

Generalized Planning as Heuristic Search

¹Javier Segovia-Aguas, ²Sergio Jiménez and ¹Anders Jonsson

¹Dept. Information and Communication Technologies, Universitat Pompeu Fabra, Spain

²VRAIN - Valencian Research Institute for Artificial Intelligence, Universitat Politècnica de València, Spain
 javier.segovia@upf.edu, serjice@dsic.upv.es, anders.jonsson@upf.edu

Abstract

Although *heuristic search* is one of the most successful approaches to classical planning, this planning paradigm does not apply straightforwardly to *Generalized Planning* (GP). *Planning as heuristic search* traditionally addresses the computation of sequential plans by searching in a grounded state-space. On the other hand GP aims at computing algorithm-like plans, that can branch and loop, and that generalize to a (possibly infinite) set of classical planning instances. This paper adapts the *planning as heuristic search* paradigm to the particularities of GP, and presents the first native heuristic search approach to GP. First, the paper defines a novel GP solution space that is independent of the number of planning instances in a GP problem, and the size of these instances. Second, the paper defines different evaluation and heuristic functions for guiding a combinatorial search in our GP solution space. Lastly the paper defines a GP algorithm, called Best-First Generalized Planning (BFGP), that implements a best-first search in the solution space guided by our evaluation/heuristic functions.

Introduction

Heuristic search is one of the most successful approaches to classical planning (Bonet and Geffner 2001; Hoffmann 2001; Helmert 2006; Richter and Westphal 2010; Lipovetzky and Geffner 2017). Unfortunately, it is not straightforward to adopt state-of-the-art search algorithms and heuristics from classical planning to *Generalized Planning* (GP). The *planning as heuristic search* approach traditionally addresses the computation of sequential plans with a grounded state-space search. Generalized planners reason however about the synthesis of algorithm-like solutions that, in addition to action sequences, contain branching and looping constructs. Furthermore, GP aims to synthesize solutions that generalize to a (possibly infinite) set of planning instances (where the domain of variables may be large), making the grounding of classical planners unfeasible (Winner and Veloso 2003; Hu and Levesque 2011; Siddharth et al. 2011; Hu and De Giacomo 2013; Bonet, Frances, and Geffner 2019; Illanes and McIlraith 2019).

This paper adapts the *planning as heuristic search* paradigm to the particularities of GP, and presents the first

native heuristic search approach to GP. Given a GP problem, that comprises an input set of classical planning instances from a given domain, our *GP as heuristic search* approach computes an algorithm-like plan that solves the full set of input instances. The contributions of the paper are three-fold:

1. *A tractable solution space for GP.* We leverage the computational models of the *Random-Access Machine* (Skiena 1998) and the *Intel x86* FLAGS register (Dandamudi 2005) to define an innovative solution space that is independent of the number of input planning instances in a GP problem, and the size of these instances (i.e. the number of state variables and their domain size).
2. *Grounding-free evaluation/heuristic functions for GP.* We define several evaluation and heuristic functions to guide a combinatorial search in our GP solution space. Evaluating these functions does not require to ground states/actions in advance, so they allow to address GP problems where state variables have large domains (e.g. integers).
3. *A heuristic search algorithm for GP.* We present the BFGP algorithm that implements a best-first search in our GP solution space. Experiments show that BFGP significantly reduces the CPU-time required to compute and validate generalized plans, compared to the classical planning compilation approach to GP (Segovia-Aguas, Jiménez, and Jonsson 2019).

Background

This section introduces the necessary notation to define our heuristic search approach to generalized planning.

Classical Planning

Let X be a set of *state variables*, each $x \in X$ with domain D_x . A *state* is a total assignment of values to the set of state variables, i.e. $\langle x_0 = v_0, \dots, x_N = v_N \rangle$ such that for every $x_i \in X$ then $v_i \in D_{x_i}$. For a variable subset $X' \subseteq X$, let $D[X'] = \times_{x \in X'} D_x$ denote its joint domain. The state space is then $S = D[X]$. Given a state $s \in S$ and a subset of variables $X' \subseteq X$, let $s|_{X'} = \langle x_i = v_i \rangle_{x_i \in X'}$ be the *projection* of s onto X' i.e. the partial state defined by the values that s assigns to the variables in X' . The *projection* of s onto X' defines the subset of states $\{s \in S, s|_{X'} \subseteq s\}$ that are consistent with the corresponding partial state.

Let A be a set of deterministic *lifted actions*. A lifted action $a \in A$ has an associated set of variables $\text{par}(a) \subseteq X$, called *parameters*, and is characterized by two functions: an *applicability function* $\rho_a : D[\text{par}(a)] \rightarrow \{0, 1\}$, and a *successor function* $\theta_a : D[\text{par}(a)] \rightarrow D[\text{par}(a)]$. Action a is applicable in a given state s iff $\rho_a(s|_{\text{par}(a)})$ equals 1, and results in a *successor state* $s' = s \oplus a$, that is built replacing the values that s assigns to variables in $\text{par}(a)$ with the values specified by $\theta_a(s|_{\text{par}(a)})$. Lifted actions generalize actions with *conditional effects*. These actions are common in GP since their state-dependent outcomes allow to adapt generalized plans to different classical planning instances.

A *classical planning instance* is a tuple $P = \langle X, A, I, G \rangle$, where X is a set of state variables, A is a set of lifted actions, $I \in S$ is an initial state, and G is a goal condition on the state variables that induces the subset of *goal states* $S_G = \{s \mid s \models G, s \in S\}$. Given P , a *plan* is an action sequence $\pi = \langle a_1, \dots, a_m \rangle$ whose execution induces a *trajectory* $\tau = \langle s_0, a_1, s_1, \dots, a_m, s_m \rangle$ such that, for each $1 \leq i \leq m$, a_i is applicable in s_{i-1} and results in the successor state $s_i = s_{i-1} \oplus a_i$. A plan π *solves* P if the execution of π in $s_0 = I$ finishes in a goal state, i.e. $s_m \in S_G$. We say π is *optimal* if $|\pi| = m$ is minimal among the plans that solve P .

Generalized Planning

Generalized planning is an umbrella term that refers to more general notions of planning (Jiménez, Segovia-Aguas, and Jonsson 2019). This work builds on top of the inductive formalism for GP, where a GP problem is defined as a set of classical planning instances with a common structure.

Definition 1 (GP problem with shared state variables). A GP problem is a finite and non-empty set of T classical planning instances $\mathcal{P} = \{P_1, \dots, P_T\}$, where each classical planning instance $P_t \in \mathcal{P}$, $1 \leq t \leq T$, shares the same sets of state variables X and lifted actions A , but may differ in the initial state and goals. Formally, $P_1 = \langle X, A, I_1, G_1 \rangle, \dots, P_T = \langle X, A, I_T, G_T \rangle$.

