

# Few-Shot Induction of Generalized Logical Concepts via Human Guidance

Mayukh Das<sup>1,\*</sup>, Nandini Ramanan<sup>2</sup>, Janardhan Rao Doppa<sup>3</sup> and Sriraam Natarajan<sup>2</sup>

<sup>1</sup>*Samsung R&D Institute India - Bangalore, Device Intelligence, Bangalore , India*

<sup>2</sup>*The University of Texas at Dallas, ECS, Dallas , Texas, USA*

<sup>3</sup>*Washington State University, EECS, Pullman, Washington, USA*

Correspondence\*:

Mayukh Das

mayukh.das@samsung.com

## 2 ABSTRACT

3 We consider the problem of learning generalized first-order representations of concepts from a  
4 small number of examples. We augment an inductive logic programming learner with two novel  
5 contributions. First, we define a distance measure between candidate concept representations  
6 that improves the efficiency of search for target concept and generalization. Second, we leverage  
7 richer human inputs in the form of advice to improve the sample-efficiency of learning. We prove  
8 that the proposed distance measure is semantically valid and use that to derive a PAC bound.  
9 Our experiments on diverse learning tasks demonstrate both the effectiveness and efficiency of  
10 our approach.

## 1 INTRODUCTION

11 We study the case of learning from *few examples*, of which *one-shot* learning is a special case (Lake et al.,  
12 2015). We consider a challenging setting – that of learning explainable, decomposable and generalizable  
13 (first-order) concepts from few examples. Plan induction becomes a special case where a generalizable  
14 plan is induced from a single (noise-free) demonstration. As an example, consider building a tower that  
15 requires learning *L-shapes* as a primitive. In our formulation, the goal is to learn a *L-shape* from a single  
16 demonstration. Subsequently, using this concept, the agent can learn to build a rectangular base (with 2  
17 *L-shapes*) from another single demonstration and so on till the tower is fully built. Concept learning has  
18 been considered as problem solving by reflection (Stroulia and Goel, 1994), mechanical compositional  
19 concepts (Wilson and Latombe, 1994), learning probabilistic programs (Lake et al., 2015), etc. While  
20 successful, they are not considered in one-shot learning except with SVM (Tax, 2001), or with a neural  
21 network (Kozerański and Turk, 2018).

22 Our work has two key differences. First, we aim to learn an ‘easily interpretable’, ‘explainable’,  
23 ‘decomposable’ and ‘generalizable’ concepts as first-order Horn clauses Horn (1951) (which are akin to  
24 If-Then rules). Second, and perhaps most important, we ‘do not assume the existence of a simulator (for  
25 plans) or employ a closed-world assumption’ to generate negative examples. Inspired by Mitchell’s [1997]  
26 observation of futility of bias-free learning, we employ domain expertise as inductive bias. The principle  
27 of structural risk minimization Vapnik (1999) shows how optimal generalization from extremely sparse  
28 observations is quite hard. The problem is harder in structured domains since most relations are false.

Thus, few-shot induction of generalized logical concepts is challenging. We employ iterative revision of first-order horn clause theories using a novel scoring metric and guidance from a human. *We emulate a ‘student’ who learns a generalized concept from an example provided by the ‘teacher’, by both reflecting as well as, asking relevant questions.*

We propose *Guided One-shot Concept Induction* (GOCI) for learning in relational domains<sup>1</sup>. GOCI builds upon an inductive logic program (ILP) learner Muggleton (1991) with two key extensions. First, a modified scoring function that explicitly computes distances between concept representations. We show the relation to *Normalized Compression Distance* (NCD) for plan induction settings. Consequently, we demonstrate that NCD is a valid distance metric. Second, we use domain knowledge from human expert as inductive bias. Unlike many advice taking systems that employ domain knowledge before training, GOCI identifies the relevant regions of the concept representation space and *actively solicits guidance* from the human expert to find the target concept in a sample-efficient manner. Overall, these two modifications allow for more effective and efficient learning using GOCI that we demonstrate both theoretically and empirically.

We make the following key contributions:

1. We derive a new distance-penalized scoring function that computes definitional distances between concepts, henceforth termed as ‘conceptual distance’.
2. We treat the human-advice as an inductive bias to accelerate learning. Our ILP learner actively solicits richer information from the human experts than mere labels.
3. Our theoretical analyses of GOCI prove that (a) our metric is indeed a valid distance, (b) NCD between plans is a special case of our metric
4. We show a PAC analysis of the learning algorithm based on Kolmogorov complexity.
5. We demonstrate the exponential gains in both sample efficiency and effectiveness of GOCI on diverse concept induction tasks with one or a few examples.

## 2 BACKGROUND AND RELATED WORK

Our approach to **Concept Learning** is closely related to Stroulia and Goel (1994)’s work which learns logical problem-solving concepts by reflection. GOCI’s scoring metric is more general and applicable to both concepts and plans and can be used for learning from a few examples. While we use discrete spatial structures as motivating examples, GOCI is not limited to discrete spaces, similar to Wilson’s Wilson and Latombe (1994) work. GOCI is also related in spirit to probabilistic (Bayesian) program induction for learning decomposable visual concepts Lake et al. (2015) which illustrates how exploiting decomposability is more effective. Our approach leverages not only decomposability but implicit relational structure as well.

### **One/few-Shot Learning and theory induction:**

Our problem setting differs from the above in that it requires learning from sparse examples (possibly one). Lake et al., (Lake et al., 2015) propose a one-shot version of Bayesian program induction of visual concepts. There is also substantial work on one/few-shot learning (both deep and shallow) in a traditional classification setting (Bart and Ullman, 2005; Vinyals et al., 2016; Wang et al., 2018), most of which either pre-train with gold-standard support example set or sample synthetic observations. We make no such assumptions about synthetic examples. ILP (Muggleton, 1991) inductively learns a logical program (first-order theory) that covers most of the positive examples and few of the negative examples by effectively employing background knowledge as search bias. In concept learning, generalization is typically performed as a search through space of candidate inductive hypotheses by (1) structuring, (2) searching and (3) constraining the

<sup>1</sup> Our algorithm can learn from one(few) example(s). We specify the number of examples in our evaluations.

space of theories. FOIL (Quinlan, 1990) is an early non-interactive learner with the disadvantage that it occasionally prunes some uncovered hypotheses. This is alleviated in systems like FOCL by introducing language-bias in form of user-defined constraints (Pazzani, 1992). With *Interactive ILP*, learner could pose questions and elicit expert advice which allows pruning large parts of search space (Sammur and Banerji, 1986; Rouveirol, 1992). To incorporate new incoming information, ILP systems with *theory revision*, incrementally refine and correct the induced theory (MUGGLETON, 1988; Sammut and Banerji, 1986). While GOCI is conceptually similar to ALEPH (Srinivasan, 2007), it learns from a few examples and actively acquires domain knowledge by interacting with human expert incrementally.

### Knowledge-Guided Learning:

Background knowledge in ILP is primarily used as search bias. Although the earliest form of knowledge injection can be found in explanation-based approaches (Shavlik and Towell, 1989), our work relates to preference-elicitation framework (Braziunas and Boutilier, 2006) which guides learning via human preferences as an inductive bias. Augmented learning with domain knowledge as an inductive bias has long been explored across various modeling formalisms, including traditional machine learning (Fung et al., 2003), probabilistic logic (Odom et al., 2015) and planning (Das et al., 2018). Our human-guided GOCI learner aims to extend these directions in the context of learning generalizable complex concepts from a few examples(including plans). Similar problem setting of concept learning from incomplete/sparse observations has also been explored in the cognitive science paradigm via explanation based inductive program synthesis (Kitzelmann and Schmid, 2006; Flener, 1997).

