
Synthesis of Causal Reactive Programs with Structured Latent State

Ria A. Das
MIT CSAIL
riadas@mit.edu

Joshua B. Tenenbaum
MIT BCS
jbt@mit.edu

Armando Solar-Lezama
MIT CSAIL
asolar@csail.mit.edu

Zenna Tavares
Columbia University
zt2297@columbia.edu

Abstract

The human ability to efficiently discover causal theories of their environments from observations is a feat of nature that remains elusive in machines. In this work, we attempt to make progress on this frontier by formulating the challenge of causal mechanism discovery from observed data as one of *program synthesis*. We focus on the domain of time-varying, Atari-like 2D grid worlds, and represent causal models in this domain using a programming language called AUTUMN. Discovering the causal structure underlying a sequence of observations is equivalent to identifying the *program* in the AUTUMN language that generates the observations. We introduce a novel program synthesis algorithm, called AUTUMNSYNTH, that approaches this synthesis challenge by integrating standard methods of synthesizing functions with an *automata synthesis* approach, used to discover the model’s latent state. We evaluate our method on a suite of AUTUMN programs designed to express the richness of the domain, and our results signal the potential of our formulation.

1 Introduction

Children are intuitive scientists [4, 5, 11]. Without being told the rules, they can figure out how a new toy or video game works—a full causal theory of which stimuli cause which changes—after just minutes of observation. Such a data-efficient and flexible ability to discover causal models has yet to be demonstrated in an artificial agent. Though there has been progress on several related fronts, current approaches fall short of the goal of human-like causal model discovery from observations in three ways. First, the choice of model representation, e.g. causal graphical models, is often not expressive enough to concisely capture the complexity of real-world causal mechanisms. Second, even when the model representation is sufficiently rich, the model learning algorithm is often unrealistically data-hungry [9, 6, 13]. Third, many framings of the real-world model discovery problem assume more information in the input specification than just raw observations [8, 3].

In this paper, we seek to overcome some of these limitations by framing the problem of discovering a causal model from observed data as one of *program synthesis*. We represent a causal model as a program in a domain-specific language (DSL), and as such are able to exploit two key advantages of programs: their ability to express complex designs very concisely, and that they can often be synthesized from small input data. In this work, we focus specifically on the domain of time-varying mechanisms in 2D Atari-like grid worlds, and design a functional reactive language called AUTUMN to succinctly represent a wide variety of interesting causal phenomena within these worlds (Figure 1). The input to an AUTUMN program is a sequence of user events (either clicking anywhere on the grid, pressing an arrow key, or no event), one per discrete time step. The output is a corresponding




Figure 1: Sequence of grid frames from the Ice program. At times 1 and 4, the user presses down (red arrow), releasing a blue water particle from the gray cloud. The water moves down to the lowest possible height, moving to the side (time 10) if necessary to reach this height. The user presses down again at time 12, and then clicks anywhere (red circle) at time 15. The click causes the sun to change color and the water to turn to ice, which *stacks* rather than tries to reach the lowest height. A down press at time 19 releases another ice particle from the cloud. Finally, a click at time 24 changes the sun color back to yellow and turns the ice back to water, which again seeks the lowest possible height.

sequence of grid frames, each a partially observed representation of the underlying AUTUMN program state. Our goal is to build a program synthesis engine which, given the output sequence of observed frames and the associated user events, produces the best program in the AUTUMN language that generates the observations.

Critically, despite the natural fit of programs as a model representation for causal discovery, the AUTUMN domain does introduce a complication not handled by most existing program synthesis algorithms. To understand this aspect, we first note that, though the overall AUTUMN synthesis problem is to construct a program that takes a sequence of inputs to a sequence of outputs, the setting is also *reactive*: The output sequence is constructed *one time step at a time* from the user event at that time and the *program state* that summarizes what has happened in the program over all the previous time steps. Thus, the program that we wish to synthesize is a function that takes as input the current program state and user event, and produces as output a *modified* program state corresponding to the next time. However, while most methods for functional program synthesis from input-output data assume that both the inputs and outputs are fully observed [2], we never directly observe the full AUTUMN program state. Instead, we only have access to the grid frames, which are partial views of the underlying program states produced by a *rendering function*. Further, while there exist synthesis approaches that reconstruct hidden elements from partially observed inputs, our problem requires that we learn *how* these latent elements are modified over time, instead of just *what* they are at one time. Specifically, a full AUTUMN program state consists of a set of *objects*, each storing a 2D position and shape along with potentially some internal (latent) data fields, as well as global latent elements that are not tied to any particular object. The values of these time-varying latent variables may pivotally impact the rendered grid objects, but they are never directly exposed themselves. As such, part of the AUTUMN synthesis challenge becomes reconstructing the appropriate values and dynamics of the *latent state* in the generating program.

For concreteness, we give an example of this additional fold in our problem formulation via an AUTUMN program called “Mario” (Figure 2). In Mario, the agent (red) moves around with arrow key presses and can collect coins (yellow). If the agent has collected a positive number of coins, on a




Figure 2: Stills (ordered but not consecutive) from the Mario model. The red agent initially cannot shoot bullets (purple), but after jumping up and collecting coins (yellow), shoots on a user click. The agent can no longer shoot when all its bullets are used up. A bullet kills the blue enemy.

click event, a bullet (purple) is released upwards from the agent’s position, and the agent’s coin count is decremented. The number of coins that the agent possesses is not displayed anywhere on the grid at any time, so the only way to write an AUTUMN program that captures this behavior is to define an *internal* or *latent variable*, which tracks the number of coins (bullets) possessed by the agent. This involves both setting the variable’s initial value, as well as learning functions that dictate when (on what stimulus) and how (increment, decrement, etc.) that value will change. Notably, though a simple exercise for a human, this identification of latent state represents an elevation of the standard program synthesis paradigm to a new regime. While much work exists on learning latent state representations, including structured representations like those in AUTUMN [7, 10, 12], the integration of dynamical latent state discovery methods with standard methods of synthesizing functions is novel, as far as we know.

