
Synthesizing Tasks for Block-based Programming*

Umair Z. Ahmed¹ Maria Christakis² Aleksandr Efremov² Nigel Fernandez²
Ahana Ghosh² Abhik Roychoudhury¹ Adish Singla²

¹National University of Singapore, {umair, abhik}@comp.nus.edu.sg,

²MPI-SWS, {maria, aefremov, nfernand, gahana, adishs}@mpi-sws.org

Abstract


Block-based visual programming environments play a critical role in introducing computing concepts to K-12 students. One of the key pedagogical challenges in these environments is in designing new practice tasks for a student that match a desired level of difficulty and exercise specific programming concepts. In this paper, we formalize the problem of synthesizing visual programming tasks. In particular, given a reference visual task T^{in} and its solution code C^{in} , we propose a novel methodology to automatically generate a set $\{(T^{\text{out}}, C^{\text{out}})\}$ of new tasks along with solution codes such that tasks T^{in} and T^{out} are *conceptually similar* but *visually dissimilar*. Our methodology is based on the realization that the mapping from the space of visual tasks to their solution codes is highly discontinuous; hence, directly mutating reference task T^{in} to generate new tasks is futile. Our task synthesis algorithm operates by first mutating code C^{in} to obtain a set of codes $\{C^{\text{out}}\}$. Then, the algorithm performs *symbolic execution* over a code C^{out} to obtain a visual task T^{out} ; this step uses the Monte Carlo Tree Search (MCTS) procedure to guide the search in the symbolic tree. We demonstrate the effectiveness of our algorithm through an extensive empirical evaluation and user study on reference tasks taken from the *Hour of Code: Classic Maze* challenge by *Code.org* and the *Intro to Programming with Karel* course by *CodeHS.com*.

1 Introduction

Block-based visual programming environments are increasingly used nowadays to introduce computing concepts to novice programmers including children and K-12 students. Led by the success of environments like Scratch [29], initiatives like *Hour of Code* by *Code.org* [24] (HOC) and online platforms like *CodeHS.com* [21], block-based programming has become an integral part of introductory computer science education. Considering HOC alone, over one billion hours of block-based programming activity has been performed so far by over 50 million unique students worldwide [24, 35].


The societal need for enhancing K-12 computing education has led to a surge of interest in developing AI-driven systems for pedagogy of block-based programming [33, 26, 27, 34, 16]. Existing works have studied various aspects of intelligent support, including providing real-time next-step hints when a student is stuck solving a task [20, 36, 18, 17, 9], giving data-driven feedback about a student’s misconceptions [31, 19, 28, 30, 35], and demonstrating a worked-out solution for a task when a student lacks the required programming concepts [37]. An underlying assumption when providing such intelligent support is that afterwards the student can practice *new similar tasks* to finally learn the missing concepts. However, this assumption is far from reality in existing systems—the programming tasks are typically hand-curated by experts/tutors, and the available set of tasks is limited. Consider HOC’s *Classic Maze* challenge [23], which provides a progression of 20 tasks: Millions of students have attempted these tasks, yet when students fail to solve a task and receive assistance, they cannot practice similar tasks, hindering their ability to master the desired concepts. We seek to tackle this pedagogical challenge by developing techniques for synthesizing new programming tasks.

*Authors listed alphabetically; Correspondence to: Ahana Ghosh <gahana@mpi-sws.org>.



(a) Visual puzzle for T^{in} (b) Solution code C^{in} (c) Visual puzzle for T^{out} (d) Solution code C^{out}

Figure 1: Illustration of our methodology for task *Maze 16* from the *Hour of Code: Classic Maze* challenge by *Code.org* [23]; the complete list of tasks with their specifications is in Fig. 6.



(a) Visual puzzle for T^{in} (b) Solution code C^{in} (c) Visual puzzle for T^{out} (d) Solution code C^{out}

Figure 2: Illustration of our methodology for task *Diagonal* from the *Intro to Programming with Karel* course by *CodeHS.com* [22]; the complete list of tasks with their specifications is in Fig. 6.

We formalize the problem of synthesizing visual programming tasks of the kind found in popular learning platforms like *Code.org* (see Fig. 1) and *CodeHS.com* (see Fig. 2). As input, we are given a reference task T^{in} , specified as a visual puzzle, and its solution code C^{in} . Our goal is to synthesize a set $\{(T^{\text{out}}, C^{\text{out}})\}$ of new tasks along with their solution codes that are *conceptually similar* but *visually dissimilar* to the input. This is motivated by the need for practice tasks that on one hand exercise the same concepts, while looking fresh in order to maintain student engagement.

When tackling the problem of synthesizing new tasks with the above desirable properties, three key challenges emerge. First, we are generating problems in a *conceptual domain* with no *well-defined procedure* that students follow to solve a task—consequently, existing work on educational problem generation in procedural domains does not apply in our setting [3, 11]. Second, the mapping from the space of visual tasks to their solution codes is highly discontinuous; hence, template-based problem generation techniques [32, 25] that rely on directly mutating the input to generate new tasks is ineffective (see Section 5 where we use this approach as a baseline). Furthermore, such a direct task-mutation approach would require access to an automated solution synthesizer; however, state-of-the-art program synthesis techniques are not yet on par with experts and their minimal solutions [5, 8, 6]. Third, the space of possible tasks and their solutions is potentially unbounded, and thus, any problem generation technique that relies on exhaustive enumeration is intractable [32, 1, 2].

To overcome these challenges, we propose a novel methodology that operates by first mutating the solution code C^{in} to obtain a set of codes $\{C^{\text{out}}\}$, and then performing symbolic execution over a code C^{out} to obtain a visual puzzle T^{out} . Mutation is efficient by creating an abstract representation of C^{in} along with appropriate constraints and querying an SMT solver [4]; any solution to this query is a mutated code C^{out} . During symbolic execution, we use Monte Carlo Tree Search (MCTS) to guide the search over the (unbounded) symbolic execution tree. We demonstrate the effectiveness of our methodology by performing an extensive empirical evaluation and user study on a set of reference tasks from the *Hour of code* challenge by *Code.org* and the *Intro to Programming with Karel* course by *CodeHS.com*. In summary, our main contributions are:

- We formalize the problem of synthesizing block-based visual programming tasks (Section 2).
- We present a novel approach for generating new visual tasks along with solution codes such that they are conceptually similar but visually dissimilar to a given reference task (Section 3).
- We demonstrate the effectiveness of our approach through an extensive empirical evaluation and user study on reference tasks from real-world programming platforms (Section 4 and Section 5).

2 Problem Formulation

The space of tasks. We define a task as a tuple $T := (T_{\text{vis}}, T_{\text{store}}, T_{\text{size}})$, where T_{vis} denotes the visual puzzle, T_{store} the available block types, and T_{size} the maximum number of blocks allowed in the

solution code. For instance, considering the task $T := T^{\text{in}}$ in Fig. 1a, T_{vis} is illustrated in Fig. 1a, $T_{\text{store}} = \{\text{move, turnL, turnR, REPEATUNTIL, IF}\}$, and $T_{\text{size}} = 4$.

The space of codes. The programming environment has a domain-specific language (DSL), which defines the set of valid codes \mathcal{C} and is shown in Fig. 4a. A code $C \in \mathcal{C}$ is characterized by several properties, such as the set $\mathcal{C}_{\text{blocks}}$ of block types in C , the number of blocks $\mathcal{C}_{\text{size}}$, the depth $\mathcal{C}_{\text{depth}}$ of the corresponding Abstract Syntax Tree (AST), and the nesting structure $\mathcal{C}_{\text{struct}}$ representing programming concepts exercised by C . For instance, considering the code $C := C^{\text{in}}$ in Fig. 1b, $\mathcal{C}_{\text{blocks}} = \{\text{move, turnL, REPEATUNTIL, IF}\}$, $\mathcal{C}_{\text{size}} = 4$, $\mathcal{C}_{\text{depth}} = 3$, and $\mathcal{C}_{\text{struct}} = \{\text{RUN}\{\text{REPEATUNTIL}\{\text{IF}\}\}\}$.

Below, we introduce two useful definitions relating the task and code space.

Definition 1 (Solution code). C is a solution code for T if the following holds: C successfully solves the visual puzzle T_{vis} , $\mathcal{C}_{\text{blocks}} \subseteq T_{\text{store}}$, and $\mathcal{C}_{\text{size}} \leq T_{\text{size}}$. \mathcal{C}_T denotes the set of all solution codes for T .

Definition 2 (Minimality of a task). Given a solvable task T with $|\mathcal{C}_T| \geq 1$ and a threshold $\delta \in \mathbb{N}$, the task is minimal if $\nexists C \in \mathcal{C}_T$ such that $\mathcal{C}_{\text{size}} < T_{\text{size}} - \delta$.

Next, we introduce two definitions formalizing the notion of conceptual similarity. Definition 3 formalizes conceptual similarity of a task T along with one solution code C . Since a task can have multiple solution codes, Definition 4 provides a stricter notion of conceptual similarity of a task T for all its solution codes. These definitions are used in our objective of task synthesis in conditions (I) and (V) below.

Definition 3 (Conceptual similarity of (T, C)). Given a reference $(T^{\text{in}}, C^{\text{in}})$ and a threshold $\delta \in \mathbb{N}$, a task T along with a solution code C is conceptually similar to $(T^{\text{in}}, C^{\text{in}})$ if the following holds: $T_{\text{store}} = T^{\text{in}}_{\text{store}}$, $|T_{\text{size}} - T^{\text{in}}_{\text{size}}| \leq \delta$, and $\mathcal{C}_{\text{struct}} = C^{\text{in}}_{\text{struct}}$.

Definition 4 (Conceptual similarity of (T, \cdot)). Given a reference $(T^{\text{in}}, C^{\text{in}})$ and a threshold $\delta \in \mathbb{N}$, a task T is conceptually similar to $(T^{\text{in}}, C^{\text{in}})$ if the following holds: $T_{\text{store}} = T^{\text{in}}_{\text{store}}$, $|T_{\text{size}} - T^{\text{in}}_{\text{size}}| \leq \delta$, and $\forall C \in \mathcal{C}_T, \mathcal{C}_{\text{struct}} = C^{\text{in}}_{\text{struct}}$.

Environment domain knowledge. We now formalize our domain knowledge about the block-based environment to measure *visual dissimilarity* of two tasks, and capture some notion of *interestingness and quality* of a task. Given tasks T and T' , we measure their visual dissimilarity by an environment-specific function $\mathcal{F}_{\text{diss}}(T_{\text{vis}}, T'_{\text{vis}}) \in [0, 1]$. Moreover, we measure generic quality of a task with function $\mathcal{F}_{\text{qual}}(T_{\text{vis}}, C) \in [0, 1]$. We provide specific instantiations of $\mathcal{F}_{\text{diss}}$ and $\mathcal{F}_{\text{qual}}$ in our evaluation.

Objective of task synthesis. Given a reference task T^{in} and a solution code $C^{\text{in}} \in \mathcal{C}_{T^{\text{in}}}$ as input, we seek to generate a set $\{(T^{\text{out}}, C^{\text{out}})\}$ of new tasks along with solution codes that are *conceptually similar* but *visually dissimilar* to the input. Formally, given parameters $(\delta_{\text{size}}, \delta_{\text{diss}}, \delta_{\text{qual}})$, our objective is to synthesize new tasks meeting the following conditions:

- (I) $(T^{\text{out}}, C^{\text{out}})$ is conceptually similar to $(T^{\text{in}}, C^{\text{in}})$ with threshold δ_{size} in Definition 3.
- (II) T^{out} is visually dissimilar to T^{in} with margin δ_{diss} , i.e., $\mathcal{F}_{\text{diss}}(T^{\text{in}}_{\text{vis}}, T^{\text{out}}_{\text{vis}}) \geq \delta_{\text{diss}}$.
- (III) T^{out} has a quality score above threshold δ_{qual} , i.e., $\mathcal{F}_{\text{qual}}(T^{\text{out}}_{\text{vis}}, C^{\text{out}}) \geq \delta_{\text{qual}}$.

In addition, depending on the use case, it is desirable that the new tasks satisfy the following criteria:

- (IV) C^{out} is different from the input solution code, i.e., $C^{\text{out}} \neq C^{\text{in}}$.
- (V) T^{out} is conceptually similar to $(T^{\text{in}}, C^{\text{in}})$ with threshold δ_{size} in Definition 4.
- (VI) T^{out} is minimal as per Definition 2 for a desired value of δ_{mini} (e.g., $\delta_{\text{mini}} = 0$ or $\delta_{\text{mini}} = 1$).

3 Our Task Synthesis Algorithm

We now present the pipeline of our algorithm (see Fig. 3), which takes as input a reference task T^{in} and its solution code C^{in} , and generates a set $\{(T^{\text{out}}, C^{\text{out}})\}$ of new tasks with their solution codes. The goal is for this set to be conceptually similar to $(T^{\text{in}}, C^{\text{in}})$, but for new tasks $\{T^{\text{out}}\}$ to be visually dissimilar to T^{in} . This is achieved by two main stages: (1) mutation of C^{in} to obtain a set $\{C^{\text{out}}\}$, and (2) symbolic execution of each C^{out} to create a task T^{out} . The first stage, presented in Section 3.1, converts C^{in} into an abstract representation restricted by a set of constraints (Fig. 3(a)), which must be satisfied by any generated C^{out} (Fig. 3(b)). The second stage, described in Section 3.2, applies symbolic execution on each code C^{out} to create a corresponding visual task T^{out} (Fig. 3(c)) while using Monte Carlo Tree Search (MCTS) to guide the search in the symbolic execution tree.



Figure 3: Stages in our task synthesis algorithm.

