



Programmable Program Synthesis

Loris D'Antoni, Qinheping Hu, Jinwoo Kim^(✉), and Thomas Reps

University of Wisconsin-Madison, Madison, USA
{ldantoni, qhu28, jkim934, treps}@wisc.edu

Abstract. Program synthesis is now a reality, and we are approaching the point where domain-specific synthesizers can now handle problems of practical sizes. Moreover, some of these tools are finding adoption in industry. However, for synthesis to become a mainstream technique adopted at large by programmers as well as by end-users, we need to design programmable synthesis frameworks that (i) are not tailored to specific domains or languages, (ii) enable one to specify synthesis problems with a variety of qualitative and quantitative objectives in mind, and (iii) come equipped with theoretical as well as practical guarantees. We report on our work on designing such frameworks and on building synthesis engines that can handle program-synthesis problems describable in such frameworks, and describe open challenges and opportunities.

1 Introduction

1.1 A Synthesis Tale

Monica, a software engineer, is trying to write a program for transforming data she has stored in an array of integer numbers. Monica needs to zero-out all the negative entries from the array (they represent irrelevant data points) and add 10 to all the positive entries (this is a normalization step needed in Monica's API). Of course, Monica is a great engineer and she could write this program herself, but since Monica knows that similar problems arise often in her company (i.e., reformatting arrays to match certain APIs), Monica decides to try out this new thing everyone is talking about: *program synthesis*.

Monica wants a tool that takes as input some examples of the desired transformation and a set of operators the program can use, and magically outputs the intended program. In fact, Monica already has an input, a unit test, that she wants to process using her newly synthesized program: $[-1, 2, 3, 10, 31, -14, -11]$, for which the output should be $[0, 12, 13, 20, 41, 0, 0]$.

Monica also knows that the final program will look like a loop that iterates over the input array *arr*, which leads her to develop the grammar in Fig. 1. Monica thinks this grammar is general enough that it will cover a reasonable range of programs for similar tasks but limited enough that it will not result in spurious programs that overfit too much to the examples.

Quickly, Monica discovers that using program synthesis is not so straightforward. There are so many different tools! And they all take different kinds of

$$\begin{aligned}
 \textit{Start} &\rightarrow x = \text{len}(\text{arr}) - 1; \text{ while } x \geq 0 \text{ do } S \\
 S &\rightarrow \text{arr}[E] = \text{arr}[E] + E \mid \text{arr}[E] = E \mid \\
 &\quad x = E \mid S; S \mid \text{if } \text{arr}[x] > 0 \text{ then } S \text{ else } S \\
 E &\rightarrow 0 \mid 1 \mid x \mid E + E \mid E - E
 \end{aligned}$$

Fig. 1. The grammar G_{ex} Monica has in mind for synthesizing programs that iterate over an input array (*Start* is the starting nonterminal). G_{ex} is general enough to cover most programs that iteratively normalize entries in an array.

inputs. After a bit more research, Monica decides to go for one of the many tools, UltraSynth™, and encodes her problem. UltraSynth is written in a C-like language and Monica has mostly programmed in Python for her job. However, Monica decides to give UltraSynth a try and after a few days of learning the ins and outs of UltraSynth, she finally manages to encode her transformation-synthesis problem in UltraSynth. To achieve her goal, Monica had to tweak a bit what the grammar looks like to provide it to UltraSynth, which only accepted grammars without unbounded recursion (i.e., without infinitely many terms) and had to encode the examples in a way that was accepted by the tool.

The time has come and Monica manages to run UltraSynth on an instance of the synthesis problem. UltraSynth outputs the program in Fig. 2b, which is correct on the example. However, this program is needlessly large and contains many unneeded operations.

```

x = len(arr)-1
while x>=0 do
  if arr[x]>0 then
    arr[x] = arr[x]+10
  else
    arr[x] = 0
  x = x-1
    
```

(a) A possible ideal solution for Monica’s problem.

```

x = len(arr)-1
while x>=0 do
  if arr[x]>0 then
    arr[x] = arr[x]+1-1...+1
  else
    arr[x] = 0
  x = x-1
    
```

(b) A solution for Monica’s problem synthesized by UltraSynth. The sequence of ± 1 s on line 3 adds up to +10.

Fig. 2. Two possible solutions for Monica’s synthesis problem.

Monica has already invested a lot of time in learning UltraSynth, so she tries to figure out a way to avoid such problematic programs. Monica astutely realizes that the needless computations in Line 3 of Fig. 2b are due to repeated applications of the minus operator. Monica would like to ask UltraSynth to synthesize the program that contains as few minus operators as possible, but UltraSynth does not support a way to “prefer” one possible program over another. To bypass this limitation, Monica decides to remove the production $E \rightarrow E - E$ in order to suppress these programs.

Monica reruns UltraSynth after removing $E \rightarrow E - E$ from the grammar, and to her surprise, UltraSynth continues to run for hours and eventually times out without providing a solution. After investigating the matter, Monica finds out that she has made a mistake and disallowed too many programs—there is no longer a valid solution to the synthesis problem because without subtraction, the variable x cannot be decremented in line 5. UltraSynth was unable to report, or even detect this simple mistake—Why is it so difficult to program a synthesizer and why can’t synthesis tools detect the simplest of mistakes?

Monica has finally had enough of synthesis. She goes back to her daily routine and just writes the 7-line piece of code that applies the transformation she intended (Fig. 2a).

1.2 Programmable Synthesis Frameworks

The story of Monica is a common one in program synthesis, where most of the recent focus has been on solving problems rather than building general algorithms, tools, and methodologies. Existing synthesis frameworks are not **programmable** as they lack at least one of the following properties:

Domain-Agnostic. Existing synthesis ideas and algorithms have been introduced with specific domains in mind and are hard to apply to arbitrary synthesis problems. The languages used to specify synthesis problems are therefore domain-specific, and often fail to abstract the logical requirements of the synthesis problem. In our example, Monica had to look for a specific tool that accepted programs of the kind she was interested in. Moreover, she was not permitted to refine the specification to add a quantitative objective she had in mind (minimizing the number of minus operators).

Solver-Agnostic. Different synthesis tools are typically not interchangeable because their underlying solvers solve different types of problems. Even when two solvers can in principle solve the same types of problems, they typically cannot be interchanged or combined because they typically use drastically different formats written in different languages (e.g., Racket [28] vs. C [27]). For example, when Monica found out that UltraSynth was not working as expected, she could not easily try another tool to see if that tool was better.

This state of affairs is unfortunate because synthesis is very general; if synthesis were easier to use, it would benefit many domains. The potential generality, which is currently held back by the need for better support for usability, underscores the need to answer the following question:

Can we make synthesis more programmable?

In this paper, we present the steps we have undertaken in the direction of making synthesis more programmable, including some of the challenges that we faced, and some of the opportunities that the work has opened.

2 An Overview of Programmable Program Synthesis

The goal of enlarging the scope of synthesis has focused our attention on the need to have a framework in which synthesis problems can be addressed. By a *framework*, we mean the conceptual underpinnings that allow one to build tools to automate the creation of solutions for problems in some domain, in this case, program-synthesis problems. The canonical example is how the theory of parsing underlies the `yacc` tool [13], which automates the construction of parsers. For instance, consider the problem that `yacc` addresses:

- An instance of a parsing problem, $\text{Parse}(L,s)$, has two parameters: L , a context-free language; and s , a string to be parsed. String s changes more frequently than language L .
- Context-free grammars are a formalism for specifying context-free languages.
- Create a tool that implements the following specification:
 - Input: a context-free grammar that describes language L .
 - Output: a parser, `yyparse()`, such that invoking `yyparse()` on s computes $\text{Parse}(L,s)$.

One consideration for building a framework is the existence of a well-defined “engine” (or collection of engines) for performing the desired task—in this case, parsing s with respect to L , once both L and s are at hand. `Yacc` supports just a single engine, which parses a string with respect to a grammar that is LALR(1). In principle, `yacc` could have been a more general tool by having it perform various tests on L to determine what grammar family L belongs to (e.g., LALR(1), LR(1), LL(1), LL(*)), and then emitting a parser that makes use of an appropriate parsing algorithm for that family, falling back on Generalized LR parsing [19] in case L is not in one of the specialized families supported.