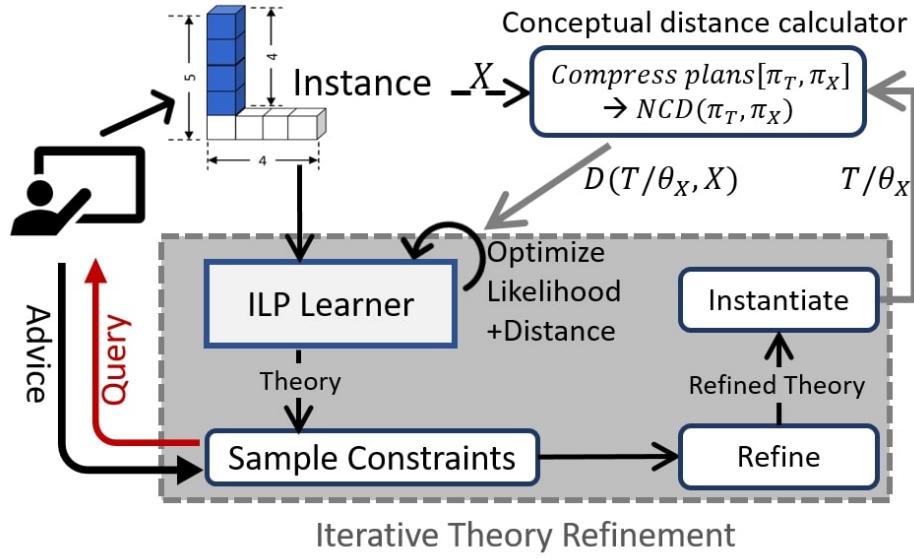
The idea of augmented learning with human guidance/knowledge has also been extensively studied in the context of evolutionary computation. Interactive Evolutionary systems (Eiben and Smith, 2015) use expert guidance to emulate a holistic fitness function that would otherwise depend on a very restricted pre-defined fitness model. The potential richness of such knowledge can be leveraged in not just evolutionary parent selection but can also optimize other parameters that leads to faster convergence, especially in mutations (Wendt et al., 2010). ILP has been shown to be conceptually similar to mutative EA in the context of program induction (Wong and Leung, 1997) and hence knowledge guided mutations are related to knowledge augmented search in ILP. Thus, in our problem setting, the interaction module that seeks human guidance to select the most useful constraints (detailed in section 3.2.3) is similar in spirit to interactive (knowledge guided) evolutionary mutation process. However, our underlying search strategy and optimization is based on ILP.

## 3 GUIDED ONE-SHOT CONCEPT INDUCTION

We are inspired by a teacher (human) and student (machine) setting in which a small number of demonstrations are used to learn generalized concepts (Chick, 2007). Intuitively, the description provided by a human teacher tends to be modular (can have distinct logical partitions), structured (entities and relations between them), and in terms of known concepts. Hence, a vectorized representation of examples is insufficient. We choose a logical representation, specifically a ‘function-free restricted form of first-order logic (FOL)’ that models structured spaces faithfully.

**Given:** A set of “facts” or assertions, i.e. a set of ground literals (or trajectories) describing 1 (or few) instance(s) of an unknown concept, availability of an expert to provide guidance and a knowledge-base of known concepts

**To Do:** Learning a representation, by inducing a first-order logic program, of the given concept that optimally generalizes the given instance(s) effectively



**Figure 1.** Highlevel overview of our GOCI framework.

107 The input to GOCI is the description of the instances(s) of a concept that the human teacher provides. An  
 108 example is, thus, conjunction of a set of ground literals (assertions). The output of GOCI is a least general  
 109 generalization (LGG) horn clause from the *input* example(s).


### 110 3.1 Concept Representation

111 Consider the following example input to the GOCI framework. The input here is an instance of the  
 112 structural concept  $\mathbb{L}$  (illustrated in Figure 2),

113 **EXAMPLE 1.** An instance in a minecraft domain, can be a  $\mathbb{L}$  with dimensions *height* = 5, *base* = 4  
 114 (Figure 2).  $\mathbb{L}(S), \text{Height}(S, 5), \text{Base}(S, 4)$ ,  $s$  is the concept identifier and may be described as conjunction  
 115 of ground literals,  
 116  $\text{Row}(A) \wedge \text{Tower}(B) \wedge \text{Width}(A, 4) \wedge \text{Height}(S, 5) \wedge \text{Base}(S, 4) \wedge \text{Contains}(S, A) \wedge \text{Contains}(S, B) \wedge$   
 117  $\text{Height}(B, 4) \wedge \text{SpRel}(B, A, \text{'NWTop'})$ ;  
 118 which denotes  $\mathbb{L}$  as composition of a ‘Row’ of  $w = 4$  and a ‘Tower’ of  $h = 4$  with appropriate literals  
 119 describing the scenario (Figure 2 left). As a special case, under partial or total ordering assumptions  
 120 among the ground literals, an input instance can represent a plan demonstration.

121 We aim to learn the optimally generalized (decomposable) representation of the concept ( $\mathbb{L}$  in the context  
 122 of the aforementioned example) referred by the one/few instances that were passed to GOCI as input.  
 123 Before further discussion on the learning such a generalized (decomposable) representation let us first  
 124 define formally what a concept representation signifies in our setting.

125 **DEFINITION 1.** Concepts in GOCI are represented as horn clause theories. A theory  $T$  is defined as,  
 126  $T = \mathcal{C}(s^k \dots) : - \bigvee [\bigwedge_{i=1}^N f_i(t_1, \dots, t_j)]$ , where the body  $\bigwedge_{i=1}^N f_i(t_1, \dots, t_j)$  is a conjunction of literals  
 127 indicating known concepts or relationships among them, such that any  $t_j \in V \cup \{s^k\} \cup C$  where  $V$  is the  
 128 set of all logical variables in the clause,  $C$  is the set of constants in the domain of any logical variable.  
 129 The head  $\mathcal{C}(s^k \dots)$  identifies a target concept, and the terms  $\{s^k\}$  are logical variables that denote the  
 130 parameters of the concept assuming there are  $k = \{1, \dots, K\}$  parameters including the identifier to the  
 131 given instance of the concept. Since a concept can be described in multiple ways (Figure 2), the final



**Figure 2.** Concept  $\mathbb{L}$  ( $base = 4$ ,  $height = 5$ ), described as composition of a **Tower** and a **Row**

theory will be a disjunction over clause bodies with the same head. A (partial) instantiation of a theory  $T$  is denoted as  $T/\theta$ .

Note that these definitions allow for the reuse of concepts, potentially in a hierarchical fashion. We believe that this is *crucial in achieving human-agent collaboration*.

EXAMPLE 2. Figure 2 illustrates an instance of the concept  $\mathbb{L}$  that can be described in multiple ways. A possible one is,

$$\begin{aligned} \mathbb{L}(s) : & -[\text{Height}(s, h_s), \text{Base}(s, w_s), \text{Contains}(s, a), \text{Contains}(s, b), \text{Row}(a), \text{Tower}(b), \\ & \text{Width}(a, w_a), \text{Height}(b, h_b), \text{Equal}(w_s, w_a), \text{Sub}(h_b, h_s, 1), \text{SpRel}(b, a, \text{"NWTop"})] \\ \vee & [\text{Height}(s, h_s), \text{Base}(s, w_s), \text{Contains}(s, a), \text{Contains}(s, b), \text{Row}(a), \text{Tower}(b), \\ & \text{Width}(a, w_a), \text{Height}(b, h_b), \text{Equal}(h_s, h_b), \text{Sub}(w_a, w_s, 1), \text{SpRel}(b, s, \text{"W"})] \end{aligned}$$

The generalization must be noted. The last argument of the  $\text{SpRel}()$  is a constant, as only this particular spatial alignment is appropriate for the concept of the  $\mathbb{L}$  structure. Although the input is a single instance (Example 1), GOCI should learn a generalized representation such as Example 2. Another interesting aspect are the additional constraints:  $\text{Equal}(X, Y)$  and  $\text{Sub}(X, Y, N)$ . While such predicates are a part of the language, they are not typically described directly in the input examples. However, they are key to generalization, since they express complex interactions between numerical (or non-numerical) parameters. Also note that the head predicate of the clause could have been designed differently as per Definition 1. For instance in case of Example 2 the head predicate could have folded in the dimensional parameters -  $\mathbb{L}(s, h_s, w_s)$ . However, the number of such dimensional parameters can vary across different concepts. Hence to maintain generality of representation format during implementation we push the dimensional parameters of the learnable concept into the body of the clause.

A specific case of our concept learning (horn clause induction) framework could be plan induction from sparse demonstrations. This can be achieved by specifying time as the last argument of both the state and action predicates. Following this definition, we can allow for plan induction as shown in our experiments. Our novel conceptual distance is clearer and more intuitive in the case of plans as can be seen later.



153 **DEFINITION 2 (Decomposable:).** *A concept  $\mathcal{C}$  is Decomposable iff it is expressed as a conjunction of*  
 154 *other concepts, and one or more additional literals to model the interactions.  $\mathcal{C} \Leftarrow (\bigwedge_i \mathcal{C}'_i) \wedge (\bigwedge_j B_j)$ .*  
 155 *Here  $\mathcal{C}'_i$  are literals that represent other concepts that are already present in the knowledge base of the*  
 156 *learner and  $B_j$  are literals that, either describe the attributes of  $\mathcal{C}'_i$  or interactions between them.*

157 Decomposable allows for an unknown concept to be constructed as a composition of other known  
 158 concepts. GOCI learns the class of decomposable concepts since it is intuitive for the “human teacher” to  
 159 describe. Decomposable concepts faithfully capture the *modular* and *structured* aspect of how humans  
 160 would understand and describe instances. It also allows for a hierarchical construction of plans.