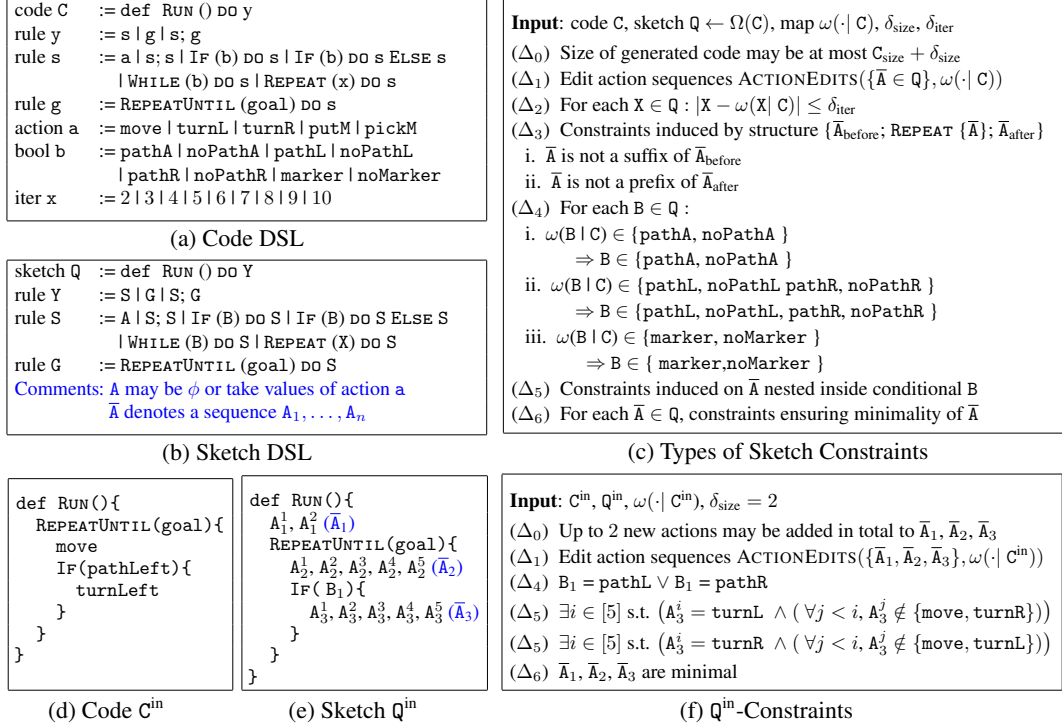


Figure 4: Illustration of key steps in Code Mutation. Fig. 4d shows code C^{in} from Fig. 1b. The code mutation stage, when applied to C^{in} , generates many output codes, including C^{out} in Fig. 1d.

3.1 Code Mutation

This stage in our pipeline mutates code C^{in} of task T^{in} such that its conceptual elements are preserved. Our mutation procedure consists of three main steps. First, we generate an abstract representation of C^{in} , called *sketch*. Second, we restrict the sketch with constraints that describe the space of its concrete instantiations. Although this formulation is inspired from work on generating algebra problems [32], we use it in the entirely different context of generating conceptually similar mutations of C^{in} . This is achieved in the last step, where we use the sketch and its constraints to query an SMT solver [4]; the query solutions are mutated codes $\{C^{\text{out}}\}$ such that $C^{\text{out}}_{\text{struct}} = C^{\text{in}}_{\text{struct}}$ (see Definition 3).

Step 1: Sketch. The sketch of code C, denoted by Q, is an abstraction of C capturing its skeleton and generalizing C to the space of conceptually similar codes. Q, expressed in the language of Fig. 4b, is generated from C with mapping Ω . In particular, the map exploits the AST structure of the code: the AST is traversed in a depth-first manner, and all values are replaced with their corresponding sketch variables, i.e., action a, bool b, and iter x are replaced with A, B, and X, respectively. In the following, we also use mapping $\omega(\cdot | C)$, which takes a sketch variable in Q and returns its value in C.

In addition to the above, we may extend a variable A to an *action sequence* \bar{A} , since any A is allowed to be empty (ϕ). We may also add an action sequence of length δ_{size} at the beginning and end of the obtained sketch. As an example, consider the code in Fig. 4d and the resulting sketch in Fig. 4e. Notice that, while we add an action sequence at the beginning of the sketch (\bar{A}_1), no action sequence is appended at the end because construct REPEATUNTIL renders any succeeding code unreachable.

Step 2: Sketch constraints. Sketch constraints restrict the possible concrete instantiations of a sketch by encoding the required semantics of the mutated codes. All constraint types are in Fig. 4c.

In particular, Δ_0 restricts the size of the mutated code within δ_{size} . Δ_1 specifies the allowed mutations to an action sequence based on its value in the code, given by $\omega(\bar{A} | C)$. For instance, this constraint could result in converting all turnLeft actions of a sequence to turnRight. Δ_2 restricts the possible values of the REPEAT counter within threshold δ_{iter} . Δ_3 ensures that the REPEAT counter is optimal, i.e., action subsequences before and after this construct are not nested in it. Δ_4 specifies the possible values of the IF condition based on its value in the code, given by $\omega(B | C)$. Δ_5 refers to constraints imposed on action sequences nested within conditionals. As an example, consider




Figure 5: Illustration of symbolic execution on C^{out} from Fig. 1d. (b) shows the initial configuration of the agent’s location and orientation as well as the status of the grid cells (*unknown*, *free*, *blocked*, *goal*). (c)–(e) show the symbolic execution steps where conditions `goal` and `pathRight` are `False`. (f) shows the step where `goal` is `True`. (g) shows the post-processing step where a puzzle $T_{\text{vis}}^{\text{out}}$ is obtained.

Δ_5 in Fig. 4f, which states that if $B_1 = \text{pathLeft}$, then the nested action sequence must have at least one `turnLeft` action, and the first occurrence of this action must not be preceded by a `move` or `turnRight`, thus preventing invalid actions within the conditional. Δ_6 ensures minimality of an action sequence, i.e., optimality of the constituent actions to obtain the desired output. This constraint would, for instance, eliminate redundant sequences such as `turnLeft`, `turnRight`, which does not affect the output, or `turnLeft`, `turnLeft`, `turnLeft`, whose output could be achieved by a single `turnRight`. All employed elimination sequences can be found in the supplementary material. The entire list of constraints applied on the solution code in Fig. 4d is shown in Fig. 4f.

Step 3: SMT query. For a sketch Q generated from code C and its constraints, we pose the following query to an SMT solver: (sketch Q , Q -constraints). As a result, the solver generates a set of instantiations, which are conceptually similar to C . In our implementation, we used the Z3 solver [7]. For the code in Fig. 4d, Z3 generated 66 mutated codes in 0.8s from an exhaustive space of 2, 997 possible codes with $\delta_{\text{size}} = 2$. One such mutation is shown in Fig. 1d.

While this approach generates codes that are devoid of most semantic irregularities, it has its limitations. Certain irregularities continue to exist in some generated codes: An example of such a code included the action sequence `move`, `turnLeft`, `move`, `turnLeft`, `move`, `turnLeft`, `move`, `turnLeft`, which results in the agent circling back to its initial location in the task space. This kind of undesirable behaviour is eliminated in the symbolic execution stage of our pipeline.

3.2 Symbolic Execution

Symbolic execution [13] is an automated test-generation technique that symbolically explores execution paths in a program. During exploration of a path, it gathers symbolic constraints over program inputs from statements along the path. These constraints are then mutated (according to a search strategy), and an SMT solver is queried to generate new inputs that explore another path.

Obtaining visual tasks with symbolic execution. This stage in our pipeline applies symbolic execution on each generated code C^{out} to obtain a suitable visual task T^{out} . The program inputs of C^{out} are the agent’s initial location/orientation and the status of the grid cells (*unknown*, *free*, *blocked*, *marker*, *goal*), which is initially *unknown*. Symbolic execution collects constraints over these from code statements. As in Fig. 5 for one path, symbolic execution generates a visual task for each path in C^{out} .

However, not all of these tasks are suitable. For instance, if the goal is reached after the first move in Fig. 1d, all other statements in C^{out} are not covered, rendering the task less suitable for this code. Naïvely, symbolic execution could first enumerate all paths in C^{out} and their corresponding tasks, and then rank them in terms of suitability. However, solution codes may have an unbounded number of paths, which leads to *path explosion*, that is, the inability to cover all paths with tractable resources.

Guiding symbolic execution using Monte Carlo Tree Search (MCTS). To address this issue, we use MCTS [14] as a search strategy in symbolic execution with the goal of generating more suitable tasks with fewer resources—we define task suitability next. Symbolic execution has been previously combined with MCTS in order to direct the exploration towards costly paths [15]. In the supplementary material, we provide an example demonstrating how MCTS could guide the symbolic execution in generating more suitable tasks.

As previously observed [12], a critical component of effectively applying MCTS is to define an evaluation function that describes the desired properties of the output, i.e., the visual tasks. Tailoring the evaluation function to our unique setting is exactly what differentiates our approach from existing work. In particular, our evaluation function, $\mathcal{F}_{\text{score}}$, distinguishes suitable tasks by assigning a score ($\in [0, 1]$) to them, which guides the MCTS search. A higher $\mathcal{F}_{\text{score}}$ indicates a more suitable task.

Task T	T _{store}	T _{size} (= C _{size})	C _{depth}	Type: Source
H1	move, turnL, turnR	5	1	HOC: <i>Maze 4</i> [23]
H2	move, turnL, turnR, REPEAT	3	2	HOC: <i>Maze 7</i> [23]
H3	move, turnL, turnR, REPEAT	5	2	HOC: <i>Maze 8</i> [23]
H4	move, turnL, turnR, REPEATUNTIL	5	2	HOC: <i>Maze 12</i> [23]
H5	move, turnL, turnR, REPEATUNTIL, IF	4	3	HOC: <i>Maze 16</i> [23]
H6	move, turnL, turnR, REPEATUNTIL, IFELSE	4	3	HOC: <i>Maze 18</i> [23]
K7	move, turnL, turnR, pickM, putM	5	1	Karel: <i>Our first</i> [22]
K8	move, turnL, turnR, pickM, putM, REPEAT	4	2	Karel: <i>Square</i> [22]
K9	move, turnL, turnR, pickM, putM, REPEAT, IFELSE	5	3	Karel: <i>One ball in each spot</i> [22]
K10	move, turnL, turnR, pickM, putM, WHILE	7	2	Karel: <i>Diagonal</i> [22]

Figure 6: Datasets for HOC and Karel tasks.

Its constituent components are: (i) $\mathcal{F}_{\text{cov}}(\mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{C}^{\text{out}}) \in \{0, 1\}$, which evaluates to 1 in the event of complete coverage of code \mathbf{C}^{out} by task $\mathbf{T}_{\text{vis}}^{\text{out}}$ and 0 otherwise; (ii) $\mathcal{F}_{\text{diss}}(\mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{T}_{\text{vis}}^{\text{in}}) \in [0, 1]$, which evaluates the dissimilarity of $\mathbf{T}_{\text{vis}}^{\text{out}}$ to $\mathbf{T}_{\text{vis}}^{\text{in}}$ (see Section 2); (iii) $\mathcal{F}_{\text{qual}}(\mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{C}^{\text{out}}) \in [0, 1]$, which evaluates the quality and validity of $\mathbf{T}_{\text{vis}}^{\text{out}}$; (iv) $\mathcal{F}_{\text{nocrash}}(\mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{C}^{\text{out}}) \in \{0, 1\}$, which evaluates to 0 in case the agent crashes into a wall and 1 otherwise; and (v) $\mathcal{F}_{\text{nocut}}(\mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{C}^{\text{out}}) \in \{0, 1\}$, which evaluates to 0 if there is a *shortcut sequence* of actions (a in Fig. 4a) smaller than $\mathbf{C}_{\text{size}}^{\text{out}}$ that solves $\mathbf{T}_{\text{vis}}^{\text{out}}$ and 1 otherwise. $\mathcal{F}_{\text{qual}}$ and $\mathcal{F}_{\text{nocut}}$ also resolve the limitations of our mutation stage by eliminating codes and tasks that lead to undesirable agent behavior. We instantiate $\mathcal{F}_{\text{score}}$ in the next section.

4 Experimental Evaluation

In this section, we evaluate our task synthesis algorithm on HOC and Karel tasks. Our implementation is publicly available.² While we give an overview of key results here, a detailed description of our setup and additional experiments can be found in the supplementary material.

4.1 Reference Tasks and Specifications

Reference tasks. We use a set of ten reference tasks from HOC and Karel, shown in Fig. 6. The HOC tasks were selected from the *Hour of Code: Classic Maze* challenge by *Code.org* [23] and the Karel tasks from the *Intro to Programming with Karel* course by *CodeHS.com* [22]. The DSL of Fig. 4a is generic in that it includes both HOC and Karel codes, with the following differences: (i) construct `WHILE`, marker-related actions `putM`, `pickM`, and conditions `noPathA`, `noPathL`, `noPathR`, `marker`, `noMarker` are specific to Karel only; (ii) construct `REPEATUNTIL` and `goal` are specific to HOC only. Furthermore, the puzzles for HOC and Karel are of different styles (see Fig. 1 and Fig. 2). For all tasks, the grid size of the puzzles is fixed to 10×10 cells (grid-size parameter $n = 10$).

Specification of scoring functions. $\mathcal{F}_{\text{qual}}(\mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{C}^{\text{out}}) \in [0, 1]$ was approximated as the sum of the normalized counts of ‘moves’, ‘turns’, ‘segments’, and ‘long-segments’ in the grid; segments and long-segments are sequences of ≥ 3 and ≥ 5 move actions respectively. More precisely, for HOC tasks, we used the following function where features are computed by executing \mathbf{C}^{out} on $\mathbf{T}_{\text{vis}}^{\text{out}}$:

$$\mathcal{F}_{\text{qual}}^{\text{HOC}}(\mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{C}^{\text{out}}) = \frac{1}{4} \left(\frac{\#\text{moves}}{2n} + \frac{\#\text{turns}}{n} + \frac{\#\text{segments}}{n/2} + \frac{\#\text{long-segments}}{n/3} \right).$$

Furthermore, in our implementation, $\mathcal{F}_{\text{qual}}(\cdot)$ value was set to 0 when $\mathcal{F}_{\text{nocrash}}(\cdot) = 0$. For Karel tasks, $\mathcal{F}_{\text{qual}}$ additionally included the normalized counts of `putM` and `pickM`, and is provided in the supplementary material. $\mathcal{F}_{\text{diss}}(\mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{T}_{\text{vis}}^{\text{in}}) \in [0, 1]$ was computed based on the dissimilarity of the agent’s initial location/orientation w.r.t. $\mathbf{T}_{\text{vis}}^{\text{in}}$, and the grid-cell level dissimilarity based on the Hamming distance between $\mathbf{T}_{\text{vis}}^{\text{out}}$ and $\mathbf{T}_{\text{vis}}^{\text{in}}$. More precisely, we used the following function:

$$\mathcal{F}_{\text{diss}}(\mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{T}_{\text{vis}}^{\text{in}}) = \frac{1}{3} \left(\text{diss}(\text{loc} \mid \mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{T}_{\text{vis}}^{\text{in}}) + \text{diss}(\text{dir} \mid \mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{T}_{\text{vis}}^{\text{in}}) + \text{diss}(\text{grid-cells} \mid \mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{T}_{\text{vis}}^{\text{in}}) \right)$$

where $\text{diss}(\text{loc} \mid \mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{T}_{\text{vis}}^{\text{in}}) \in \{0, 1\}$, $\text{diss}(\text{dir} \mid \mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{T}_{\text{vis}}^{\text{in}}) \in \{0, 1\}$, and $\text{diss}(\text{grid-cells} \mid \mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{T}_{\text{vis}}^{\text{in}}) \in [0, 1]$ (after the Hamming distance is normalized with a factor of $\frac{2}{n^2}$).