Another aspect illustrated by `yacc` is that the parameters to the problem have different “binding times”. In this case, string s changes more frequently than language L —i.e., L is bound early, and s is bound late. The framework implementation can exploit the known value of the early-bound parameter to create a more efficient implementation. In the case of `yacc`, it compiles L to tables used by a table-driven LALR(1) parsing algorithm.

2.1 Why Isn’t Existing Work in Synthesis Programmable?

There do exist synthesis tools (mostly, solver-aided languages [27, 28]) that allow one to control some aspects of a synthesis problem in a programmable fashion. However, the nature of existing synthesis tools also forces an association between how a synthesis problem is written and how it is solved. For instance, in Sect. 1, the fact that solvers are tightly coupled to some specification language prevented Monica from trying out a different tool after UltraSynth produced an inadequate answer. The current state of program-synthesis tools is depicted in Fig. 3.




Fig. 3. Program synthesis today, where the lack of separation between specification and solver causes a user to have to encode a problem multiple times to use different tools.

This situation is in direct conflict with the principles articulated at the end of Sect. 1, namely, that a user should be able to program the various aspects of a synthesis problem using a formalism that is both (i) *domain-agnostic* and (ii) *solver-agnostic*. The first property addresses generality: the formalism should be powerful enough to capture a wide variety of synthesis problems (e.g., SQL, regular expressions, and imperative programs). The second property opens the door for synthesis-problem specifications to be fed—possibly after a compilation/translation step—to different specialized solvers, or to multiple solvers with different capabilities.

Another example that one may consider a programmable framework is Syntax-Guided Synthesis (SYGUS) [1], which is a successful synthesis framework targeted at expressions. The defining characteristic of SYGUS, compared to other synthesis approaches such as solver-aided languages, is that it allows one to write synthesis problems in a completely logical format.

Example 1. Consider the simple problem of synthesizing the maximum max of two input variables, x and y . There are two parts to a SYGUS problem: a syntactic part, written as a context free grammar such as the example G_S below:

$$G_S ::= Start \rightarrow x \mid y \mid Start + Start \mid \text{if } x < y \text{ then } Start \text{ else } Start$$

and the specification part, which is written as a Boolean formula ψ_S :

$$\psi_S \equiv \forall x, y. max(x, y) \geq x \wedge max(x, y) \geq y \wedge (max(x, y) = x \vee max(x, y) = y)$$

A SYGUS problem is simply the pair $sy = (G_S, \psi_S)$, where a solution to the SYGUS problem sy is a term $t \in L(G)$ such that ψ_S holds. For example, the following term is a solution for the function max in the problem we just described:

$$\text{if } x < y \text{ then } y \text{ else } x.$$

The advantage of such a logic-based formalism is that it achieves a separation from solver and specification, which allows SYGUS to be *solver-agnostic*. Several different SYGUS solvers have been developed (e.g., [4, 7, 21, 26]), many of which use drastically different internal algorithms that have different strengths for solving different kinds of problems. Moreover, a user of SYGUS need not consider the differing input languages or characteristics of these solvers, and instead can encode their problem just once in the SYGUS format to have access to all the different solvers.

While SYGUS achieves—and shows the benefits of—solver-agnosticity, it *fails to achieve domain-agnosticity* because the framework is targeted specifically at expressions. For example, consider the scenario from Sect. 1: Monica would be unable to encode her problem in SYGUS, because the grammar G_{ex} in Fig. 1 contains a production with a while loop, and loops, which require a custom semantics, cannot be expressed in any decidable theory—a key restriction of SYGUS. SYGUS also does not allow one to express intent outside of the behavioral specification ψ , which would have prevented Monica from trying to optimize the program obtained from UltraSynth in Fig. 2b.

All in all, the current state of program synthesis is an unsatisfactory mess, as depicted in Fig. 3. There are multiple non-interoperable solvers with different input languages, targeting different synthesis domains with varying degrees of overlap. SYGUS, by virtue of solver-agnosticity, provides a unified approach to synthesizing expressions, which forms the basis of multiple solvers. However, while SYGUS is a bright spot, it fails to be general: it does not cope with (1) the *variety of domains* used in synthesis, required to deal with arbitrary languages (e.g., SQL, regular expressions, and imperative programs), and (2) the *variety of collateral considerations* that arise for different domains (e.g., types, quantitative objectives, and probabilities).

2.2 What Does a Programmable Synthesis Framework Look Like?

Our vision of programmable synthesis can be summed up as follows:

programmable synthesis
==

easily instantiable, domain-agnostic, solver-agnostic synthesis framework.

In contrast with Fig. 3, what we would like to have is depicted in Fig. 4, where both user and solver work with a unified general format, regardless of domain or solving technique. Such an approach would allow one to specify a synthesis problem once and for all, without having to worry about the underlying solving strategy. To achieve this goal, it is necessary to distill out the essence of many program-synthesis problems into a specification formalism that is ground in formal methods (e.g., automata and logic) and is agnostic to any specific domain of application. This degree of abstraction also opens the opportunity to lift certain synthesis algorithms and ideas to a higher level that makes these algorithms reusable across different tools. Our framework can then interface to




Fig. 4. Programmable program synthesis, where a synthesis problem with arbitrary constraints can be written once and for all in a general format, which can then be dispatched to compatible solvers.

different solving tools (backend solvers) in a way that allows one to easily swap one solver for another, or to use multiple solvers in tandem. If our vision is achieved, the capabilities that would be available to tool designers—discussed in greater detail in Sect. 5—would allow synthesis tools to be created that have the kind of flexibility that Monica expected and needed in Sect. 1.1.

Let us now be more concrete about the requirements for such a framework for synthesis. Following the pattern for yacc given above, a framework for synthesis could follow a similar scheme:

- An instance of a synthesis problem $\text{Synthesize}(\mathcal{L}, \llbracket \cdot \rrbracket_{\mathcal{L}}, \varphi)$ has three parameters: \mathcal{L} , a formal language; $\llbracket \cdot \rrbracket_{\mathcal{L}}$, a semantics to ascribe to \mathcal{L} ; and φ , a behavioral specification for some desired member of \mathcal{L} . The behavioral specification φ changes more frequently than \mathcal{L} and $\llbracket \cdot \rrbracket_{\mathcal{L}}$.
- Let F_{syntax} and $F_{\text{semantics}}$ be appropriate formalisms for specifying \mathcal{L} and $\llbracket \cdot \rrbracket_{\mathcal{L}}$, respectively.
- Create a tool that implements the following specification:
 - Input: an F_{syntax} specification of a language’s syntax, and an $F_{\text{semantics}}$ specification of the language’s semantics.
 - Output: a function $\text{Synth}_{\mathcal{L}, \llbracket \cdot \rrbracket_{\mathcal{L}}}(\cdot)$ that takes φ as input and computes $\text{Synthesize}(\mathcal{L}, \llbracket \cdot \rrbracket_{\mathcal{L}}, \varphi)$.

To be even more concrete, F_{syntax} could be a regular-tree grammar [5],¹ and $F_{\text{semantics}}$ would be defined over the grammar in a compositional manner, production by production. What we have called collateral considerations (types, quantitative objectives, probabilities, etc.) would be handled as part of the F_{syntax} or $F_{\text{semantics}}$ specifications, depending on the issue at hand. For instance, constraints on program behavior, such as refinement types [24], minimizing/bounding evaluation resources usage [11, 15], and probabilistic behavior

¹ The grammar would also be equipped with production-by-production pretty-printing rules to specify how to convert a tree to its textual representation.

[22], are semantic concerns that would be part of $F_{\text{semantics}}$. Other considerations would be part of F_{syntax} , such as bounds on the use of syntactic constructs [12], or the use of probabilistic generative models of syntactic structures [3, 17]. For instance, for these two issues, one could weight the productions of the grammar with values from a semiring, and place a (possibly learned) distribution on the productions, respectively.

The scheme in the box above would allow us to meet the goals of being both domain-agnostic and solver-agnostic,² as long as (i) the formalisms for F_{syntax} and $F_{\text{semantics}}$ are sufficiently powerful to qualify as “domain-agnostic,” and (ii) specifications in these formalisms can be analyzed and broken down into components that can be farmed out to existing solvers (or perhaps to new implementations of the kinds of algorithms used in existing solvers).

