Compliant Conditions for Polynomial Time Approximation of Operator Counts

Tathagata Chakraborti[1] **Sarath Sreedharan**[1] **Sailik Sengupta**[1]
T. K. Satish Kumar[2] **Subbarao Kambhampati**[1]

Department of Computer Science[1]
Arizona State University
{tchakra2, ssreedh3, sailiks, rao}@asu.edu

Department of Computer Science[2]
University of Southern California
tkskwork@gmail.com

Abstract

We develop a computationally simpler version of the operator count heuristic for a particular class of domains. The contribution of this abstract is thus threefold, we (1) propose an efficient closed form approximation to the operator count heuristic; (2) leverage compressed sensing techniques to obtain an integer approximation in polynomial time; and (3) discuss the relationship of the proposed formulation to existing heuristics and investigate properties of domains where such approaches are useful.

The OP−COUNT Heuristic

Domain Model. The domain is described by a set of variables $f \in \mathcal{F}$ which can assume values from a (finite) domain $D(f) \subseteq \mathbb{N}$. A state is given by the particular assignment of values to these variables: $\mathbb{S} = \{f = v \mid v \in D(f)\ \forall f \in \mathcal{F}\}$. The value of variable f in state $\mathbb{S}$ is referred to as $\mathbb{S}(f)$. The action model $\mathcal{A}$ consists of operators $a = \langle C_a, E_a \rangle$ where C_a is the cost of the action, and $E_a = \{\langle f, v_o, v_n \rangle \mid f \in \mathcal{F}; v_o, v_n \in \{-1\} \cup D(f)\}$ is the set of effects. The transition function $\delta(\cdot)$ determines the next state after the application of action a to state $\mathbb{S}$ as -

$$\delta(a, \mathbb{S}) = \perp \text{ if } \exists \langle f, v_o, v_n \rangle \in E_a \text{ s.t. } v_o \neq -1 \wedge v_o \neq \mathbb{S}(f);$$
$$= \{f = v_n \forall \langle f, v_o, v_n \rangle \in E_a;\ \text{else } f = \mathbb{S}(f)\} \text{ otherwise.}$$

Plans and Operator Counts. A planning problem is a tuple $\Pi = \langle \mathcal{F}, \mathcal{A}, \mathbb{I}, \mathbb{G} \rangle$, where $\mathbb{I}, \mathbb{G}$ are the initial and (partial) goal states respectively. The solution to the planning problem is a *plan* $\pi = \langle a_1, a_2, \ldots \rangle$, $\pi(i) = a_i \in \mathcal{A}$ such that $\delta(\pi, \mathbb{I}) \models \mathbb{G}$, where the cumulative transition function is given by $\delta(\pi, \mathbb{S}) = \delta(\langle a_2, a_3, \ldots \rangle, \delta(a_1, \mathbb{S}))$. The cost of the plan is given by $C(\pi) = \sum_{a \in \pi} C_a$ and an *optimal plan* π^* is such that $C(\pi^*) \leq C(\pi)\ \forall \pi$. The operator count for an action a given a plan π is given by $\lambda(a, \pi) = |\{i \mid a = \pi(i)\}|$ and the total operator count of the plan $\lambda(\pi) = |\pi|$.

Compliant Variables. We define compliant variables as those that whenever they occur as a precondition of an action, they must also be an effect, and vice versa. Thus, $f \in \mathcal{F}$ is *compliant* iff $\forall a \in \mathcal{A}, \langle f, v_o, v_n \rangle \in E_a \implies v_o \neq -1 \wedge v_n \neq -1$; f is referred to as *rogue* otherwise. Let $\Phi \subseteq \mathcal{F}$ be the set of all compliant variables, and the set of compliant variables whose values are specified in the goal be $\phi \subseteq \Phi$, henceforth referred to as goal compliant conditions.

The State Transformation Equation. Let $|\phi| = m$ and $|\mathcal{A}| = n$. Consider an $m \times n$ matrix $\mathbf{M}$ whose ij^{th} element $M_{ij} \in \mathbb{Z}$ is the numerical change in $f_i \in \phi$ produced by action $a_j \in \mathcal{A}$, i.e. $M_{ij} = v_n - v_o$; $\langle f_i, v_o, v_n \rangle \in E_{a_j}$. Also, let $\mathbf{D}$ be a vector of size m whose i^{th} entry d_i is the change in a goal compliant $f \in \phi$ from the current state to the final state, i.e. $d_i = v_g - v_c; v_g = f_i \in \mathbb{G}, v_c = f_i \in \mathbb{S}$; and let $\mathbf{x}$ be a vector of size n, whose i^{th} element is $x_i \in \mathbb{N}$. Then the following equality holds:

$$\mathbf{Mx} = \mathbf{D} \tag{1}$$

The integer solution $\mathbf{x}^*$ to this system of linear equations with the least $|x^*|$ gives a lower bound on the operator counts required to solve the planning problem, i.e. $|x^*| \leq |\pi^*|$. We can compute a real-valued approximation to this in closed form, by minimizing the l_2-*norm* $||x||_2^2$ using the Lagrangian multiplier method as follows -

$$L(\mathbf{x}) = \frac{1}{2}||\mathbf{Qx}||^2 + \lambda^T(\mathbf{D} - \mathbf{Mx}) \tag{2}$$

$$\implies \mathbf{x}^* = \mathbf{Q}^{-2}\mathbf{M}^T(\mathbf{MQ}^{-2}\mathbf{M}^T)^{-1}\mathbf{D} \tag{3}$$

