} - \frac{x_2}{2} < 0$	$-3 + 2x_1 + x_1^2 + \frac{1}{2}x_2^2 > 0$
CAV13-1 [11]	$-1 + \frac{x_2^2}{2} - \frac{y}{3} + \frac{xy}{3} - \frac{y^2}{4} < 0$	$436.45(x^2 - 2y^2 - 4) + \frac{1}{2} \leq 0$ $-14629.26 + 2983.44x_3 + 10972.97x_3^2 +$ $297.62x_2 + 297.64x_2x_3 + 0.02x_2x_3^2 + 9625.61x_2^2 -$ $1161.80x_2^2x_3 + 0.01x_2^2x_3^2 + 811.93x_2^3 +$ $105x^4 + x^2(140y^2 + 24y(5z + 7) + 35z(3z + 8)) +$ $2(70y^3z + 5y^2(12z^2 + 21z + 28) - 14y(6z^3 + 5z^2 +$ $10) - 35(3z^4 + 8z^2 + 4z - 9)) < 14x(20x^2(z + 1) +$ $10y^2(z + 2) - 3y(4z^2 - 5z + 4) - 20z(z^2 + 2))$
CAV13-2 [11]		$2745.14x_2^4 - 10648.11x_1 + 3101.42x_1x_3 +$ $8646.17x_1x_3^2 + 511.84x_1x_2 - 1034x_1x_2x_3 +$ $0.02x_1x_2x_3^2 + 9233.66x_1x_2^2 + 1342.55x_1x_2^2x_3 -$ $138.70x_1x_1^2x_3^2 + 11476.61x_1^2 - 3737.70x_1^2x_3 +$ $4071.65x_1^2x_3^2 - 2153.00x_1x_2x_2 + 373.14x_1^2x_2x_3 +$ $7616.18x_1^2x_2^2 + 8950.77x_1^3 + 1937.92x_1^3x_3 -$ $64.07x_1^3x_2 + 4827.25x_1^4 > 0$
CAV13-3 [11]	$-1 + \frac{2vc_1}{99} < 0$	$-1.3983vc_1 + 69.358 > 0$
Sharper-1 [33]	$2 + y < y^2$	$34y^2 - 68y - 102 \geq 0$
Sharper-2 [33]	$y > 0$	$8y + 4x^2 > 0$
ICAR16-2 [14]	$x_1 < x_2$	$-x_1 + x_2 > 0$
CAV13-4 [11]	$2xa + 4ya > 5$	$716.77 + 1326.74(ya) + 1.33(ya)^2 + 433.90(ya)^3 +$ $668.16(xa) - 155.86(xa)(ya) + 317.29(xa)(ya)^2 +$ $222.00(xa)^2 + 592.39(xa)^2(ya) + 271.11(xa)^3 > 0$
TACAS16 [15]	$15x^2 < 4 + 20y$	$y > 1.8 \vee (0.59 \leq y \leq 1.8 \wedge -1.35 \leq x \leq 1.35) \vee$ $(0.09 \leq y < 0.59 \wedge -0.77 \leq x \leq 0.77) \vee$ $(y \geq 0 \wedge -0.3 \leq x \leq 0.3)$

often produces interpolants of simpler forms, particularly for examples CAV13-2, CAV13-4 and TACAS16. Such a simplicity benefits from both the rounding effect and the form of interpolant (i.e., a single polynomial inequality) that we tend to construct.

Bottleneck of Efficiency and Potential Solutions. The current implementation of NIL works promisingly for small examples; it does not scale to interpolation problems with large numbers of common variables, as reported in Table 1. The bottleneck stems from quantifier eliminations performed within every iteration of the learning process, for entailment checking and generating counterexamples. We pose here several potential solutions that are expected to significantly reduce computational efforts: (i) substitute general purpose QE procedure that perform CAD by the so-called variant quantifier-elimination (VQE) algorithm [21], which features singly-exponential complexity in the number of variables. This however requires a careful inspection of whether our problem meets the geometric conditions imposed by VQE; (ii) incorporate relaxation schemes, e.g., Lagrangian relaxation and sum-of-squares decompositions [34], and complement with QE only when the relaxation fails to produce desired results.




Fig. 5. ϵ -Face: introducing perturbations (with ϵ up to 0.5) in the Face example. The synthesized interpolant is resilient to any ϵ -perturbation in the radii satisfying $-0.5 \leq \epsilon \leq 0.5$.