Who benefits from such a framework? The existence of a domain- and solver-agnostic framework benefits two parties: (i) users of synthesis tools such as Monica, and (ii) designers of synthesis tools, such as the team behind UltraSynth. Both scenarios can be illustrated by making an analogy with LLVM [20]—which provides an intermediate representation for compilation that is similarly both domain- and solver-agnostic. Users of LLVM, which are front-end language designers, benefit from two facts: (i) that the LLVM IR is rich enough to support the range of features their language might have, and (ii) that once their language is compiled down into LLVM IR, the entire library of LLVM IR optimizations is accessible to them. Similarly, a programmable synthesis framework benefits users in two ways: (i) by supporting the full range of features that may be required for a synthesis problem, and (ii) by putting multiple solvers within reach for problems written in the framework. Additionally, a well-defined framework also facilitates *reuse* of problem components: for example, Monica can reuse G_{ex} for synthesizing other array transformations.

On the other hand, backend optimization designers of LLVM benefit from the fact that once their optimization is written in LLVM, all LLVM users may easily access those optimizations if need be. Similarly, tool designers for a programmable synthesis framework rest easy knowing that once their tool supports the framework, those who need it will find it accessible and easy to use—regardless of what internal techniques they decide to use. Note that while the framework intends to be general, tools that interface with the framework can choose to be selective in the problems they support—it is up to the users, or perhaps the framework designers, to match a problem with an appropriate solver (similar to how language designers mix and match backend optimizations for their language in LLVM). In addition, advances at the framework-level—such as

² We also acknowledge that even the scheme given above, which was modeled on the one for `yacc`, is open to revision. In particular, the additional degree of parameterization for synthesis (\mathcal{L} , $\llbracket \cdot \rrbracket_{\mathcal{L}}$, and φ) opens the door for a variety of alternatives, based on different “binding times” for \mathcal{L} , $\llbracket \cdot \rrbracket_{\mathcal{L}}$, and φ . For instance, a solver that uses different abstract domains as part of a refinement-based search strategy [29] would have \mathcal{L} and φ fixed, but vary $\llbracket \cdot \rrbracket_{\mathcal{L}}$. Similarly, when one has quantitative syntactic objectives [12], the solver would carry out its search with φ fixed, \mathcal{L} varying, and $\llbracket \cdot \rrbracket_{\mathcal{L}}$ induced as \mathcal{L} changes.

the development of meta-algorithms, as illustrated in Sect. 4.2—instantly benefit all tools that support the framework.

This Paper. New technical challenges, as well as new opportunities, come along with our broader goals. In this paper, we present some of the work that we have done toward building the kind of framework sketched out above.

Specifying Programmable Synthesis Problems (Sect. 3). Semantics-guided synthesis (SEMGUS) is our proposed framework that allows a user to provide both the syntax and semantics for the constructs in the language over which programs are to be synthesized. We show how SEMGUS can easily be extended with quantitative objectives for specifying when a synthesized program is “good” according to a certain metric—e.g., the program should be of minimal size or should maximize a certain outcome.

Solving Programmable Synthesis Problems (Sect. 4). We present solvers that can tackle problems specified in the SEMGUS framework. We also present a meta-solver that can be combined with other SEMGUS solvers to support quantitative objectives. Because our framework does not impose solver-specific restrictions on how synthesis problems are programmed, our solvers can prove unrealizability—i.e., whether a synthesis problem has no solution—of many complex synthesis problems with infinite search spaces.

These steps are just the beginning of what we expect to be a multi-year journey into designing a framework that achieves our goals, and solvers for such a framework. We discuss some of the open challenges and opportunities in Sect. 5.

3 Programmable-Synthesis Specifications

Designing synthesis frameworks that are *programmable* requires one to formally abstract the essence of how one specifies different program-synthesis problems. While we do not claim to have developed a completely unified framework that can capture all synthesis problems yet, in this section we present two ideas for programming many practical synthesis problems: (i) SEMGUS, a framework that uses logic and formal methods to make the search space and specifications of all synthesis problems easy to program in arbitrary domains (Sect. 3.1), and (ii) an extension of SEMGUS that allows one to specify *quantitative objectives* over the syntactic structure of a synthesized program (Sect. 3.2).

3.1 Semantics-Guided Synthesis

Existing work on program synthesis [1] typically identifies two main components to a synthesis problem: (i) a *search space* of candidate programs, which is in essence a small programming language, and (ii) a *behavioral specification*, which describes what the synthesized program should do. A programmable synthesis framework must represent (at the very least) these two components in a domain- and solver-agnostic way. Take the syntax-guided synthesis (SYGUS) framework,

for example: SYGUS achieves solver-agnosticity by representing the search space as a regular tree grammar, and the specification as a Boolean formula in a decidable background theory.

Then why is SYGUS, and this particular combination of representations, unable to achieve domain-agnosticity? The syntactic component of SYGUS—the grammar—actually does achieve some degree of domain-agnosticity, in the sense that one is free to define a language of one’s own. However, SYGUS requires that the specified grammar be contained within a fixed background theory, which are terms with a pre-defined fixed and standardized *semantics*. While this design choice makes the solutions to SYGUS problems easy to verify (using an SMT solver), it limits the *programmability of the search space*.

For example, let us reconsider the example in Sect. 1. If Monica attempted to write her example as a SYGUS problem, she would have been unable to use loops because loops are not part of the supported background theory. What if Monica wanted a solution that operates over a DSL, or had some pre-defined components that she wanted to use (like $len(arr)$)? What if Monica wanted to synthesize regular expressions, or some other programs with relatively non-standard semantics?

One can intuitively understand these scenarios as synthesis problems over different programming languages (search spaces)—a DSL, library functions, regular expressions. To support different programming languages, a synthesis framework needs more than the ability to accept a syntax, it needs the ability to accept a semantics for a language as well. Therefore, developing a programmable synthesis framework capable of supporting all these scenarios requires designing a solver-agnostic way of specifying the semantics of such arbitrary programming languages. SYGUS has shown that regular tree grammars are an effective formalism for programming the syntax of a search space; we extend this with a formalism to program the semantics of the search space as well, which, to achieve true domain-agnosticity, need not be constrained to a fixed background theory.

Semantics as Constrained Horn Clauses. Our solution to this challenge is the Semantics-Guided Synthesis (SEMGUS) framework [14], which allows users to customize the syntax and semantics of the search space. To see how SEMGUS supports programmable semantics, let us consider the production $Start \rightarrow \text{while } x \geq 0 \text{ do } S$ from Fig. 1 as an example. This production is a while loop, and part of the semantics for a term produced by this production can be expressed using the inference rule below³ (where Γ represents a state that maps variables to integer values):

$$\frac{\llbracket x \geq 0 \rrbracket(\Gamma) = \text{True} \quad \llbracket s \rrbracket(\Gamma) = \Gamma_1 \quad \llbracket \text{while } x \geq 0 \text{ do } s \rrbracket(\Gamma_1) = \Gamma_2}{\llbracket \text{while } x \geq 0 \text{ do } s \rrbracket(\Gamma) = \Gamma_2} \quad (1)$$

Such semantics are supported in the SEMGUS framework by expressing the inference rule in Eq. (1) as a Constrained Horn Clause (CHC). CHCs are logical formulas, and more precisely, they are implications where one is only allowed

³ A similar rule must be added for the case in which the guard evaluates to false.

to have a single relation in the conclusion, and a conjunction of relations along with one constraint in the premise:

Definition 1 (Constrained Horn Clauses). A Constrained Horn Clause is a first-order formula of the form

$$\forall \vec{x}, \vec{x}_1, \dots, \vec{x}_n. (\phi \wedge R_1(\vec{x}_1) \wedge \dots \wedge R_n(\vec{x}_n)) \implies H(\vec{x}),$$

where ϕ is a constraint over some background theory that may contain variables from $\vec{x}, \vec{x}_1, \dots, \vec{x}_n$, and R_1, \dots, R_n and H are uninterpreted relations.