The representations of GP solutions range from *programs* (Winner and Veloso 2003; Segovia-Aguas, Jiménez, and Jonsson 2019) and *generalized policies* (Martín and Geffner 2004; Francès Medina et al. 2019), to *finite state controllers* (Bonet, Palacios, and Geffner 2010; Segovia-Aguas, Jiménez, and Jonsson 2016), *formal grammars* or *hierarchies* (Nau et al. 2003; Ramirez and Geffner 2016; Segovia-Aguas, Jiménez, and Jonsson 2017). Each representation has its own expressiveness capacity, as well as its own validation/computation complexity. We can however define a common condition under which a generalized plan is considered a solution to a GP problem. First, let us define $\text{exec}(\Pi, P) = \langle a_1, \dots, a_m \rangle$ as the sequential plan produced by the execution of a generalized plan Π , on a classical planning instance $P = \langle X, A, I, G \rangle$.

Definition 2 (GP solution). A generalized plan Π *solves* a GP problem $\mathcal{P} = \{P_1, \dots, P_T\}$ iff, for every classical planning instance $P_t \in \mathcal{P}$, $1 \leq t \leq T$, it holds that $\text{exec}(\Pi, P_t)$ solves P_t .

We define the cost of a GP solution as the sum of the costs of the sequential plans that are produced for each of




Figure 1: Classical planning instances for sorting the content of two six-element lists by swapping the list elements.

the classical planning instances in a GP problem. Formally, $\text{cost}(\Pi, \mathcal{P}) = \sum_{P_t \in \mathcal{P}} |\text{exec}(\Pi, P_t)|$. A GP solution Π is optimal if $\text{cost}(\Pi, \mathcal{P})$ is minimal among the generalized plans that solve \mathcal{P} .

Example. Figure 1 shows the initial state and goal of two instances, $P_1 = \langle X, A, I_1, G_1 \rangle$ and $P_2 = \langle X, A, I_2, G_2 \rangle$, for sorting two six-element lists. Both instances share the set of state variables $X = \{x_1, x_2, x_3, x_4, x_5, x_6\}$ that encode six integers. They also share the set of lifted actions A , with $\frac{6 \times 5}{2}$ $\text{swap}(x_i, x_j)$ actions that swap the content of two list positions $i < j$. An example solution plan for P_1 is $\pi_1 = \langle \text{swap}(x_1, x_6), \text{swap}(x_2, x_3), \text{swap}(x_2, x_4) \rangle$ while $\pi_2 = \langle \text{swap}(x_1, x_3), \text{swap}(x_4, x_6) \rangle$ is a solution for P_2 . The set $\mathcal{P} = \{P_1, P_2\}$ is an example GP problem. Sorting instances with different list sizes is made possible by including enough state variables to represent the elements of the longest list, and an extra state variable for the list length; e.g. $\text{length} = l$ implies that list positions $[0, l - 1]$ are valid.

Planning Programs

In this work we represent GP solutions as *planning programs* (Segovia-Aguas, Jiménez, and Jonsson 2019). The execution of a planning program on a classical planning instance is a deterministic procedure that requires no variable matching; this is a useful property when implementing native heuristics for the space of GP solutions.

Formalization. A *planning program* is a sequence of n instructions $\Pi = \langle w_0, \dots, w_{n-1} \rangle$, where each instruction $w_i \in \Pi$ is associated with a *program line* $0 \leq i < n$ and is either:

- A *planning action* $w_i \in A$.
- A *goto instruction* $w_i = \text{go}(i', !y)$, where i' is a program line $0 \leq i' < i$ or $i + 1 < i' < n$, and y is a proposition.
- A *termination instruction* $w_i = \text{end}$. The last instruction of a program Π is always $w_{n-1} = \text{end}$.

The execution model for a planning program is a *program state* (s, i) , i.e. a pair of a planning state $s \in S$ and program counter $0 \leq i < n$. Given a program state (s, i) , the execution of a programmed instruction w_i is defined as:

- If $w_i \in A$, the new program state is $(s', i + 1)$, where $s' = s \oplus w_i$ is the *successor* when applying w_i in s .
- If $w_i = \text{go}(i', !y)$, the new program state is $(s, i + 1)$ if y holds in s , and (s, i') otherwise¹. Proposition y can

¹We adopt the convention of jumping to line i' whenever y is *false*, inspired by jump instructions in the *Random-Access Machine* that jump when a register equals zero.

be the result of an arbitrary expression on state variables, e.g. high-level state features (Lotinac et al. 2016).

- If $w_i = \text{end}$, program execution terminates.

To execute a planning program Π on a planning instance $P = \langle X, A, I, G \rangle$, the initial program state is set to $(I, 0)$, i.e. the initial state of P and the first program line of Π . A program Π *solves* P iff execution terminates in a program state (s, i) that satisfies the goal condition, i.e. $w_i = \text{end}$ and $s \in S_G$. If a planning program *fails* to solve the instance, the only possible sources of failure are: (i) *incorrect program*, i.e. execution terminates in a program state (s, i) that does not satisfy the goal condition, i.e. $(w_i = \text{end}) \wedge (s \notin S_G)$; (ii) *inapplicable program*, i.e. executing action $w_i \in A$ fails in program state (s, i) since w_i is not applicable in s ; (iii) *infinite program*, i.e. execution enters into an infinite loop that never reaches an end instruction.

The space of planning programs. Given a maximum number of program lines n , where the last instruction is $w_{n-1} = \text{end}$, and defined the propositions of goto instructions as $(x = v)$ atoms, where $x \in X$ and $v \in D_x$. The space of possible planning programs is compactly represented with three bit-vectors

1. The *action vector* of length $(n-1) \times |A|$, indicating whether action $a \in A$ appears on line $0 \leq i < n-1$.
2. The *transition vector* of length $(n-1) \times (n-2)$, indicating whether $\text{go}(i', *)$ appears on line $0 \leq i < n-1$.
3. The *proposition vector* of length $(n-1) \times \sum_{x \in X} |D_x|$, indicating if $\text{go}(*, !\langle x = v \rangle)$ appears on line $0 \leq i < n-1$.

A planning program is encoded as the concatenation of these three bit-vectors. The length of the resulting vector is:

$$(n-1) \left(|A| + (n-2) + \sum_{x \in X} |D_x| \right). \quad (1)$$

This encoding allows us to quantify the similarity of two planning programs (their *Hamming distance*) and more importantly, to implement a combinatorial search in the space of planning programs with n lines. In fact, this is the encoding leveraged by the classical planning compilation approach to GP (Segovia-Aguas, Jiménez, and Jonsson 2019). Equation 1 reveals an important scalability limitation of this solution space that limits the applicability of the cited compilation to instances of small size: the number of planning programs with n lines depends on the number, and domain size, of the state variables.