161 **EXAMPLE 3.** *Following the Minecraft structure described in 2, note how  $\mathbb{L}$  is described with*  
 162 *already known concepts  $\mathcal{C}'_1 = \text{Row}()$  and  $\mathcal{C}'_2 = \text{Tower}()$  and the other literals such as*  
 163  *$\text{Height}(b, h_b), \text{SpRel}(b, a, \text{”NWTop”}), \dots \in \{B_j\}$ , i.e. they describe the parameters of the known*  
 164 *concepts or interactions between them. Note that, known concepts in the knowledge base could have*  
 165 *been manually coded in by experts or learned previously and are essentially represented in the same way.*  
 166 *For instance  $\text{Row}()$  can be encoded as recursive the clause program representing a composition of one*  
 167 *block and one unit shorter row,*


$$\text{Row}(r) : -[\text{Width}(r, w_r), \text{Block}(a), \text{Row}(b), \text{Width}(b, w_b), \text{SpRel}(a, b, \text{”East”}), \text{Sub}(w_b, w_r, 1)] \\ \bigvee [\text{Width}(r, w_r), \text{Equal}(w_r, 1), \text{Block}(a)]$$

168  *$\text{Tower}()$  could also be defined in the knowledge base in the same way. When the optimally general*  
 169 *representation of the concept  $\mathbb{L}$  is learned that is persisted in the knowledge base as well, such that more*  
 170 *complex concepts can be represented by decomposing into  $\mathbb{L}$  and other known concepts.*

171 An obvious question that arises here is why  $\{B_j\} \not\subseteq \{\mathcal{C}'\}$ ? i.e. why can the other literals not be treated  
 172 similarly as a part of the knowledge base of known concepts? Ideally, that would be correct. However, that  
 173 will also cause infinite levels of concept definitions which cannot be implemented in practice. Additionally,  
 174 following the paradigm of a student-teacher scenario, it can always be assumed that the student has prior  
 175 understanding of many concepts from outside the current system. Thus we can safely assume, without loss  
 176 of generality, that set of literals  $\{B_j\}$  are implicitly understood and defined as a part of the framework  
 177 itself. This argument applies to the semantics of the ‘constraint predicates’ (described later) as well.

178 Finally, before we discuss the details of the learning methodology let us briefly look into a motivating,  
 179 and presently relevant, real-world scenario that represents our problem setting.

180 **EXAMPLE 4.** *Consider a Decision Support AI system for resource planning and management in hospitals*  
 181 *as illustrated in Figure 3. The AI agent forecasts the need for increased resources in the Infectious Diseases*  
 182 *(ID) ward, given the early signs of an outbreak of some disease such as Covid-19 or Ebola etc. and a*  
 183 *potential spike in ID ward admissions. However, as noted by the administrators and/or physicians there*  
 184 *isn’t enough budget to procure additional resources for ID ward. But the problem is quite critical and*  
 185 *needs to be solved. So the human teacher (administrators in this case) teaches the AI agent the concept of*  
 186 *“Divert”-ing resources from Cancer ward since cancer ward admissions are usually stable and does not*  
 187 *have spikes. The AI agent is hence expected to learn a generalized representation of the concept of divert*



**Figure 3.** A motivating real-world scenario for concept induction. The concept learnt by the AI agent is “Divert()”

188 such that it may be applied later for other wards or for other tasks and furthermore in a “decomposable”  
 189 fashion. For instance, “Divert” may be learned as a clause program such as,

$$\text{Divert}(R, \text{qty}_R) : \neg \text{To}(R, \text{Loc}_{\text{dest}}), \text{AcquireFrom}(R, \text{Loc}_{\text{source}}^1, \text{qty}_1, ),$$

$$\text{AcquireFrom}(R, \text{Loc}_{\text{source}}^2, \text{qty}_2), \text{AssignTo}(R, \text{Loc}_{\text{dest}}, \text{qty}_{\text{dest}}),$$

$$\text{sum}(\text{qty}_1, \text{qty}_2, \text{qty}_{\text{dest}})$$

190 Obviously the above representation assumes that concepts such as “AcquireFrom()” are known concepts,  
 191 either implicitly defined inside the learning framework or its explicit representation has been learned and  
 192 persisted inside the knowledge base in the past.

193 The above example is solely to motivate the potential impact of our problem setting and the proposed  
 194 solution. For an explanation of different components and aspects of GOCI we refer to the much simpler  
 195 and unambiguous structural example outlines earlier (ℒ).

## 196 3.2 Methodology

### 197 3.2.1 Search

198 ILP systems perform a greedy search through the space of possible theories. Space is typically defined,  
 199 declaratively, by a set of mode definitions Muggleton (1995) that guide the search. We start with the most  
 200 specific clause (known as a bottom clause) (Srinivasan, 2007) from the ground assertions and successively  
 201 add/modify literals that might improve a rule that best explains the domain. Typically, the best theory  
 202 is the one that covers the most positive and least negative examples. Thus, it optimizes the likelihood  
 203 of a theory  $T$  based on the data. We start with a bottom clause and variablize the statements via anti-  
 204 substitution. Variabilization of  $T$  is denoted by  $\theta^{-1} = \{a/x\}$  where  $a \in \text{consts}(T)$ ,  $x \notin \text{vars}(T)$ . That  
 205 is, anti-substitution  $\theta^{-1}$ , is a mapping from occurrences of ground terms in  $T$  to new or existing logical  
 206 variables.

207 **Evaluation Score:** We redesign the ILP scoring (e.g., ALEPH’s compression heuristics) as:

- 208 • The user-provided advice forces the learner to learn longer theory, hence the *search space can be*  
 209 *exponentially large*. Thus, modes alone are not sufficient as the search bias.
- 210 • There is only one (a few) positive training example(s) to learn from and *many possible rules can*  
 211 *accurately match the training example*. Coverage based scores fail.

212 Most learners optimize some form of likelihood. For a candidate theory  $T$ , likelihood given data  $\mathbb{D}$  is,  
 213  $LL(T) = \log P(\mathbb{D}|T)$  (i.e., coverage). To elaborate further, in most classification tasks in discrete domains  
 214 (with categorical/ordinal feature and target variables), goodness of fit of candidate models is achieved via  
 215 the measure of how well the candidate models explain (or cover) the given data i.e. a good model is the one  
 216 that will predict positive class for maximum possible positive examples and for minimum possible negative  
 217 examples. This measure is expressed as likelihood of the data given a candidate model. In GOCI, we have  
 218 one (at most few) positive example(s). Coverage will not suffice. Hence, we define a modified objective as  
 219 follows.

$$T^* = \arg \min_{T \in \tau} (-LL(T) + D(T/\theta_X, X)) \quad (1)$$

220 where  $T^*$  is the optimal theory,  $\tau$  is the set of all candidate theories, and  $D$  is the conceptual distance  
 221 between the instantiated candidate theory  $T/\theta_X$  and the original example  $X$ . Recall that a theory  $\mathbb{T}$  is a  
 222 disjunction of horn clause bodies (or conjunction of clauses).

223

### 224 3.2.2 Distance metric

225 Conceptual distance,  $D(T/\theta_X, X)$ , is a penalty in our objective. The key idea is that any learned  
 226 first-order horn clause theory must recover the given instance by equivalent substitution. However, syntactic  
 227 measures, such as edit distance, are not sufficient since changing even a *single literal*, especially, literals  
 228 that indicate inter-concept relations, could potentially result in a completely different concept. For instance,  
 229 in blocks-world, the difference between a block being in the middle of a row and one at the end of the row  
 230 can be encoded by changing one literal. Hence, a more sophisticated semantic distance such as conceptual  
 231 distance is necessary Friend et al. (2018). However, such distances require deeper understanding of the  
 232 domain and its structure.

233 Our solution is to employ *inter-plan distances*. Recall that the concepts GOCI can induce are  
 234 decomposable and, hence, are equivalent to parameterized planning tasks. One of our key contributions  
 235 is to exploit this equivalence by using a domain-independent planner to find grounded *plans* for both  
 236 the theory learned at a particular iteration  $i$ ,  $T_i$  and the instance given as input,  $X$ . We then compute the  
 237 Normalized Compression Distance (NCD) between the plans.