²https://github.com/adishs/neurips2020_synthesizing-tasks_code

Task T^{in}	Code Mutation				Symbolic Execution			Fraction of T^{out} with criteria		
	2: # $C_{\Delta=0}^{\text{out}}$	3: # $C_{\Delta=0,1}^{\text{out}}$	4: # $C_{\Delta=\text{all}}^{\text{out}}$	5: Time	6: # C^{out}	7: # T^{out}	8: Time	9: (V)	10: (VI) $_{\delta_{\text{mini}}=1}$	11: (VI) $_{\delta_{\text{mini}}=0}$
H1	3, 159	112	64	0.6s	28	272	68s	1.00	1.00	1.00
H2	8, 991	594	138	1.7s	48	428	61s	1.00	1.00	1.00
H3	798, 255	13, 122	720	13.3s	196	1, 126	60s	0.90	0.98	0.90
H4	5, 913	152	108	1.0s	44	404	167s	1.00	1.00	0.50
H5	2, 997	294	66	0.8s	46	444	348s	0.98	0.59	0.27
H6	1, 728	294	54	0.6s	48	480	347s	0.80	0.45	0.07
K7	96, 875	150	122	1.3s	122	1, 196	61s	1.00	1.00	1.00
K8	484, 875	4, 506	990	11.6s	469	4, 506	63s	1.00	1.00	1.00
K9	8.595×10^6	60, 768	888	11.3s	432	4, 258	185s	0.92	0.92	0.88
K10	132.625×10^6	19, 328	1, 404	17.1s	532	5, 032	158s	1.00	1.00	1.00

Figure 7: Results on HOC and Karel tasks; details are provided in Section 4.

Next, we define the evaluation function $\mathcal{F}_{\text{score}}(T^{\text{out}}, C^{\text{out}}, T^{\text{in}}, C^{\text{in}}) \in [0, 1]$ used by MCTS:

$$\mathcal{F}_{\text{score}}(T^{\text{out}}, C^{\text{out}}, T^{\text{in}}, C^{\text{in}}) = \underbrace{\mathbb{1}(\mathcal{F}_{\text{qual}}(T_{\text{vis}}^{\text{out}}, C^{\text{out}}) \geq \delta_{\text{qual}}, \mathcal{F}_{\text{nocrash}}(T_{\text{vis}}^{\text{out}}, C^{\text{out}}) = 1, \mathcal{F}_{\text{nocut}}(T_{\text{vis}}^{\text{out}}, C^{\text{out}}) = 1)}_{(i)} \cdot \underbrace{[\alpha_1 \mathcal{F}_{\text{cov}}(T_{\text{vis}}^{\text{out}}, C^{\text{out}}) + \alpha_2 \mathcal{F}_{\text{qual}}(T_{\text{vis}}^{\text{out}}, C^{\text{out}}) + \alpha_3 \mathcal{F}_{\text{diss}}(T_{\text{vis}}^{\text{out}}, T_{\text{vis}}^{\text{in}})]}_{(ii)}$$

where $\mathbb{1}$ is an indicator function and each constant $\alpha = 1/3$. Component (ii) in the above function supplies the gradients for guiding the search in MCTS; Component (i) is applied at the end of the MCTS run to pick the output. More precisely, the *best* task (i.e. the one with the highest $\mathcal{F}_{\text{score}}$ value) is picked only from the pool of generated tasks which have $\mathcal{F}_{\text{score}}(\cdot) > 0$ and satisfy $\mathcal{F}_{\text{cov}}(\cdot) = 1$.

Specification of task synthesis and MCTS. As per Section 2, we set the following thresholds for our algorithm: (i) $\delta_{\text{size}} = 2$, (ii) $\delta_{\text{diss}} = 0.33$, and (iii) $\delta_{\text{qual}} = 0.2$ for codes with `WHILE` or `REPEATUNTIL`, and 0.05 otherwise. We run MCTS 10 times per code, with each run generating one task. We set the maximum iterations of a run to 2 million (M) and the exploration constant to 2 [14]. Even when considering a tree depth of $2n$ ($= 20$), there are millions of leaves for difficult tasks H5 and H6, reflecting the complexity of task generation. For each code C^{out} , we generated 10 different visual tasks. To ensure sufficient diversity among the tasks generated for the same code, we introduced a measure $\mathcal{F}_{\text{diversity}}$. This measure, not only ensures visual task dissimilarity, but also ensures sufficient diversity in entire symbolic paths during generation (for details, see supplementary material).

4.2 Results

Performance of task synthesis algorithm. Fig. 7 shows the results of our algorithm. The second column illustrates the enormity of the unconstrained space of mutated codes; we only impose size constraint Δ_0 from Fig. 4c. We then additionally impose constraint Δ_1 resulting in a partially constrained space of mutated codes (column 3), and finally apply all constraints from Fig. 4c to obtain the final set of generated codes (column 4). This reflects the systematic reduction in the space of mutated codes by our constraints. Column 5 shows the total running time for generating the final codes, which denotes the time taken by Z3 to compute solutions to our mutation query. As discussed in Section 3.1, few codes with semantic irregularities still remain after the mutation stage. The symbolic execution stage eliminates these to obtain the reduced set of valid codes (column 6). Column 7 shows the final number of generated tasks and column 8 is the average time per output task (i.e., one MCTS run).

Analyzing output tasks. We further analyze the generated tasks based on the objectives of Section 2. All tasks satisfy properties (I)–(III) by design. Objective (IV) is easily achieved by excluding generated tasks for which $C^{\text{out}} = C^{\text{in}}$. For a random sample of 100 of the generated tasks per reference task, we performed manual validation to determine whether objectives (V) and (VI) are met. The fraction of tasks that satisfy these objectives is listed in the last three columns of Fig. 7. We observe that the vast majority of tasks meet the objectives, even if not by design. For H6, the fraction of tasks satisfying (VI) is low because the corresponding codes are generic enough to solve several puzzles.

Deep dive into an MCTS run. To offer more insight into the task generation process, we take a closer look at an MCTS run for task H5, shown in Fig. 8. Fig. 8a illustrates the improvement in various components of $\mathcal{F}_{\text{score}}$ as the number of MCTS iterations increases. Best tasks at different iterations are shown in Fig. 8b, 8c, 8d. As expected, the more the iterations, the better the tasks are.




Figure 8: Illustration of a single MCTS run on C^{out} from Fig. 1d obtained from solution code of task H5 by mutation. (a) shows the temporal trends of different feature values in $\mathcal{F}_{\text{score}}$ averaged over a time window of 100 steps. (b)–(d) show the best, i.e., highest scoring, tasks generated up to times 2×10^2 , 2×10^4 , and 2×10^6 respectively. $T_{\text{vis}}^{\text{out}}$ shown in Fig. 1c is the puzzle produced in (d).

Remarks. We also ran the mutation stage by enumerating the programs within size constraints and then post-checking other constraints without Z3. This implementation leads to a run-time increase by a factor of 10 to 100 for different tasks. So, Z3 seems to be very effective by jointly considering all the constraints. As a search method, although MCTS seems computationally expensive, the actual run-time and memory footprint of an MCTS run depend on the unique traces explored (i.e., unique symbolic executions done)—this number is typically much lower than the number of iterations, also see discussion in the supplementary material. Considering the MCTS output in Figs. 8c, 8d, to obtain a comparable evaluation score through a random search, the corresponding number of unique symbolic executions required is at least 10 times more than executed by MCTS. We note that while we considered one I/O pair for Karel tasks, our methodology can be easily extended to multiple I/O pairs by adapting techniques designed for generating diverse tasks.

5 User Study and Comparison with Alternate Methods

In this section, we evaluate our task synthesis algorithm with a user study focusing on tasks H2, H4, H5, and H6. We developed an online app³, which uses the publicly available toolkit of Blockly Games [10] and provides an interface for a participant to practice block-based programming tasks for HOC. Each “practice session” of the study involves three steps: (i) a reference task $T^{\text{in}} \in \{H2, H4, H5, H6\}$ is shown to the participant along with its solution code C^{in} , (ii) a new task T^{out} is generated for which the participant has to provide a solution code, and (iii) a post-survey asks the participant to assess the visual dissimilarity of the two tasks on a 4-point Likert scale as used in [25]. Details on the app interface and questionnaire are provided in the supplementary material. Participants for the study were recruited through Amazon Mechanical Turk. We only selected four tasks due to the high cost involved in conducting the study (about 1.8 USD per participant). The number of participants and their performance are documented in Fig. 9.

Baselines and methods evaluated. We evaluated four different methods, including three baselines (SAME, TUTOR, MUTASK) and our algorithm (SYNTASK). SAME generates tasks such that $T^{\text{in}} = T^{\text{out}}$. TUTOR produces tasks that are similar to T^{in} and designed by an expert. We picked similar problems from the set of 20 *Classic Maze* challenge [23] tasks exercising the same programming concepts: *Maze 6, 9* for H2, *Maze 11, 13* for H4, *Maze 15, 17* for H5, and *Maze 19* for H6.

MUTASK generated tasks by directly mutating the grid-world of the original task, i.e., by moving the agent or goal by up to two cells and potentially changing the agent’s orientation. A total of 18, 20, 15, and 17 tasks were generated for H2, H4, H5, and H6, respectively. Fig. 10 shows two output tasks for H4 and illustrates the challenge in directly mutating the input task, given the high discontinuity in mapping from the space of tasks to their codes. For H4, a total of 14 out of 20 new tasks were structurally very different from the input.

SYNTASK uses our algorithm to generate tasks. We picked the generated tasks from three groups based on the size of the code mutations from which they were produced, differing from the reference solution code by $+\delta_{\text{size}}$ for $\delta_{\text{size}} \in \{0, 1, 2\}$. For H2 and H4, we randomly selected 5 tasks from each group, for a total of 15 new tasks per reference task. For H5 and H6, we selected 10 tasks from the first group ($\delta_{\text{size}} = 0$) only, due to their complexity stemming from nested constructs in their codes. We observed that TUTOR tasks for H5, H6 were also of $\delta_{\text{size}} = 0$, i.e., $C_{\text{size}}^{\text{out}} = C_{\text{size}}^{\text{in}}$. All the generated tasks picked for SYNTASK adhere to properties (I)–(VI) in Section 2.

³<https://www.teaching-blocks.cc/>

Method	Total participants					Fraction of tasks solved					Time spent in secs					Visual dissimilarity				
	H-	H2	H4	H5	H6	H-	H2	H4	H5	H6	H-	H2	H4	H5	H6	H-	H2	H4	H5	H6
SAME	96	24	24	24	24	.94	.92	1.00	.96	.88	89	60	59	93	145	1.07	1.12	1.04	1.00	1.12
TUTOR	170	48	48	49	25	.90	.90	.92	.88	.92	121	107	113	118	169	2.90	2.81	2.79	2.96	3.16
MUTTASK	278	72	79	60	67	.68	.76	.71	.65	.60	219	135	299	219	215	2.17	2.36	2.33	1.95	1.99
SYNTASK	197	59	57	40	41	.89	.92	.89	.92	.83	144	85	183	130	189	2.63	2.41	2.42	2.68	3.20

Figure 9: User study results for HOC tasks (**H**-represents all tasks in the study); see Section 5.




Figure 10: MUTTASK applied to H4. T_{vis}^{in} and C^{in} are shown in (a) and (d). (b)–(c) illustrate two tasks T_{vis}^{out} obtained via small mutations of T_{vis}^{in} . (e) is the smallest solution code for (b) and is structurally similar to C^{in} . (f) is the smallest solution code for (c) and is drastically different from C^{in} .

Results on task solving. In terms of successfully solving the generated tasks, SAME performed best (mean success = 0.94) in comparison to TUTOR (mean = 0.90), SYNTASK (mean = 0.89), and MUTTASK (mean = 0.68)—this is expected given the tasks generated by SAME. In comparison to TUTOR, the performance of SYNTASK was not significantly different ($\chi^2 = 0.04, p = 0.83$); in comparison to MUTTASK, SYNTASK performed significantly better ($\chi^2 = 28.74, p < e^{-8}$). The complexity of the generated tasks is also reflected in the average time that participants spent on solving them. As shown in Fig. 9, they spent more time solving the tasks generated by MUTTASK.

Results on visual task dissimilarity. Visual dissimilarity was measured on a Likert scale ranging from 1–4, 1 being highly similar and 4 highly dissimilar. Comparing the dissimilarity of the generated tasks w.r.t. the reference task, we found that the performance of SAME was worst (mean dissimilarity = 1.07), while that of TUTOR was best (mean = 2.90). SYNTASK (mean = 2.63) performed significantly better than MUTTASK (mean = 2.17), yet slightly worse than TUTOR. This is because TUTOR generates tasks with additional distracting paths and noise, which can also be done by our algorithm (although not done for this study). Moreover, for H2, which had no conditionals, the resulting codes were somewhat similar, and so were the generated puzzles. When excluding H2 from the analysis, the difference between SYNTASK (mean = 2.72) and TUTOR (mean = 2.93) was not statistically significant. A detailed distribution of the responses can be found in the supplementary material.

Remarks. SAME’s performance in terms of tasks solved is below 1.00, possibly because participants overlooked the solution of Step 1, unaware they will be receiving the same task in Step 2, and the app did not allow them to go back to Step 1. This user study provides a proof-of-concept; more elaborate studies are needed to fully reach the motivational goal of teaching K-12 students, and evaluate the long term impact on students’ concept learning. As additional studies, it would be important to understand the sensitivity of user study results w.r.t. the Likert scale definition; another possibility is to use pairwise comparisons in eliciting user evaluations.