A Tractable Solution Space for GP

This section defines a scalable GP solution space of candidate *planning programs*, that is independent of the size of the input planning instances in a GP problem (i.e. the number of state variables and their domain size). Our approach is to extend the input classical planning instances of a GP problem with a general, and compact, feature language. This way the different classical planning instances in a given GP problem share a common set of features, even if their sets of states variables are actually different.

Our general and compact feature language leverages the computational model of the *Random-Access Machine* (RAM), an abstract computation machine that is polynomial-equivalent to a Turing machine (Boolos, Burgess, and Jeffrey 2002). The RAM machine enhances the *counter machine* with indirect memory addressing (Minsky 1961), and it allows augmented instruction sets and auxiliary *dedicated* registers, like the FLAGS register of the *Intel x86 microprocessors* (Dandamudi 2005).

Classical Planning with Pointers

We extend the classical planning model with a RAM machine of $|Z| + 2$ registers: $|Z|$ *pointers* that reference the original planning state variables, plus the *zero* and *carry* FLAGS (whose joint value will represent our common set of features for GP). Given a classical planning instance $P = \langle X, A, I, G \rangle$, the *extended classical planning instance* with a RAM machine of $|Z| + 2$ registers is defined as $P'_Z = \langle X'_Z, A'_Z, I'_Z, G \rangle$, where:

- The new set of **state variables** X'_Z comprises:
 - The state variables X of the original planning instance.
 - Two *Boolean variables* $Y = \{y_z, y_c\}$, that play the role of the *zero* and *carry* FLAGS register.
 - The *pointers* Z , a non-empty set of extra state variables with finite domain $[0, \dots, |X| - 1]$.
- The new set of actions A'_Z includes:
 - The **planning actions** A' that result from abstracting each original planning action $a \in A$ into its corresponding action scheme. The abstraction is a two-step procedure that requires at least as many pointers as the largest arity of an action in A : (i), each action parameter is replaced with a pointer in Z (ii), its applicability/successor function is rewritten to refer to the state variables pointed by the action parameters. For instance, the $\frac{6 \times 5}{2}$ $\text{swap}(x_i, x_j)$ actions from the example of Figure 1, are abstracted into the action scheme $\{\text{swap}(z_1, z_2) \mid z_1, z_2 \in Z\}$ using two pointers, $Z = \{z_1, z_2\}$.
 - The **RAM actions** that implement the following set of RAM instructions $\{\text{inc}(z_1), \text{dec}(z_1), \text{cmp}(z_1, z_2), \text{cmp}(*z_1, *z_2), \text{set}(z_1, z_2) \mid z_1, z_2 \in Z\}$ over the pointers in Z . Respectively, these RAM instructions increment/decrement a pointer, compare two pointers (or their content), and set the value of a pointer z_1 to another pointer z_2 . Each RAM action also updates the FLAGS $Y = \{y_z, y_c\}$, according to the result of the corresponding RAM instruction (denoted here by res):

$$\begin{aligned} y_z &:= (\text{res} == 0), \\ y_c &:= (\text{res} > 0), \\ \text{inc}(z_1) &\implies \text{res} := z_1 + 1, \\ \text{dec}(z_1) &\implies \text{res} := z_1 - 1, \\ \text{cmp}(z_1, z_2) &\implies \text{res} := z_1 - z_2, \\ \text{cmp}(*z_1, *z_2) &\implies \text{res} := *z_1 - *z_2, \\ \text{set}(z_1, z_2) &\implies \text{res} := z_2. \end{aligned}$$

The number of actions of a classical planning instance, extended with a RAM machine of $|Z| + 2$ registers, is

$$|A'_Z| = 2|Z|^2 + |A'|. \quad (2)$$

The *increment/decrement* instructions induce $2|Z|$ actions, the *set* instructions over pointers induce $|Z|^2 - |Z|$ actions, and *comparison* instructions induce $|Z|^2 - |Z|$ actions; they can compare pointers or their content, but we only consider a single parameter ordering for symmetry breaking, i.e. $\text{cmp}(z_1, z_2)$ but not $\text{cmp}(z_2, z_1)$.

- The new **initial state** I'_Z is the initial state of the original planning instance, but extended with all *pointers* set to zero and *FLAGS* set to `False`. The **goals** are the same as those of the original planning instance.

Theorem 1. *Given a classical planning instance P , its extension P'_Z , with a RAM machine of $|Z| + 2$ registers, preserves the completeness and soundness of P .*

Proof. \Rightarrow : Let $\pi = \langle a_1, \dots, a_m \rangle$ be a plan that solves P . For each action a_i with parameter set $\text{par}(a_i) \subseteq X$, A' contains an action scheme a'_i that replaces the parameters in $\text{par}(a_i)$ with pointers in Z . For each such pointer $z \in Z$, repeatedly apply RAM actions $\text{inc}(z)$, or $\text{dec}(z)$, until it references the associated state variable in $\text{par}(a_i)$, and then apply a'_i . The resulting plan π' has exactly the same effect as π on the original planning state variables in X , and since the goal condition of P'_Z is the same as that of P , it follows that π' solves P'_Z .

\Leftarrow : Let $\pi' = \langle a'_1, \dots, a'_m \rangle$ be a plan that solves P'_Z . Identify each action in A' among those of π' , and execute π' to identify the assignment of variables to pointers when applying each action in A' . Construct a plan π corresponding to the subsequence of actions in A' from π' , replacing each action scheme $a'_i \in A'$ by an original action $a_i \in A$ and choosing as parameters of a_i the state variables referenced by the pointers of a'_i at the moment of execution. Hence a_i has the same effect as a'_i on the state variables in X , implying that π has the same effect as π' on X . Since the goal condition of P is the same as that of P'_Z , it follows that π solves P . \square

Generalized Planning with Pointers

Our extension of classical planning instances with a RAM machine of $|Z| + 2$ registers allows us to define the following minimalist feature language.

Definition 3 (The feature language). *We define the feature language $\mathcal{L} = \{\neg y_z \wedge \neg y_c, y_z \wedge \neg y_c, \neg y_z \wedge y_c, y_z \wedge y_c\}$ of the four possible joint values for the pair of Boolean variables $Y = \{y_z, y_c\}$.*

We say that our feature language \mathcal{L} is general, because it is independent of the number (and domain size) of the planning state variables. Different classical planning instances, that are extended with a RAM machine of $|Z| + 2$ registers, share the features language \mathcal{L} even if they actually have different sets of planning state variables. In this regard, we extend our previous GP problem definition to a set of classical planning instances that, as in (Hu and De Giacomo 2011), they share the same features and actions, but may actually have different state variables.