238 **NCD:** Goldman & Kuter Goldman and Kuter (2015) proved that NCD is arguably the most robust inter-plan  
 239 distance metric. NCD is a reasonable approximation of *Normalized Information Distance*, which is not  
 240 computable Vitányi et al. (2009). Let the plans for  $T_i/\theta_X$  and  $X$  be  $\pi_T$  and  $\pi_X$ . To obtain NCD, we execute  
 241 string compression (lossy or lossless) on each of the plans as well as the concatenation of the two plans to  
 242 recover the compressed strings  $C_T$ ,  $C_X$ , and  $C_{T,X}$  respectively. NCD between the plans can be computed  
 243 as,

$$NCD(\pi_T, \pi_X) = \frac{C_{T,X} - \min(C_T, C_X)}{\max(C_T, C_X)} \quad (2)$$



The conceptual distance between a theory  $T$  and  $X$  is the NCD between the respective plans,  $D(T/\theta_X, X) = NCD(\pi_T, \pi_X)$ . This entire computation is performed by the *Conceptual distance calculator* as shown in Figure 1.

**Observations (1)** Conceptual distance as a penalty term in the  $LL$  score ensures that the learned theory will correctly recover the given example/demonstration. **(2)**  $D(T/\theta_X, X)$  generalizes to the *Kolmogorov-Smirnov* statistic between two target distributions if we induce probabilistic logic theories. We prove these insights theoretically.

### 3.2.3 Human Guidance:

The search space in ILP is provably infinite. Typically language-bias (modes) and model assumptions (closed world) are used to prune the search space. However, it is still intractable with one (or few) examples. So, we employ human expert guidance as constraints that can directly refine an induced theory, acting as a strong inductive bias. Also, we are learning decomposable concepts (see Def. 2). This allows us to exploit another interesting property. Constraints can now be applied over the attributes of the known concepts that compose the target concept, or over the relations between them. Thus, GOCI directly includes such constraints in the clauses as literals (see Ex. 2). Though such constraint literals come from the pre-declared language, they are not directly observed in the input example(s). So an ILP learner will fail to include such literals.

If the human inputs (constraints) are provided upfront before learning, it can be wasteful/irrelevant. More importantly, it places an additional burden on the human. To alleviate this, *our framework explicitly queries for human advice on the relevant constraint literals, which are most useful*. Let  $\mathbb{U}$  be a predefined library of constraint predicates in the language, and let  $\mathcal{U}() \in \mathbb{U}$  be a relevant constraint literal. Human advice  $\mathcal{A}$  can be viewed as a preference over the set of relevant constraints  $\{\mathcal{U}()\}$ . If  $\mathcal{U}_{\mathcal{A}}$  denotes the preferred set of constraints, then we denote the theory having a preferred constraint literal in the body of a clause as  $\tau_{\mathcal{A}}$ . (For instance, as per Example 2 GOCI queries ‘which of the two sampled constraints  $\text{Sub}(\mathbf{h}_b, \mathbf{h}_s, 1)$  &  $\text{Greater}(\mathbf{h}_b, \mathbf{h}_s)$  is more useful’. Human could prefer  $\text{Sub}(\mathbf{h}_b, \mathbf{h}_s, 1)$ , since it subsumes the other.) The scoring function now becomes:

$$T^* = \arg \min_{T \in \tau} (-LL(T) + D(T/\theta_X, X)) : \tau \subseteq \{\tau_{\mathcal{A}}\} \quad (3)$$

Thus, we are optimizing the constrained form of the same objective as Equation 1 which aims to prune the search space. This is inspired by advice elicitation approaches Odom et al. (2015). While our framework can incorporate different forms of advice, we focus on preference over constraints on the logical variables. The formal algorithm, described next, illustrates how we achieve this via an iterative greedy refinement (Figure 1, query-advice loop shown in left).

### 3.3 The Goci Algorithm

Algorithm 1 outlines the GOCI framework. It initializes a theory  $T_0$ , by variablizing the ‘bottom clause’ obtained from  $X$  and background knowledge [lines 3 & 5]. Then it performs a standard ILP search (described earlier) to propose a candidate theory [line 6]. This is followed by the guided refinement steps, where constraint literals are sampled (parameter tying guides the sampling) and the *human teacher* is queried for preference over them, such that the candidate theory can be modified using preferred constraints [lines 7–9]. The function  $NCD()$  performs the computation of the conceptual distance, by first grounding

**Algorithm 1** Guided One-shot Concept Induction


---

```

1: procedure GOCI(Instance  $X$ )
2:   Initialize: Set Iteration  $\ell \leftarrow 1$ 
3:   Initialize: Bootstrap theory  $T_0 \leftarrow X/\theta^{-1}$ 
4:   repeat
5:     Use  $T_{\ell-1}$  as initial model
6:     Candidate theory  $T_\ell \leftarrow \text{SEARCH}(T \in \tau|T_{\ell-1})$ 
7:     Sample applicable constraints  $\mathcal{U} \in \mathbb{U}$ 
8:      $\mathcal{U}_A \leftarrow \text{QUERY}(\text{human}, \mathcal{U})$ 
9:      $T' \leftarrow T_\ell \oplus \mathcal{U}_A$   $\triangleright \forall \mathcal{U}_A \in \mathcal{A}$ 
10:     $D_\ell(T'/\theta_X, X) \leftarrow \text{NCD}(\pi_{T'/\theta_X}, \pi_X)$ 
11:    Score  $\mathcal{S}_\ell \leftarrow (-LL(T') + D(T'/\theta_X, X))$ 
12:    if  $\mathcal{S}_\ell < \mathcal{S}_{\ell-1}$  then  $\triangleright$  minimize
13:      Retain  $T'$ : Update  $T_\ell = T'$ 
14:    end if
15:  until  $\ell \leq L$  OR  $T_\ell = T_{\ell-1}$ 
16: end procedure

```

---

**Figure 4.** Difference in evaluation of a concept instance across different learning paradigms

the current modified candidate theory  $T'$  with the same parameter values as the input example  $X$ , then generating grounded plans and finally calculating the normalized compression distance between the plan strings (as shown in Figure 1 and Equation 2) [line 10]. The distance-penalized negative log-likelihood score is estimated and minimized to find the best theory at the current iteration [lines 11–14], which is then used as the initial model in the next iteration. This process is repeated either until convergence (no change in induced theory) or maximum iteration bound ( $L$ ).

**Connection to plan induction:** Evaluation, both in traditional ML and ILP, generally predicts the value of  $\hat{y}_X$  for a test instance  $X$  represented as a fixed (ML) or arbitrary (ILP) length feature vector. In GOCI, however, the notion of evaluation of an instance  $X$  depends on successful construction of a valid/correct plan  $\pi_X$  (Figure 4). Thus while learning, most research aim to maximize coverage of positive instances  $E^+$  ( $\max P(\hat{y}_x = \text{true} | y_x \in E^+)$ ) and minimize coverage of negatives  $E^-$ , (i.e.  $\min P(\hat{y}_x = \text{true} | y_x \in E^-)$ ). GOCI evaluates a candidate concept representation by allowing the agent to realize that concept - by computing a valid plan for the goal/task implied by the instance  $x$ . This is akin to plan induction, since we are learning parameterized plan for realizing the concept as a surrogate for the concept itself. Additionally, planning has long been shown to be conceptually same as logic programming (Preiss and Shai, 1989) and hence induction of logic programs (theories) is the same as plan induction where the examples are trajectories (plan traces) in this case.

300

### 3.4 Theoretical analysis

#### 3.4.1 Validity of Distance Metric

NCD  $\delta(x, y)$  between 2 strings  $x$  and  $y$  is provably a valid distance metric Vitányi et al. (2009) -  $\delta(x, y) = \frac{\max K(x|y), K(y|x)}{\max K(x), K(y)}$ , where  $K(x)$  is the Kolmogorov complexity of a string  $x$  and  $K(x|y)$  is the conditional Kolmogorov complexity of  $x$  given another string  $y$ . NCD is a computable approximation of the same [ $D(x, y) \approx \delta(x, y)$ ]. Thus, we just verify if  $\delta$  is a correct conceptual distance measure. Let  $T_Y$  and  $T_Z$  be 2 theories, with same parameterizations (i.e., same heads). Let  $T_Y/\theta$  and  $T_Z/\theta$  be their groundings with identical parameter values  $\theta$ . Our learned theories are equivalent to planning tasks. Assuming access to a planner  $\Pi()$  which returns  $Y = \Pi(T_Y/\theta)$  and  $Z = \Pi(T_Z/\theta)$ , the two plan strings w.r.t the instantiations of concepts represented by  $T_Y$  and  $T_Z$  respectively.