6 Conclusions and Outlook

We developed techniques for a critical aspect of pedagogy in block-based programming: Automatically generating new tasks that exercise specific programming concepts, while looking visually dissimilar to input. We demonstrated the effectiveness of our methodology through an extensive empirical evaluation and user study on reference tasks from popular programming platforms. We believe our techniques have the potential to drastically improve the success of pedagogy in block-based visual programming environments by providing tutors and students with a substantial pool of new tasks. Beyond the application domain of programming education, our methodology can be used for generating large-scale datasets consisting of tasks and solution codes with desirable characteristics—this can be potentially useful for training neural program synthesis methods.

There are several promising directions for future work, including but not limited to: Learning a policy to guide the MCTS procedure (instead of running vanilla MCTS); automatically learning the constraints and cost function from a human-generated pool of problems; and applying our methodology to other programming environments (e.g., Python problems).

Broader Impact

This paper develops new techniques for improving pedagogy in block-based visual programming environments. Such programming environments are increasingly used nowadays to introduce computing concepts to novice programmers, and our work is motivated by the clear societal need of enhancing K-12 computing education. In existing systems, the programming tasks are hand-curated by tutors, and the available set of tasks is typically very limited. This severely limits the utility of existing systems for long-term learning as students do not have access to practice tasks for mastering the programming concepts.

We take a step towards tackling this challenge by developing a methodology to generate new practice tasks for a student that match a desired level of difficulty and exercise specific programming concepts. Our task synthesis algorithm is able to generate 1000's of new similar tasks for reference tasks taken from the *Hour of Code: Classic Maze* challenge by *Code.org* and the *Intro to Programming with Karel* course by *CodeHS.com*. Our extensive experiments and user study further validate the quality of the generated tasks. Our task synthesis algorithm could be useful in many different ways in practical systems. For instance, tutors can assign new practice tasks as homework or quizzes to students to check their knowledge, students can automatically obtain new similar tasks after they failed to solve a given task and received assistance, and intelligent tutoring systems could automatically generate a personalized curriculum of problems for a student for long-term learning.

Acknowledgments and Disclosure of Funding

We would like to thank the anonymous reviewers for their helpful comments. Ahana Ghosh was supported by Microsoft Research through its PhD Scholarship Programme. Umair Z. Ahmed and Abhik Roychoudhury were supported by the National Research Foundation, Singapore and National University of Singapore through its National Satellite of Excellence in Trustworthy Software Systems (NSOE-TSS) project under the National Cybersecurity R&D (NCR) Grant award no. NRF2018NCR-NSOE003-0001.

References

- [1] Umair Z. Ahmed, Sumit Gulwani, and Amey Karkare. Automatically generating problems and solutions for natural deduction. In *IJCAI*, pages 1968–1975, 2013.
- [2] Chris Alvin, Sumit Gulwani, Rupak Majumdar, and Supratik Mukhopadhyay. Synthesis of geometry proof problems. In *AAAI*, pages 245–252, 2014.
- [3] Erik Andersen, Sumit Gulwani, and Zoran Popovic. A trace-based framework for analyzing and synthesizing educational progressions. In *CHI*, pages 773–782, 2013.
- [4] Clark W. Barrett and Cesare Tinelli. Satisfiability modulo theories. In *Handbook of Model Checking*, pages 305–343. Springer, 2018.
- [5] Rudy Bunel, Matthew J. Hausknecht, Jacob Devlin, Rishabh Singh, and Pushmeet Kohli. Leveraging grammar and reinforcement learning for neural program synthesis. In *ICLR*, 2018.
- [6] Xinyun Chen, Chang Liu, and Dawn Song. Execution-guided neural program synthesis. In *ICLR*, 2018.
- [7] Leonardo de Moura and Nikolaj Bjørner. Z3: An efficient SMT solver. In *TACAS*, pages 337–340, 2008.
- [8] Jacob Devlin, Rudy Bunel, Rishabh Singh, Matthew J. Hausknecht, and Pushmeet Kohli. Neural program meta-induction. In *Advances in Neural Information Processing Systems*, pages 2080–2088, 2017.
- [9] Aleksandr Efremov, Ahana Ghosh, and Adish Singla. Zero-shot learning of hint policy via reinforcement learning and program synthesis. In *EDM*, 2020.
- [10] Blockly Games. Games for tomorrow's programmers. <https://blockly.games/>.

- [11] Sumit Gulwani. Example-based learning in computer-aided STEM education. *Communications of the ACM*, 57(8):70–80, 2014.
- [12] Bilal Kartal, Nick Sohre, and Stephen J. Guy. Data driven Sokoban puzzle generation with Monte Carlo tree search. In *AIIDE*, 2016.
- [13] James C. King. Symbolic execution and program testing. *Communications of the ACM*, 19:385–394, 1976.
- [14] Levente Kocsis and Csaba Szepesvári. Bandit based Monte-Carlo planning. In *ECML*, pages 282–293, 2006.
- [15] Kasper Luckow, Corina S Păsăreanu, and Willem Visser. Monte Carlo tree search for finding costly paths in programs. In *SEFM*, pages 123–138, 2018.
- [16] John H. Maloney, Kylie Peppler, Yasmin Kafai, Mitchel Resnick, and Natalie Rusk. Programming by choice: Urban youth learning programming with Scratch. In *SIGCSE*, pages 367–371, 2008.
- [17] Samiha Marwan, Nicholas Lytle, Joseph Jay Williams, and Thomas W. Price. The impact of adding textual explanations to next-step hints in a novice programming environment. In *ITiCSE*, pages 520–526, 2019.
- [18] Benjamin Paaßen, Barbara Hammer, Thomas W. Price, Tiffany Barnes, Sebastian Gross, and Niels Pinkwart. The continuous hint factory – Providing hints in vast and sparsely populated edit distance spaces. *Journal of Educational Data Mining*, 2018.
- [19] Chris Piech, Jonathan Huang, Andy Nguyen, Mike Phulsuksombati, Mehran Sahami, and Leonidas J. Guibas. Learning program embeddings to propagate feedback on student code. In *ICML*, pages 1093–1102, 2015.
- [20] Chris Piech, Mehran Sahami, Jonathan Huang, and Leonidas J. Guibas. Autonomously generating hints by inferring problem solving policies. In *L@S*, pages 195–204, 2015.
- [21] CodeHS platform. CodeHS.com: Teaching Coding and Computer Science. <https://codehs.com/>.
- [22] CodeHS platform. Intro to Programming with Karel the Dog. <https://codehs.com/info/curriculum/introkarel>.
- [23] Code.org platform. Hour of Code: Classic Maze Challenge. <https://studio.code.org/s/hourofcode>.
- [24] Code.org platform. Hour of Code Initiative. <https://hourofcode.com/>.
- [25] Oleksandr Polozov, Eleanor O’Rourke, Adam M. Smith, Luke Zettlemoyer, Sumit Gulwani, and Zoran Popovic. Personalized mathematical word problem generation. In *IJCAI*, 2015.
- [26] Thomas W. Price and Tiffany Barnes. Position paper: Block-based programming should offer intelligent support for learners. In *B&B*, pages 65–68, 2017.
- [27] Thomas W. Price, Yihuan Dong, and Dragan Lipovac. iSnap: Towards intelligent tutoring in novice programming environments. In *SIGCSE*, pages 483–488, 2017.
- [28] Thomas W. Price, Rui Zhi, and Tiffany Barnes. Evaluation of a data-driven feedback algorithm for open-ended programming. *EDM*, 2017.
- [29] Mitchel Resnick, John Maloney, Andrés Monroy-Hernández, Natalie Rusk, Evelyn Eastmond, Karen Brennan, Amon Millner, Eric Rosenbaum, Jay Silver, Brian Silverman, et al. Scratch: Programming for all. *Communications of the ACM*, 52(11):60–67, 2009.
- [30] Reudismam Rolim, Gustavo Soares, Loris D’Antoni, Oleksandr Polozov, Sumit Gulwani, Rohit Gheyi, Ryo Suzuki, and Björn Hartmann. Learning syntactic program transformations from examples. In *ICSE*, pages 404–415, 2017.

- [31] Rishabh Singh, Sumit Gulwani, and Armando Solar-Lezama. Automated feedback generation for introductory programming assignments. In *PLDI*, pages 15–26, 2013.
- [32] Rohit Singh, Sumit Gulwani, and Sriram K. Rajamani. Automatically generating algebra problems. In *AAAI*, 2012.
- [33] Lisa Wang, Angela Sy, Larry Liu, and Chris Piech. Learning to represent student knowledge on programming exercises using deep learning. *EDM*, 2017.
- [34] David Weintrop and Uri Wilensky. Comparing block-based and text-based programming in high school computer science classrooms. *TOCE*, 18(1):1–25, 2017.
- [35] Mike Wu, Milan Mosse, Noah Goodman, and Chris Piech. Zero shot learning for code education: Rubric sampling with deep learning inference. In *AAAI*, 2019.
- [36] Jooyong Yi, Umair Z. Ahmed, Amey Karkare, Shin Hwei Tan, and Abhik Roychoudhury. A feasibility study of using automated program repair for introductory programming assignments. In *ESEC/FSE*, 2017.
- [37] Rui Zhi, Thomas W. Price, Samiha Marwan, Alexandra Milliken, Tiffany Barnes, and Min Chi. Exploring the impact of worked examples in a novice programming environment. In *SIGCSE*, pages 98–104, 2019.

A List of Appendices

In this section, we provide a brief description of the content provided in the appendices of the paper.

- Appendix B shows all the 10 reference tasks from Fig. 6. For each task, we also illustrate our methodology as was done in Fig. 1 and Fig. 2.
- Appendix C expands on the user study analysis. (Section 5)
- Appendix D provides additional details on the code mutation stage. (Section 3.1)
- Appendix E demonstrates how MCTS could guide the symbolic execution in generating more suitable tasks. (Section 3.2)
- Appendix F provides additional details and results about experiments. (Section 4)

B Illustration of Our Methodology for the HOC and Karel Dataset

In this section, we illustrate our methodology for all the 10 reference tasks from Fig. 6. For each reference task ($T^{\text{in}}, C^{\text{in}}$), we picked one output ($T^{\text{out}}, C^{\text{out}}$) from the entire set of generated outputs.




Figure 11: Task H1 – Illustration of our methodology.




Figure 12: Task H2 – Illustration of our methodology.




Figure 13: Task H3 – Illustration of our methodology.






Figure 14: Task H4 – Illustration of our methodology.




(a) Visual puzzle for T^{in} (b) Solution code C^{in} (c) Visual puzzle for T^{out} (d) Solution code C^{out}
 Figure 15: Task H5 – Illustration of our methodology (same as Fig. 1 and shown for completeness).




(a) Visual puzzle for T^{in} (b) Solution code C^{in} (c) Visual puzzle for T^{out} (d) Solution code C^{out}

Figure 16: Task H6 – Illustration of our methodology.




(a) Visual puzzle for T^{in} (b) Solution code C^{in} (c) Visual puzzle for T^{out} (d) Solution code C^{out}

Figure 17: Task K7 – Illustration of our methodology.




(a) Visual puzzle for T^{in} (b) Solution code C^{in} (c) Visual puzzle for T^{out} (d) Solution code C^{out}

Figure 18: Task K8 – Illustration of our methodology.



(a) Visual puzzle for T^{in} (b) Solution code C^{in} (c) Visual puzzle for T^{out} (d) Solution code C^{out}

Figure 19: Task K9 – Illustration of our methodology.



(a) Visual puzzle for T^{in} (b) Solution code C^{in} (c) Visual puzzle for T^{out} (d) Solution code C^{out}

Figure 20: Task K10 – Illustration of our methodology (same as Fig. 2 and shown for completeness).

C User Study: Additional Details and Results

In this section, we discuss additional details of our user study and expand on the key results provided in Section 5. For our user study, we used 4 reference tasks (H2, H4, H5, and H6). We developed an web app where participants, recruited through Amazon Mechanical Turk, were asked to solve tasks generated by four algorithms, SAME, TUTOR, MUTTASK, or SYNTASK (described in Section 5). Next, we describe the interface of our app and details of the questionnaire. The web app is publicly available (see Footnote 3).

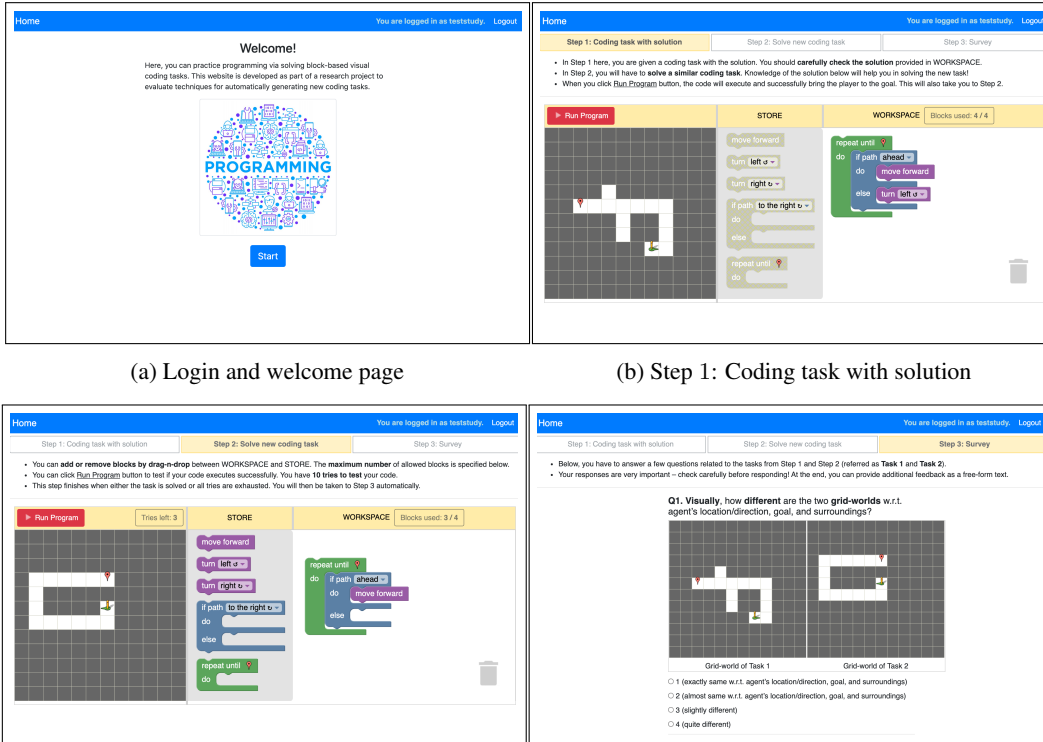


Figure 21: Interface of the app.