To solve the AUTUMN latent state learning problem, we first clarify that we can reasonably parse the visible elements of the program state—the object shapes and positions—directly from the grid frames, but cannot do the same for *invisible* elements, which include internal object-specific and global fields like `numCoins`. We address this challenge by making a key insight: The invisible state in AUTUMN programs can be viewed as finite state automata (FSMs), where the labels on the automata transitions are simply AUTUMN predicates (over the visible *and* invisible state). As such, we design our synthesis algorithm for AUTUMN programs, called AUTUMNSYNTH, to be a fusion of standard methods of functional synthesis with techniques drawn from the largely orthogonal area of *automata synthesis*. At a high level, AUTUMNSYNTH first attempts to synthesize small component functions in the full generating program without constructing new invisible state. Upon failure of this attempt, the method uses a novel automata synthesis algorithm to enrich the original program state with new, compact latent elements, which then enable the previous functional synthesis step to succeed.

We give a fuller description of the AUTUMNSYNTH algorithm, which includes important other steps related to entity extraction from scenes and object tracking, in Section 3. Prior to those details, we first provide more context about the AUTUMN language in Section 2. Finally, in Section 4, we describe the results of evaluating our synthesis algorithm on a benchmark dataset constructed to highlight the diversity of the AUTUMN domain. We note that, in the rest of the paper, we will use the term *latent state* to specifically refer to *invisible state*, though, technically, the full program state is latent since it is all viewed through a rendering function.

2 The AUTUMN Language

The AUTUMN language was designed to concisely express a rich variety of causal mechanisms in interactive 2D grid worlds. These mechanisms range from distillations of real-world, everyday causal phenomena, such as water interacting with a sink, plants growing upon exposure to sunlight, or an egg breaking upon being dropped, to video game-inspired domains such as Atari’s Space Invaders or Mario. In this section, we describe the key structure of an AUTUMN program that enables it to support this variety, as well as is necessary to understand our synthesis algorithm.

Every AUTUMN program is composed of four parts (Figure 3). The first part defines the grid dimensions and background color. The second part defines *object types*, which are simply structs which define an object *shape*, or a list of 2D positions each associated with a color, as well as a set of *internal fields*, which store additional information about the object (e.g. a Boolean `healthy` field may store an indicator of the object’s health). The third part defines *object instances*, which are

concrete instantiations of the object types defined previously, as well as *latent variables*, which are values with type `int`, `string`, or `bool`. Object instances and latent variables are defined using a primitive AUTUMN language construct called `initnext`, which defines a *stream* of values over time via the syntax `var = init expr1 next expr2`. The initial value of the variable (`expr1`) is set with `init`, and the value at later time steps is defined using `next`. The `next` expression (`expr2`) is re-evaluated at each subsequent time step to produce the new value of the variable at the present time. Further, the previous value of a variable may be accessed using the primitive `prev`, e.g. `prev var`. Indeed, the `next` expression frequently utilizes the `prev` primitive to express dependence on the past. For example, the definition of the Mario object in the example program from the introduction is `mario = init (Mario (Position 7 15)) next (moveDownNoCollision (prev mario))`, indicating that later values of Mario should move down one unit from the previous value whenever that is possible without collision.

Finally, the fourth segment of an AUTUMN program defines what we call *on-clauses*, which are expressed via the high-level form

```
on event
  update_function,
```

where `event` is a predicate (Boolean expression) and `update_function` is a variable update of the form `var = expr`, or multiple such updates. An *on-clause* represents an *override* of the default modification to a variable that is defined in the `next` clause. In particular, when the event predicate evaluates to true, the new value of the variable `var` at that specific time is computed by evaluating the associated `update_function` instead of the standard `next` expression. Each *on-clause* may contain multiple update statements for different variables, and a single program may contain multiple *on-clauses*. In the latter scenario, the *on-clauses* are evaluated sequentially, with the effect that later *on-clauses* may update a variable in a way that composes with updates from earlier *on-clauses*, or completely overrides it.

3 The AUTUMNSYNTH Algorithm

The AUTUMN synthesis problem that we focus on is one of passive discovery: given a sequence of images and associated user events, what is the underlying model of the world that generated this observed data? Synthesizing the correct AUTUMN program involves determining the latter three of the four program components described in the previous section: object types, object instance and latent variable definitions, and *on-clauses* (the first program component is known because the grid

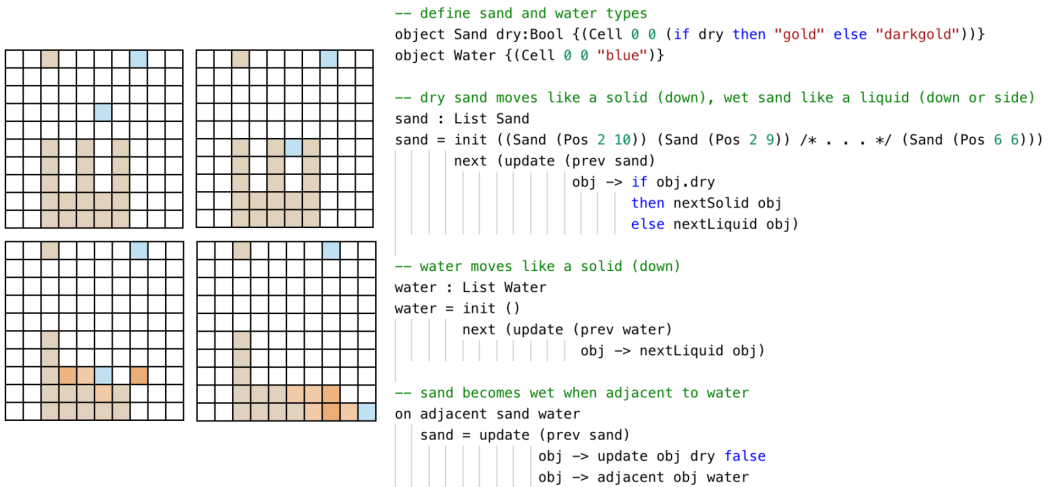


Figure 3: AUTUMN program describing dynamics of the sandcastle model. Note that we have omitted for clarity the definition of the two top buttons and the active button variable, which tracks which was the last button clicked, as well as the grid dimensions and background color.