**PROPOSITION 1 (Valid Conceptual Distance).** *Normalized information distance  $\delta(Y, Z)$  is a valid and sound conceptual distance measure between  $T_Y$  and  $T_Z$ , i.e.,  $\delta(Y, Z) = 0$  iff the concepts represented by  $T_Y$  and  $T_Z$  are equivalent.*

**PROOF SKETCH FOR PROPOSITION 1.** Let  $T_Y$  and  $T_Z$  be 2 induced consistent first-order Horn clause theories, that may or may not represent the same concept. Let  $\theta$  be some substitution. Now let  $T_Y/\theta$  and  $T_Z/\theta$  be the grounded theories under the same substitution. This is valid since we are learning horn clause theories with the same head, that indicates the target concept being learned. As explained in the manuscript a theory is equivalent to a planning task. We assume access to a planner  $\Pi()$ , we get plan strings  $Y = \Pi(T_Y/\theta)$  and  $Z = \Pi(T_Z/\theta)$  with respect to the planing tasks  $T_Y/\theta$  and  $T_Z/\theta$ .

Friend et al.,(2018) proved that *Conceptual Distance* is the step distance between 2 consistent theories in a cluster network  $(\mathbb{T}, \rightleftharpoons, \sim)$ , where  $\mathbb{T}$  is the class of consistent theories,  $\rightleftharpoons$  is the definitional equivalence relation (equivalence over bidirectional concept extensions) and  $\sim$  implies symmetry relation. We have shown in the paper that, given the class of concepts we focus on, *a concept is a planning task*.

Let there be a theory  $T^*$ , which represents the optimal generalization of a concept  $\mathcal{C}$ . If step distance  $\langle T_Y, T^* \rangle = 0$  in a cluster network and  $\langle T_Z, T^* \rangle = 0$  then  $\langle T_Y, T_Z \rangle = 0$ , i.e. they represent the same concept  $\mathcal{C}$  and they are definitionally equivalent  $T_Y \rightleftharpoons T_Z$ . Thus both  $T_Y/\theta$  and  $T_Z/\theta$  will generate the same set of plans as  $T^*$ , since they will denote the same planning tasks (By structural induction). Thus,

$$T_Y \rightleftharpoons T_Z \iff [\Pi(Y) \cap \Pi(Z) = \Pi(Y) = \Pi(Z)] \quad (4)$$

upto equivalence of partial ordering in planning. let  $\pi^*$  be a minimum length plan in a set of plans  $\Pi()$ . Let  $y$  and  $z$  be strings indicating plans  $\pi^*(Y)$  and  $\pi^*(Z)$  ignoring partial order. If  $\Pi(Y) = \Pi(Z)$  then  $\pi^*(Y) = \pi^*(Z)$ . Hence, the conditional Kolmogorov complexities  $K(y|z)$  and  $K(z|y)$  will both be set to 0, if the strings  $x$  and  $y$  are equivalent (ignoring partial ordering). This is based on the principle that if they are equivalent, then a Universal prefix-Turing machine will recover one string given the other in 0 steps.

$$\therefore \frac{\max(K(y|z), K(z|y))}{\max(K(y), K(z))} = 0 = \delta(Y, Z)$$

**PROPOSITION 2 (Generalization to Kolmogorov-Smirnov).** *In generalized probabilistic logic, following Vitányi Vitányi (2013),  $\delta(Y, Z)$  corresponds to two-sample **Kolmogorov-Smirnov statistic** between two random variables  $T_Y/\theta$  and  $T_Z/\theta$  with distributions  $P_{T_Y}$  and  $P_{T_Z}$  respectively [ $v(T_Y, T_Z) =$*

336  $\sup_{\theta \in \mathcal{F}} |F_{T_Y}(\theta) - F_{T_Z}(\theta)|$ , where  $F_T()$  is the cumulative distribution function for  $P_T$  and  $\sup_{\theta \in \mathcal{F}}$  is the  
 337 supremum operator. In a deterministic setting,  $\delta$  is a special case of the Kolmogorov-Smirnov statistic  $v$ ,  
 338  $\delta(Y, Z) \preceq v(F_{T_Y}, F_{T_Z})$ .

339 PROOF SKETCH FOR PROPOSITION 2. This can be proved by considering the connection between  
 340 NID and the distributions induced by the concept classes we are learning. NID is defined as  $\delta(x, y) =$   
 341  $\frac{\max K(x|y), K(y|x)}{\max K(x), K(y)}$ , where,  $K(a|b)$  is the conditional Kolmogorov complexity of a string  $a$ , given  $b$ . There  
 342 is no provable equivalence between Kolmogorov complexity and traditional notions of probability  
 343 distributions.

344 However, if we consider a *reference* universal semi-computable semi-probability mass function  $\mathbf{m}(x)$ ,  
 345 then there is a provable equivalence  $-\log \mathbf{m}(x) = K(x) \pm O(1)$ . Similarly for conditional Kolmogorov  
 346 complexity, by Conditional Coding Theorem,  $-\log \mathbf{m}(y|x) = K(y|x) \pm O(1)$  Vitányi (2013). By  
 347 definition,

$$m(y|x) = \sum_{j \geq 1} 2^{-K(j)-c_j} P_j(y|x)$$

348 where  $c_j > 0$  are constants and  $P_j(y|x)$  is the lower semi-computable conditional. A lower semi-  
 349 computable semi-probability conditional mass function is based on the string generating complexity of a  
 350 universal prefix Turing machine. Thus  $m(y|x)$  is greater than all the lower semi-computable. Note that our  
 351 compressed plans are equivalent to a string generated by Universal Turing Machines. The conditional case  
 352 implies, if a compressed plan string  $x$  is given as an auxiliary prefix tape, how complex it is to generate  
 353 compressed string  $y = \theta$ .

354 Given 2 grounded theories  $T_Y/\theta$  and  $T_Z/\theta$ , let  $P_{T_Y/\theta}$ ,  $P_{T_Z/\theta}$  be the respective distributions when  
 355 learning probabilistic logic rules. Now let us define the semantics of a distribution  $P_{T/\theta}$  in our case:  $P_{T/\theta} =$   
 356  $P(\pi(T/\theta))$ , i.e. distribution over the plan strings, which can be considered as lower semi-computable  
 357 probability based on coding theory. We know,

$$\sum_{j \geq 1} 2^{-K(j)-c_j} P_j(y) \approx F(y|x) \quad (5)$$

358 where  $F(y)$  is the cumulative distribution. So, NID  $\delta(Y, Z)$  now becomes,  $\delta(Y, Z) = \frac{\max(K(y|z), K(z|y))}{\max(K(y), K(z))}$   
 359 We know that  $\max(K(y), K(z))$  is a normalizer. Thus,  $\delta(Y, Z) < \max(K(y|z), K(z|y))$

$$\begin{aligned}
\max(K(y|z), K(z|y)) &= \max(-\log m(y|z), -\log m(z|y)) \\
&= \max\left(-\log \frac{m(y, z)}{m(z)}, -\log \frac{m(y, z)}{m(y)}\right) \\
&= \max([-\log m(y, z) + \log m(z)], [-\log m(y, z) + \log m(y)])
\end{aligned}$$

Under partial ordering max yields supremum

$$\begin{aligned}
&\approx \sup |\log m(y) - \log m(z)| \\
&\approx \sup |\log F(y) - \log F(z)| \\
&\approx \sup |F(y) - F(z)| \text{ log is monotonic}
\end{aligned}$$

**Significance of Propositions 1 & 2:** Proposition 1 outlines how our proposed NCD-based metric is a valid conceptual distance. It is well understood that the true measure of conceptual distance is not straightforward and is subject to the semantic interpretation of the domain itself. But designing a unique distance metric based on the semantics of every domain limits the generality of any learning system. So NCD acts as a surrogate ‘conceptual distance’. It is based on the notion that “if two concepts are fundamentally same the complexity of optimal action plans to realize the concepts should also be fundamentally same”. NCD (or NID) essentially measures the difference in generative complexities of two plans. Also note that other types of distances that are limited to a syntactic level such as edit distance (or literal distance) will fail to capture the similarity or diversity between concept representations since the same concept can be represented with more than one theories which may vary in one or more literals.

Proposition 2 on the other hand, proves that our proposed metric is not limited to our specific scenario. It positions our work in the context of known statistical distance metrics and establishes its credibility as a valid solution. It proves how in a non-deterministic setting, i.e. probabilistic logic formulation, our proposed metric generalizes to Kolmogorov-Smirnov statistic.

### 3.4.2 PAC Learnability

PAC analysis of GOCI follows from GOLEM for function-free *horn* clause induction Muggleton and Feng (1990). Let initial hypothesis space be  $\mathcal{H}_0$  and the final be  $\mathcal{H}^*$  (s.t.  $T^* \in \mathcal{H}^*$ ).