In SEMGUS, search spaces are represented as regular tree grammars, where productions have associated semantics. In Eq. (1), the semantics of a term $x \geq 0$ is represented using the semantic function $\llbracket \cdot \rrbracket$. SEMGUS, assumes that each non-terminal N appearing in the grammar has a corresponding logical relation sem_N , which we refer to as the *semantic relation*, that represents the behavior of the semantic function $\llbracket \cdot \rrbracket$ in Eq. (1). For example, the expression $\llbracket s \rrbracket(\Gamma) = \Gamma_1$ from Eq. (1) can be translated into the relation $\text{sem}_S(\langle s, \Gamma \rangle, \Gamma_1)$.

Example 2 (Semantic Rules as CHCs). The following CHC captures how one would express in SEMGUS the semantics of the production $\text{Start} \rightarrow \text{while } x \geq 0 \text{ do } S$ shown in Eq. (1):

$$\frac{\Gamma[x] \geq 0 \quad \text{sem}_S(\langle s, \Gamma \rangle, \Gamma_1) \quad \text{sem}_{\text{Start}}(\langle \text{while } x \geq 0 \text{ do } s, \Gamma_1 \rangle, \Gamma_2)}{\text{sem}_{\text{Start}}(\langle \text{while } x \geq 0 \text{ do } s, \Gamma \rangle, \Gamma_2)} \quad (2)$$

One can read Eq. (2) as the following implication:

$$\text{sem}_S(\langle s, \Gamma \rangle, \Gamma_1) \wedge \text{sem}_{\text{Start}}(\langle \text{while } x \geq 0 \text{ do } s, \Gamma_1 \rangle, \Gamma_2) \wedge \Gamma[x] \geq 0 \implies \text{sem}_{\text{Start}}(\langle \text{while } x \geq 0 \text{ do } s, \Gamma \rangle, \Gamma_2) \quad (3)$$

Equation (3) is a CHC where $\text{sem}_{\text{Start}}$ and sem_S are relations, and $\Gamma[x] \geq 0$ corresponds to the first-order constraint ϕ .

SEMGUS allows one to specify multiple such CHCs⁴ for each *production* in the grammar. CHCs are the logical formalism of choice for expressing these semantics in a language-agnostic way, which are an intuitive and expressive format.

The SEMGUS Framework. Once a user has understood how to define a semantics for their grammar, a SEMGUS problem then can be specified simply as a synthesis problem over a grammar equipped with such a semantics.

Definition 2 (SEMGUS). A SEMGUS problem over a theory \mathcal{T} is a tuple $\text{sem} = (G_{\llbracket \cdot \rrbracket}, \psi(x, f(x)))$, where:

- G is a regular tree grammar equipped with the semantics $\llbracket \cdot \rrbracket$,

⁴ The ability to define multiple semantic rules for a production is useful for productions such as while loops, which are commonly equipped with two rules that describe looping and loop termination.

- $\psi(x, f(x))$ is a Boolean formula over the theory \mathcal{T} , that serves as the behavioral specification,
- f is a free second-order variable that serves as the function to be synthesized.

A **solution** to the SEMGUS problem sem is a term $s \in L(G_{\llbracket \cdot \rrbracket})$ such that $\psi(x, \llbracket s \rrbracket(x))$ holds.

Example 3 (Monica’s Synthesis Problem in SEMGUS). Consider the synthesis problem Monica had in Sect. 1. Let $G_{ex\llbracket \cdot \rrbracket}$ be the grammar G_{ex} from Fig. 1, equipped with semantic rules such as the one defined in Eq. (2). Let $E = \{-1, 2, 3, 10, 31, -14, -11\}$, the input array Monica considered for her task. Let $\psi(arr, f(arr))$ be a formula over the theory of arrays and CLIA describing what it means for the program f to be correct on an input arr :

$$\psi(arr, f(arr)) \equiv \bigwedge_{0 \leq i < len(arr)} f(arr)[i] = ITE(arr[i] > 0, arr[i] + 10, 0).$$

Then $sem_{ex} = (G_{ex\llbracket \cdot \rrbracket}, \bigwedge_{arr \in E} \psi(arr, f(arr)))$ is a SEMGUS problem defined over a background theory of arrays and CLIA—the behavioral specification requires that the final program satisfies all the examples in E .⁵ Moreover, sem_{ex} is written in a completely logical format, and is thus not tied to a specific tool like UltraSynth and can be dispatched to multiple backend solvers (assuming tooling) as Monica pleases.

The ability to customize the semantics for a language in a framework allows that framework to support a plethora of different synthesis problems. One can define synthesis problems over regular expressions, domain-specific languages, imperative languages, or any other language that has a semantics definable as CHCs within the framework, all of which can be tested using different solvers utilizing different strategies.

Example 4 (Regular Expressions Synthesis in SEMGUS). Synthesis problems over regular expressions can be expressed succinctly in SEMGUS. The grammar of regular expressions can be captured with the following grammar, where c is a character and ϕ the empty set:

$$R \rightarrow c \mid \epsilon \mid \phi \mid R + R \mid R \cdot R \mid R^*$$

Using CHCs, one can also naturally express the semantics of terms $r \in L(R)$. For example, the semantics of Kleene star can be given as the following two CHCs:

$$\frac{}{\text{sem}_R(r^*, \epsilon)} \quad \frac{\text{sem}_R(r, s_1) \quad \text{sem}_R(r^*, s_2) \quad s = s_1 s_2}{\text{sem}_R(r^*, s)}$$

⁵ In this example, one could have used a formula simply describing the input/output examples instead of a more complex logical formula. We chose the latter option to illustrate how the behavioral specification can involve terms in interesting theories—e.g., CLIA and arrays.

The rules are based on the expansion $r^* \rightarrow \epsilon + r \cdot r^*$: the first rule lets r^* accept ϵ , and the second rule accepts a string s by finding two substrings s_1, s_2 , such that s_1 is accepted by r , s_2 is accepted by r^* , and the concatenation $s_1 \cdot s_2$ is equal to s . The specification of the problem can then use expressions of the form $\text{sem}_R(r, s)$ and $\neg \text{sem}_R(r, s)$ to denote whether an example s is positive or negative, respectively.

3.2 Adding Quantitative Syntactic Objectives

In the example discussed in Sect. 1.1, the original synthesis problem Monica posed to the solver was under-constrained and caused the underlying tool to synthesize an undesirable solution that contained unnecessary operations. While the logical-specification mechanism is powerful, it can only capture the functional requirements of the synthesis problem—e.g., the program should perform correctly on a given set of input/output examples. When multiple possible programs can satisfy the specification, a programmable synthesis framework should provide a way to prefer one to the other—i.e., the user of the framework should be able to describe a quantitative objective. In this section, we show how the formal foundations of SEMGUS (i.e., the use of grammars and logic) allow us to easily extend the framework to incorporate quantitative objectives over the syntax of the synthesized program. The ideas we present were originally described in the context of SYGUS [12]; here we show how they can also be applied to SEMGUS.

Adding Quantitative Objectives Using Weighted Grammars. Recall that a SEMGUS problem is given along with a *regular tree grammar* specifying the search space. In our running example, Monica would like to synthesize a program that has few occurrences of the minus operator. A natural way to express this intent is allowing Monica to tag productions involving such an operator with a cost, let’s say 1. Our quantitative extension of SEMGUS builds on this intuition and allows users to add weights/costs to productions in the grammar. This extension leads to a well-studied formalism, *weighted tree grammars*, keeping the SEMGUS framework general. Intuitively, a *weighted tree grammar* is a grammar in which each production p has an associated weight/cost $\mu(p)$.

Intuitively, the weight of a derivation tree is the *sum* of the weights of all productions.⁶ For simplicity, in this paper, we assume that the domain of weights is the natural numbers, and that their sum is the usual application of the $+$ -operator. We use $w_G(t)$ to denote the weight of a term t with respect to the weighted grammar G .

With the weights specified by the weighted grammars, users can specify quantitative objectives as constraint objectives and optimization objectives. A constraint objective is a predicate $\omega(v)$ over a numerical variable v ; we say that

⁶ Weights have to come equipped with operators that tell us how to combine weights of individual productions to obtain the weights of terms. Formally, the weights must be from a semiring; we refer the reader to the original work on this topic [12] for details.

a term t satisfies the constraint objective if $\omega(w_G(t))$ holds. An optimization objective is a flag $\text{OPT} \in \{\text{True}, \text{False}\}$ indicating whether we want to minimize the weight of the solution.