dimensions are given and we currently evaluate on just white-background models). For simplicity, we elect to define each object in our synthesized program with the trivial `next` expression, `prev obj`, and instead express default object behavior using the “default” on-clause

```
on true
  update_function.
```

When this on-clause is first in the list of on-clauses, it acts as the default behavior that takes place when later on-clauses evaluate to false and hence do not override it. Using this equivalent form in our synthesized programs simplifies our discussion by allowing us to describe synthesis as just involving determining object types, initial object and latent values, and on-clauses.

The AUTUMNSYNTH algorithm, as an end-to-end synthesis algorithm taking images as input, consists of four distinct steps, each producing a new representation of the input sequence. These steps are

1. **perception**, in which objects are parsed from the observed grid frames;
2. **object tracking**, which involves assigning each object in a frame to either (1) an object in the subsequent frame, deemed to be its transformed image in the next time, or (2) no object, indicating that the object was removed in the next time;
3. **update function synthesis**, in which AUTUMN expressions, called update functions, describing each object-object mapping from Step 2 are determined; and
4. **event synthesis**, in which AUTUMN predicates associated with each update function from Step 3 are sought, and new latent state in the form of automata is constructed upon search failure.

We describe each of these steps below, with greatest space given to the fourth step of event and latent state synthesis, since that procedure represents the most novel aspect of our work.

3.1 Step 1: Perception

In the perception step, each frame in the observation sequence is parsed into a set of object variables. Each object variable is characterized by a *shape*, which is a list of 2D grid positions relative to (0, 0) that are each associated with a color, as well as an *origin*, which indicates the location of the object in the grid frame (the positions of the shape are translated by the origin to obtain the final rendering of each object). We use two different object parsing algorithms, each of which produces a different representation. We perform the rest of the synthesis procedure atop both of these parsing results one at a time, and take the output program from the first parsing using which the procedure succeeds.

The first parsing algorithm (“multi-cell”) is based on a breadth-first search pixel crawler, which identifies groups of adjacent cells with the same color as multi-celled objects. This approach currently supports only uniform-colored objects instead of individual objects composed of multiple colors. Extending this algorithm to handle more diverse object renderings is an area of future work. Object types are extracted from the union of the parsed object sets over all frames by finding shapes that contain the same 2D positions, though not necessarily the same colors. Shapes that support multiple colors are described by object types that have a custom field `(color, string)`, which allows individual instances of the object type to specify a particular color. The second parsing algorithm (“single-cell”) simply identifies each colored cell in a frame as an individual object, with the set of object types being the set of single-celled shapes each with a particular fixed color.

3.2 Steps 2 and 3: Object Tracking and Update Function Synthesis

Together, the second and third steps in the synthesis procedure answer the question, “What does each object do at each time step?” Concretely, this means identifying the *update function* undergone by each object in each frame to produce the object’s rendering in the subsequent frame. The first element of answering this question is *object tracking* (Step 2), which involves assigning each object in a frame either to (1) an object in the subsequent frame, which is considered to be the transformed image of the object after the time step, or (2) no object, which means that the object has been removed after the time step. Multiple objects may not map to the same object in the subsequent frame, and further, objects in the subsequent frame without a pre-image in the previous frame are deemed to have been just *added* to the program in the current time. The algorithm that performs this mapping is based upon a heuristic that embodies the following prior assumption about object motion in AUTUMN:

Objects are unlikely to move very far in a single time step. As such, the tracking algorithm performs assignments based on a proximity metric that tries to maximally assign objects in one frame to their closest objects (with the same object type) in the next frame.

Having determined which objects in a frame become which objects in the next, the *update function synthesis* procedure (Step 3) computes an AUTUMN expression, the update function, that describes every object-object, object-null, and null-object mapping. The null object simply represents a non-existent object, so an object-null mapping indicates object removal and a null-object mapping indicates object addition. To identify a matching update function, the procedure simply enumerates through a space of update function expressions, such as `obj = moveLeft obj` or `obj = nextLiquid obj`. Since it is often the case that there are multiple update functions that correctly describe a single mapping, we define a simple heuristic (based on how common an update function is across objects of a type) to select one update function from a set of correct options. At the end of the update function synthesis procedure, the synthesized update functions may be visualized in a matrix depiction, which we call the *update function matrix*. In the update function matrix, the rows represent `object_id`'s, where objects are assigned the same `object_id` if one is transformed into the other over time, and the columns display the update function undergone by each object at each time.

While the above procedure is effective in some scenarios, we have actually found that the ambiguity in deciding which update function should be chosen for each object-object mapping has made using a slightly *relaxed* version of our heuristic more successful. In this case, rather than producing a single update function matrix as output, the update function synthesis procedure actually produces a *list* of possible matrices, each with a different combination of update functions across cells. We perform the rest of the synthesis procedure with each matrix until the first successful output program is produced.

3.3 Step 4: Event Synthesis

By this stage in the synthesis process, the object types, the object instance definitions, and the update functions undergone by each object at every time have been identified. Remaining to be synthesized are the *event predicates* associated with the update functions in on-clauses, and potentially *latent variables* that are necessary for the appropriate events to exist. At a high level, the event synthesis step answers the simple question, “Why does each object do what it does in each time step?”