**PROPOSITION 3 (Sample complexity).** *Following Valiant (1984) and Mooney (1994), with probability  $(1 - \delta)$ , the sample complexity of inducing the optimal theory  $T^*$  is:*

$$n^* = \mathcal{O}\left(\frac{1}{\epsilon} \left[ d^L j^i \ln((t f m)) + \ln\left(\frac{1}{\delta}\right) \right]\right) \quad (6)$$

where  $\epsilon$  is the regret,  $n^*$  - sample complexity of  $\mathcal{H}^*$ ,  $i$  is the maximum depth of a variable in a clause and  $j$  is the maximum arity.  $m$  - number of distinct predicates,  $t$  is the number of terms,  $p$  is the place and  $d$  is the distance of the current revision from the last known consistent theory, and  $L$  is the upper bound on the number of refinement steps (iterations).

**PROOF SKETCH FOR PROPOSITION 3.** *In our learning setting the learned theory will always have non-zero uncertainty.* To understand the properties, we build upon the PAC analysis for recursive *rlgg* (Relative Least General Generalization) approach for function-free *Horn* clause learning shown by Muggleton & Feng (1990) in GOELM. With some restrictions, it applies here as well. Let,  $n$ : denote the sample size and



387  $\mathcal{H}$ : the hypothesis space. Our approach can be considered as an *rlgg* approach with refinement steps. Note  
 388 that constraint predicates that refine the clauses are not part of  $\mathcal{K}$ .

389 To begin with, we are interested in regret bounds for the initially learned hypothesis by the ILP learner  
 390  $\mathcal{H}_0$ , before refinement. We know from Valiant (1984), that with probability  $(1 - \delta)$  the sample complexity  
 391  $n$  for  $\mathcal{H}_0$  is,

$$n \geq \frac{1}{\epsilon} \left( \ln(|\mathcal{H}_0|) + \ln\left(\frac{1}{\delta}\right) \right) \quad (7)$$

392 where  $\epsilon$  is the regret. Now, our ILP learner induces *ij*-Determinate clauses Muggleton and Feng (1990),  
 393 where  $i$  is the maximum depth of the clause and  $j$  is the maximum arity. In our problem setting it can be  
 394 proven that  $|\mathcal{H}_0| = \mathcal{O}((tpm)^{j^i})$ , where  $m$  is the number of distinct predicates in the language.  $t$  is the  
 395 number of terms, and  $p$  is the place (for details about place refer Muggleton and Feng (1990)). Also note  
 396 that, (if  $j$  &  $i$  is bounded : then  $j^i \leq c$ ). Mooney (1994) shows that for theory refinement/revision sample  
 397 complexity is expressed as ,

$$n^* = \mathcal{O} \left( \frac{1}{\epsilon} \left[ d^k \ln(|\mathcal{H}_0| + d + m) + \ln\left(\frac{1}{\delta}\right) \right] \right) \quad (8)$$

398 where distance  $d$  to be the number of single literal changes in a single refinement step and  $k$  is the number  
 399 of refinement/revision iterations. In Algorithm 1, we observe that at each iteration  $\ell \leq L$ , updates are w.r.t.  
 400 the preferred constraint predicates  $\mathcal{U}_A \in \mathbb{U}$ . Thus we know that  $k = L$ . Substituting  $|\mathcal{H}_0| = (tfm)^{j^i}$  and  
 401  $j^i = c$  constant) in Equation 8 and ignoring the additive terms  $d + m$  since  $(tfm)^{j^i} \gg d + m$ , we get,

$$n = \mathcal{O} \left( \frac{1}{\epsilon} \left[ d^L c \ln(tpm) + \ln\left(\frac{1}{\delta}\right) \right] \right) \quad (9)$$

402 **PROPOSITION 4 (Refinement distance).**  *$d$  is upper bounded by the expected number of literals that*  
 403 *can be constructed out of the library of constraint predicates with human advice  $\mathbb{E}_{\sim \mathcal{A}}[|\mathbb{U}|]$  and lower*  
 404 *bounded by the conceptual distance between theory learned at two consecutive iterations since we adopt*  
 405 *a greedy approach. If  $Pr_{\mathcal{A}}(\mathcal{U})$  denotes the probability of a constraint predicate being preferred then,*  
 406  *$|D_\ell - D_{\ell-1}| \leq d \leq \sum_{i=1}^{2^{(|\mathbb{U}|-1)} \times {}^tP_q} Pr_{\mathcal{A}}(\mathcal{U}_i)$  where  $2^{(|\mathbb{U}|-1)} \times {}^tP_q$  is the maximum possible number of*  
 407 *constraint literals and  $q$  is the maximum arity of the constraints. In case of only pairwise constraints  $q = 2$ .*

408 **PROOF SKETCH FOR PROPOSITION 4.** The proof is straightforward and hence we present it in brief.  
 409 In our setting to show that,

$$|D_\ell - D_{\ell-1}| \leq d \leq \sum_{i=1}^{2^{(|\mathbb{U}|-1)} \times {}^tP_q} P_{\mathcal{A}}(\mathcal{U}_i) \quad (10)$$

410 (where  $2^{(|\mathbb{U}|-1)} \times {}^tP_q$  is the maximum number of constraint literals possible, since  $\mathbb{U}$  is the library of  
 411 constraint predicates) consider that the number of constraint predicates that can be picked up at any  
 412 iteration is  $2^{(|\mathbb{U}|-1)}$ . To form constraint literals we need to tie arguments to existing logical variables in  
 413 the current theory. We have defined  $t$  to be the number of terms in the existing theory. Let  $q$  be the max  
 414 arity of a constraint thus total possible number of constraint literals are  $2^{(|\mathbb{U}|-1)} \times {}^tP_q$ . So if the distribution  
 415 induced on the constraint literals by human advice  $\mathcal{A}$  be  $P_{\mathcal{A}}$ , then  $\sum_{i=1}^{2^{(|\mathbb{U}|-1)} \times {}^tP_q} P_{\mathcal{A}}(\mathcal{U}_i)$  is the expected  
 416 number of literals added given the advice. Now this is the upper bound of  $d$ . Again  $d$  should at least be the

conceptual distance between the new theory after constraint addition and the last consistent theory. Note  $d$  and conceptual distance  $D$  is not the same. Thus it is the difference between the NCD of last theory to original example and current updated theory to the original example  $|D_\ell - D_{\ell-1}|$ .

Observe that, if at each layer  $\ell \leq L$  we add constraint predicates  $\mathcal{U}_\ell$  then at layer  $\ell$ ,  $d = |\{\mathcal{U}\}^\ell| \leq 2^{mtP_q}$  (assuming  $q$  is maximum arity of the constraint predicates). Also, as per our greedy refinement framework, at each layer  $\ell$ , distance new theory  $\mathbb{T}_\ell$  should at least be the change in conceptual distance.

**Significance of Propositions 3 & 4:** Propositions 3 & 4 aim to illustrate what the general sample complexity would be for a theory refinement based RLGG clause learner such as GOCI and how the conceptual distance controls the complexity by establishing bounds on the refinement distance. Furthermore the complexity is also subject to the maximum refinement iterations, which in turn is affected by human guidance. Thus we establish the theoretical connection between the two dimensions of the contribution of this work.

**PROPOSITION 5 (Advice complexity).** *From Equations 6 and 8, at convergence  $\ell = L$ , we get  $\frac{n^* - |X|}{L}$  examples, on an average, for a concept  $\mathcal{C}$  to be PAC learnable using GOCI.*

The proof is quite straightforward and hence we just discuss the brief idea behind it. Our input is sparse (one or few instances). GOCI elicits advice over constraints to acquire additional information. Let  $|X|$  be the number of input examples. We query the human once at each layer and hence the maximum query budget is  $L$ . Given that the sample size is  $|X|$ , each query to the human must acquire information about at least  $\frac{n^* - |X|}{L}$  examples, on an average, for our a concept  $\mathcal{C}$  to be PAC learnable using our approach.

## 4 EVALUATION

We next aim to answer the following questions explicitly:

(Q1) Is GOCI effective in ‘one-shot’ concept induction?

(Q2) How sample efficient is GOCI compared to baselines?

(Q3) What is the relative contribution of the novel scoring function vs. human guidance towards performance?

Our framework extends a Java version of Aleph Srinivasan (2007). We modified the scoring function with NCD penalty computed via a customized SHOP2 planner Nau et al. (2003). We added constraint sampling and human guidance elicitation iteratively (Algorithm 1).