6 Taming Perturbations in Parameters

An interpolant synthesized by the SVM-based technique features inherent robustness due to the way optimal-margin classifier is defined (Sect. 2). That is, the validity of such an interpolant is not easily perturbed by changes (in the feature space) in the formulae to be interpolated. It is straightforward in NIL to deal with interpolation problems under explicitly specified perturbations, which are treated as constraints over fresh variables. An example named ϵ -Face is depicted in Fig. 5, which perturbs $\langle \phi, \psi \rangle$ in the Face example as $\phi \hat{=} -0.5 \leq \epsilon_1 \leq 0.5 \wedge ((x+4)^2 + y^2 - (1+\epsilon_1)^2 \leq 0 \vee (x-4)^2 + y^2 - (1+\epsilon_1)^2 \leq 0)$ and $\psi \hat{=} -0.5 \leq \epsilon_2 \leq 0.5 \wedge x^2 + y^2 - 64 \leq 0 \wedge (x+4)^2 + y^2 - (3+\epsilon_2)^2 \geq 0 \wedge (x-4)^2 + y^2 - (3+\epsilon_2)^2 \geq 0$. The synthesized interpolant over common variables of ϕ and ψ is $\frac{x^4}{139} + \frac{x^3y}{268} + x^2 \left(\frac{y^2}{39} - \frac{11}{36} \right) + x \left(-\frac{y^3}{52} - \frac{y^2}{157} - \frac{y}{52} - \frac{1}{116} \right) + \frac{y^4}{25} - \frac{y^3}{182} + \frac{2y^2}{19} - \frac{y}{218} + 1 < 0$ which is hence resilient to any ϵ -perturbation in the radii satisfying $-0.5 \leq \epsilon \leq 0.5$, as illustrated in Fig. 5(b).

7 Conclusions

We have presented a unified, counterexample-guided method named NIL for generating polynomial interpolants over the general quantifier-free theory of nonlinear arithmetic. Our method is based on classification techniques with space transformations and kernel tricks as established in the community of machine-learning. We proved the soundness of NIL and proposed sufficient conditions for its completeness and convergence. The applicability and effectiveness of our technique are demonstrated experimentally on a collection of representative benchmarks from the literature, including those extracted from program verification. Experimental results indicated that our method suffices to address more interpolation tasks, including those with perturbations in parameters, and in many cases synthesizes simpler interpolants compared with existing approaches.

For future work, we would like to improve the efficiency of NIL by substituting the general purpose quantifier-elimination procedure with alternative methods previously discussed in Sect. 5. An extension of our approach to cater for the combination of nonlinear arithmetic with EUFs, by resorting to predicate-abstraction techniques [23], will be of particular interest. Additionally, we plan to investigate the performance of NIL over different classification techniques, e.g., the widespread regression-based methods [19], though SVMs are expected to be more competent concerning the robustness and predictability, as also observed in [37].