To synthesize events, we first define a finite set of AUTUMN predicates, which roughly embodies a prior about what types of events are likely to be triggers of changes in the grid world. We call these predicates *atomic events*, because we ultimately enumerate both through the events themselves as well as *conjunctions* and *disjunctions* of those atoms when searching for a matching event. The atomic event set includes *global events*, including user events like `clicked`, `clicked obj1`, and `leftPress` as well as object contact events like `intersects obj1 obj2` and `adjacent obj1 obj2`, among other forms. These stand in contrast to the other type of event in the atomic event set, called an *object-specific event*, which takes different values for *distinct object_id*'s in addition to distinct times. These events are effectively implemented as functions in a filter operation; for example, the event `obj.color == ‘red’` is true for an object if the object is contained in the filtered list




Figure 4: The latent state automaton for `numCoins` learned in the Mario example. Note that the learned automaton is the finite version of a more general Mario program, in which Mario can collect an uncountably infinite number of coins. Since we feed our synthesis procedure a finite trace, we cannot expect to learn the infinite automaton directly. In future work, however, we aim to capture this generalization via an *abstraction* step.

```
filter (obj -> (obj.color == 'red')) objects,
```

where `objects` denotes the set of all objects at the current time. We note that while the evaluation of a global event over time consists of a single vector of true/false values (one per time), the full evaluation of an object-specific event consists of a set of such vectors, one per distinct `object_id`.

Next, we describe the set of update functions for which we must find associated events. In our setting, we make the assumption that objects that belong to the same object type are all controlled by the same set of on-clauses. This means that if two objects both undergo the update `moveLeft` and the objects have the same object type, then a single event (on-clause) caused both of them to undergo the update. In contrast, if two objects undergo `moveLeft` and belong to different object types, we must synthesize a different event associated with each one, since a different on-clause caused each object type’s update. Thus, we synthesize events by enumerating through the object types, and finding an event for each distinct update function that appears across objects of that type.

Lastly, for each update function under consideration, we construct what is called an *update function trajectory*, which is a set of vectors $v \in \{-1, 0, 1\}^T$ that describes the times when the update function took place versus did not take place (T is the length of the observation sequence). There is one vector for each `object_id` with the object type under consideration. Each vector position is 1 if the update function took place at that time for that `object_id`, 0 if it did not take place, and -1 if it *may* have taken place but could have been *overridden* by another update function. This third scenario is interesting, and arises because we structure synthesized AUTUMN programs so on-clauses with update functions that are more frequent in the observed sequence are ordered before on-clauses with less frequent update functions. Thus, those later on-clauses will always override the earlier ones. With respect to event search, an event is a match for an update function if it is true for every time and `object_id` for which the update function trajectory vector is 1, and false whenever it is 0. The event may be either true or false when the corresponding update function trajectory value is -1 .

Notably, if the number of unique vectors in an update function trajectory is 1, then the matching event may be a global event, because there is no variance based on object-specific features. Otherwise, if there is more than one unique vector in the trajectory, then the matching event must be an object-specific event, since the evaluated vector depends on the particular `object_id`. It is possible that a matching event may not be found in either of these cases, which signals that we must enrich the program state with new elements that were not used in the original event space. For simplicity, in the following section about latent state synthesis, we focus only on the case where the unmatched update function trajectory contains a single unique vector. This setting is called *global latent state synthesis*; the alternative setting, called *object-specific latent state synthesis*, is a straightforward extension.

3.3.1 Discovery of Latent State: Automata Synthesis

The input to the automata synthesis step is an update function trajectory composed of a single vector $v \in \{-1, 0, 1\}^T$. The goal of the procedure is to design the simplest latent state automaton that enables us to write a latent-state-based event predicate that matches v .

This problem statement is best illustrated with an example. Consider the Mario program from the introduction. As discussed, the bullet addition function does not have an event that perfectly matches its trajectory of occurrence versus non-occurrence (update function trajectory). Indeed, the event `clicked` co-occurs with every occurrence of a bullet addition, but there are times when `clicked` is true but no bullet addition takes place. These are the times when the number of coins possessed by Mario is zero, but there is no way to express this event using the existing program state representation, in which the only variables are the object variables. The latent state automaton described below alleviates this expressiveness limitation:

```
% initialize latent variable
numCoins : Int
numCoins = init 0 next (prev numCoins)
% define dynamics of variable
on intersects (prev mario) (prev coins)
    numCoins = (prev numCoins) + 1
on clicked && ((prev numCoins) > 0)
    numCoins = (prev numCoins) - 1
```


The automaton diagram associated with this AUTUMN description is shown in Figure 4. With this new latent variable `numCoins`, the event

```
clicked && ((prev numCoins) > 0)
```

suddenly is a perfect match for the update function trajectory vector v for bullet addition, so that we can write the following on-clause:

```
on clicked && ((prev numCoins) > 0)
    bullets = addObj (prev bullets) (Bullet (prev mario).origin).
```

Having developed a concrete example, we return to the challenging part of the problem, which is actually synthesizing the `numCoins` latent variable (along with its associated on-clauses) from the observed data alone. We do this by first framing our problem with respect to the classic formulation of automata synthesis given input-output examples. Classically, the problem of inductive automata synthesis is to determine the minimum-state automaton that accepts a given set of accepted input strings (positive examples) and rejects a given set of rejected input strings (negative examples) [1]. In our scenario, these positive and negative input “strings” may be determined from the sequence of program states (one per time) corresponding to the observation sequence. In particular, we consider the set of *prefixes* (sub-arrays starting from the first position) of the program state sequence that have, as their last element, a program state where the *optimal co-occurring event* is true. The optimal co-occurring event is defined to be the event that co-occurs with the update function in question, and has the minimum number of false positive times, i.e. times when the event is true but the update function does not occur. In the Mario example, this co-occurring event is `clicked`. We then partition the set of program state sequence prefixes into those that end with a program state in which the update function took place and those in which it did not take place. The former set is the set of positive examples and the latter is the set of negative examples in our automata synthesis problem.