### 4.1 Experimental Design

We compare GOCI with Aleph with no enhancements. We focus on the specific task of “one-shot concept induction”, with a single input example for each of the several types of concepts and report aggregated precision. We consider precision because preference queries are meant to eliminate false positives in our case. To demonstrate general robustness of GOCI, beyond one-shot case, we experimented with varying sample sizes for each concept type and show learning curves for the same. We perform an ablation study to show the relative contribution of two important components of GOCI: (a) novel scoring metric and (b) human guidance; *i.e.*, we compare against two more baselines (*ILP+Score* and *ILP+Guidance*). For every domain, we consider ten different types of concepts (ten targets) and aggregate results over five runs.

Note that human guidance was obtained from distinct human experts for every run. The expertise level of all the advice providers was reasonably at par since they were chosen from the same pool of candidates with

Domain	Approach	Avg. Precision	#Queries
Minecraft	GOCI	0.85	$5.5 \pm 3$
	ILP	0.35	-
Assembly	GOCI	0.65	$16.5 \pm 4$
	ILP	0.2	-
ChEBI	GOCI	0.615	$13.1 \pm 2.13$
	ILP	0.45	-
Barman	GOCI	0.7	$10.5 \pm 5.4$
	ILP	0.51	-

**Table 1.** Results for one-shot concept learning.


zero visibility and knowledge of our proposed framework. However, for all the human advice providers we assumed a basic level of knowledge in geometry or fundamentals of logic and reasoning. Additionally, we also explained each of the experimental domains to the human participants to create a similar level of awareness about the domains among all of them.

**Domains** We employ 4 domains with varying complexity. Note that we have selected the below domains based on multiple considerations. The domain encoding need to be such that target concepts can be learned in a modular fashion (i.e. decomposable). Thus the first two domains are structure construction domains either spatial (Minecraft) or chemical/molecular (CheBI). Spatial structures are implicitly modular (such as the  $\mathbb{L}$ -structure in Figure 2). Chemical entities, molecules/compounds/complexes are similarly modular as well. The last two domains are fundamentally planning domains. However, they are also compositional in nature i.e. any planning goal is a composition task. For instance, machine structure in ‘Assembly’ domain and cocktails etc. ‘Barman’ domain. So these two domains do not just demonstrate learning modular/decomposable concepts but they also illustrate the plan induction feature of GOCI.

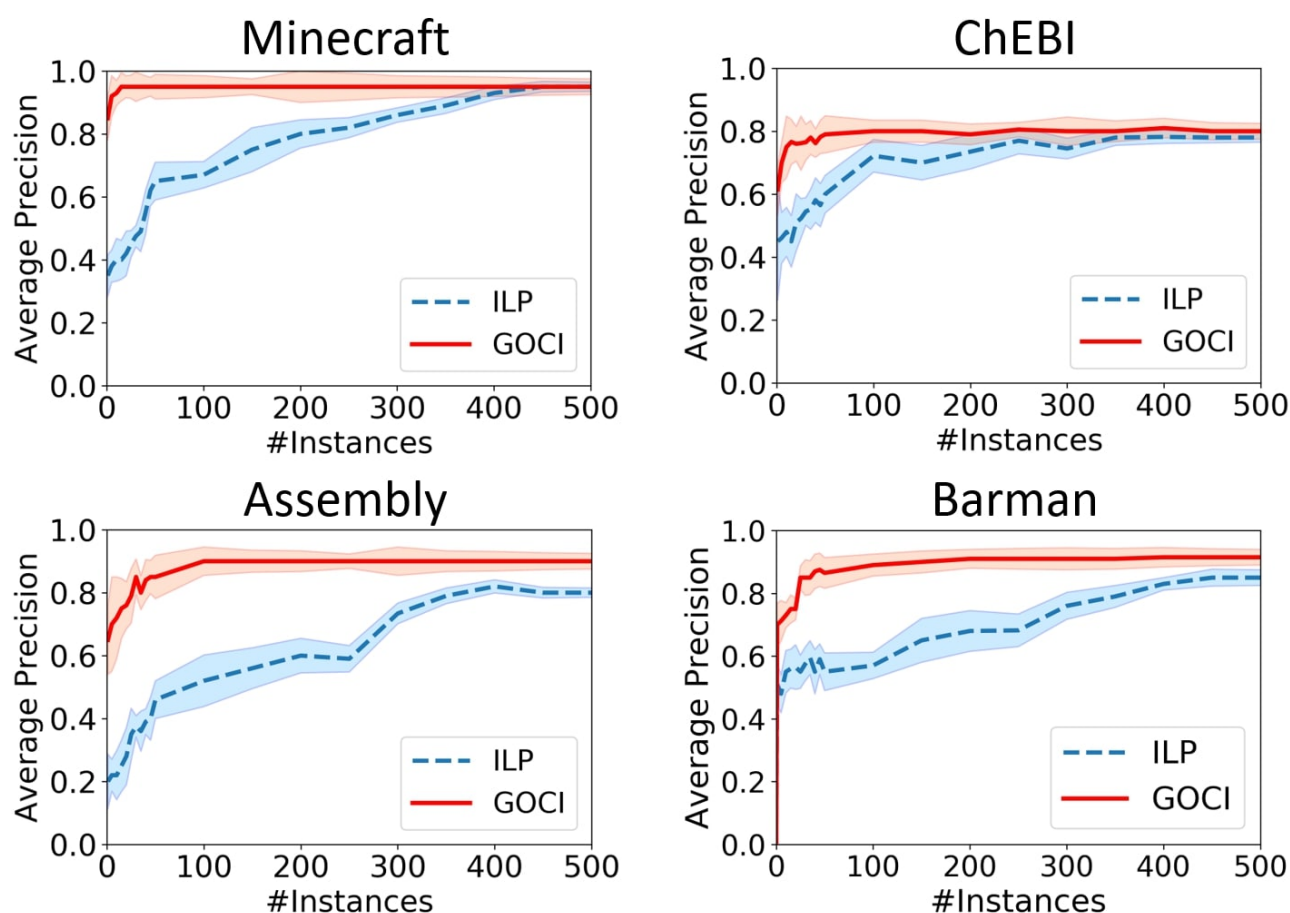
1. **Minecraft (Spatial Structures):** The goal is to learn discrete spatial concepts in a customized Narayan-Chen et al. (2019) Project Malmo platform for Minecraft. Dialogue data in Malmo is available online, and we converted them into a logical representation. All structures are in terms of discrete atomic unit blocks (cubes). Figure 5 shows examples of some spatial structures that GOCI was able to learn.
2. **Chemical Entities of Biological Interest (ChEBI):** ChEBI Degtyarenko et al. (2007) is a compound database containing important structural features and activity-based information, for classification of chemicals, such as - (1) Molecular structure, (2) Biological role etc. We model the Benzene molecule prediction task as molecular-compositional concepts. The data has predicates such as SingleBond, DoubleBond, HasAtom etc.
3. **Assembly (planning domain):** Assembly is a planning domain, where different mechanical structure concepts are compositions of different parts and resources. Input is a conjunction of ground literals indicating ground plan demonstration (assuming total ordering).
4. **Barman (planning domain):** A standard planning domain where a bartender is supposed to follow certain recipes and sequence of techniques to create cocktails. The different cocktails are decomposable concepts in this setting.

## 4.2 Experimental Results

**[Effective One-shot (Q1)]** Table 1 shows the performance of GOCI on one-shot concept learning tasks as compared to standard ILP. GOCI significantly outperforms ILP across all domains answering (Q1) affirmatively. Also, note that GOCI is very ‘query’ efficient as observed from the small average number of queries posed in the case of each domain. Note that in the case of CheBI, the number of queries is the




**Figure 5.** Instances of spatial concepts in Minecraft. (Left) Upright Tee, (Middle) Upright L, (Right) Orthogonal overlapping Ls



**Figure 6.** Learning curves for varying sample size to compare the sample-efficiency of GOCI and ILP. Top 2 plots are with respect to structural composition domains-Minecraft & ChEBI and the bottom 2 are for planning domains Assembly & Barman (best viewed in color).

488 highest among all the domains. This can be attributed to that fact that CheBI is a domain which requires  
 489 a certain degree of understanding of fundamental chemistry (chemical bonds and their types, molecules,  
 490 atoms etc.). Thus some of the human participants required more iterations (consequently more queries) to  
 491 converge to the most relevant set of constraint literals, given the difference of their prior understanding of  
 492 school chemistry.



**Figure 7.** Results of ablation study on Minecraft domain. Relative contribution of our distance-penalized score vs. human guidance.

Query efficiency is an important consideration in any learning paradigm that leverages human guidance, since controlling the cognitive load on the human expert is critical. So, in general, the observed average query numbers being reasonably low across all domains corroborates our theoretical advice complexity (Section 3.4.2).