Here $\mathbf{Q}$ is a $n \times n$ matrix of action costs whose ij^{th} entry $Q_{ij} = C_{a_i}$ if $i = j$; 0 otherwise (for unit cost domains) $\mathbf{Q}$ is an identity matrix and $\mathbf{x}^* = \mathbf{M}^T(\mathbf{MM}^T)^{-1}\mathbf{D}$ The most costly operation here is the calculation of the pseudo inverse, which can be done in $\approx \mathcal{O}(n^{2.3})$ time. Further, $\mathbf{M}$ is problem independent, and hence the factor $\mathbf{Z} = \mathbf{Q}^{-2}\mathbf{M}^T(\mathbf{MQ}^{-2}\mathbf{M}^T)^{-1}$ can be *precomputed* given an action model. Thus it follows that we can readily use $||\mathbf{QZD}||$ as a heuristic for state-space search. Note that this formulation can also determine infeasibility of goal reachability immediately (in domains where actions are not reversible this is extremely useful in the course of search) when the system in unsolvable, as shown in Algorithm 1.

Sparse coding. We would have ideally wanted an integer solution to Eqn 1, but this makes the problem computationally intractable. The real-valued approximation of minimizing l_2-norm mentioned above can be improved further. For this, we note that in most cases $n \gg m$ and also $n \gg |\mathbf{x}^*|$ due to the combinatorial explosion during grounding of domains. To exploit this knowledge about the sparsity of $|\mathbf{x}^*|$, we draw upon compressed sensing techniques (Candès and Wakin 2008) to enforce sparsity. To this end, we suggest

Algorithm 1 Using OP−COUNT Heuristic for State-Space Search

procedure PRE-COMPUTE(Π)
  Compute $\mathbf{M}$, $\mathbf{Q}$
  Convert $\mathbf{M}$ to row echelon form $\rightarrow$ $\mathbf{T}$ is the transformation matrix, r is the rank
  $\mathbf{Y} \leftarrow \mathbf{M}[1:r, :]$, $\mathbf{Z} \leftarrow \mathbf{Q}^{-2}\mathbf{Y}^T(\mathbf{Y}\mathbf{Q}^{-2}\mathbf{Y}^T)^{-1}$

procedure $h(\mathbb{S}) =$ OP−COUNT($\mathbb{S}$, $\mathbb{G}$)
  Compute $\mathbf{D} = \mathbf{G} - \mathbf{S}$
  Compute $T^d = \mathbf{T} \times \mathbf{D}$ and $\tau = \mathbf{T}^d[1:r]$
  if $t_i^d \neq 0 \; \forall i \geq r+1$ **then** *No solution!*
  else return $\lceil \mathbf{Q} \times \mathbf{Z} \times \tau \rceil$

Domains	%-compliance	l_1-MILP	l_1-LP	$\omega - l_1$-LP	OP−COUNT
GED	34.29%	55.48%	55.48%	75.76%	55.48%
Blocks-3ops	31.25%	47.80%	47.80%	23.60%	52.60%
Blocks-4ops	19.64%	67.71%	67.71%	35.42%	67.71%
Visitall	-	-	-	-	-
GED	25.49%	37.61%	34.02%	53.36%	48.32%
Blocks-3ops	31.25%	47.80%	47.80%	23.60%	52.60%
Blocks-4ops	19.64%	67.71%	67.71%	35.42%	67.71%
Visitall	21.75%	28.41%	28.41%	44.37%	100.00%
Blocks-3ops	48.13%	28.68%	28.68%	44.38%	32.32%
Blocks-4ops	42.86%	56.25%	56.25%	12.50%	64.58%
8-puzzle	88.89%	33.33%	40.00%	46.67%	40.00%

minimization of l_1-*norm* (l_1-*LP*) or weighted l_1-*norm* (ω-l_1-*LP*) (Candès, Wakin, and Boyd 2008) (with the added constraint $\mathbf{x} \succeq 0$ to enforce positive integer solutions). For ω-l_1-*LP*, we empirically observe that rounding up the individual operator counts produce a more informed heuristic (even though it loses out on admissibility). Thus, we arrive at a more informed polynomial time proxy for integer solutions.

Evaluations.[1] The table shows the evaluation of the proposed heuristics across a total of 83 problems from five well-known unit cost planning domains. Each entry in the table represents the percentage difference in the initial state heuristic value and the optimal plan length averaged across the problems in each domain. The %-compliance column shows the average number of goal compliant predicates in the problems. Rows 1-3 show the performance of our heuristic on the original domains ('-' indicates that the heuristics could not be computed due to absence of any goal complaint variables). Rows 3-6 show the performance in domains where the %-compliance was increased (this was done by identifying instances in the action model where variables assume a don't care condition, i.e. a value of -1, and replacing it with appropriate values as entailed by domain axioms). Finally, rows 6-9 show the performance of our heuristics in problems with more completely specified goals (which results in higher percentage compliance). As expected, our heuristic performs better as %-compliance increases across a particular domain. The performance of l_1 LP and ω-l_1 LP highlights the usefulness of compressed sensing techniques in obtaining better integer approximations to the MILP.