Definition 4 (GP problem with shared features). *A GP problem with shared features is a finite and non-empty set of*

T classical planning instances $\mathcal{P} = \{P_1, \dots, P_T\}$, where instances share the same set of actions A'_Z , but may differ in the state variables, initial state, and goals. Formally, $P_1 = \langle X'_{1_Z}, A'_Z, I'_{1_Z}, G_1 \rangle, \dots, P_T = \langle X'_{T_Z}, A'_Z, I'_{T_Z}, G_T \rangle$.

We leverage our minimalist feature language \mathcal{L} to define a tractable solution space for GP; we represent GP solutions as *planning programs* where goto instructions are exclusively conditioned on a feature in \mathcal{L} . Limiting the conditions of goto instructions to any of the four features in \mathcal{L} reduces the number of planning programs with n lines, specially when state variables have large domains; the *proposition vector* required to encode a planning program becomes now a vector of only $(n - 1) \times 4$ bits, one bit for each of the four features in \mathcal{L} . Equation 1 simplifies then to:

$$(n - 1) (|A'_Z| + (n - 2) + 4). \quad (3)$$

As illustrated by Equation 3, the size of this new solution space for GP is independent of the number (and domain size) of the planning state variables. Equation 2 already showed that $|A'_Z|$ no longer grows with the number/domain of the planning state variables, and that instead, it grows with the number of pointers $|Z|$. This means that this GP solution space scales to GP problems where state variables have large domains (e.g. integers). The solution space is still incomplete in the sense that either the given bound n on the maximum number of program lines, or the maximum number of pointers available $|Z|$, may be too small to accommodate a solution to a given GP problem.

Example. Figure 2 shows two example planning programs found by our BFGP algorithm for GP: (left) reversing a list; and (right) sorting a list. In both cases the programs generalize, no matter the list length or its content. In *reverse*, line 0 sets the pointer j to the last list element. Then, line 1 swaps the element pointed by i (initially set to zero) and the element pointed by j , pointer j is decremented, pointer i is incremented, and this sequence of instructions is repeated until the condition on line 5 becomes false, i.e. when pointers i and j meet, which means that reversing the list is finished. In *sorting*, pointers i and j are used for inner (lines 5-7) and outer (lines 8-11) loops respectively, and *min* to point to the minimum value in the inner loop (lines 3-4); $\neg y_z \wedge \neg y_c$ on line 2 represents whether the content of j is less than the content of *min*, while $y_z \wedge \neg y_c$ on line 7 represents whether $j == \text{length}$ (and respectively $i == \text{length}$ on line 11).

Generalized Planning as Heuristic Search

This section describes our heuristic search approach to GP; a best-first search in our space of candidate planning programs with at most n program lines, and $|Z|$ pointers available. First, we describe the evaluation and heuristic functions and then, we describe the search algorithm.

Evaluation and Heuristic Functions

We exploit two different sources of information to guide the search in the space of candidate planning programs:

- *The program structure.* These are evaluation functions, computed in linear time, traversing the bit-vector representation of a planning program II.

0. set(j, tail)	0. set(min, i)
1. swap(*i, *j)	1. cmp(*j, *min)
2. dec(j)	2. goto(5, ¬(¬y _z ∧ ¬y _c))
3. inc(i)	3. set(min, j)
4. cmp(j, i)	4. swap(*i, *min)
5. goto(1, ¬(¬y _z ∧ ¬y _c))	5. inc(j)
6. end	6. cmp(length, j)
	7. goto(1, ¬(y _z ∧ ¬y _c))
	8. inc(i)
	9. set(j, i)
	10. cmp(length, i)
	11. goto(0, ¬(y _z ∧ ¬y _c))
	12. end

Figure 2: *Generalized plans*: (left) for reversing a list; (right) for sorting a list with the *selection-sort* algorithm.

- $f_1(\Pi)$, the number of *goto* instructions in Π .
- $f_2(\Pi)$, the number of *undefined* program lines in Π .
- $f_3(\Pi)$, the number of repeated actions in Π .
- *The empirical performance of the program.* These functions assess the performance of a planning program Π on a GP problem \mathcal{P} , executing Π on each of the classical planning instances $P_t \in \mathcal{P}$, $1 \leq t \leq T$:
 - $h_4(\Pi, \mathcal{P}) = n - PC^{MAX}$, where PC^{MAX} is the maximum program line that is eventually reached after executing Π on all the classical planning instances in \mathcal{P} .
 - $h_5(\Pi, \mathcal{P}) = \sum_{P_t \in \mathcal{P}} h_5(\Pi, P_t)$, where

$$h_5(\Pi, P_t) = \sum_{x \in X_t} (v_x - G_t(x))^2.$$

Here, $v_x \in D_x$ is the value eventually reached, for the state variable $x \in X_t$, after executing Π on the classical planning instance $P_t \in \mathcal{P}$, and $G_t(x)$ is the value for this same variable as specified in the goals of P_t . Computing $h_5(\Pi, P_t)$ requires then that the goal condition of the classical planning instance P_t is specified as a partial state. Note also that for Boolean variables the squared difference becomes a simple difference.

- $f_6(\Pi, \mathcal{P}) = \sum_{P_t \in \mathcal{P}} |exec(\Pi, P_t)|$, is the cost of a GP solution, as defined in the *Background* section of the paper. That is $exec(\Pi, P_t)$ is the sequence of actions induced from executing the planning program Π on the planning instance P_t .

All these functions are *cost functions* (i.e. smaller values are preferred). Functions $h_4(\Pi, \mathcal{P})$ and $h_5(\Pi, \mathcal{P})$ are *cost-to-go heuristics*; they provide an estimate on how far a program is from solving the given GP problem. Functions $h_4(\Pi, \mathcal{P})$, $h_5(\Pi, \mathcal{P})$, and $f_6(\Pi, \mathcal{P})$ aggregate several costs that could be expressed as a combination of different functions, e.g. *sum*, *max*, *average*, *weighted average*, etc.

Example. We illustrate how our evaluation/heuristic functions work on the program $\Pi = 0. \text{ swap}(*i, *j) 1. \text{ inc}(i) 2. \text{ dec}(j) 3. \dots 4. \dots 5. \text{ end}$, where only lines 0-2 are specified because lines 3-4 are not programmed yet. The value of the evaluation functions for this partially

specified program is $f_1(\Pi) = 0$, $f_2(\Pi) = 5 - 3 = 2$, $f_3(\Pi) = 0$. Given the GP problem $\mathcal{P} = \{P_1, P_2\}$ that comprises the two classical planning instances illustrated in Figure 1, and pointers i and j starting at the first and last memory register, respectively, we can compute h_4 and h_5 to evaluate how far Π is from solving the GP problem of sorting lists, and the accumulated cost f_6 . In this case $h_4(\Pi, \mathcal{P}) = 5 - 3 = 2$, $h_5(\Pi, \mathcal{P}) = 22$ and $f_6(\Pi, \mathcal{P}) = 3 + 3 = 6$.