C.1 App Interface and Questionnaire

Our online app was developed using the publicly available toolkit of Blockly Games [10] and provides an interface for a participant to practice Block-based programming tasks for HOC. Participants were familiarized with the tasks by a small tutorial given before they logged-in to the app. Each participant was randomly assigned a reference task, an algorithm (out of the four chosen for evaluation), and a particular task generated based on the chosen algorithm. These elements constituted a “practice session” for a participant. Each session consisted of three steps. In Step 1, the reference task along with its solution code was shown to the participant (Fig. 21b). In Step 2, the participant was asked to solve a new task (Fig. 21c). The new task was generated by one of the four algorithms: SAME, TUTOR, MUTTASK, or SYNTASK. To solve this new task, the participant was given 10 tries. If they successfully solved the task or ran out of tries, they were directed to Step 3 of the practice session, the survey (Fig. 21d). Here, they were presented with a question on the visual dissimilarity of the tasks from Step 1 and Step 2, which was to be answered on a 4-point Likert scale as used in [25]:

- Score 1: indicates that the tasks are visually exactly same
- Score 2: indicates that the tasks are visually similar
- Score 3: indicates that the tasks are visually slightly different
- Score 4: indicates that the tasks are visually very different




Figure 22: Distribution of participant responses on visual dissimilarity of tasks, rated on a 4-point Likert scale (1: exactly same, 2: almost same, 3: slightly different, 4: very different)

C.2 Results on Visual Task Dissimilarity

The distribution of participant responses (on a 4-point scale) on the visual dissimilarity of the presented tasks is shown in Fig. 22; the aggregate statistics are provided in Fig. 9.

Comparing the dissimilarity of the generated tasks w.r.t. all the reference tasks H-, we found that the performance of SAME was worst (mean dissimilarity = 1.07) in comparison to MUTTASK (mean = 2.17), SYNTASK (mean = 2.63), and TUTOR (mean = 2.90). The poor performance of SAME, with the majority of scores being 1, is expected. Furthermore, we also see in Fig. 22a that MUTTASK has a greater percentage of lower scores of 1 and 2, and this is because the baseline directly mutates the task puzzle. Comparing the performance of SYNTASK to the baselines, we find that SYNTASK performed significantly better than MUTTASK ($\chi^2 = 38.81, p < e^{-7}$). However, its performance was worse than that of TUTOR and this difference was significant ($\chi^2 = 12.20, p = 0.0053$).

When analyzing the differences in the performance of SYNTASK w.r.t. TUTOR, we find that TUTOR performs better primarily because of the following two reasons: (i) some of the tasks generated by TUTOR have additional distracting paths / noise, and (ii) for simpler tasks without conditionals like H2, more powerful code mutations were used. Next, we performed additional analysis by limiting to two complex tasks with nested conditionals, H5 and H6. On these two tasks, the mean scores of the methods MUTTASK, SYNTASK, and TUTOR were 1.97, 2.94, and 3.03 respectively. Furthermore, we find that SYNTASK’s performance continued to be significantly better than MUTTASK ($\chi^2 = 64.33, p < e^{-13}$). But, the difference in the performance of SYNTASK and TUTOR was not statistically significant ($\chi^2 = 2.68, p = 0.44$).

In general, the performance of our task synthesis algorithm can be further improved by allowing for more powerful code mutations in tasks without conditionals (such as H2) and by adding more variability in the output tasks by incorporating techniques discussed in Appendix F.4.

D Code Mutation: Additional Details

This section describes the mutation stage of the task synthesis algorithm. In particular, we describe in detail the constraints applied on sketch $Q = \Omega(C)$. Note that we denote an empty action as ϕ . Our implementation of the code mutation stage, using the Z3 solver [7], is publicly available (see Footnote 2).

Constraint (Δ_1): ACTIONEDITS. ACTIONEDITS returns the values that all action sequences can take, based on their values in the reference code, C^{in} . Consider the action sequence \bar{A} in sketch Q . The function $\omega(\bar{A} | C^{in})$ returns the value of \bar{A} in C^{in} . The types of constraints returned by ACTIONEDITS are:

1. Local \bar{A} constraints. These constraints describe the values that one \bar{A} can take w.r.t. $\omega(\bar{A} | C^{in})$. It has the following rules:

- move action $\in \omega(\bar{A} | C^{in})$, would imply that the corresponding action in \bar{A} must be move.
- Set of turnLeft, turnRight actions $\in \omega(\bar{A} | C^{in})$, would imply that the corresponding actions in \bar{A} will either have all the ‘turn’ actions changed to turnLeft, or to turnRight, or remain the same, or be flipped, i.e. turnLeft \rightarrow turnRight and turnRight \rightarrow turnLeft.

- Set of `pickMarker`, `putMarker` actions $\in \omega(\bar{A} \mid C^{in})$, would imply that the corresponding actions in \bar{A} will either have all the ‘marker’ actions changed to `pickMarker`, or to `putMarker`, or remain the same, or be flipped i.e. `pickMarker` \rightarrow `putMarker` and `putMarker` \rightarrow `pickMarker`.
- Additional actions (up to δ_{size}) to \bar{A} can either be appended before the existing actions (in $\omega(\bar{A} \mid C^{in})$) or after, but not both. In our experiments we set $\delta_{size} = 2$.

These constraints are listed under “Local \bar{A} constraints” in Fig. 23h and Fig. 24h.

2. Global \bar{A} constraints. These constraints apply on all the \bar{A} ’s in sketch Q. They allow only one of all the \bar{A} ’s to have additional actions added to them. These constraints are listed under “Global \bar{A} constraints” in Fig. 23h and Fig. 24h.

Constraint (Δ_6): Action sequence is minimal. These constraints describe the sequences that invalidate minimality of code. The constraints ensure that certain sequences of actions do not occur in \bar{A} . The detailed list of sequences for the two example codes described, which invalidate \bar{A} if they occur, are given in Fig. 23g and Fig. 24g.

D.1 Details of Code Mutation for HOC

Here, we build on the generic DSL presented in Fig. 4 for HOC codes. We describe the HOC-DSL, Sketch DSL, and constraints for HOC codes (shown in Fig. 23a, Fig. 23b, and Fig. 23c, respectively). It is to be noted that the HOC DSL does not contain ‘marker’ based actions/conditionals: (action) `pickMarker`, (action) `putMarker`, (conditional) `marker`, (conditional) `noMarker`. It does not allow few more conditionals: `noPathA`, `noPathL` and `noPathR`. It also does not contain the `WHILE` construct. We consider a concrete example (continued from the example presented in Fig. 4d), in Fig. 23d, illustrate its sketch in Fig. 23e, and its constraints in Fig. 23f. We elaborate its \bar{A} minimality constraints and `ACTIONEDITS` in Fig. 23g and Fig. 23h, respectively.

A brief description of elimination sequences to ensure \bar{A} -minimality. We identify four sequences in HOC codes, which must be removed to ensure code minimality. These are: `turnLeft`, `turnRight` or `turnRight`, `turnLeft`, which do not lead to any change in the output; `turnLeft`, `turnLeft`, `turnLeft`, which can be replaced by a single `turnRight`; and finally `turnRight`, `turnRight`, `turnRight`, which can be replaced by a single `turnLeft`. It is to be noted that we also eliminate variants of the above sequences, that contain ϕ , but effectively reduce to these four sequences only, given that ϕ denotes an empty action.

D.2 Details of Code Mutation for Karel

Here, we present a modified/detailed form of Fig. 4 for Karel codes in particular. We describe the Karel-DSL, Sketch DSL, and constraints for Karel codes (shown in Fig. 24a, Fig. 24b, and Fig. 24c, respectively). It is to be noted that Karel-DSL does not contain the `REPEATUNTIL` construct. We consider a concrete example (same as the Karel solution code presented in Fig. 2b), in Fig. 24d, illustrate its sketch in Fig. 24e, and its constraints in Fig. 24f. We elaborate its \bar{A} minimality constraints and `ACTIONEDITS` in Fig. 24g and Fig. 24h, respectively.

A brief description of elimination sequences to ensure \bar{A} -minimality. In addition to the four sequences considered in HOC codes, we identify sixteen more sequences in Karel codes, which must be removed to ensure code minimality. These are: `pickMarker`, `putMarker` or `putMarker`, `pickMarker` or `pickMarker`, `pickMarker` or `putMarker`, `putMarker`, which do not lead to any change in the output or leads to a crash in the Karel grid (only one marker is allowed per grid-cell); `turnLeft`, `pickMarker`, `turnRight` or `turnRight`, `pickMarker`, `turnLeft` or `turnLeft`, `putMarker`, `turnRight` or `turnRight`, `putMarker`, `turnLeft`, which bring about the same output without the ‘turn’ actions; `pickMarker`, `turnLeft`, `putMarker` or `putMarker`, `turnLeft`, `pickMarker` or `pickMarker`, `turnRight`, `putMarker` or `putMarker`, `turnRight`, `pickMarker`, which bring about the same output without the ‘marker’ actions; and finally `pickMarker`, `turnLeft`, `pickMarker` or `pickMarker`, `turnRight`, `pickMarker` or `putMarker`, `turnLeft`, `putMarker` or `putMarker`, `turnRight`, `putMarker`, which leads to a crash in the Karel grid (only one marker is allowed per grid-cell). It is to be noted that we also eliminate variants of the above sequences, that contain ϕ , but effectively reduce to these basic sequences only, given that ϕ denotes an empty action.




Figure 23: Illustration of Code Mutation for HOC on the solution code for task *Maze 16* from the *Hour of Code: Classic Maze* challenge by *Code.org* [23]; We build on Fig. 4 here.

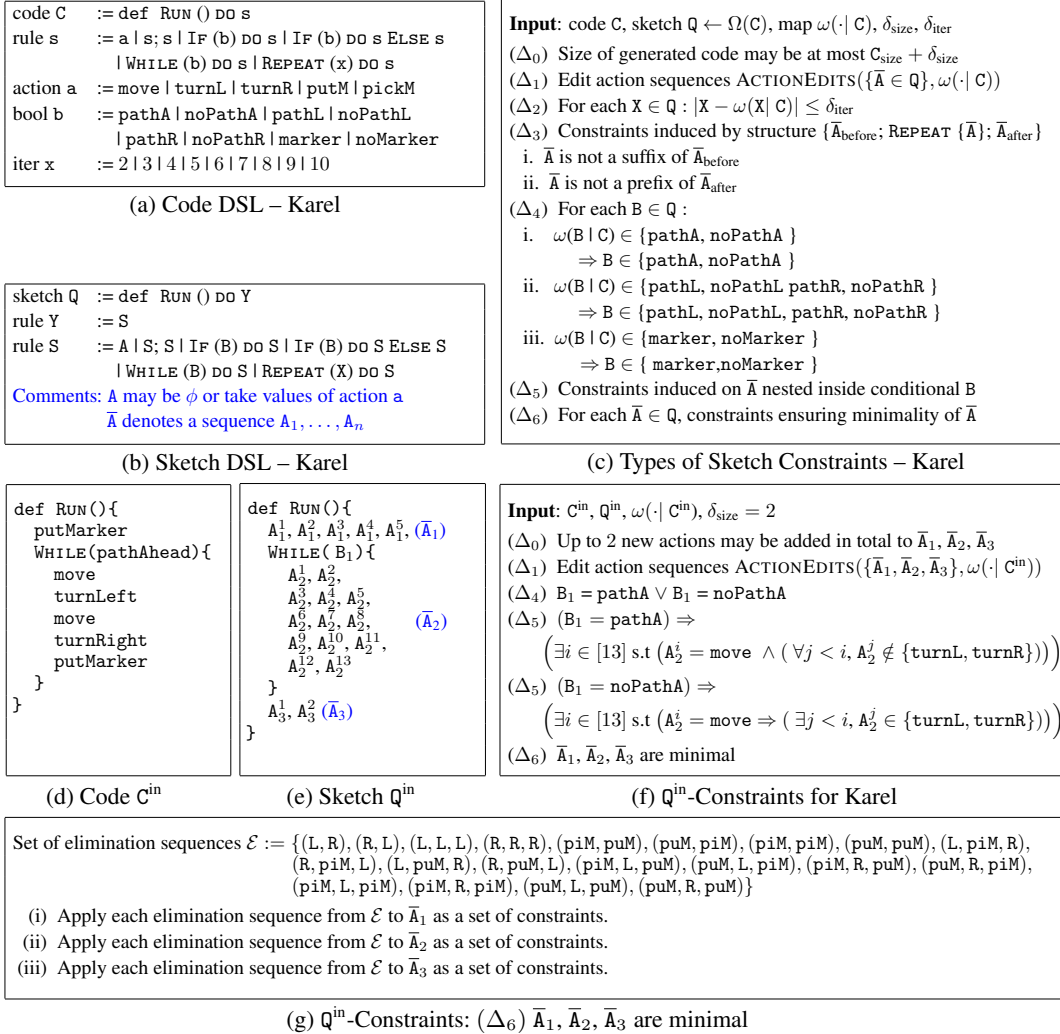


Figure 24: (a)–(g): Illustration of Code Mutation for Karel task *Diagonal* from the *Intro to Programming with Karel* course by *CodeHS.com* [22]; We present the Karel variant of Fig. 4 here.

Shorthand notation for actions: move \rightarrow M, turnL \rightarrow L, turnR \rightarrow R, pickM \rightarrow piM, putM \rightarrow puM, $\phi \rightarrow$ empty-action.