Discussion and Related Work

Relation to Existing Heuristics. The proposed heuristic has close associations with both heuristics on state change equations and operator counts (Pommerening et al. 2014; Bonet, Van Den Briel, and others 2014; Van Den Briel et al. 2007). Specifically, compliant conditions capture the net change criteria very succinctly and are thus extremely useful where such properties are relevant. Another interesting connection to existing work is with respect to graph-plan based heuristics (Blum and Furst 1997), except here we are relaxing preconditions instead of delete effects.

Compliance. Our approach works better in domains that have many goal compliant conditions, e.g. in manufacturing domains (Nau, Gupta, and Regli 1995) or in puzzles like Sudoku (Babu et al. 2010). Thus goal completion strategies and semantic preserving actions have a direct effect on the quality of the heuristic. Intermediate representations such as transition normal form (TNF) (Pommerening and Helmert 2015) should be investigated in this context.

Landmarks. Our purpose here is not to compete with the most sophisticated heuristics of today but to motivate a special case that can be computed extremely efficiently. We discussed the simplest version of this formulation here, but it can be easily extended to incorporate more informative features like *landmarks* (Porteous, Sebastia, and Hoffmann 2001). A landmark constraint is added by simply subtracting the corresponding net change from $\mathbb{D}$: $d_i \leftarrow d_i - k_a \times (x_n - x_o)$ if $\langle d_i, x_o, x_n \rangle \in E_a$ and $a \in \mathcal{A}$ is an action landmark with cardinality k_a; and the closed form solution remains valid. In fact in terms of plan recognition with operator counts, observations are landmarks and the same approach applies. This demonstrates the flexibility of our approach.

Resource Constrained Interaction. The approach is especially relevant in the context of multi-agent interactions constrained by usage $\pi^\alpha(\eta)$ of a shared resource η by a plan π^α of an agent α. For example, in an adversarial setting, if an agent α_2 wanted to stop α_1 from executing its plan, all it needs to do is to ensure that $\exists \eta$ s.t. $\pi^{\alpha_1}(\eta) + \pi^{\alpha_2}(\eta) > |\eta|$. Similarly, in a cooperative setting, if agent α_2 wanted to ensure that α_1's plan succeeds, it would need to make sure that $\forall \eta \; \pi^{\alpha_1}(\eta) + \pi^{\alpha_2}(\eta) \leq |\eta|$. In fact, as resource variables are compliant, our approach may provide quick estimates of an agent's intent without computing the entire plan.

Acknowledgment. This research is supported in part by the ONR grants N00014-13-1-0176, N00014-13-1-0519 and N00014-15-1-2027, and ARO grant W911NF-13-1-0023.

References

[Babu et al. 2010] Babu, P.; Pelckmans, K.; Stoica, P.; and Li, J. 2010. Linear systems, sparse solutions, and sudoku. *Signal Processing Letters, IEEE* 17(1):40–42.

[Blum and Furst 1997] Blum, A. L., and Furst, M. L. 1997. Fast planning through planning graph analysis. *Artificial intelligence* 90(1):281–300.

[Bonet, Van Den Briel, and others 2014] Bonet, B.; Van Den Briel, M.; et al. 2014. Flow-based heuristics for optimal planning: Landmarks and merges. In *ICAPS*.

[Candès and Wakin 2008] Candès, E. J., and Wakin, M. B. 2008. An introduction to compressive sampling. *Signal Processing Magazine, IEEE* 25(2):21–30.

[Candès, Wakin, and Boyd 2008] Candès, E. J.; Wakin, M. B.; and Boyd, S. P. 2008. Enhancing sparsity by reweighted l_1 minimization. *Journal of Fourier analysis and applications*.

[Nau, Gupta, and Regli 1995] Nau, D. S.; Gupta, S. K.; and Regli, W. C. 1995. Ai planning versus manufacturing-operation planning: A case study. In *IJCAI*.

[Pommerening and Helmert 2015] Pommerening, F., and Helmert, M. 2015. A normal form for classical planning tasks. In *ICAPS*, 188–192.

[Pommerening et al. 2014] Pommerening, F.; Röger, G.; Helmert, M.; and Bonet, B. 2014. Lp-based heuristics for cost-optimal planning. In *ICAPS*.

[Porteous, Sebastia, and Hoffmann 2001] Porteous, J.; Sebastia, L.; and Hoffmann, J. 2001. On the extraction, ordering, and usage of landmarks in planning. In *ECP*, 37–48.

[Van Den Briel et al. 2007] Van Den Briel, M.; Benton, J.; Kambhampati, S.; and Vossen, T. 2007. An lp-based heuristic for optimal planning. In *CP*. Springer. 651–665.

[1] More extensive evaluations will be available at http://bit.ly/1sA5wM7.