Best First Search for Generalized Planning

Here we provide the implementation details of our BFGP algorithm for generalized planning. This algorithm implements a *Best-First Search* (BFS) in the space of planning programs with n program lines, and a RAM machine of $|Z| + 2$ registers.

The BFGP algorithm starts a BFS with an empty planning program, that is represented by the concatenation of the bit-vectors defined in the previous section, with all their bits set to `False`. To generate a tractable set of successor nodes, child nodes in the search tree are restricted to programs that result from programming the bits that correspond to the PC^{MAX} line (i.e. the maximum line reached after executing the current program on the classical planning instances in \mathcal{P}). The planning program of a search node is then at Hamming distance 1 from its parent, when programming a planning action, or at Hamming distance 2, when programming a *goto* instruction. This procedure for successor generation guarantees that duplicate successors are not generated.

BFGP is a *frontier search* algorithm, meaning that, to reduce memory requirements, BFGP stores only the open list of generated nodes, but not the closed list of expanded nodes (Korf et al. 2005). BFS sequentially expands the best node in a priority queue (aka *open list*) sorted by an evaluation/heuristic function. When evaluating the *empirical performance* of a program, e.g. computing heuristics $h_i(\Pi, \mathcal{P})$, if the execution of Π on a given instance $P_t \in \mathcal{P}$ fails, this means that the search node corresponding to the planning program Π is a dead-end, and hence it is not added to the open list. If the planning program Π solves all the instances $P_t \in \mathcal{P}$, then search ends, and Π is a valid solution for the GP problem \mathcal{P} . BFGP can compute optimal solutions when run in *anytime mode*. In this case we can use $f_6(\Pi, \mathcal{P})$ to rank GP solutions (e.g. to prefer a sorting program with smaller computational complexity).

Evaluation

This section evaluates the empirical performance of our *GP as heuristic search* approach. All experiments are performed in an Ubuntu 20.04 LTS, with AMD® Ryzen 7 3700x 8-core processor $\times 16$ and 32GB of RAM. We make fully available the compilation source code, evaluation scripts, and used benchmarks (Segovia-Aguas, Jiménez, and Jonsson 2021), so any experimental data reported in the paper is fully reproducible.

Benchmarks

We report results in eight domains of two kinds (*numerical series* and *memory manipulation*). Numerical series compute the n^{th} term of a math sequence. *Memory manipulation*

Domain	$n, Z $	f_1				f_2				f_3			
		Time	Mem.	Exp.	Eval.	Time	Mem.	Exp.	Eval.	Time	Mem.	Exp.	Eval.
T. Sum	5, 2	0.24	4.2	4.8K	5.8K	0.30	3.8	6.4K	6.4K	0.13	3.9	2.2K	2.8K
Corridor	7, 2	3.04	6.6	12.4K	26.7K	0.41	3.8	2.2K	2.2K	3.66	4.2	24.2K	25.9K
Reverse	7, 3	82	61	0.28M	0.57M	181	4.0	0.95M	0.95M	170	23	0.75M	0.84M
Select	7, 3	198	110	0.83M	1.10M	27	3.9	0.12M	0.12M	76.49	11	0.34M	0.38M
Find	7, 3	195	175	0.50M	1.36M	271	4.0	1.46M	1.46M	86	12	0.41M	0.45M
Fibonacci	8, 3	496	922	2.48M	6.79M	1,082	3.9	11.8M	11.8M	TO	-	-	-
Gripper	8, 4	TO	-	-	-	3,439	4.1	19.9M	19.9M	TO	-	-	-
Sorting	9, 3	TO	-	-	-	TO	-	-	-	3,143	711	19.5M	22.9M
Average		162	213	0.68M	1.64M	714	3.9	4.89M	4.89M	580	128	3.51M	4.09M

Domain	$n, Z $	h_4				h_5				f_6			
		Time	Mem.	Exp.	Eval.	Time	Mem.	Exp.	Eval.	Time	Mem.	Exp.	Eval.
T. Sum	5, 2	0.10	3.8	1.6K	1.6K	0.09	3.9	1.2K	1.4K	0.45	4.8	7.4K	7.4K
Corridor	7, 2	1.09	3.9	5.3K	5.4K	5.29	4.1	30.3K	31.3K	6.16	7.6	35.3K	35.3K
Reverse	7, 3	205	4.2	0.88M	0.88M	1.46	4.2	4.9K	6.3K	369	230	1.57M	1.65M
Select	7, 3	0.80	3.9	3.0K	4.2K	94	5.7	0.34M	0.35M	255	155	1.06M	1.14M
Find	7, 3	415	4.4	1.76M	1.76M	140	7.0	0.58M	0.59M	423	244	1.76M	1.77M
Fibonacci	8, 3	TO	-	-	-	1,500	120	11.3M	11.8M	TO	-	-	-
Gripper	8, 4	TO	-	-	-	83	5.5	0.34M	0.35M	TO	-	-	-
Sorting	9, 3	TO	-	-	-	TO	-	-	-	TO	-	-	-
Average		125	4.0	0.53M	0.53M	260	21	1.80M	1.88M	211	128	0.89M	0.92M

Table 1: We report the number of program lines n , and pointers $|Z|$ per domain, and for each evaluation/heuristic function, CPU (secs), memory peak (MBs), and the numbers of expanded and evaluated nodes. TO stands for Time-Out (> 1 h of CPU).

domains include *Reverse* for reversing the content of a list, *Select* for choosing the minimum value in the list, *Find* for counting the number of occurrences of a specific value in a list, *Corridor* to move to any arbitrary location in a corridor starting from another arbitrary location, and *Gripper* to move all balls from one room to another. In the numerical series domains, the planning instances of a GP problem represent the computation of the first ten terms of the series. In the memory manipulation domains, the planning instances have initial random content, with list sizes ranging from two to twenty-one (besides *grripper* where all balls are initially in room A). The results of arithmetical operations is bounded to 10^2 in the synthesis of GP solutions, and to 10^9 in the validation of GP solutions.

All domains include the base RAM instructions $\{\text{inc}(z_1), \text{dec}(z_1), \text{cmp}(z_1, z_2), \text{cmp}(*z_1, *z_2), \text{set}(z_1, z_2) \mid z_1, z_2 \in Z\}$. In addition, each domain also contains regular planning action schemes, that do not affect the FLAGS.