References

1. Bennett, K.P., Bredensteiner, E.J.: Duality and geometry in SVM classifiers. In: ICML 2000, pp. 57–64 (2000)
2. Bishop, C.M.: Pattern Recognition and Machine Learning, pp. 326–328. Springer, New York (2006)
3. Boser, B.E., Guyon, I., Vapnik, V.: A training algorithm for optimal margin classifiers. In: COLT 1992, pp. 144–152 (1992)
4. Bourbaki, N.: Topological Vector Spaces. Elements of Mathematics. Springer, Heidelberg (1987). <https://doi.org/10.1007/978-3-642-61715-7>
5. Chang, C., Lin, C.: LIBSVM: a library for support vector machines. ACM TIST **2**(3), 27:1–27:27 (2011)
6. Chen, M., Wang, J., An, J., Zhan, B., Kapur, D., Zhan, N.: NIL: learning nonlinear interpolants (full version). http://lcs.ios.ac.cn/~chenms/papers/CADE-27_FULL.pdf
7. Cimatti, A., Griggio, A., Sebastiani, R.: Efficient interpolant generation in satisfiability modulo theories. In: Ramakrishnan, C.R., Rehof, J. (eds.) TACAS 2008. LNCS, vol. 4963, pp. 397–412. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78800-3_30
8. Clarke, E., Grumberg, O., Jha, S., Lu, Y., Veith, H.: Counterexample-guided abstraction refinement. In: Emerson, E.A., Sistla, A.P. (eds.) CAV 2000. LNCS, vol. 1855, pp. 154–169. Springer, Heidelberg (2000). https://doi.org/10.1007/10722167_15
9. Collins, G.E.: Quantifier elimination for real closed fields by cylindrical algebraic decomposition. In: Brakhage, H. (ed.) GI-Fachtagung 1975. LNCS, vol. 33, pp. 134–183. Springer, Heidelberg (1975). https://doi.org/10.1007/3-540-07407-4_17
10. Craig, W.: Linear reasoning. A new form of the Herbrand-Gentzen theorem. J. Symb. Log. **22**(3), 250–268 (1957)
11. Dai, L., Xia, B., Zhan, N.: Generating non-linear interpolants by semidefinite programming. In: Sharygina, N., Veith, H. (eds.) CAV 2013. LNCS, vol. 8044, pp. 364–380. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-39799-8_25
12. Dathathri, S., Arechiga, N., Gao, S., Murray, R.M.: Learning-based abstractions for nonlinear constraint solving. In: IJCAI 2017, pp. 592–599 (2017)
13. D’Silva, V., Kroening, D., Purandare, M., Weissenbacher, G.: Interpolant strength. In: Barthe, G., Hermenegildo, M. (eds.) VMCAI 2010. LNCS, vol. 5944, pp. 129–145. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-11319-2_12

14. Gan, T., Dai, L., Xia, B., Zhan, N., Kapur, D., Chen, M.: Interpolant synthesis for quadratic polynomial inequalities and combination with *EUF*. In: Olivetti, N., Tiwari, A. (eds.) IJCAR 2016. LNCS (LNAI), vol. 9706, pp. 195–212. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-40229-1_14
15. Gao, S., Zufferey, D.: Interpolants in nonlinear theories over the reals. In: Chechik, M., Raskin, J.-F. (eds.) TACAS 2016. LNCS, vol. 9636, pp. 625–641. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49674-9_41
16. Gilbert, S.: A nullstellensatz and a positivstellensatz in semialgebraic geometry. *Math. Ann.* **207**(2), 87–97 (1974)
17. Graf, S., Saidi, H.: Construction of abstract state graphs with PVS. In: Grumberg, O. (ed.) CAV 1997. LNCS, vol. 1254, pp. 72–83. Springer, Heidelberg (1997). https://doi.org/10.1007/3-540-63166-6_10
18. Gulavani, B.S., Chakraborty, S., Nori, A.V., Rajamani, S.K.: Automatically refining abstract interpretations. In: Ramakrishnan, C.R., Rehof, J. (eds.) TACAS 2008. LNCS, vol. 4963, pp. 443–458. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78800-3_33
19. Hastie, T., Tibshirani, R., Friedman, J.: *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. SSS, 2nd edn. Springer, New York (2009). <https://doi.org/10.1007/978-0-387-84858-7>
20. Henzinger, T.A., Jhala, R., Majumdar, R., McMillan, K.L.: Abstractions from proofs. In: POPL 2004, pp. 232–244 (2004)
21. Hong, H., Din, M.S.E.: Variant quantifier elimination. *J. Symb. Comput.* **47**(7), 883–901 (2012)
22. Jha, S., Gulwani, S., Seshia, S.A., Tiwari, A.: Oracle-guided component-based program synthesis. In: ICSE 2010, pp. 215–224 (2010)
23. Jhala, R., Podelski, A., Rybalchenko, A.: Predicate abstraction for program verification. In: Clarke, E., Henzinger, T., Veith, H., Bloem, R. (eds.) *Handbook of Model Checking*, pp. 447–491. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-10575-8_15
24. Jung, Y., Lee, W., Wang, B.-Y., Yi, K.: Predicate generation for learning-based quantifier-free loop invariant inference. In: Abdulla, P.A., Leino, K.R.M. (eds.) TACAS 2011. LNCS, vol. 6605, pp. 205–219. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-19835-9_17
25. Kapur, D., Majumdar, R., Zarba, C.G.: Interpolation for data structures. In: FSE 2006, pp. 105–116 (2006)
26. Kovács, L., Voronkov, A.: Interpolation and symbol elimination. In: Schmidt, R.A. (ed.) CADE 2009. LNCS (LNAI), vol. 5663, pp. 199–213. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-02959-2_17
27. Krajíček, J.: Interpolation theorems, lower bounds for proof systems, and independence results for bounded arithmetic. *J. Symb. Log.* **62**(2), 457–486 (1997)
28. Kupferschmid, S., Becker, B.: Craig interpolation in the presence of non-linear constraints. In: Fahrenberg, U., Tripakis, S. (eds.) FORMATS 2011. LNCS, vol. 6919, pp. 240–255. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-24310-3_17
29. Lang, S.: *Introduction to Diophantine Approximations: New Expanded Edition*. Springer, New York (2012)
30. McMillan, K.L.: Interpolation and SAT-based model checking. In: Hunt, W.A., Somenzi, F. (eds.) CAV 2003. LNCS, vol. 2725, pp. 1–13. Springer, Heidelberg (2003). https://doi.org/10.1007/978-3-540-45069-6_1