This definition of positive and negative input strings may be understood by considering the fact that, if there existed a latent state automaton that fit this specification, then the event

```
co_occurring_event && (latent_var in [/* accepting state labels */])
```

would be a perfect match for the update function. This is because the co-occurring event is true during a set of false positive times with respect to the update function trajectory, and the latent automaton is in rejecting states at exactly those times (since those times correspond to the rejected program state prefixes). Thus, finding such an automaton would mean we would have an event that matches the update function under consideration.

To synthesize automata, we first recognize that a single automaton may provide the latent structure necessary to define the matching predicate for *multiple* distinct update functions, instead of just one. We thus employ a heuristic algorithm that first *groups* update functions into sets for which one automaton will be sought. The algorithm then starts with a small automaton (i.e. small number of states), and tries to determine *transition events* that result in the appropriate program state sequences being accepted and rejected, for each update function in the group. If no such transition event can be found that perfectly partitions the input sets into positive and negative, then the number of states in the automaton is increased by *splitting* one of the existing states into two states. This process of searching for transition events and dividing a state into two on failure is repeated until success.

4 Experiments

We curated a benchmark suite of 26 AUTUMN programs to evaluate our synthesis algorithm, and stills from a subset of these programs are displayed in Figure 6 in the Appendix. These programs span three general domain categories—physics-inspired, video-game-inspired, and miscellaneous—as well as a range of difficulty levels. We summarize statistics about the synthesis procedure for each of these models in Table 1 in the Appendix. In addition, we look in depth at the synthesized output for one program, Water Plug, below.

Prior to these results, we note some additional details about our evaluation procedure. As this work is still ongoing, we currently declare a successful outcome for the synthesis procedure if it produces an output program that is *consistent* with the given observation sequence. In other words, we declare

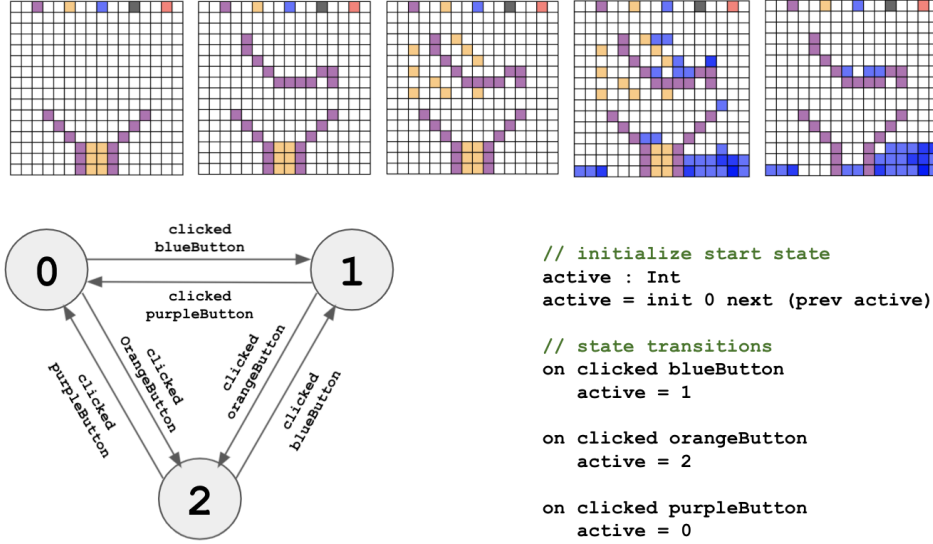


Figure 5: The latent state automaton learned for the Water Plug program. The last clicked of the top-row purple, orange, and blue buttons dictates which colored cell is added upon clicking free positions. *Top*: A sequence of frames (with time jumps) from the model. *Left*: Diagram of the learned state machine. *Right*: The AUTUMN description of the structure (we renamed the raw synthesized output variable names with more meaningful names for simplicity of exposition).

success even if the synthesized program does not exactly match the ground-truth generating program. In practice, we find that many of the synthesized programs are *almost* exactly the ground-truth program, except in cases of ambiguity in the observation sequence. For example, it is often the case that multiple events are a match for triggering an update function, such that an output program using any of them would technically match the observations. The way to disambiguate between multiple matches is to check the hypothesis program’s evaluation against additional, independent input sequences, or against a forecast of the existing input sequence into the future. These strategies remain part of our future work. Using our definition of a solution, our algorithm synthesizes 24 out of the 26 models.

In addition, the in-progress nature of our submission manifests in the fact that our evaluations across different models are not yet fully standardized. Specifically, we allow modification of the event and update function search spaces between models to try and minimize scenarios of ambiguity, along with other low-level modifications to the algorithm. Standardization of these variations are currently in progress and a priority for future work; see the Appendix for further evaluation details. Finally, while we take care to emphasize to the reader that our results are still preliminary and should not be the basis of overzealous speculation, the fact that the complex dynamics, and especially latent state automata, that underlie these models is being captured in some form by our procedure is exciting.

4.1 Water Plug

In Water Plug (Figure 5), clicking on an empty (uncolored) cell adds a single-celled object to that position that is either purple, orange, or blue. The color of the added square depends on which of the purple, orange, or blue buttons in the top row was last clicked. Added purple and orange squares maintain their positions, but blue squares move according to the `nextLiquid` library function. Clicking the top black button removes all the non-button orange objects from the grid, and clicking the top red button removes all the non-button objects from the grid. The input to synthesis for the Water Plug example is a sequence of 42 grid frames and user events. The synthesized output is 50 lines long, with 10 on-clauses, some of which represent state transitions in the learned automaton diagrammed in Figure 5. The synthesized source code may be found in the Appendix.