Example 5. Recall that in the example introduced in Sect. 1, Monica wants to avoid redundant occurrences of the minus ($-$) operator. To express this intent in SEMGUS, Monica can utilize the following weighted grammar.

$$\begin{aligned} \text{Start} &\rightarrow x = \text{len}(\text{arr}) - 1; \text{ while } x \geq 0 \text{ do } S \\ S &\rightarrow \text{arr}[E] = \text{arr}[E] + E \mid \text{arr}[E] = E \mid \\ &\quad x = E \mid S; S \mid \text{if } \text{arr}[x] > 0 \text{ then } S \text{ else } S \\ E &\rightarrow 0 \mid 1 \mid x \mid E + E \mid E - E/1 \end{aligned}$$

In the weighted grammar, only the rule $E \rightarrow E - E$ is assigned the weight 1. All other rules are assigned the weight 0 (omitted for readability). The weight of a term t with respect to this grammar is the number of occurrence of the minus operator in t . If Monica wants to restrict the number of occurrences of the minus operators to be less than 5, she can use the constraint objective $\omega(v) = v < 5$. Furthermore, if she wants to minimize the occurrences of the minus operator, she can set the flag OPT to **True**.

To summarize, a SEMGUS problem with quantitative syntactic objectives is a tuple $\text{sem} = (W_{[\cdot]}, \psi(x, f(x)), \omega, \text{OPT})$ where $W_{[\cdot]}$ is a weighted grammar with a corresponding semantics, ψ is a Boolean formula like before, ω is the constraint objective, and the flag OPT is the optimization objective. The goal is to find a solution that not only satisfies the specification ψ , but also the quantitative objective ω , and is of minimal cost if OPT is set to **True**.

Quantitative syntactic objectives are useful in applications such as programming by examples [10] and program repair [6], where it is desirable to produce small programs with fewer constants, because such programs are more likely to generalize to examples and test cases outside of the set of examples given by the user. When allowing real-valued weights, syntactic objectives can be also used to find the most likely solution with respect to a given probability distribution. We can assign productions weights that represent their probabilities; the weight of a candidate solution is its likelihood.

4 Programmable-Synthesis Solvers

While a programmable synthesis framework as discussed in Sect. 3 is certainly desirable, it is of little practical use if one is unable to solve the problems that are written in such a framework. In this section, we show that SEMGUS problems can be solved practically. We first describe two general solving techniques for SEMGUS (Sect. 4.1) and then present new algorithmic solving techniques enabled by the SEMGUS framework (Sect. 4.2).

4.1 General Solving Procedures for *SEMGUS* Problems

We start off by presenting two solving procedures for general *SEMGUS* problems we implemented as a tool, rooted in strategies commonly used in existing program synthesizers: enumeration (used in the tool *MESSY-Enum*) and constraint solving (used in the tool *MESSY*). Specifically, we will be considering *SEMGUS*-with-examples problems: *SEMGUS* problems where the specification is given in terms of a finite set of examples E . An algorithm for solving *SEMGUS*-with-examples problems can be combined with counterexample-guided inductive synthesis (CEGIS) [27], which generates counterexamples in case a synthesized answer does not meet the general specification, to iteratively increase the example set E and eventually obtain a correct program.

MESSY-Enum: A Basic Enumerator for *SEMGUS* Problems. Because *SEMGUS* also relies on a grammar to specify the syntax of valid terms, like *SYGUS*, one can employ a simple enumerator that generates terms of increasing size from the grammar and test the enumerated terms against the behavioral specification. With *SEMGUS*, a term (representing a program) cannot be executed directly, because the semantics to ascribe to it has been specified in the semantic specification. However, because the semantics is specified with CHCs, the term can be executed with a level of interpretation supplied by an off-the-shelf CHC solver. Therefore, *MESSY-Enum* employs an off-the-shelf CHC solver such as [18] to check if the CHCs are consistent with the specification.⁷

Concretely, given a term t_e to test, one can use the following CHC to check whether t_e meets the specification:

$$\frac{\bigwedge_{e_i \in E} \mathbf{sem}_{\text{start}}(\langle e_i, t_e \rangle, o_i)}{\text{Realizable}} \text{Query} \quad (4)$$

The *Query* rule in Eq. (4) exactly encodes the specification as a CHC: it asks whether the semantics of t_e computed by $\mathbf{sem}_{\text{start}}$ is consistent with the set of input-output examples E . If so, the conclusion *Realizable* is provable using the existing set of CHCs—i.e., t_e is a solution to the synthesis problem.

Because we cannot directly execute candidate terms and instead rely on CHC solvers (which may be treated as a blackbox), it is difficult to employ common enumeration optimizations, such as behavioral equivalence caching, or equality saturation. Developing an enumeration-based solver capable of utilizing these ideas would require generating an explicit and efficiently executable interpreter from the given semantics, which is an interesting research challenge and future direction that we discuss in Sect. 5.

⁷ One can treat CHC solving as akin to a proof search, where the objective is to prove that a specific query holds (in this case, *Realizable* from Eq. (4)) using the provided CHCs.

MESSY: SEMGUS Problem Solving as CHC-Solving. MESSY-Enum uses a CHC solver to check whether an enumerated term t_e is consistent with the specification or not—however, CHC solvers are also capable of automatically searching for terms that satisfy the specification, as well. Our next solver, MESSY, takes advantage of this fact by expressing *both* the syntax of the search space and the semantics using CHCs. Once the entire search space is modeled this way, one can then slightly modify the Query rule to accommodate this change and directly use a CHC solver to solve the entire SEMGUS problem. In essence, Messy reduces solving the SEMGUS problem into finding a configuration of variables for which the set of CHC rules (containing syntax, semantics, and specification) is valid—similar to how constraint-based methods in existing synthesizers reduce the synthesis problem to one of solving a set of constraints.

Example 6 (MESSY Encoding). We show how the syntax and semantics used in the production $Start \rightarrow \text{while } x \geq 0 \text{ do } S$ from Fig. 1 can be captured using CHCs. This production states that one can obtain a syntactically valid term $\text{while } x \geq 0 \text{ do } s \in L(Start)$ for the nonterminal $Start$, given a valid term $s \in L(S)$. Equation (5) encodes this idea as a CHC using the *syntax relations* syn_S , and syn_{Start} , which capture whether the supplied arguments are valid terms that may be derived from the corresponding nonterminals S , and $Start$.

$$\frac{\text{syn}_S(s)}{\text{syn}_{Start}(\text{while } x \geq 0 \text{ do } s)} \quad (5)$$

Because the syntax relations provide a way to guarantee that a term t is a valid term in the syntax of a SEMGUS problem, one can rewrite the Query rule from Eq. (4) to use this relation instead of an explicitly enumerated term t_e .

$$\frac{\text{syn}_{Start}(t) \quad \bigwedge_{e_i \in E} \text{sem}_{Start}(\langle e_i, t \rangle, o_i)}{\text{Realizable}} \quad \text{Query} \quad (6)$$

The new Query rule in Eq. (6) has the term t as a free variable—i.e., proving *Realizable* amounts to finding a term $t \in L(Start)$ that is consistent on the input-output examples. A CHC solver presented with this rule, in tandem with the syntax and semantic rules, will then attempt to find a configuration of t such that *Realizable* holds. If the solver can prove that the premises of Equation Eq. (4) hold, then the term t is a solution to the SEMGUS problem.

One of the advantages of using such a CHC-based method is when dealing with cases where there is no answer to the synthesis problem, i.e., when there exists no t such that *Realizable* holds. In this case, the SEMGUS problem contains *no* answer satisfying the specification within its search space; we say that such a problem is *unrealizable*. Proving unrealizability is something that many existing solvers fail to consider, but is important: for example, Monica would not have had to wait for several hours after modifying the grammar in Sect. 1 if her solver had been able to show that the problem was unrealizable.

4.2 Meta Algorithms for Solving SEMGUS Problems

Now that we have shown how to build solvers for general SEMGUS problems (that do not involve quantitative objectives), we turn to ‘meta’-algorithms for solving SEMGUS problems, which are ‘meta’ in the sense that they (i) may be used atop any general SEMGUS solver, (ii) generate modified SEMGUS problems (rather than solutions) that can be easier to solve than the original SEMGUS problem or can be used to solve SEMGUS problems with quantitative objectives. The key component behind these meta-algorithms is the customizability of the search-space description in SEMGUS.