- *Triangular Sum* and *Fibonacci* include the action scheme $\text{add}(*z_1, *z_2)$ for adding the referenced value of a pointer to the content of the other pointer.
- *Select* and *Find* only require the base RAM instructions. *Reverse* and *Sorting* include also the $\text{swap}(*z_1, *z_2)$ action scheme. *Corridor* includes $\text{left}(*z_1)$ and $\text{right}(*z_1)$ schemes for moving left and right in the corridor, respectively, and *Gripper* includes the action schemes $\text{pick}(*z_1)$, $\text{drop}(*z_1)$, $\text{moveAB}()$, and $\text{moveBA}()$.

Synthesis of GP Solutions

Table 1 summarizes the results of BFGP with our six different evaluation/heuristic functions (best results in bold): i) f_2 and h_5 exhibited the best coverage; ii) there is no clear dominance of a *structure* evaluation function, f_2 has the best

Dom.	BFGP(f_1, h_5) / BFGP(h_5, f_1)			
	T.	M.	Exp.	Eval.
T. Sum	0.2/ 0.1	4.3/ 3.8	2.8K/ 1.1K	4.8K/ 1.4K
Corr.	3.2/4.5	6.6/ 5.9	6.5K/26.0K	21.6K/27.5K
Rev.	63/ 1.4	52/ 4.7	81.9K/ 3.7K	0.3M/ 7.7K
Sel.	203/ 80	110/ 7.0	0.6M/ 0.3M	0.9M/ 0.3M
Find	313/ 162	176/ 14	0.9M/ 0.7M	1.5M/ 0.7M
Fibo.	528/ 22	828/ 32	1.4M/ 75K	5.3M/ 0.2M
Grip.	TO/ 6.9	-/ 10	-/ 5.8K	-/ 37.3K
Sort.	TO/ 713	-/ 730	-/ 4.4M	-/ 4.5M

Table 2: For each domain we report, CPU time (secs), memory peak (MBs), num. of expanded and evaluated nodes. TO means time-out (> 1 h of CPU). Best results in bold.

memory consumption while f_3 is the only structural function that solves *Sorting*; iii) the *performance*-based function h_5 dominates h_4 and f_6 .

Interestingly, the base performance of BFGP with a single evaluation/heuristic function is improved combining both structural and cost-to-go information; we can guide the search of BFGP with a cost-to-go heuristic function and break ties with a structural evaluation function, and vice versa. Table 2 shows the performance of $BFGP(f_1, h_5)$ and its reversed configuration $BFGP(h_5, f_1)$ which actually resulted in the overall best configuration solving all domains.

Figure 3 shows the programs computed by $BFGP(h_5, f_1)$. The state variables in *Triangular Sum* are two pointers (a and b) and two memory registers ($v[0], v[1]$). The program for *Triangular Sum* starts with $a = 0, b = 0, v[0] = 0$ and $v[1] = n$; first it increases b by one, then iterates adding the content of $v[1]$ to $v[0]$, and decreasing $v[1]$ until it becomes 0. The *Corridor* domain has a pointer i , which points to a register with the current

location, and a constant pointer g_i that accesses a memory register with the goal location. The solution moves the agent right in the corridor until the goal location is surpassed, then it moves the agent left until the goal is reached.

The *Reverse*, *Select* and *Find* domains have three pointers each, and one of these pointers is a constant pointer indicating the vector size. The *Reverse* solution is the same illustrated in Figure 2, that was already explained above. *Select* uses pointer a to iterate over the vector, compares each position with the content of b , and assigns a to b if the content of a is smaller than the content of b . The *Find* program requires a counter to accumulate the number of occurrences of the target element to be found; the *Find* program compares the content of a with a target content, and if they are equal, the counter is increased by one.

In *Fibonacci* there is a vector that stores the first n numbers of the Fibonacci sequence. The constant pointer n indicates the vector size, and pointers b and c are used to iterate over the vector (pointer c acts as F_n , while b plays the role of F_{n-1} and F_{n-2}). The solution program for *Gripper* only uses the left arm of the robot to pick a ball, then moves the robot to the next room, drops the ball, and moves the robot back to the initial room, repeating the entire process until the last object is processed. *Sorting* is solved by $BFGP(h_5, f_1)$ with a succinct but more complex implementation of the selection-sort algorithm, where the first register is used as the selected element, and that requires many back and forth iterations until pointer i reaches the first element.

Validation of the GP Solutions

Table 3 reports the CPU time, and peak memory, yield when running the solutions synthesized by $BFGP(h_5, f_1)$ on a validation set. All the solutions synthesized by $BFGP(h_5, f_1)$ were successfully validated. In the validation set the domain of the planning state variables equals the integer bounds (10^9). The solution for *Triangular Sum* is validated over 44,709 instances, the last one corresponding to the 44,720th term in the sequence, i.e. the integer 999,961,560. *Fibonacci* has a validation set of 33 instances, ranging from the 11th Fibonacci term to the 43rd, i.e. the integer 701,408,733. The solutions for *Reverse*, *Select*, and *Find* domains, are validated on 50 instances each, with vector sizes ranging from 1K to 50K, and random integer elements bounded by 10^9 . *Corridor* validation instances range from corridor lengths of 12 to 1,011, and *Gripper* is validated on 1,000 instances with up to 1,011 balls. Validation instances in *Sorting* are vectors with sizes ranging from 12 to 31, and with random integer elements bounded by 10^9 . The largest CPU-times and memory peaks correspond to the configuration that implements the detection of *infinite programs*, which requires saving states to detect whether they are revisited during execution. Skipping this mechanism allows to validate non-infinite programs faster (Segovia-Aguas, Jiménez, and Jonsson 2020).

Comparison with GP Compilation

Table 4 compares the performance of $BFGP(h_5, f_1)$, in terms of CPU-time, with PP, which stands for the compilation-based approach for GP (Segovia-Aguas,

Dom.	Inst.	Time _∞	Mem _∞	Time	Mem
T. Sum	44,709	1,066.74	53MB	574.08	47MB
Corr.	1,000	0.23	5.0MB	0.15	4.7MB
Rev.	50	37.96	5.2GB	2.70	0.3GB
Sel.	50	144.75	19.6GB	2.29	33MB
Find	50	114.55	19.6GB	2.12	33MB
Fibo.	33	0.00	4.2MB	0.00	3.9MB
Grip.	1,000	2.71	0.1GB	1.65	0.1GB
Sort.	20	272.06	15.2GB	52.04	3.8MB

Table 3: Validation set, CPU-time (secs) and memory peak for program validation, with/out *infinite program* detection.

Domain	PP in sec.	BFGP(h_5, f_1) in sec.
Triangular Sum	0.85	0.1
Corridor	-	4.5
Reverse	87.86	1.4
Select	204.20	80
Find	274.86	162
Fibonacci	3,570	22
Gripper	1	6.9
Sorting	-	713

Table 4: Computing CPU-time (secs) for solving domains in the GP compilation approach (PP) and $BFGP(h_5, f_1)$.