31. McMillan, K.L.: An interpolating theorem prover. In: Jensen, K., Podelski, A. (eds.) TACAS 2004. LNCS, vol. 2988, pp. 16–30. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24730-2_2
32. McMillan, K.L.: Quantified invariant generation using an interpolating saturation prover. In: Ramakrishnan, C.R., Rehof, J. (eds.) TACAS 2008. LNCS, vol. 4963, pp. 413–427. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78800-3_31
33. Okudono, T., Nishida, Y., Kojima, K., Suenaga, K., Kido, K., Hasuo, I.: Sharper and simpler nonlinear interpolants for program verification. In: Chang, B.-Y.E. (ed.) APLAS 2017. LNCS, vol. 10695, pp. 491–513. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-71237-6_24
34. Parrilo, P.A.: Semidefinite programming relaxations for semialgebraic problems. *Math. Program.* **96**(2), 293–320 (2003)
35. Pudlák, P.: Lower bounds for resolution and cutting plane proofs and monotone computations. *J. Symb. Log.* **62**(3), 981–998 (1997)
36. Rybalchenko, A., Sofronie-Stokkermans, V.: Constraint solving for interpolation. In: Cook, B., Podelski, A. (eds.) VMCAI 2007. LNCS, vol. 4349, pp. 346–362. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-69738-1_25
37. Sharma, R., Nori, A.V., Aiken, A.: Interpolants as classifiers. In: Madhusudan, P., Seshia, S.A. (eds.) CAV 2012. LNCS, vol. 7358, pp. 71–87. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-31424-7_11
38. Sofronie-Stokkermans, V.: Interpolation in local theory extensions. In: Furbach, U., Shankar, N. (eds.) IJCAR 2006. LNCS (LNAI), vol. 4130, pp. 235–250. Springer, Heidelberg (2006). https://doi.org/10.1007/11814771_21
39. Sofronie-Stokkermans, V.: On interpolation and symbol elimination in theory extensions. In: Olivetti, N., Tiwari, A. (eds.) IJCAR 2016. LNCS (LNAI), vol. 9706, pp. 273–289. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-40229-1_19
40. Solar-Lezama, A., Rabbah, R.M., Bodík, R., Ebcioğlu, K.: Programming by sketching for bit-streaming programs. In: PLDI 2005, pp. 281–294 (2005)
41. Strzeboński, A.W.: Real root isolation for exp-log functions. In: ISSAC 2008, pp. 303–314 (2008)
42. Strzeboński, A.W.: Real root isolation for tame elementary functions. In: ISSAC 2009, pp. 341–350 (2009)
43. Strzeboński, A.W.: Cylindrical decomposition for systems transcendental in the first variable. *J. Symb. Comput.* **46**(11), 1284–1290 (2011)
44. Tarski, A.: *A Decision Method for Elementary Algebra and Geometry*. University of California Press, Berkeley (1951)
45. Vladimir, V.: Pattern recognition using generalized portrait method. *Autom. Remote Control* **24**, 774–780 (1963)
46. Yorsh, G., Musuvathi, M.: A combination method for generating interpolants. In: Nieuwenhuis, R. (ed.) CADE 2005. LNCS (LNAI), vol. 3632, pp. 353–368. Springer, Heidelberg (2005). https://doi.org/10.1007/11532231_26
47. Zhang, J., Feng, Y.: Obtaining exact value by approximate computations. *Sci. China Ser. A Math.* **50**(9), 1361 (2007)