A Meta Solver for Quantitative Objectives. We first present an algorithm for solving SEMGUS problems with quantitative objectives [12]—i.e., where productions in the grammar have weights. We assume, for simplicity, that the only quantitative objective is to find the program of *least cost* that satisfies the specification. The idea of the algorithm is to iteratively reduce the SEMGUS problem with a quantitative objective to a sequence of SEMGUS problems without quantitative objectives, which are used to iteratively find a solution that has least cost—i.e., at each step of the sequence the cost of the solution is improved.

The algorithm operates as follows. Initially, we are given a SEMGUS problem sem with a weighted grammar W (we omit the semantic information for brevity) and with the minimization objective OPT set to *true*.⁸ The first step of the algorithm is to construct an unweighted grammar G^W by merely erasing all the weights in W . We can now use any SEMGUS solver to solve the resulting SEMGUS problem and obtain a term t_0 . This term will have a weight c according to the weighted grammar W , but it might not be the term of least cost that satisfies the specification. Our algorithm therefore tries to find out whether a solution with a lower weight exists, and accordingly constructs an (unweighted) grammar $G_{<c}^W$ such that a term t is accepted by the grammar $G_{<c}^W$ if and only if the weight of t according to W is less than c . When the weights are natural numbers, this construction is always possible [12]. We now have again an unweighted grammar, and we can use a SEMGUS solver to solve the resulting problem. This procedure can be repeated until no better solution exists.

Example 7. Consider the weighted grammar W we presented in Example 5. In particular, let us focus our attention on the following subset of productions that involve non-zero weights:

$$E \rightarrow 0 \mid 1 \mid x \mid E + E \mid E - E/1$$

The grammar $G_{<3}^W$, which accepts all terms of weight less than 3 is as follows:

$$\begin{aligned} E &\rightarrow E_2 \mid E_1 \mid E_0 \\ E_2 &\rightarrow E_1 - E_0 \mid E_0 - E_1 \\ E_1 &\rightarrow E_0 - E_0 \\ E_0 &\rightarrow 0 \mid 1 \mid x \mid E_0 + E_0 \end{aligned}$$

⁸ For simplicity, we assume no further quantitative objectives are present, but the general case can be handled using similar ideas [12].

Intuitively, each non-terminal E_i produces all and only terms with exactly i minus operators.

The meta solver for quantitative objectives shows how using a solver-agnostic specification formalism—i.e., grammars—enables algorithms that operate at the specification level and can be reused across multiple solvers.

Underapproximating Semantics with SEMGUS. The previous section showed how the programmability of the search-space syntax (i.e., the grammar) allows us to design meta-algorithms to solve SEMGUS problems involving quantitative objectives. In this section, we show how the programmability of the search-space *semantics* can be used to build meta-algorithms that can make synthesis faster. The key idea is to generate “simpler” variants of the original SEMGUS problem that use an *underapproximating semantics*, where an underapproximating semantics is defined as a subset of the original semantics that must be precise on the subset on which it is defined.

Definition 3. *For a grammar G equipped with a semantics $\llbracket \cdot \rrbracket$, we say $\llbracket \cdot \rrbracket^p$ underapproximates $\llbracket \cdot \rrbracket$ on G , or that $\llbracket \cdot \rrbracket^p$ is an underapproximating semantics for G with respect to $\llbracket \cdot \rrbracket$, if for every term $t \in L(G)$, every state Γ , and every value v on which $\llbracket \cdot \rrbracket^p$ is defined, $\llbracket t \rrbracket^p(\Gamma, v) = \llbracket t \rrbracket(\Gamma, v)$.*

One easy way to underapproximate a semantics is to simply “eliminate” certain operators from a grammar by not defining semantic rules for them. However, the concept of underapproximation need not be bounded to eliminating operators from a grammar—it may have a fully semantic meaning instead, for example, a bound on the number of possible loop iterations. The key intuition is that underapproximation is sound for use in synthesis—if a term t is the answer to a synthesis problem *sy*, *sy* actually does not need to contain any syntax or semantics outside of what is used to define and compute t . (In contrast, overapproximation is sound for proving unrealizability.)

Example 8. Recall, once again, the synthesis problem Monica has in Sect. 1. The grammar G_{ex} of Fig. 1 contains a while loop, which has a complex semantics that can be expensive to compute and, most importantly, allows nonterminating behavior. Most existing synthesizers [27, 28] explicitly prohibit nontermination by only considering finitely many unrollings for loops (because most answers to a synthesis problem will indeed terminate).

Fortunately, Monica knows that on her example $[-1, 2, 3, 10, 31, -14, -11]$, the loop should iterate no more than 7 times to process every element of the array. Monica may then choose to supply the synthesizer with an underapproximating semantics that limits the number of loop iterations to 7, which could greatly reduce the amount of computation the synthesizer must perform—for example, a naive enumerator might get stuck on a nonterminating loop when using the precise semantics, while terminating quickly when using the underapproximating semantics. Such a semantics can be expressed easily by adding a loop counter c to the semantics of loops given in Eq. (2), yielding the following CHC:

$$\frac{c \geq 0 \quad \Gamma[x] \geq 0 \quad \text{sem}_S(\langle s, \Gamma, c-1 \rangle, \Gamma_1) \quad \text{sem}_{\text{Start}}(\langle \text{while } x \geq 0 \text{ do } s, \Gamma_1, c-1 \rangle, \Gamma_2)}{\text{sem}_{\text{Start}}(\langle \text{while } x \geq 0 \text{ do } s, \Gamma, c \rangle, \Gamma_2)} \quad (7)$$

Setting $c = 7$ in the Query rule now ensures that loops run at most 7 iterations.

Abstract Semantics with SEMGUS. Similar to how we used underapproximating semantics to find solutions to a SEMGUS problem, abstract (overapproximating) semantics can be used to prove that a SEMGUS problem is unrealizable.

Definition 4. For a grammar G equipped with a semantics $\llbracket \cdot \rrbracket$, we say $\llbracket \cdot \rrbracket^\#$ is an abstract semantics for G with respect to $\llbracket \cdot \rrbracket$ if there exists an abstraction function α and a concretization function γ , such that for all $t \in L(G)$, if $\llbracket t \rrbracket(\Gamma, v)$ holds, then $\llbracket t \rrbracket^\#(\alpha(\Gamma), \alpha(v))$ holds, and $\Gamma \in \gamma(\alpha(\Gamma))$, $v \in \gamma(\alpha(v))$, i.e., α and γ form a Galois connection.

In contrast to underapproximating semantics, abstract semantics are sound when used to prove *unrealizability*—i.e., that a synthesis problem has no solution that satisfies its specification within its search space. Consider the use of abstract interpretation in program analysis: abstract interpretation is most often used to prove that a program cannot reach a certain set of bad states, while often being unable to guarantee that a program will produce a specific value, due to lack of precision. Similarly, an abstract semantics will often be unable to guarantee that a synthesized program satisfies the specification, due to lack of precision—but it can guarantee that all programs in the search space will *never* be able to produce a certain set of values, which can be used to prove unrealizability.

Example 9. Consider the scenario from Sect. 1, in which Monica removed subtraction from her grammar in an attempt to simplify the synthesized program. The removal of subtraction made the problem unrealizable—and UltraSynth ran for hours on end because it could not prove that this was the case. While proving unrealizability can be very difficult in general, a solver capable of reasoning about abstract domains and semantics could have utilized an (abstract) semantic rule such as Eq. (8) below:

$$\frac{\text{sem}_E(\langle e_1, \Gamma \rangle, \{pos\}) \quad \text{sem}_E(\langle e_2, \Gamma \rangle, \{pos\})}{\text{sem}_E(\langle e_1 + e_2, \Gamma \rangle, \{pos\})} \quad (8)$$

Equation (8) is defined on the abstract domain $\{pos, zero, neg\}$ —corresponding to positive, zero, and negative values—and captures the fact that the sum of two positive numbers will always be positive. This rule will be able to prove that G_{ex} without subtraction will *never* be able to modify x to a negative value, and thus that no program in the search space will terminate (leading to unrealizability).