Jiménez, and Jonsson 2019), that computes planning programs with the LAMA-2011 setting (first solution) to solve the classical planning problems that result from the compilation. The *Gripper* domain is the only one where it is easier to compute a generalized plan following the classical planning compilation. In addition, the compilation-based approach reported no solution for the *Corridor* and *Sorting* domains. What is more, the compilation approach (PP) is unable to perform the previous validation; off-the-shelf planner cannot handle state variables with the reported domain sizes.

Related Work

This work is the first full native heuristic search approach to GP, with a search space, evaluation/heuristic functions, and search algorithm, that are specially targeted to GP. Our *GP as heuristic search* approach is related to the classical planning compilation for GP (Segovia-Aguas, Jiménez, and Jonsson 2019). The paper showed that this compilation presented an important drawback; the size of the classical planning instance output by the compilation grows exponentially with the number and domain size of the state variables. In practice, this drawback limits the applicability of the cited compilation to planning instances of small size.

A related trend in GP represents generalized plans as lifted policies that are computed via a FOND planner (Bonet, Frances, and Geffner 2019; Illanes and McIlraith 2019). Instantiating a lifted policy for solving a specific classical planning instance requires a matching process over the objects in the instance. This instantiation requirement makes it difficult to define native heuristic functions for searching in the space of possible GP solutions with a classic heuristic search algorithm. On the other hand, our approach naturally suits a classic best-first search framework; the execu-

0. inc(b)	0. right(*i)	0. set(j, tail)	0. inc(a)
1. add(*a,*b)	1. cmp(*i,*gi)	1. swap(*i,*j)	1. cmp(*b,*a)
2. dec(*b)	2. goto(0, ¬(¬y _z ∧ y _c))	2. dec(j)	2. goto(4, ¬(¬y _z ∧ y _c))
3. goto(1, ¬(y _z ∧ ¬y _c))	3. left(*i)	3. inc(i)	3. set(b, a)
4. end	4. cmp(*i,*gi)	4. cmp(j, i)	4. cmp(tail, a)
	5. goto(1, ¬(y _z ∧ ¬y _c))	5. goto(1, ¬(¬y _z ∧ ¬y _c))	5. goto(0, ¬(y _z ∧ ¬y _c))
	6. end	6. end	6. end
0. cmp(*target,*a)	0. inc(c)	0. pick(*left-arm)	0. set(i, tail)
1. goto(3, ¬(y _z ∧ ¬y _c))	1. inc(c)	1. moveAB()	1. swap(*i,*j)
2. inc(accumulator)	2. add(*c,*b)	2. drop(*left-arm)	2. cmp(*j,*i)
3. inc(a)	3. inc(b)	3. inc(left-arm)	3. goto(6, ¬(¬y _z ∧ y _c))
4. cmp(tail, a)	4. add(*c,*b)	4. moveBA()	4. inc(i)
5. goto(0, ¬(y _z ∧ ¬y _c))	5. cmp(n, c)	5. cmp(left-arm, last-obj)	5. goto(0, ¬(y _z ∧ y _c))
6. end	6. goto(0, ¬(y _z ∧ ¬y _c))	6. goto(0, ¬(y _z ∧ ¬y _c))	6. dec(i)
	7. end	7. end	7. goto(1, ¬(y _z ∧ ¬y _c))
			8. end

Figure 3: Solutions computed by $BFGP(h_5, f_1)$ for Tr. Sum, Corridor, Reverse, Select, Find, Fibonacci, Gripper and Sorting.

tion of a planning program on a classical planning instance is a matching-free process that enables the definition of effective evaluation/heuristic functions.

Our feature language is related to the ones used in *Qualitative numeric planning* (Srivastava et al. 2011; Bonet and Geffner 2020; Illanes and McIlraith 2019). *Qualitative numeric planning* leverages propositional variables to abstract the value of numeric state variables. In our work the value of FLAGS $Y = \{y_z, y_c\}$ depends on the last executed action. Considering that only RAM instructions update the variables in Y , we have an observation space of $2^{|Y|} \times 2^{|Z|}$ state observations implemented with only $|Y|$ Boolean variables. The four joint values of $\{y_z, y_c\}$ can model a large space of observations, e.g. $=, \neq, <, >, \leq, \geq$ on variable pairs. Our state observations may also refer to pointer contents; in the sorting program of Figure 2, $\text{cmp}(*j, *min)$ is conditioned on the value of the variables referenced by two pointers, while $\text{cmp}(\text{length}, j)$ is conditioned on the pointers themselves.

Last but not least, our GP solution representation can be understood as a program encoded in a minimalist programming language (i.e. with a lean instruction set). Furthermore, many GP benchmarks have solutions that are known to be polynomial algorithms. Previous work on generalized planning (Segovia-Aguas, Jiménez, and Jonsson 2019; Illanes and McIlraith 2019; Bonet et al. 2020) was already aware of the connection between GP and *program synthesis* (Gulwani et al. 2017; Lee et al. 2018; Alur et al. 2018). We believe this work builds a stronger connection between these two closely related areas.

Conclusion

We presented the first native heuristic search approach for GP. Since we are approaching GP as a classic tree search, a wide landscape of effective techniques, coming from heuristic search and classical planning, may improve the base performance of our approach. We mention some of the more promising ones. Large open lists can be handled more effec-

tively splitting them in several smaller lists (Helmert 2006). *Delayed duplicate detection* could be implemented to manage large closed lists with magnetic disk memory (Korf 2008). Further, once a search node is cancelled (e.g. because $h_i(\Pi, \mathcal{P})$ identified that the *planning program* fails on a given instance), any program equivalent to this node should also be cancelled, e.g. any program that can be built with transpositions of the causally-independent instructions.

BFGP starts from the empty program but nothing prevents us from starting search from a partially specified planning program. In fact, local search approaches have already shown successful for planning (Gerevini, Saetti, and Serina 2003) and program synthesis (Solar-Lezama 2009; Gulwani et al. 2017). SATPLAN planners exploit multiple-thread computing to parallelize search in solution spaces with different bounds (Rintanen 2012). This same idea could be applied to multiple searches for GP solutions with different program sizes. Last but not least, our goal-driven heuristic $h_5(\Pi, P_t)$ builds on top of the *Euclidean distance*; better estimates may be obtained by building on top of better informed planning heuristics (Hoffmann 2003; Helmert 2006; Francès et al. 2017).

Acknowledgments

Javier Segovia-Aguas is supported by TAILOR, a project funded by EU H2020 research and innovation programme no. 952215, an ERC Advanced Grant no. 885107, and grant TIN-2015-67959-P from MINECO, Spain. Sergio Jiménez is supported by the *Ramon y Cajal* program, RYC-2015-18009, the Spanish MINECO project TIN2017-88476-C2-1-R. Anders Jonsson is partially supported by Spanish grants PID2019-108141GB-I00 and PCIN-2017-082.