Local \bar{A} constraints that define the values that each action sequence can take:

- (i) $(A_1^1 = M \vee A_1^1 = L \vee A_1^1 = R \vee A_1^1 = \text{piM} \vee A_1^1 = \text{puM} \vee A_1^1 = \phi)$
 $\wedge(A_1^2 = M \vee A_1^2 = L \vee A_1^2 = R \vee A_1^2 = \text{piM} \vee A_1^2 = \text{puM} \vee A_1^2 = \phi)$
 $\wedge(A_1^3 = \text{piM} \vee A_1^3 = \text{puM})$
 $\wedge(A_1^4 = M \vee A_1^4 = L \vee A_1^4 = R \vee A_1^4 = \text{piM} \vee A_1^4 = \text{puM} \vee A_1^4 = \phi)$
 $\wedge(A_1^5 = M \vee A_1^5 = L \vee A_1^5 = R \vee A_1^5 = \text{piM} \vee A_1^5 = \text{puM} \vee A_1^5 = \phi)$
- (ii) $\wedge(A_1^1 \neq \phi \vee A_1^2 \neq \phi) \Rightarrow (A_1^4 = \phi \wedge A_1^5 = \phi) \wedge$
 $\wedge(A_1^4 \neq \phi \vee A_1^5 \neq \phi) \Rightarrow (A_1^1 = \phi \wedge A_1^2 = \phi)$
- (iii) $(A_2^1 = M \vee A_2^1 = L \vee A_2^1 = R \vee A_2^1 = \text{piM} \vee A_2^1 = \text{puM} \vee A_2^1 = \phi)$
 $\wedge(A_2^2 = M \vee A_2^2 = L \vee A_2^2 = R \vee A_2^2 = \text{piM} \vee A_2^2 = \text{puM} \vee A_2^2 = \phi)$
 $\wedge(A_2^3 = M)$
 $\wedge(A_2^4 = M \vee A_2^4 = L \vee A_2^4 = R \vee A_2^4 = \text{piM} \vee A_2^4 = \text{puM} \vee A_2^4 = \phi)$
 $\wedge((A_2^5 = L \wedge A_2^9 = L) \vee (A_2^5 = R \wedge A_2^9 = R) \vee (A_2^5 = L \wedge A_2^9 = R) \vee (A_2^5 = R \wedge A_2^9 = L))$
 $\wedge(A_2^6 = M \vee A_2^6 = L \vee A_2^6 = R \vee A_2^6 = \text{piM} \vee A_2^6 = \text{puM} \vee A_2^6 = \phi)$
 $\wedge(A_2^7 = M)$
 $\wedge(A_2^8 = M \vee A_2^8 = L \vee A_2^8 = R \vee A_2^8 = \text{piM} \vee A_2^8 = \text{puM} \vee A_2^8 = \phi)$
 $\wedge(A_2^{10} = M \vee A_2^{10} = L \vee A_2^{10} = R \vee A_2^{10} = \text{piM} \vee A_2^{10} = \text{puM} \vee A_2^{10} = \phi)$
 $\wedge(A_2^{11} = \text{piM} \vee A_2^{11} = \text{puM})$
 $\wedge(A_2^{12} = M \vee A_2^{12} = L \vee A_2^{12} = R \vee A_2^{12} = \text{piM} \vee A_2^{12} = \text{puM} \vee A_2^{12} = \phi)$
 $\wedge(A_2^{13} = M \vee A_2^{13} = L \vee A_2^{13} = R \vee A_2^{13} = \text{piM} \vee A_2^{13} = \text{puM} \vee A_2^{13} = \phi)$
- (iv) $((A_2^1 \neq \phi \vee A_2^2 \neq \phi) \Rightarrow (A_2^4 = \phi \wedge A_2^6 = \phi \wedge A_2^8 = \phi \wedge A_2^{10} = \phi \wedge A_2^{12} = \phi \wedge A_2^{13} = \phi))$
 $\wedge((A_2^4 \neq \phi) \Rightarrow (A_2^1 = \phi \wedge A_2^2 = \phi \wedge A_2^6 = \phi \wedge A_2^8 = \phi \wedge A_2^{10} = \phi \wedge A_2^{12} = \phi \wedge A_2^{13} = \phi))$
 $\wedge((A_2^6 \neq \phi) \Rightarrow (A_2^1 = \phi \wedge A_2^2 = \phi \wedge A_2^4 = \phi \wedge A_2^8 = \phi \wedge A_2^{10} = \phi \wedge A_2^{12} = \phi \wedge A_2^{13} = \phi))$
 $\wedge((A_2^8 \neq \phi) \Rightarrow (A_2^1 = \phi \wedge A_2^2 = \phi \wedge A_2^4 = \phi \wedge A_2^6 = \phi \wedge A_2^{10} = \phi \wedge A_2^{12} = \phi \wedge A_2^{13} = \phi))$
 $\wedge((A_2^{10} \neq \phi) \Rightarrow (A_2^1 = \phi \wedge A_2^2 = \phi \wedge A_2^4 = \phi \wedge A_2^6 = \phi \wedge A_2^8 = \phi \wedge A_2^{12} = \phi \wedge A_2^{13} = \phi))$
 $\wedge((A_2^{12} \neq \phi \vee A_2^{13} \neq \phi) \Rightarrow (A_2^1 = \phi \wedge A_2^2 = \phi \wedge A_2^4 = \phi \wedge A_2^6 = \phi \wedge A_2^8 = \phi \wedge A_2^{10} = \phi \wedge A_2^{13} = \phi))$
- (v) $(A_3^1 = M \vee A_3^1 = L \vee A_3^1 = R \vee A_3^1 = \text{piM} \vee A_3^1 = \text{puM} \vee A_3^1 = \phi)$
 $\wedge(A_3^2 = M \vee A_3^2 = L \vee A_3^2 = R \vee A_3^2 = \text{piM} \vee A_3^2 = \text{puM} \vee A_3^2 = \phi)$

Global \bar{A} constraints that allow actions to be added to either of $\bar{A}_1, \bar{A}_2, \bar{A}_3$:

- (i) $(A_1^1 \neq \phi \vee A_2^2 \neq \phi \vee A_3^3 \neq \phi \vee A_1^4 \neq \phi)$
 $\Rightarrow (A_2^1 = \phi \wedge A_2^2 = \phi \wedge A_2^3 = \phi \wedge A_2^4 = \phi \wedge A_2^5 = \phi \wedge A_2^6 = \phi \wedge A_2^7 = \phi \wedge A_2^8 = \phi \wedge A_2^9 = \phi \wedge A_2^{10} = \phi \wedge A_2^{11} = \phi \wedge A_2^{12} = \phi \wedge A_2^{13} = \phi \wedge A_3^1 = \phi \wedge A_3^2 = \phi)$
- (ii) $(A_2^1 \neq \phi \vee A_2^2 \neq \phi \vee A_2^4 \neq \phi \wedge A_2^6 \neq \phi \wedge A_2^8 \neq \phi \wedge A_2^{10} \neq \phi \wedge A_2^{12} \neq \phi \wedge A_2^{13} \neq \phi)$
 $\Rightarrow (A_1^1 = \phi \wedge A_1^2 = \phi \wedge A_1^3 = \phi \wedge A_1^4 = \phi \wedge A_1^5 = \phi \wedge A_1^6 = \phi \wedge A_1^7 = \phi \wedge A_1^8 = \phi \wedge A_1^9 = \phi \wedge A_1^{10} = \phi \wedge A_1^{11} = \phi \wedge A_1^{12} = \phi \wedge A_1^{13} = \phi)$
- (iii) $(A_3^1 \neq \phi \vee A_3^2 \neq \phi)$
 $\Rightarrow (A_1^1 = \phi \wedge A_1^2 = \phi \wedge A_1^3 = \phi \wedge A_1^4 = \phi \wedge A_1^5 = \phi \wedge A_1^6 = \phi \wedge A_1^7 = \phi \wedge A_1^8 = \phi \wedge A_1^9 = \phi \wedge A_1^{10} = \phi \wedge A_1^{11} = \phi \wedge A_1^{12} = \phi \wedge A_1^{13} = \phi)$

(h) Q^{in} -Constraints: (Δ_1) ACTIONEDITS($\{\bar{A}_1, \bar{A}_2, \bar{A}_3\}, \omega(\cdot | C^{\text{in}})$)

Figure 24: (h): Illustration of Code Mutation for Karel task *Diagonal* from the *Intro to Programming with Karel* course by *CodeHS.com* [22]; We present the Karel variant of Fig. 4 here.

E Symbolic Execution: Additional Details

In this section, we provide an example demonstrating how MCTS could guide the symbolic execution in generating more suitable tasks, see Fig. 25.

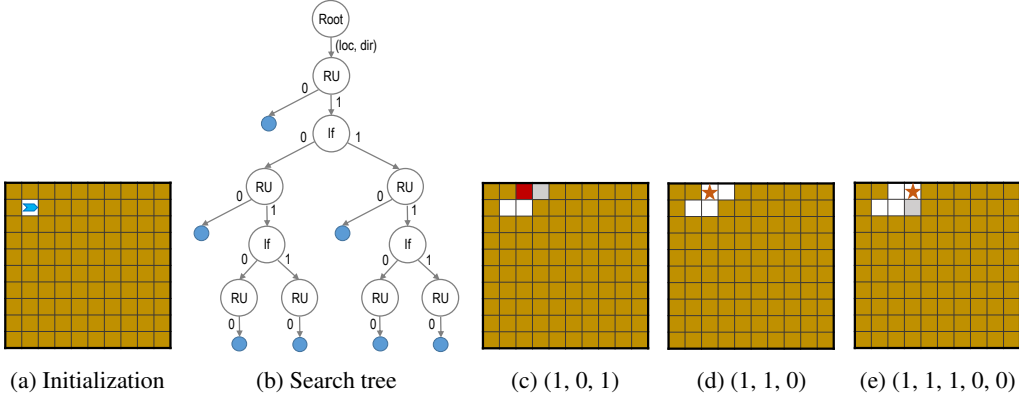


Figure 25: Search strategy for guiding the symbolic execution; see the text below for details.

In Fig. 25b, we consider the MCTS search tree for the code C^{out} from Fig. 1d obtained from the solution code of task H5 by mutation. At the top of the tree, we have a “Root” node where the initial location/direction of the agent is picked from among the available choices. Node “RU” corresponds to the block REPEATUNTIL; the child “1” from a node “RU” corresponds to unrolling of the loop (i.e., goal is false). Node “If” corresponds to the block IF; the child “1” from a node “If” corresponds to executing the code inside (i.e., pathR is true). For the purpose of this demonstration, we limited the unrolling of the loops to a maximum of 3; for our experiments, this depth is $2n(= 20)$ as discussed in Appendix F. Furthermore, in this demonstration, the initial location/direction for the agent are fixed as shown in Fig. 25a; for our experiments, we consider $5 \times 4(= 20)$ initial configurations which are picked by MCTS resulting in a large branching factor at the root of the tree.

Figs. 25c, 25d, and 25e illustrate the symbolic execution output for three different paths in the search tree as discussed below:

- Fig. 25c corresponds to the path (“RU”:1, “If”:0, “RU”: 1). This path results in agent crashing the wall. As discussed in Section 4, $\mathcal{F}_{\text{qual}}(\cdot)$ score is set to 0 when $\mathcal{F}_{\text{nocrash}}(\cdot) = 0$ and hence this tree path evaluates to low $\mathcal{F}_{\text{score}}(\cdot)$ score.
- Fig. 25d corresponds to the path (“RU”:1, “If”:1, “RU”: 0). This path results in two moves and two turns.
- Fig. 25e corresponds to the path (“RU”:1, “If”:1, “RU”: 1, “If”:0, “RU”: 0). This path results in three moves and two turns, and is the higher scoring path compared to other two paths.

MCTS is extremely effective as a search strategy and quickly learns to pick paths with high scores. For the same code C^{out} used in Fig. 25 above, Fig. 8a shows the temporal trends of different feature values in $\mathcal{F}_{\text{score}}$ averaged over a time window of 100 steps in our experiments. For further details, we refer the reader to [14] for an overview of the MCTS procedure, and to [15] where MCTS is used with symbolic execution to direct the exploration towards costly paths.

F Experimental Evaluation: Additional Details and Results

This section elaborates the experimental setup and results described in Section 4. We begin by providing details on the MCTS procedure in the symbolic execution stage of our task synthesis algorithm. Furthermore, we provide insights on generating multiple tasks for a single code. We also illustrate some example output tasks which violate criteria (V) and (VI), and mutated codes which got pruned in the symbolic execution stage (see Figure 7). Our implementation of the symbolic execution stage with MCTS procedure is publicly available (see Footnote 2).

F.1 Specification of MCTS for Single Run and Additional Results

In this section, we elaborate on the specification of MCTS, briefly discussed in Section 4. We begin by discussing choices that effect the scalability and run-time of MCTS procedure.

Choice of initial location and direction. When doing symbolic execution using MCTS, the procedure begins by picking an initial location and direction for the agent (see Fig. 5(b)). Given a grid-size n , and four initial directions (*north, east, south, west*) of the agent to choose from, we get a total of $4n^2$ choices for our initial configuration of the grid-puzzle. In the implementation used for generating the results, we restricted the set of initial choices to 20 (by choosing only 5 initial locations of the grid including four corners and centre, and 4 directions)—this aids in exploration by limiting the branching factor at the root of the MCTS’s tree.

Tree depth for symbolic execution. The depth of the symbolic tree depends on the nature of the corresponding solution code. For codes without the REPEATUNTIL or WHILE constructs (H1, H2, H3, K7, K8 and K9), the tree-depth is bounded. But, for more complex codes (H4, H5, H6, and K10), the tree depth is unbounded. In our implementation, we limited the unrolling of the loops to a maximum of $2n(= 20)$. To get an insight into the complexity of the problem, we note that with 20 initial choices and a depth of $2n$, there are over 40 million leaves in the symbolic tree for codes H5 and H6 which contain conditionals nested inside REPEATUNTIL or WHILE constructs.⁴ Next, we describe the details of our evaluation function for MCTS.