Unrealizability is a property that is ignored by many current synthesizers, but it is a very important property nonetheless. One practical way to think about

unrealizability is as a sanity check, like a type system: the fact that a synthesis problem provided by an end user is unrealizable means that the synthesis problem is malformed in the sense that the user has got some of their specifications wrong. Similar scenarios happen daily with ordinary programming, and we expect them to happen with synthesis as well—thus, it is desirable that synthesizers be able to detect these problems, and report them early on if possible, without running indefinitely, as in Sect. 1. Unrealizability also has applications in computing optimal solutions, as in Sect. 4.2: unrealizability given a grammar with a lower weight bound ensures that the current solution is optimal.

5 The Future of Programmable Synthesis and SEMGUS

We hope we have convinced the reader that synthesis could use more programmability, and that SEMGUS addresses many of the programmability issues of existing synthesis work. But what lies ahead? How can we make programmable synthesis truly practical? In this section, we first outline some of the steps we are undertaking to answer this question (Sect. 5.1).

More importantly, we would also like to emphasize that *the vision of programmable program synthesis can only be realized through a community effort*. We will conclude this section with ideas to involve the synthesis community to help us realize our vision (Sect. 5.2).

5.1 What Are We Working on Next?

In this section, we present some of the directions our group is pursuing in extending SEMGUS to richer objectives and building better solvers for it. We also describe some open problems related to SEMGUS.

Interfacing Existing Program Synthesizers with SEMGUS. The bulk of our discussion in Sect. 3 was about achieving domain-agnosticity by building upon the ideas that SYGUS used in achieving solver-agnosticity. However, there also exist synthesis tools that are already domain-agnostic; most notably, solver-aided languages such as Sketch [27], Rosette [28], MiniKanren [8], and Prose [25]. While these tools are not solver-agnostic, they can in principle be used as SEMGUS solvers by virtue of their domain-agnosticity.

To use such existing tools as SEMGUS solvers, one must develop a compiler of sorts to translate a SEMGUS problem (written in the logical format from Sect. 3) to the specific front-end language of the tool. This task is not trivial for a number of reasons. First, each of these tools implement restrictions on the types of synthesis problems they accept; these restrictions are what enables their fast algorithms. For example, Rosette, Sketch, and MiniKanren only support finite search spaces (i.e., finite grammars), and this fact is encoded in different ways for different tools (e.g., by imposing bounds on the search depth or by imposing syntactic bounds on the search space). Second, some of these solvers implicitly use limited semantics—e.g., Sketch limits how many times a loop can

be executed. Third, some of these solvers require special inputs that are useful to guide the synthesis engine—e.g., Prose requires the user to provide a semantics for each operator in the input language as well as an *inverse semantics* that can be executed *backwards*; the inverse semantics is used to perform efficient enumeration.

Soundly compiling SEMGUS problems to these tools requires one to modify the original problems to fit these restrictions. Thankfully, the flexibility of SEMGUS comes to our aid! In Sect. 4.2, we have described ideas for transforming SEMGUS problems using restricted grammars or underapproximating semantics. These transformations are sound for synthesis—i.e., a solution to the transformed synthesis problem, which satisfies the restrictions of a particular external tool, is still a solution to the original problem—and thus can be used to interface with external tools. We are currently working on automating such translations.

The case of Prose is particularly interesting in that it requires inverse semantics, which are not immediately available from a SEMGUS problem. However, because SEMGUS semantics are expressed logically as CHCs, one can automatically invert these semantics starting from the CHCs—we are currently developing a tool that performs this inversion automatically and uses the inverse semantics to interface with Prose.

Other more specialized solvers, such as those for synthesizing regular expressions [23], could also be interfaced with our framework, with the limitation that they will only be able to handle specific problems. The more general question here is: how can we determine whether a specific SEMGUS problem is compatible with a specialized solver? We are working on designing “theories” that describe specific semantics for which specialized solvers exist. For example, if one were to use SEMGUS to work with regular expressions, they could import the regular-expression theory, which by design would enable compatibility with certain solvers. Note that this approach is still solver-agnostic because any general SEMGUS solver would still be able to use this problem definition.

Lifting Existing Synthesis Algorithms to Work with SEMGUS. While interfacing existing synthesizers with SEMGUS is one straightforward way of creating SEMGUS solvers, we envision that higher efficiency can be achieved by designing solvers that take advantage of the structure of SEMGUS problems. Is it possible to lift algorithms (not tools) that have previously been successful with SYGUS or other synthesizers up into SEMGUS?

For example, consider the problem of building an efficient enumeration algorithm for SEMGUS, an algorithmic technique that is now successfully employed in most SYGUS solvers [2, 4, 21]. The success of enumeration has been driven by a number of clever ideas for efficiently pruning the search space of relevant programs. An example was mentioned in Sect. 4.1, where we discussed the challenges with employing strategies such as behavioral equivalence caching or equality saturation on SEMGUS due to the lack of an executable semantics—i.e., in SEMGUS, evaluating a term on an input requires a costly call to a CHC solver. We are currently building an enumeration algorithm for SEMGUS that addresses this limitation. Our algorithm first synthesizes an executable interpreter from

the SEMGUS problem semantics, and then uses this executable interpreter to guide the search. To scale, our approach must handle other challenges, which we are also working on—e.g., discovering which operators have a semantics that is associative or commutative can help us avoid enumerating equivalent terms.

While the generality of SEMGUS is an obstacle to adapting some well-known algorithms, the same generality also helps SEMGUS provide a natural interface to express other algorithms, such as program synthesis using abstraction refinement [29]. The approach taken here is to synthesize programs that work on an abstract domain, and repeatedly refine the abstract domain until a program is found that is correct under the concrete semantics. This approach, in a sense, uses a meta-algorithm that can be expressed naturally in SEMGUS, as discussed in Sect. 4.2. We believe that SEMGUS will naturally be able to express many such meta-algorithms, and further accelerate the development of new meta-algorithms.

Supporting Richer Specifications. Beyond the basic specification mechanisms, SEMGUS already supports syntactic quantitative objectives through weighted grammars (Sect. 3.2). To capture the breadth of specifications appearing in modern synthesis applications, the SEMGUS framework will have to evolve over time. While we are investigating a number of complex objectives that will require extensions to the framework (e.g., probabilistic specifications), in the following paragraph we describe a specification mechanism the current SEMGUS framework can already capture for free: types.

Consider the problem of synthesizing a program that meets a given time complexity (or asymptotic resource usage in general) [11, 16]. In existing work, such bounds are specified (and proven correct) using a dependent type system. The solver uses the type system to guide the search, by enumerating only terms that satisfy a certain type. We observe that the SEMGUS framework is already able to capture such type-based specifications! In particular, types are a form of *static* semantics that can be associated with terms and, in most cases, typing rules can be encoded as CHCs, similarly to how one encodes semantic rules. For example, the following dependent type rule can be captured using a CHC where each typing judgment $t : \textit{type}$ is described using a relation $r(t, \textit{type})$.

$$\frac{a : \{\mathbf{Int} \mid \varphi_a(v)\} \quad b : \{\mathbf{Int} \mid \varphi_b(v)\} \quad + : x : \mathbf{Int} \rightarrow y : \mathbf{Int} \rightarrow \{\mathbf{Int} \mid v = x + y\}}{a + b : \{\mathbf{Int} \mid v = x + y \wedge \varphi_a(x) \wedge \varphi_b(y)\}} \quad (9)$$

5.2 What Can the Synthesis Community Do?

As we mentioned at the beginning of this section:

The vision of programmable program synthesis can only be realized through a community effort.

We discuss problems the community can help with in this concluding section.

A Broader Scope for Synthesis. The scope and potential of synthesis is very broad, in fact even broader than what has been discussed in this paper. An invited paper by Gulwani began [9]

Program Synthesis is the task of discovering an executable program from user intent expressed in the form of some constraints.