References

- [1] Dana Angluin. Learning regular sets from queries and counterexamples. *Inf. Comput.*, 75(2):87–106, November 1987.
- [2] Kevin Ellis, Maxwell I. Nye, Yewen Pu, Felix Sosa, Josh Tenenbaum, and Armando Solar-Lezama. Write, execute, assess: Program synthesis with a REPL. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 9165–9174, 2019.
- [3] Bernd Finkbeiner, Felix Klein, Ruzica Piskac, and Mark Santolucito. Synthesizing functional reactive programs. In *Proceedings of the 12th ACM SIGPLAN International Symposium on Haskell*, Haskell 2019, page 162–175, New York, NY, USA, 2019. Association for Computing Machinery.
- [4] Alison Gopnik. Scientific thinking in young children: Theoretical advances, empirical research, and policy implications. *Science*, 337(6102):1623–1627, 2012.
- [5] Alison Gopnik and Laura Schulz. Mechanisms of theory formation in young children. *Trends in cognitive sciences*, 8(8):371–377, 2004.
- [6] Mohammadhosein Hasanbeig, Natasha Yogananda Jeppu, Alessandro Abate, Tom Melham, and Daniel Kroening. DeepSynth: Automata Synthesis for Automatic Task Segmentation in Deep Reinforcement Learning, 2021.
- [7] Brenden M Lake, Ruslan Salakhutdinov, and Joshua B Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.
- [8] Parthasarathy Madhusudan. Synthesizing reactive programs. In Marc Bezem, editor, *Computer Science Logic, 25th International Workshop / 20th Annual Conference of the EACSL, CSL 2011, September 12-15, 2011, Bergen, Norway, Proceedings*, volume 12 of *LIPICs*, pages 428–442. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2011.
- [9] Gary Marcus. Deep learning: A critical appraisal, 2018.
- [10] Feras A Saad, Marco F Cusumano-Towner, Ulrich Schaechtle, Martin C Rinard, and Vikash K Mansinghka. Bayesian synthesis of probabilistic programs for automatic data modeling. *Proceedings of the ACM on Programming Languages*, 3(POPL):1–32, 2019.
- [11] Laura Schulz. The origins of inquiry: Inductive inference and exploration in early childhood. *Trends in cognitive sciences*, 16(7):382–389, 2012.
- [12] Joshua B Tenenbaum, Charles Kemp, Thomas L Griffiths, and Noah D Goodman. How to grow a mind: Statistics, structure, and abstraction. *science*, 331(6022):1279–1285, 2011.
- [13] Rodrigo Toro Icarte, Ethan Waldie, Toryn Klassen, Rick Valenzano, Margarita Castro, and Sheila McIlraith. Learning reward machines for partially observable reinforcement learning. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.

Appendix

A Full Results

Model Name	Input Length (Frames)	On-Clause Count	Contains Latent State?	Synthesized?
PARTICLES	22	2	No	Yes
ANTS	24	3	No	Yes
CHASE	42	7	No	Yes
ICE	27	10	No	Yes
LIGHTS	24	2	No	Yes
MAGNETS	53	7	No	Yes
SPACE INVADERS	42	11	No	Yes
SOKOBAN	25	7	No	Yes
DISEASE	22	7	Yes	Yes
GROW	40	-	Yes	No
SANDCASTLE	32	-	Yes	No
BULLETS	54	18	Yes	Yes
GRAVITY I	19	9	Yes	Yes
GRAVITY II	24	14	Yes	Yes
GRAVITY III	27	32	Yes	Yes
GRAVITY IV	48	18	Yes	Yes
COUNT I	22	6	Yes	Yes
COUNT II	39	10	Yes	Yes
COUNT III	69	14	Yes	Yes
COUNT IV	109	18	Yes	Yes
DOUBLE COUNT I	156	10	Yes	Yes
DOUBLE COUNT II	94	18	Yes	Yes
WIND	21	9	Yes	Yes
PAINT	27	10	Yes	Yes
WATER PLUG	42	10	Yes	Yes
MARIO	81	19	Yes	Yes

The failure of the synthesis procedure on SANDCASTLE and GROW is a result of a failure of the heuristic used for automata synthesis. Devising improved heuristic algorithms for automata synthesis is an ongoing area of future work. We also note that we selected input observation sequences for each evaluation model manually, in a way that we knew was compatible with our heuristic-based automata synthesis algorithm. We have not yet fully studied the degree of manual input curation necessary for our synthesis procedure to succeed, though this (along with algorithm modifications to alleviate this limitation) remains an important priority for future work. Lastly, we emphasize that the purpose of the benchmark suite evaluation is to give some early sense of the capabilities of this (still in progress)




Figure 6: Stills from a selection of AUTUMN programs in the evaluation suite. From top-left to bottom-right: water interacting with a sink, a clone of Space Invaders, plants growing under sunlight and water, magnets, ants foraging for food, water destroying a sandcastle, an alternative gravity simulation, and snow falling in windy conditions.

algorithmic approach, rather than a full picture of its limitations. We intend to evaluate the latter using externally sourced benchmarks, including the grid-based games of the GVGAI competition, and use those results to guide further development of the algorithm before the conference submission for this project.

B Additional Synthesis Details

We mention two details omitted due to space constraints in the main paper: Instead of completely raw images, we currently send a slightly more processed version of the images to the synthesis engine as input, namely a list of 2D pixel positions with colors. The significance of this representation is that, if two objects overlap at one pixel, the synthesizer does not need to figure out from that pixel’s color and transparency value (all Autumn renderings are partially transparent) that there are really two overlapping colors there. Instead, the input will already include two elements with the same x-y coordinates and color, e.g. $\{(x, y, color), (x, y, color)\}$. This detangling of pixels with overlap into their individual components can be trivially performed by storing a mapping between all RGBA values formed via overlaps of a finite number of colors, and the lists of colors that compose them. We will implement this procedure in the final version of the algorithm.