Details on the evaluation function $\mathcal{F}_{\text{score}}$. Our evaluation function $\mathcal{F}_{\text{score}}(\mathbf{T}^{\text{out}}, \mathbf{C}^{\text{out}}, \mathbf{T}^{\text{in}}, \mathbf{C}^{\text{in}}) \in [0, 1]$ measures the suitability of a generated task. A higher $\mathcal{F}_{\text{score}}$ indicates a more suitable task. We describe the elements of our evaluation function in greater detail here. We defined it in Section 4 and present it here again for completeness:

$$\mathcal{F}_{\text{score}}(\mathbf{T}^{\text{out}}, \mathbf{C}^{\text{out}}, \mathbf{T}^{\text{in}}, \mathbf{C}^{\text{in}}) = \underbrace{\mathbb{1}(\mathcal{F}_{\text{qual}}(\mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{C}^{\text{out}}) \geq \delta_{\text{qual}}, \mathcal{F}_{\text{nocrash}}(\mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{C}^{\text{out}}) = 1, \mathcal{F}_{\text{nocut}}(\mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{C}^{\text{out}}) = 1)}_{2a} \cdot \underbrace{[\alpha_1 \mathcal{F}_{\text{cov}}(\mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{C}^{\text{out}}) + \alpha_2 \mathcal{F}_{\text{qual}}(\mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{C}^{\text{out}}) + \alpha_3 \mathcal{F}_{\text{diss}}(\mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{T}_{\text{vis}}^{\text{in}})]}_{2b} \quad (2)$$

where $\mathbb{1}$ is an indicator function and each constant $\alpha = 1/3$. It is to be noted that component 2b in Eq.2 supplies the gradients for guiding the search in MCTS. At the end of the MCTS run (containing 2 million iterations), the *best* task (i.e, the one with the highest $\mathcal{F}_{\text{score}}$ value) is picked only from the pool of generated tasks which satisfy $\mathcal{F}_{\text{cov}}(\cdot) = 1, \mathcal{F}_{\text{score}}(\cdot) > 0$. We discuss each constituent function of $\mathcal{F}_{\text{score}}$ next.

Task quality component of evaluation function. $\mathcal{F}_{\text{qual}}(\mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{C}^{\text{out}}) \in [0, 1]$ evaluates the quality and validity of \mathbf{T}^{out} . Its is defined as a linear combination of the normalized counts of certain features of $\mathbf{T}_{\text{vis}}^{\text{out}}$ when \mathbf{C}^{out} is executed. As certain elements differ in the two task types, HOC and Karel, we define the features differently for each. More precisely, for HOC tasks, we have:

$$\mathcal{F}_{\text{qual}}^{\text{HOC}}(\mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{C}^{\text{out}}) = \frac{1}{4} \left(\frac{\#\text{moves}}{2n} + \frac{\#\text{turns}}{n} + \frac{\#\text{segments}}{n/2} + \frac{\#\text{long-segments}}{n/3} \right)$$

where the individual features are defined as

⁴The actual memory footprint of the MCTS procedure is small given that the symbolic tree is dynamically constructed during a run and the actual number of leaves explored is much lesser. In fact, the average runtime per output task (i.e., one MCTS run) as reported in Column 8 of Fig. 7 is achieved on a laptop machine with 2.8 GHz Quad-Core Intel Core i7 processor and 16 GB RAM.

- #moves: This refers to the count of ‘moves’.
- #turns: This refers to the count of ‘turns’.
- #segments: This refers to the number of consecutive sequence (≥ 3) of ‘moves’.
- #long-segments: This refers to the number of longer consecutive sequence (≥ 5) of ‘moves’.

For Karel tasks, we additionally have two marker based features to define the quality of the task i.e,

$$\mathcal{F}_{\text{qual}}^{\text{Karel}}(\mathbb{T}_{\text{vis}}^{\text{out}}, \mathbb{C}^{\text{out}}) = \frac{3}{4} \cdot \frac{1}{4} \left(\frac{\#\text{moves}}{2n} + \frac{\#\text{turns}}{n} + \frac{\#\text{segments}}{n/2} + \frac{\#\text{long-segments}}{n/3} \right) + \frac{1}{4} \cdot \frac{1}{2} \left(\frac{\#\text{pick-markers}}{n} + \frac{\#\text{put-markers}}{n} \right)$$

where the additional features are defined as

- #pick-markers: This refers to the count of ‘pick-marker’ activity.
- #put-markers: This refers to the count of ‘put-marker’ activity.

While we have used high-level features of $\mathbb{T}_{\text{vis}}^{\text{out}}$ to define the quality of a task, one could also embed more specific domain knowledge in defining these features to obtain more interesting/complex tasks.

Task dissimilarity component of the evaluation function. $\mathcal{F}_{\text{diss}}(\mathbb{T}_{\text{vis}}^{\text{out}}, \mathbb{T}_{\text{vis}}^{\text{in}}) \in [0, 1]$ evaluates the visual dissimilarity of $\mathbb{T}_{\text{vis}}^{\text{out}}$ w.r.t. $\mathbb{T}_{\text{vis}}^{\text{in}}$. We define it as a linear combination of the dissimilarity features as follows:

$$\mathcal{F}_{\text{diss}}(\mathbb{T}_{\text{vis}}^{\text{out}}, \mathbb{T}_{\text{vis}}^{\text{in}}) = \frac{1}{3} \left(\text{diss}(\text{loc} \mid \mathbb{T}_{\text{vis}}^{\text{out}}, \mathbb{T}_{\text{vis}}^{\text{in}}) + \text{diss}(\text{dir} \mid \mathbb{T}_{\text{vis}}^{\text{out}}, \mathbb{T}_{\text{vis}}^{\text{in}}) + \text{diss}(\text{grid-cells} \mid \mathbb{T}_{\text{vis}}^{\text{out}}, \mathbb{T}_{\text{vis}}^{\text{in}}) \right)$$

where the individual features are defined as

- $\text{diss}(\text{loc} \mid \mathbb{T}_{\text{vis}}^{\text{out}}, \mathbb{T}_{\text{vis}}^{\text{in}}) \in \{0, 1\}$ measures the dissimilarity in the agent’s initial location in the task-puzzles $\mathbb{T}_{\text{vis}}^{\text{out}}$ and $\mathbb{T}_{\text{vis}}^{\text{in}}$.
- $\text{diss}(\text{dir} \mid \mathbb{T}_{\text{vis}}^{\text{out}}, \mathbb{T}_{\text{vis}}^{\text{in}}) \in \{0, 1\}$ measures the dissimilarity in the agent’s initial direction in the task-puzzles $\mathbb{T}_{\text{vis}}^{\text{out}}$ and $\mathbb{T}_{\text{vis}}^{\text{in}}$.
- $\text{diss}(\text{grid-cells} \mid \mathbb{T}_{\text{vis}}^{\text{out}}, \mathbb{T}_{\text{vis}}^{\text{in}}) \in [0, 1]$ measures the grid-cell level dissimilarity in the task-puzzles $\mathbb{T}_{\text{vis}}^{\text{out}}$ and $\mathbb{T}_{\text{vis}}^{\text{in}}$. This is computed as the normalized Hamming distance w.r.t. the two grid-worlds (i.e., number of cells which are different, multiplied with a normalization factor of $\frac{2}{n^2}$).

Deep dive into an MCTS run for Karel. Analogous to the example provided in Section 4, we take a closer look at an MCTS run for the Karel task K10, shown in Fig. 26. Fig. 26a and Fig. 26b illustrate the improvement in various components of $\mathcal{F}_{\text{score}}$ as the number of MCTS iterations increases. Best tasks at different iterations are shown in Fig. 26c and Fig. 26d. As expected, the more the iterations, the better the tasks which are generated.

F.2 Specification of MCTS for Multiple Runs with Diversity and Additional Results

Our task synthesis algorithm can also generate multiple tasks for a *single* code, with sufficient diversity. To achieve this, we modify the evaluation function $\mathcal{F}_{\text{score}}$ guiding the MCTS search. We introduce a diversity measure $\mathcal{F}_{\text{diversity}}$ which measures the diversity between the generated tasks.

More concretely, when generating a new $(k + 1)^{\text{th}}$ task, we capture the diversity score w.r.t the tasks generated in the previous k-runs of MCTS $\{\mathbb{T}_{\text{vis},1}^{\text{out}}, \dots, \mathbb{T}_{\text{vis},k}^{\text{out}}\}$.

Modified evaluation function $\mathcal{F}_{\text{score}}$. Our evaluation function with the new diversity component is given below with each $\alpha = 1/4$. We have $\mathcal{F}_{\text{score}}(\mathbb{T}_{\text{vis}}^{\text{out}}, \mathbb{C}^{\text{out}}, \mathbb{T}_{\text{vis}}^{\text{in}}, \mathbb{C}^{\text{in}}, \{\mathbb{T}_{\text{vis},1}^{\text{out}}, \dots, \mathbb{T}_{\text{vis},k}^{\text{out}}\}) :=$

$$\underbrace{\mathbb{1}(\mathcal{F}_{\text{qual}}(\mathbb{T}_{\text{vis}}^{\text{out}}, \mathbb{C}^{\text{out}}) \geq \delta_{\text{qual}}, \mathcal{F}_{\text{nocrash}}(\mathbb{T}_{\text{vis}}^{\text{out}}, \mathbb{C}^{\text{out}}) = 1, \mathcal{F}_{\text{nocut}}(\mathbb{T}_{\text{vis}}^{\text{out}}, \mathbb{C}^{\text{out}}) = 1)}_{3a} \cdot \underbrace{\left[\alpha_1 \mathcal{F}_{\text{cov}}(\mathbb{T}_{\text{vis}}^{\text{out}}, \mathbb{C}^{\text{out}}) + \alpha_2 \mathcal{F}_{\text{qual}}(\mathbb{T}_{\text{vis}}^{\text{out}}, \mathbb{C}^{\text{out}}) + \alpha_3 \mathcal{F}_{\text{diss}}(\mathbb{T}_{\text{vis}}^{\text{out}}, \mathbb{T}_{\text{vis}}^{\text{in}}) + \alpha_4 \mathcal{F}_{\text{diversity}}(\mathbb{T}_{\text{vis}}^{\text{out}}, \mathbb{C}^{\text{out}} \mid \{\mathbb{T}_{\text{vis},1}^{\text{out}} \dots \mathbb{T}_{\text{vis},k}^{\text{out}}\}) \right]}_{3b} \quad (3)$$

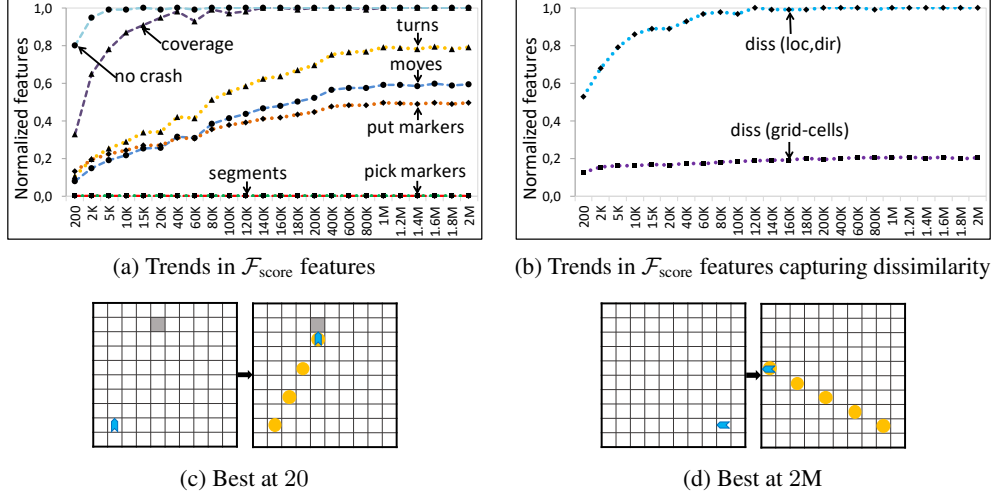


Figure 26: Illustration of a single MCTS run on C^{out} from Fig. 2d obtained from solution code of task K10 by mutation. (a, b) show the temporal trends of different feature values in $\mathcal{F}_{\text{score}}$ averaged over a time window of 100 steps. (c, d) show the best, i.e., highest scoring, tasks generated up to times 2×10^1 and 2×10^6 respectively.

Diversity score of tasks $\mathcal{F}_{\text{diversity}}$. Here, we describe the diversity component of the evaluation function. $\mathcal{F}_{\text{diversity}}(\mathbf{T}_{\text{vis}}^{\text{out}}, C^{\text{out}} | \{\mathbf{T}_{\text{vis},1}^{\text{out}}, \dots, \mathbf{T}_{\text{vis},k}^{\text{out}}\}) \in [0, 1]$ operates on a pool of generated tasks, and computes a diversity score for the new task w.r.t the tasks in the pool. Initially, the pool of tasks generated is empty. In this case, we have $\mathcal{F}_{\text{diversity}}(\mathbf{T}_{\text{vis}}^{\text{out}}, C^{\text{out}} | \{\}) = 1$.

After one run of MCTS, we have one task in the task pool $\{\mathbf{T}_{\text{vis},1}^{\text{out}}\}$. We define the diversity score for a subsequent task $\mathbf{T}_{\text{vis}}^{\text{out}}$ as follows. First, if $\mathbf{T}_{\text{vis}}^{\text{out}} = \mathbf{T}_{\text{vis},1}^{\text{out}}$ then we set $\mathcal{F}_{\text{diversity}}(\mathbf{T}_{\text{vis}}^{\text{out}}, C^{\text{out}} | \{\mathbf{T}_{\text{vis},1}^{\text{out}}\}) = 0$. Otherwise, the diversity score is given by

$$\mathcal{F}_{\text{diversity}}(\mathbf{T}_{\text{vis}}^{\text{out}}, C^{\text{out}} | \{\mathbf{T}_{\text{vis},1}^{\text{out}}\}) = \frac{1}{4} \left(\text{diss}(\text{loc} | \mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{T}_{\text{vis},1}^{\text{out}}) + \text{diss}(\text{dir} | \mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{T}_{\text{vis},1}^{\text{out}}) \right. \\ \left. + \text{diss}(\text{grid-cells} | \mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{T}_{\text{vis},1}^{\text{out}}) + \text{diss}(\text{symbolic-paths} | \mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{T}_{\text{vis},1}^{\text{out}}) \right)$$

where the individual features are defined as

- $\text{diss}(\text{loc} | \mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{T}_{\text{vis},1}^{\text{out}}) \in \{0, 1\}$ measures the dissimilarity in the agent’s initial location in the task-puzzles $\mathbf{T}_{\text{vis}}^{\text{out}}$ and $\mathbf{T}_{\text{vis},1}^{\text{out}}$.
- $\text{diss}(\text{dir} | \mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{T}_{\text{vis},1}^{\text{out}}) \in \{0, 1\}$ measures the dissimilarity in the agent’s initial direction in the task-puzzles $\mathbf{T}_{\text{vis}}^{\text{out}}$ and $\mathbf{T}_{\text{vis},1}^{\text{out}}$.
- $\text{diss}(\text{grid-cells} | \mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{T}_{\text{vis},1}^{\text{out}}) \in [0, 1]$ measures the grid-cell level dissimilarity in the task-puzzles $\mathbf{T}_{\text{vis}}^{\text{out}}$ and $\mathbf{T}_{\text{vis},1}^{\text{out}}$. This is computed as the normalized Hamming distance w.r.t. the two grid-worlds (i.e., number of cells which are different, multiplied with a normalization factor of $\frac{2}{n^2}$).
- $\text{diss}(\text{symbolic-paths} | \mathbf{T}_{\text{vis}}^{\text{out}}, \mathbf{T}_{\text{vis},1}^{\text{out}}) \in [0, 1]$ measures the dissimilarity in the symbolic-paths used in the generation of $\mathbf{T}_{\text{vis}}^{\text{out}}$ and $\mathbf{T}_{\text{vis},1}^{\text{out}}$. This is computed as the normalized “edit distance” between these paths.