However, we feel that this viewpoint is actually somewhat narrow. We believe that insight on many problems can be obtained via the “lens” of synthesis: for many computing tasks, the goal is to produce some artifact to which some semantics is attached, and the process of producing that artifact can be thought of as a synthesis problem. For instance, in an AI planning problem, the artifact is a plan—i.e., Monica from Sect. 1 is a robot, and the sought-for program must navigate her from point A to point B (e.g., minimizing power consumption and time, while avoiding collisions and satisfying other safety guarantees). Closer to home for the CAV community, inside many tools for statically checking assertions in programs (such as SLAM or BLAST), the key component is one that creates an abstracted model of the program that is sufficiently precise to show that an assertion violation is not possible. Among the artifacts that may need to be synthesized are inductive invariants, abstract transformers, function summaries, and interface specifications. Thus, we conclude by offering the following wider definition of synthesis, which connects this broader outlook with the semantics-based perspective that we have presented in this paper:

Synthesis is the task of discovering a syntactic object—selected from some formalism in which each syntactic object has a rigorously defined semantics—from an “intent” expressed in the form of some kind of constraint.

We believe that the issues discussed in this paper will be increasingly important if synthesis is to be applied successfully to the creation of artifacts that have semantics, but are not programs *per se*.

The generality of our framework can bridge the gap between the many applications of synthesis, and we hope that the community will engage in our work by modeling their synthesis problems in SEMGUS, and by adapting their solvers to work with SEMGUS. Such contributions will result in new benchmarks and solvers, contributing to the programmability and effectiveness of SEMGUS.

Standardization and Competitions. We believe that the idea of a programmable synthesis framework, and SEMGUS, the start of such a framework, represents a step forward in program synthesis. Similarly to what happened with SYGUS, SEMGUS must be standardized, other researchers should build solvers for it, and these solvers should compete annually in SEMGUS competitions.

We hope that this paper will encourage readers to experiment with and advance the ideas presented here, in three ways: First, we hope that the generality of the framework will make it easy for people to use it on various problems, which in turn will make it easy to collect large and diverse sets of benchmarks that

will make the design of new solvers focused and effective. Second, we hope that researchers will build new algorithms and techniques that are general and can solve problems built in this framework. Third, we hope to soon create a yearly competition that will foster further interest in building general synthesizers for our framework. More than anything, this paper is a call-to-arms—an invitation to help broaden the scope and abilities of program synthesis, toward an era where Monica uses synthesizers just as much as Python during her daily work.

Acknowledgments. Supported, in part, by a gift from Rajiv and Ritu Batra; by multiple Facebook Research Awards; by a Microsoft Faculty Fellowship; by NSF under grants 1420866, 1763871, and 1750965; by ONR under grants N00014-17-1-2889 and N00014-19-1-2318; and a grant from the Korea Foundation for Advanced Studies. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors, and do not necessarily reflect the views of the sponsoring entities.

References

1. Alur, R., et al.: Syntax-guided synthesis. In: Formal Methods in Computer-Aided Design (FMCAD), pp. 1–8. IEEE (2013)
2. Alur, R., Radhakrishna, A., Udupa, A.: Scaling enumerative program synthesis via divide and conquer. In: Legay, A., Margaria, T. (eds.) TACAS 2017, Part I. LNCS, vol. 10205, pp. 319–336. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-662-54577-5_18
3. Amodio, M., Chaudhuri, S., Reps, T.W.: Neural attribute machines for program generation. CoRR, abs/1705.09231 (2017)
4. Barke, S., Peleg, H., Polikarpova, N.: Just-in-time learning for bottom-up enumerative synthesis. Proc. ACM Program. Lang. 4(OOPSLA.), 1–29 (2020)
5. Comon, H., et al.: Tree automata techniques and applications (2007). <http://www.grappa.univ-lille3.fr/tata>. Accessed 12 October 2007
6. D’Antoni, L., Samanta, R., Singh, R.: QLOSE: program repair with quantitative objectives. In: Chaudhuri, S., Farzan, A. (eds.) CAV 2016. LNCS, vol. 9780, pp. 383–401. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-41540-6_21
7. ESolver. <https://github.com/abhishekudupa/sygyus-comp14>
8. Feldman, M.Q., Wang, Y., Byrd, W.E., Guimbretière, F., Andersen, E.: Towards answering “Am I on the right track?” Automatically using program synthesis. In Proceedings of the 2019 ACM SIGPLAN Symposium on SPLASH-E, SPLASH-E 2019, pp. 13–24, New York, NY, USA. Association for Computing Machinery (2019)
9. Gulwani, S.: Dimensions in program synthesis. In: PPDP (2010)
10. Gulwani, S.: Programming by examples: applications, algorithms, and ambiguity resolution. In: Olivetti, N., Tiwari, A. (eds.) IJCAR 2016. LNCS (LNAI), vol. 9706, pp. 9–14. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-40229-1_2
11. Hu, Q., Cyphert, J., D’Antoni, L., Reps, T.: Synthesis with asymptotic resource bounds. In: CAV (2021)
12. Hu, Q., D’Antoni, L.: Syntax-guided synthesis with quantitative syntactic objectives. In: Chockler, H., Weissenbacher, G. (eds.) CAV 2018, Part I. LNCS, vol. 10981, pp. 386–403. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-96145-3_21

13. Johnson, S.: YACC: Yet another compiler-compiler. Technical Report Computer Science Technical report 32, Bell Laboratories (1975)
14. Kim, J., Hu, Q., D'Antoni, L., Reps, T.: Semantics-guided synthesis. *Proc. ACM on Program. Lang.* **5**(POPL), 1–32 (2021)
15. Knoth, T., Wang, D., Polikarpova, N., Hoffmann, J.: Resource-guided program synthesis. In: *PLDI*, pp. 253–268 (2019)
16. Knoth, T., Wang, D., Polikarpova, N., Hoffmann, J.: Resource-guided program synthesis. In: *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation*, pp. 253–268 (2019)
17. Kobayashi, N., Sekiyama, T., Sato, I., Unno, H.: Toward neural-network-guided program synthesis and verification. *CoRR*, abs/2103.09414 (2021)
18. Komuravelli, A., Gurfinkel, A., Chaki, S.: SMT-based model checking for recursive programs. *Formal Methods Syst. Des.* **48**(3), 175–205 (2016)
19. Lang, B.: Deterministic techniques for efficient non-deterministic parsers. In: Loeckx, J. (ed.) *ICALP 1974*. LNCS, vol. 14, pp. 255–269. Springer, Heidelberg (1974). https://doi.org/10.1007/978-3-662-21545-6_18
20. Lattner, C., Adve, V.: LLVM: a compilation framework for lifelong program analysis and transformation. In: *Proceedings of the 2004 International Symposium on Code Generation and Optimization (CGO 2004)*, Palo Alto, California, March (2004)
21. Lee, W.: Combining the top-down propagation and bottom-up enumeration for inductive program synthesis. *Proc. ACM Program. Lang.* **5**(POPL), 1–28 (2021)
22. Nori, A.V., Ozair, S., Rajamani, S.K., Vijaykeerthy, D.: Efficient synthesis of probabilistic programs. In: *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation*, pp. 208–217 (2015) *Language Design and Implementation*, pp. 208–217 (2015)
23. Pan, R., Hu, Q., Xu, G., D'Antoni, L.: Automatic repair of regular expressions. *Proc. ACM Program. Lang.* **3**(OOPSLA), 1–29 (2019)
24. Polikarpova, N., Kuraj, I., Solar-Lezama, A.: Program synthesis from polymorphic refinement types. In: *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation*, pp. 522–538 (2016)
25. Polozov, O., Gulwani, S.: Flashmeta: a framework for inductive program synthesis. In *Proceedings of the 2015 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2015*, part of SPLASH 2015, Pittsburgh, PA, USA, 25–30 October 2015, pp. 107–126 (2015)
26. Reynolds, A., Deters, M., Kuncak, V., Tinelli, C., Barrett, C.: Counterexample-guided quantifier instantiation for synthesis in SMT. In: Kroening, D., Păsăreanu, C.S. (eds.) *CAV 2015, Part II*. LNCS, vol. 9207, pp. 198–216. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-21668-3_12
27. Solar-Lezama, A.: Program sketching. *Int. J. Softw. Tools Technol. Transf.* **15**, 475–495 (2012). <https://doi.org/10.1007/s10009-012-0249-7>
28. Torlak, E., Bodik, R.: Growing solver-aided languages with Rosette. In: *Proceedings of the 2013 ACM International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software*, pp. 135–152 (2013)
29. Wang, X., Dillig, I., Singh, R.: Program synthesis using abstraction refinement. *PACMPL* **2**(POPL), 63:1-63:30 (2018)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