References

- Alur, R.; Singh, R.; Fisman, D.; and Solar-Lezama, A. 2018. Search-based program synthesis. *Communications of the ACM* 61(12): 84–93.
- Bonet, B.; De Giacomo, G.; Geffner, H.; Patrizi, F.; and

- Rubin, S. 2020. High-Level Programming via Generalized Planning and LTL Synthesis. In *KR*, volume 17, 152–161.
- Bonet, B.; Frances, G.; and Geffner, H. 2019. Learning features and abstract actions for computing generalized plans. In *AAAI*, volume 33, 2703–2710.
- Bonet, B.; and Geffner, H. 2001. Planning as heuristic search. *Artificial Intelligence* 129: 5–33.
- Bonet, B.; and Geffner, H. 2020. Qualitative Numeric Planning: Reductions and Complexity. *JAIR* 69: 923–961.
- Bonet, B.; Palacios, H.; and Geffner, H. 2010. Automatic Derivation of Finite-State Machines for Behavior Control. In *AAAI*, volume 24.
- Boolos, G. S.; Burgess, J. P.; and Jeffrey, R. C. 2002. *Computability and logic*. Cambridge university press.
- Dandamudi, S. P. 2005. Installing and using nasm. *Guide to Assembly Language Programming in Linux*.
- Francès, G.; et al. 2017. *Effective planning with expressive languages*. Ph.D. thesis, Universitat Pompeu Fabra.
- Francès Medina, G.; Corrêa, A. B.; Geissmann, C.; and Pommerening, F. 2019. Generalized potential heuristics for classical planning. In *IJCAI*.
- Gerevini, A.; Saetti, A.; and Serina, I. 2003. Planning through stochastic local search and temporal action graphs in LPG. *JAIR* 20: 239–290.
- Gulwani, S.; Polozov, O.; Singh, R.; et al. 2017. Program synthesis. *Foundations and Trends in Programming Languages*.
- Helmert, M. 2006. The Fast Downward Planning System. *Journal of Artificial Intelligence Research* 26: 191–246.
- Hoffmann, J. 2001. FF: The fast-forward planning system. *AI magazine* 22(3): 57–57.
- Hoffmann, J. 2003. The Metric-FF Planning System: Translating "Ignoring Delete Lists" to Numeric State Variables. *JAIR* 20: 291–341.
- Hu, Y.; and De Giacomo, G. 2011. Generalized planning: Synthesizing plans that work for multiple environments. In *IJCAI*, volume 22, 918–923.
- Hu, Y.; and De Giacomo, G. 2013. A Generic Technique for Synthesizing Bounded Finite-State Controllers. In *ICAPS*.
- Hu, Y.; and Levesque, H. J. 2011. A Correctness Result for Reasoning about One-Dimensional Planning Problems. In *IJCAI*, 2638–2643.
- Illanes, L.; and McIlraith, S. A. 2019. Generalized planning via abstraction: arbitrary numbers of objects. In *AAAI*, volume 33, 7610–7618.
- Jiménez, S.; Segovia-Aguas, J.; and Jonsson, A. 2019. A review of generalized planning. *KER* 34: 1–28.
- Korf, R. E. 2008. Linear-time disk-based implicit graph search. *Journal of the ACM* 55(6): 1–40.
- Korf, R. E.; Zhang, W.; Thayer, I.; and Hohwald, H. 2005. Frontier search. *Journal of the ACM* 52(5): 715–748.
- Lee, W.; Heo, K.; Alur, R.; and Naik, M. 2018. Accelerating search-based program synthesis using learned probabilistic models. *ACM SIGPLAN Notices*.
- Lipovetzky, N.; and Geffner, H. 2017. Best-first width search: Exploration and exploitation in classical planning. In *AAAI*, volume 31.
- Lotinac, D.; Segovia-Aguas, J.; Jiménez, S.; and Jonsson, A. 2016. Automatic generation of high-level state features for generalized planning. In *IJCAI*, 3199–3205.
- Martín, M.; and Geffner, H. 2004. Learning generalized policies from planning examples using concept languages. *Applied Intelligence* 20: 9–19.
- Minsky, M. L. 1961. Recursive unsolvability of Post's problem of "tag" and other topics in theory of Turing machines. *Annals of Mathematics*.
- Nau, D. S.; Au, T.-C.; Ilghami, O.; Kuter, U.; Murdock, J. W.; Wu, D.; and Yaman, F. 2003. SHOP2: An HTN planning system. *JAIR* 20: 379–404.
- Ramirez, M.; and Geffner, H. 2016. Heuristics for Planning, Plan Recognition and Parsing. *arXiv:1605.05807*.
- Richter, S.; and Westphal, M. 2010. The LAMA planner: Guiding cost-based anytime planning with landmarks. *JAIR* 39: 127–177.
- Rintanen, J. 2012. Planning as satisfiability: Heuristics. *Artificial intelligence* 193: 45–86.
- Segovia-Aguas, J.; Jiménez, S.; and Jonsson, A. 2016. Hierarchical Finite State Controllers for Generalized Planning. In *IJCAI*, 3235–3241.
- Segovia-Aguas, J.; Jiménez, S.; and Jonsson, A. 2017. Generating context-free grammars using classical planning. In *IJCAI*.
- Segovia-Aguas, J.; Jiménez, S.; and Jonsson, A. 2019. Computing programs for generalized planning using a classical planner. *Artificial Intelligence* 272: 52–85.
- Segovia-Aguas, J.; Jiménez, S.; and Jonsson, A. 2020. Generalized Planning with Positive and Negative Examples. In *AAAI*, volume 34, 9949–9956.
- Segovia-Aguas, J.; Jiménez, S.; and Jonsson, A. 2021. Best First Generalized Planning. doi:10.5281/zenodo.4600833. <https://doi.org/10.5281/zenodo.4600833>.
- Siddharth, S.; Neil, I.; Shlomo, Z.; and Tianjiao, Z. 2011. Directed Search for Generalized Plans Using Classical Planners. In *ICAPS*, volume 21.
- Skiena, S. S. 1998. *The algorithm design manual: Text*, volume 1. Springer Science & Business Media.
- Solar-Lezama, A. 2009. The sketching approach to program synthesis. In *Asian Symposium on Programming Languages and Systems*.
- Srivastava, S.; Zilberstein, S.; Immerman, N.; and Geffner, H. 2011. Qualitative numeric planning. In *AAAI*, volume 25.
- Winner, E.; and Veloso, M. 2003. DISTILL: Learning Domain-Specific Planners by Example. In *ICML*, 800–807.