After we have run k MCTS runs, the pool of tasks generated is populated ($\{\mathbf{T}_{\text{vis},1}^{\text{out}} \dots \mathbf{T}_{\text{vis},k}^{\text{out}}\}$) and the diversity score of the subsequent task is computed as follows:

$$\mathcal{F}_{\text{diversity}}(\mathbf{T}_{\text{vis}}^{\text{out}}, C^{\text{out}} | \{\mathbf{T}_{\text{vis},1}^{\text{out}}, \dots, \mathbf{T}_{\text{vis},k}^{\text{out}}\}) = \min_{i \in [k]} \mathcal{F}_{\text{diversity}}(\mathbf{T}_{\text{vis}}^{\text{out}}, C^{\text{out}} | \{\mathbf{T}_{\text{vis},i}^{\text{out}}\})$$

As the pool of tasks grow, it becomes more constrained to generate a task that is diverse from all the tasks in the pool. In general, there is a limit to the number of tasks per code that can be generated—eventually, the new task will have $\mathcal{F}_{\text{diversity}}$ to be 0 or $\mathcal{F}_{\text{score}}$ to be 0. As stated in Section 4, for each code C^{out} , we generated up to 10 different visual tasks using this process.

Illustration of output tasks using diversity. We illustrate our diverse task generation process on both HOC and Karel tasks. Fig. 27 shows the 10 diverse tasks generated for code shown in Fig. 1d, while Fig. 28 shows the 6 diverse tasks generated for the code shown in Fig. 2d.

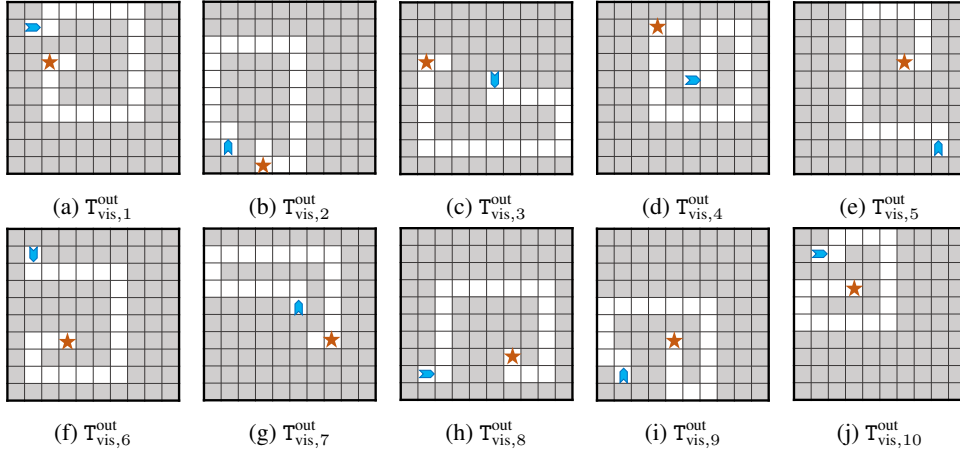


Figure 27: Illustration of task diversity on C^{out} from Fig. 1d which was obtained from solution code of task H5 by mutation. All the 10 diverse tasks generated are shown here.

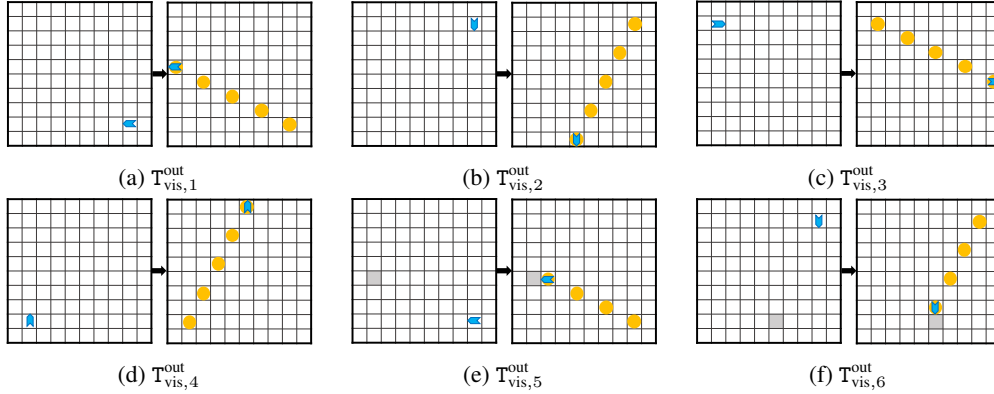


Figure 28: Illustration of task diversity on C^{out} from Fig. 2d which was obtained from solution code of task K10 by mutation. First 6 diverse tasks are shown here.

F.3 Insights into Results of Fig. 7

In this section, we provide further insights into the final results presented Fig. 7. In particular, we illustrate few limitations of the codes and tasks we generate (w.r.t task-synthesis objectives defined in Section 2):

- **Limitations of mutated codes generated:** Column 4 of Fig. 7 lists the number of codes generated by our mutation stage. However this set ($\#C_{\Delta=\text{all}}^{\text{out}}$) continues to have few semantic irregularities (as discussed in Section 3.1). Some of these irregularities are illustrated in Fig. 29. These codes are pruned by the symbolic execution stage of our algorithm, and the revised set of codes ($\#C^{\text{out}}$) is listed in Column 6 of Fig. 7.
- **Tasks which violate task-synthesis objective (V):** Out of the final set of tasks generated (shown in Column 7 of Fig. 7), some of them violate objective (V). The fraction of the tasks which satisfy this particular objective are listed in Column 9 of Fig. 7. In Fig. 30 we illustrate two examples which violate the objective, in output tasks for H5 and H6.
- **Tasks which violate task-synthesis objective (VI) $_{\delta_{\text{mini}}=1}$:** Column 10 of Fig. 7 lists the fraction of the tasks generated that satisfy this particular objective, on task-minimality. In Fig. 31, we illustrate two examples of tasks which violate this minimality constraint with $\delta_{\text{mini}} = 1$.

<pre>def RUN(){ move turnRight move turnRight move }</pre>	<pre>def RUN(){ turnLeft REPEAT(5){ move turnRight } }</pre>	<pre>def RUN(){ REPEAT(4){ move } turnLeft move turnLeft REPEAT(6){ move } }</pre>	<pre>def RUN(){ REPEATUNTIL(goal){ move turnLeft move turnLeft } }</pre>
(a) H1: C^{out}	(b) H2: C^{out}	(c) H3: C^{out}	(d) H4: C^{out}
<pre>def RUN(){ REPEATUNTIL(goal){ turnRight turnRight move IF(pathRight){ turnRight } } }</pre>	<pre>def RUN(){ REPEATUNTIL(goal){ IF(pathAhead){ move turnLeft } ELSE{ turnLeft } } }</pre>	<pre>def RUN(){ REPEAT(5){ pickMarker move turnRight putMarker } }</pre>	<pre>def RUN(){ putMarker WHILE(pathAhead){ move turnLeft move turnLeft putMarker } }</pre>
(e) H5: C^{out}	(f) H6: C^{out}	(g) K8: C^{out}	(h) K10: C^{out}

Figure 29: Illustration of codes with semantic irregularities in Column $C_{\Delta=all}^{out}$ of Fig. 7. (a)–(f) show the semantic irregularities in mutated codes generated for HOC tasks. All of these codes lead to circular paths in output task, and are pruned out in the symbolic execution stage, by the \mathcal{F}_{nocut} component of the \mathcal{F}_{score} measure. In particular, consider the semantic irregularity presented in (g). This code corresponding to the reference task K8, has redundant marker activity. With each iteration of REPEAT, the pickMarker and putMarker actions occur consecutively, leading to no change in marker activity in the output. (h) illustrates a mutated code for reference task K10 where if the WHILE is executed more than once, putMarker activity occurs in the same location consecutively leading to a crash (as only one marker is allowed per grid-cell); if WHILE is executed only once, the corresponding task generated has a short-cut to the goal.

	<pre>def RUN(){ REPEATUNTIL(goal){ move IF(pathRight){ turnRight turnRight } } }</pre>		<pre>def RUN(){ REPEATUNTIL(goal){ IF(pathAhead){ move move turnRight } ELSE{ turnRight } } }</pre>
(a) H5: T_{vis}^{out}	(b) H5: C^{out}	(c) H6: T_{vis}^{out}	(d) H6: C^{out}

Figure 30: Illustration of violation of task-synthesis objective (V) by generated output tasks for reference tasks H5 and H6. (a, b) show the irregularity in generated task (T_{vis}^{out}, C^{out}) for H5 which leads to only a straight path in the visual-puzzle. The corresponding code renders the IF construct redundant. (c, d) illustrate a similar irregularity in the generated task (T_{vis}^{out}, C^{out}) for H6. Here, IFELSE construct is not needed for the optimal solution.

	<pre>def RUN(){ move turnRight REPEATUNTIL(goal){ move IF(pathRight){ turnRight } } }</pre>		<pre>def RUN(){ turnRight move REPEATUNTIL(goal){ IF(pathAhead){ move } ELSE{ turnRight } } }</pre>
(a) H5: T_{vis}^{out}	(b) H5: C^{out}	(c) H6: T_{vis}^{out}	(d) H6: C^{out}

Figure 31: Illustration of violation of task-synthesis objective (V1), with $\delta_{mini} = 1$ by generated output tasks for reference tasks H5 and H6. (a, b) show the violation of task minimality, for the task (T_{vis}^{out}, C^{out}) generated for H5. In particular, the first two actions in the code shown in (b) are unnecessary. Similarly, (c, d) show the violation of task minimality, for the task (T_{vis}^{out}, C^{out}) generated for H6.

E.4 Adding More Variability to Output Tasks

Here, we propose a few simple extensions to our task generation algorithm allowing us to add more variability in the visual puzzles. We show that high variability can be achieved through suitable post-processing and pre-processing of the generated tasks (happening after and before the symbolic execution process, respectively). We describe each of these strategies next.

Different grid-sizes. A simple strategy to enhance task variability is by altering the grid-size parameter n . Fig. 32 illustrates the tasks generated with different values of the grid-size parameter n . For instance, this strategy is employed in generating multiple input-output pairs for Karel tasks in *Intro to Programming with Karel* course by *CodeHS.com* [22].




Figure 32: Tasks generated for C^{out} from Fig. 1d when varying grid-size (n).

Post-processing: Using distractor paths. One of the ways to add variability to the tasks is by adding *distractor-paths*, to the generated tasks. These paths are added after the symbolic execution stage is complete, and a basic task has been generated. Fig. 33a shows the task generated without any post-processing of the output of symbolic execution, carried out on C^{out} from Fig. 1d. Fig. 33b, and Fig. 33c illustrate two different post-processing patterns that yield tasks with greater variability, for the same code.

Pre-processing: Using different initializations of grid cells. We could also add task variability by initializing the grid-cells differently, before they are subjected to symbolic execution for task synthesis. We have a set of fixed grid-patterns, which when chosen as initializations of the grid-world yield very different looking output tasks. Fig. 34a shows the task generated for C^{out} from Fig. 1d without any pre-processing on the grid-cells. Fig. 34b and Fig. 34c show two different tasks obtained using different grid-initializations, for the same code.




Figure 33: Illustration of the (post-processing) distractor path strategy to increase task variability on tasks generated from C^{out} in Fig. 1d which was obtained from solution code of task H5 by mutation. (a) illustrates the basic task obtained after symbolic execution. (b, c) show two different distractor paths added after the symbolic execution stage, yielding visually very different tasks.

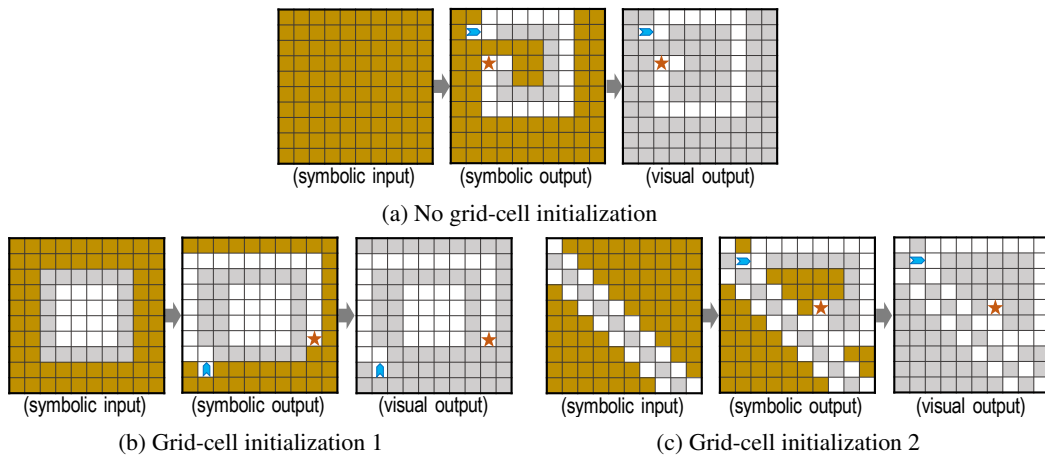


Figure 34: Illustration of the (pre-processing) grid-cell initialization strategy to increase task variability on tasks generated from C^{out} in Fig. 1d which was obtained from solution code of task H5 by mutation. (a) illustrates the basic task obtained after symbolic execution. (b, c) show two different grid-cell initializations, yielding visually very different tasks.