Lastly, since our enumeration procedure through the event predicate space is currently just blind enumeration, we only compute conjunctions and disjunctions of these atoms up to a depth of 2 atoms for performance reasons. Consequently, we occasionally seed the event space with events that themselves have contain conjunctions and disjunctions, so that combinations with greater depth may be effectively constructed with just depth-2 enumeration. This is necessary for models where the events have more than 2 atoms (i.e. this is part of the search space modifications discussed in Section 4). We intend to use more sophisticated enumeration procedures, e.g. SyGuS solvers, to overcome these performance bottlenecks in the final version of the algorithm.

C Sample Synthesized Programs

Below are raw outputs from the synthesis procedure (hence their lack of formatting/readability), with the exception of the all-capitals comments in `/* */` that were human-annotated. Note that the latent variables are all of type `int`; in the main paper, we implicitly renamed these values with strings or permuted some state labels (with no change to the underlying automaton dynamics) for human readability purposes.

C.1 Water Plug Output

```
1 (program
2   (= GRID_SIZE 16)
3   (= background "white")
4   /* OBJECT TYPE DEFINITIONS */
5   (object ObjType1 (: field1 Int) (list (Cell 0 0 "purple" )))
6   (object ObjType2 (: field1 Int) (list (Cell 0 0 "orange" )))
7   (object ObjType3 (: field1 Int) (list (Cell 0 0 "blue" )))
8   (object ObjType4 (list (Cell 0 0 "black" )))
9   (object ObjType5 (list (Cell 0 0 "red" )))
10
11  /* OBJECT INSTANCE DEFINITIONS */
12  (: addedObjType1List (List ObjType1))
13  (: addedObjType2List (List ObjType2))
14  (: addedObjType3List (List ObjType3))
15  (: obj22 ObjType4)
16  (: obj23 ObjType5)
17
18  (: addedObjType4List (List ObjType4))
19  (: addedObjType5List (List ObjType5))
20
21  (= addedObjType1List (initnext (list (ObjType1 1 (Position 6 15)) (
    ObjType1 1 (Position 6 14)) (ObjType1 1 (Position 6 13)) (ObjType1
    1 (Position 5 12)) (ObjType1 1 (Position 4 11)) (ObjType1 1 (
    Position 3 10)) (ObjType1 1 (Position 9 15)) (ObjType1 1 (Position
    9 14)) (ObjType1 1 (Position 9 13)) (ObjType1 1 (Position 10 12))
    (ObjType1 1 (Position 11 11)) (ObjType1 1 (Position 12 10)) (
    ObjType1 2 (Position 2 0))) (prev addedObjType1List)))
22  (= addedObjType2List (initnext (list (ObjType2 1 (Position 7 15)) (
    ObjType2 1 (Position 8 15)) (ObjType2 1 (Position 7 14)) (ObjType2
    1 (Position 8 14)) (ObjType2 1 (Position 7 13)) (ObjType2 1 (
    Position 8 13)) (ObjType2 2 (Position 5 0))) (prev
    addedObjType2List)))
23  (= addedObjType3List (initnext (list (ObjType3 0 (Position 8 0))) (
    prev addedObjType3List)))
24  (= obj22 (initnext (ObjType4 (Position 11 0)) (prev obj22)))
25  (= obj23 (initnext (ObjType5 (Position 14 0)) (prev obj23)))
26
27  (= addedObjType4List (initnext (list) (prev addedObjType4List)))
28  (= addedObjType5List (initnext (list) (prev addedObjType5List)))
29
30  /* LATENT VARIABLE (ACTIVE BUTTON INDICATOR) */
31  (: globalVar1 Int)
32  (= globalVar1 (initnext 1 (prev globalVar1)))
33
34  (: time Int)
35  (= time (initnext 0 (+ time 1)))
36
37  (on true
38  (= addedObjType1List (removeObj addedObjType1List (--> obj (& (clicked
    (prev obj23)) (== (.. obj field1) 1))))))
39  (on (& (== (prev globalVar1) 1) (& clicked (isFree click))) (let ((=
    addedObjType1List (addObj addedObjType1List (ObjType1 1 (Position
    .. click x) (.. click y))))))
40  (on true
41  (= addedObjType2List (removeObj addedObjType2List (--> obj (& (| (
    clicked (prev obj22)) (clicked (prev obj23))) (== (.. obj field1)
    1))))))
42  (on (& (== (prev globalVar1) 2) (& clicked (isFree click))) (let ((=
    addedObjType2List (addObj addedObjType2List (ObjType2 1 (Position
    .. click x) (.. click y))))))
43  (on true
44  (= addedObjType3List (updateObj addedObjType3List (--> obj (nextLiquid
    (prev obj))) (--> obj (& true (== (.. obj field1) 1))))))
```

```

45 (on (& (== (prev globalVar1) 3) (& clicked (isFree click))) (let ((=
    addedObjType3List (addObj addedObjType3List (ObjType3 1 (Position
    .. click x) (.. click y))))))
46 (on true
47 (= addedObjType3List (removeObj addedObjType3List (--> obj (& (clicked
    (prev obj23)) (== (.. obj field1) 1))))))
48 (on (clicked (filter (--> obj (== (.. obj id) 20)) (prev
    addedObjType2List)))
49 (= globalVar1 2))
50 (on (clicked (filter (--> obj (== (.. obj id) 21)) (prev
    addedObjType3List)))
51 (= globalVar1 3))
52 (on (clicked (filter (--> obj (== (.. obj id) 13)) (prev
    addedObjType1List)))
53 (= globalVar1 1)))

```

C.2 Mario Output

```

1 (program
2   (= GRID_SIZE 16)
3   (= background "white")
4   (object ObjType1 (: color String) (list (Cell 0 0 color)))
5   (object ObjType2 (list (Cell -1 0 "darkorange" ) (Cell 0 0 "
    darkorange" ) (Cell 1 0 "darkorange" )))
6   (object ObjType3 (list (Cell -1 -1 "blue" ) (Cell -1 0 "blue" ) (
    Cell 0 -1 "blue" ) (Cell 0 0 "blue" ) (Cell 1 -1 "blue" ) (Cell
    1 0 "blue" )))
7
8   (: addedObjType1List (List ObjType1))
9   (: addedObjType2List (List ObjType2))
10  (: obj8 ObjType3)
11
12  (: addedObjType3List (List ObjType3))
13
14  (= addedObjType1List (initnext (list (ObjType1 "red" (Position 7
    15)) (ObjType1 "gold" (Position 4 12)) (ObjType1 "gold" (
    Position 7 4)) (ObjType1 "gold" (Position 11 6))) (prev
    addedObjType1List)))
15  (= addedObjType2List (initnext (list (ObjType2 (Position 4 13)) (
    ObjType2 (Position 8 10)) (ObjType2 (Position 11 7))) (prev
    addedObjType2List)))
16  (= obj8 (initnext (ObjType3 (Position 7 1)) (prev obj8)))
17
18  (= addedObjType3List (initnext (list) (prev addedObjType3List)))
19
20  /* LATENT VARIABLE: ENEMY DIRECTION */
21  (: globalVar2 Int)
22  (= globalVar2 (initnext 1 (prev globalVar2)))
23
24  /* LATENT VARIABLE: NUM COINS */
25  (: globalVar1 Int)
26  (= globalVar1 (initnext 2 (prev globalVar1)))
27
28  (: time Int)
29  (= time (initnext 0 (+ time 1)))
30
31  (on true (= addedObjType1List (updateObj addedObjType1List (--> obj
    (moveDownNoCollision (prev obj))) (--> obj (intersects (list "red
    ") (map (--> obj (.. obj color)) (list (prev obj)))))))
32  (on true (= addedObjType1List (updateObj addedObjType1List (--> obj
    (moveUp (prev obj))) (--> obj (intersects (list "mediumpurple") (
    map (--> obj (.. obj color)) (list (prev obj)))))))

```

```

33 (on true (= addedObjType1List (updateObj addedObjType1List (--> obj
    (moveRightNoCollision (prev obj))) (--> obj (& right (intersects (
      list "red") (map (--> obj (... obj color)) (list (prev obj))))))))))
34 (on true (= addedObjType1List (updateObj addedObjType1List (--> obj
    (moveLeftNoCollision (prev obj))) (--> obj (& left (intersects (
      list "red") (map (--> obj (... obj color)) (list (prev obj))))))))))
35 (on true (= addedObjType1List (updateObj addedObjType1List (--> obj
    (move (prev obj) 0 -4)) (--> obj (& (intersects (list "red") (map
      (--> obj (... obj color)) (list (prev obj))) (& up (intersects (
        map (--> obj (... (moveDownNoCollision obj) origin)) (list (prev
          obj))) (map (--> obj (... obj origin)) (list (prev obj))))))))))
36 (on true (= addedObjType1List (removeObj addedObjType1List (--> obj
    (! (intersects (prev obj8) (list (prev obj))) (& (intersects (list
      (prev obj)) (filter (--> obj (== (... obj color) "red")) (prev
        addedObjType1List))) (intersects (list "gold") (map (--> obj (...
          obj color)) (list (prev obj))))))))))
37 (on (& (in (prev globalVar1) (list 1 3)) clicked) (let ((=
    addedObjType1List (addObj addedObjType1List (ObjType1 "
    mediumpurple" (uniformChoice (map (--> obj (... obj origin)) (
      filter (--> obj (== (... obj id) 1)) (prev addedObjType1List))))))
    )))
38 (on true (= addedObjType1List (updateObj addedObjType1List (--> obj
    (moveDown (prev obj))) (--> obj (intersects (map (--> obj (
      moveDown obj)) (list (prev obj))) (filter (--> obj (== (... obj
        color) "gold")) (prev addedObjType1List))))))
39 (on true (= addedObjType1List (updateObj addedObjType1List (--> obj
    (moveRight (prev obj))) (--> obj (& right (intersects (map (-->
      obj (moveRight obj)) (list (prev obj))) (filter (--> obj (== (...
        obj color) "gold")) (prev addedObjType1List))))))
40 (on true (= addedObjType1List (updateObj addedObjType1List (--> obj
    (moveLeft (prev obj))) (--> obj (& left (intersects (map (--> obj
      (moveLeft obj)) (list (prev obj))) (filter (--> obj (== (... obj
        color) "gold")) (prev addedObjType1List))))))
41 (on true (= obj8 (moveRight (prev obj8))))
42 (on (& (== (prev globalVar2) 1) (& (... (prev obj8) alive) (! (
    intersects (prev obj8) (prev addedObjType1List)))) (= obj8 (
    moveLeft (prev obj8))))
43 (on (intersects (prev obj8) (prev addedObjType1List)) (= obj8 (
    removeObj (prev obj8))))
44 (on (== (... (... (prev obj8) origin) x) 2)
45 (= globalVar2 2))
46 (on (== (... (... (prev obj8) origin) x) 13)
47 (= globalVar2 1))
48 (on (& (intersects (filter (--> obj (== (... obj color) "red")) (prev
    addedObjType1List)) (filter (--> obj (== (... obj color) "gold"))
    (prev addedObjType1List))) (= (prev globalVar1) 2))
49 (= globalVar1 1))
50 (on (& (intersects (filter (--> obj (== (... obj color) "red")) (prev
    addedObjType1List)) (filter (--> obj (== (... obj color) "gold"))
    (prev addedObjType1List))) (= (prev globalVar1) 1))
51 (= globalVar1 3))
52 (on (& clicked (= (prev globalVar1) 3))
53 (= globalVar1 1))
54 (on (& clicked (= (prev globalVar1) 1))
55 (= globalVar1 2))

```