**[Sample Efficiency (Q2)]** In figure 6, we observe that GOCI converges within significantly smaller sample size across all domains, thus, supporting our theoretical claims in Section 3.4. In ChEBI, though, quality of planner encoding might explain mildly lower-precision yet GOCI does perform significantly better than vanilla ILP learner. In ChEBI, we see that the sample efficiency is not vastly distinct. One of the possible reasons could be the sub-optimal encoding of the planning domain language, which is necessary for NCD computation, for this task. If we can improve the planner setup for this domain then we will likely be able to observe enhanced performance.

**[Relative contribution (Q3)]** Figure 7 validates our intuition that both components (scoring function and human-guidance) together make GOCI a robust one-shot (sample-efficient) concept induction framework. Though human guidance, alone, is able to enhance the performance of a vanilla ILP learner in sparse samples, yet it is not sufficient for optimal performance. In contrast, although the advantage of our novel distance-penalized scoring metric is marginal in sparse samples, it is essential for optimal performance at convergence.

### 4.3 Discussion

The most important conclusion from the experiments is that when available, the guidance along with the novel score leads to a jump-start, better slope and in some cases, asymptotically sample efficient with a fraction of the number of instances needed than merely learning from data.

Another important aspect to note here is that our experimental setup did not attempt to ensure in any way that the quality of guidance provided by the human participants is optimal. The formulation of the objective function, itself, in GOCI is designed to handle sub-optimal human advice implicitly, in a seamless manner. The two primary features in the design that make GOCI robust to advice quality are as follows,



- 519 1. As explained earlier and shown in Equation 3, human advice and conceptual distance deals with two  
520 distinct aspects of the search process. Human advice controls the size and nature of the search space  
521 while conceptual distance ensures the quality of the candidates. Advice and distance have a balancing  
522 effect on each other and thus, it is our novel conceptual distance that makes GOCI robust to bad advice.
- 523 2. Also, the nature of human advice in our setting is of choosing the most useful set of ‘constraint  
524 predicates’ among the set of candidate constraints. Now the candidates are generated by GOCI in a  
525 conservative fashion selecting only the ones that are logically valid for the theory learned at the current  
526 iteration of revision. Thus human experts have very little option of choosing an invalid or extremely  
527 unlikely constraint predicate.

528 Our ablation study in Figure 7 also supports our analysis. On closer inspection, we see that it is due to our  
529 novel distance penalized scoring function (ILP+Score) that ensures convergence to an optimal solution.  
530 Human advice (ILP+Guidance) contributes to sample efficiency.

## 5 CONCLUSIONS

531 We developed a human-in-the-loop one-shot concept learning framework in which the agent learns a  
532 generalized representation of a concept as FOL rules, from a single (few) positive example(s). We make  
533 two specific contributions – deriving a new distance measure between concepts and allowing for richer  
534 human inputs than mere labels, solicited actively by the agent. Our theoretical and experimental analyses  
535 show the promise of GOCI method. An exhaustive evaluation involving richer human inputs including  
536 varying levels of expertise and analyzing our claim that learning performance of GOCI is robust to expertise  
537 levels (which should only affect query efficiency) is an immediate future research objective. Integration  
538 with hierarchy learning also remains an interesting direction for future research.

## 6 ADDITIONAL REQUIREMENTS

### CONFLICT OF INTEREST STATEMENT

539 The authors declare that the research was conducted in the absence of any commercial or financial  
540 relationships that could be construed as a potential conflict of interest.

### AUTHOR CONTRIBUTIONS

541 MD and SN contributed equally to the ideation. MD and NR led the empirical evaluation. MD, NR, SN  
542 and JRD contributed nearly equally to the manuscript preparation.

### FUNDING

543 MD & SN gratefully acknowledge the support of CwC Program Contract W911NF-15- 1-0461 with the  
544 US Defense Advanced Research Projects Agency (DARPA) and the Army Research Office (ARO). SN &  
545 NR gratefully acknowledge AFOSR award FA9550-18-1-0462. Any opinions, findings and conclusion or  
546 recommendations are those of the authors and do not necessarily reflect the view of the DARPA, ARO,  
547 AFOSR or the US government.

### ACKNOWLEDGMENTS

548 The authors acknowledge the support of members of STARLING lab for the discussions.

549 This manuscript has been released as a pre-print at <http://arxiv.org> as part of the StarAI 2020  
550 workshop, Das et al. (2020)

## DATA AVAILABILITY STATEMENT

551 The datasets ANALYZED for this study can be found in following repository respectively:

- 552 • *Minecraft*: <http://juliahr.cs.illinois.edu/Minecraft/>
- 553 • *CheBI*: <ftp://ftp.ebi.ac.uk/pub/databases/chebi/>
- 554 • *Assembly*: <https://github.com/mayukhdas/ConceptLearning1Shot/tree/master/code>
- 555 • *Barman*: <https://github.com/mayukhdas/ConceptLearning1Shot/tree/master/code>

## REFERENCES

- 556 Bart, E. and Ullman, S. (2005). Single-example learning of novel classes using representation by similarity.  
 557 In *BMVC*
- 558 Braziunas, D. and Boutilier, C. (2006). Preference elicitation and generalized additive utility. In *AAAI*
- 559 Chick, H. L. (2007). Teaching and learning by example. *Mathematics: Essential research, essential*  
 560 *practice*
- 561 Das, M., Odom, P., Islam, M. R., Doppa, J. R., Roth, D., and Natarajan, S. (2018). Preference-Guided  
 562 Planning: An Active Elicitation Approach. In *AAMAS*
- 563 Das, M., Ramanan, N., Doppa, J. R., and Natarajan, S. (2020). One-shot induction of generalized logical  
 564 concepts via human guidance. In *StarAI-20 Workshop*. <https://arxiv.org/abs/1912.07060>
- 565 Degtyarenko, K., De Matos, P., Ennis, M., Hastings, J., Zbinden, M., McNaught, A., et al. (2007). Chebi: a  
 566 database and ontology for chemical entities of biological interest. *Nucleic acids research*
- 567 Eiben, A. E. and Smith, J. E. (2015). *Interactive Evolutionary Algorithms* (Berlin, Heidelberg: Springer  
 568 Berlin Heidelberg). 215–222. doi:10.1007/978-3-662-44874-8\_14
- 569 Flener, P. (1997). Inductive logic program synthesis with dialogs. In *ILP*
- 570 [Dataset] Friend, M., Khaled, M., Lefever, K., and Székely, G. (2018). Distances between formal theories
- 571 Fung, G. M., Mangasarian, O. L., and Shavlik, J. W. (2003). Knowledge-based support vector machine  
 572 classifiers. In *NIPS*
- 573 Goldman, R. P. and Kuter, U. (2015). Measuring plan diversity: Pathologies in existing approaches and a  
 574 new plan distance metric. In *AAAI*
- 575 Horn, A. (1951). On sentences which are true of direct unions of algebras. *JSL*
- 576 Kitzelmann, E. and Schmid, U. (2006). Inductive synthesis of functional programs: An explanation based  
 577 generalization approach. *JMLR*
- 578 Kozerawski, J. and Turk, M. (2018). Clear: Cumulative learning for one-shot one-class image recognition.  
 579 In *CVPR*
- 580 Lake, B. M., Salakhutdinov, R., and Tenenbaum, J. B. (2015). Human-level concept learning through  
 581 probabilistic program induction. *Science*
- 582 Mitchell, T. (1997). *Machine Learning* (McGraw-Hill Higher Education)
- 583 Mooney, R. J. (1994). A Preliminary PAC Analysis of Theory Revision. *Computational Learning Theory*  
 584 *and Natural Learning Systems*
- 585 MUGGLETON, S. (1988). Machine invention of first-order predicates by inverting resolution. In *Proc. of*  
 586 *5th Machine Learning Conference*
- 587 Muggleton, S. (1991). Inductive logic programming. *New generation computing*
- 588 Muggleton, S. (1995). Inverse entailment and prolog. *New generation computing*
- 589 Muggleton, S. and Feng, C. (1990). Efficient induction of logic programs. In *ALT*
- 590 Narayan-Chen, A., Jayannavar, P., and Hockenmaier, J. (2019). Collaborative dialogue in Minecraft. In  
 591 *ACL*

- 592 Nau, D., Au, T.-C., Ilghami, O., Kuter, U., Murdock, J. W., Wu, D., et al. (2003). Shop2: An htn planning  
593 system. *JAIR*
- 594 Odom, P., Khot, T., Porter, R., and Natarajan, S. (2015). Knowledge-based probabilistic logic learning. In  
595 *AAAI*
- 596 Pazzani, M. J. (1992). An information-based approach to integrating empirical and explanation-based  
597 learning. *ILP*
- 598 Preiss, K. and Shai, O. (1989). Process planning by logic programming. *Robotics and Computer-Integrated*  
599 *Manufacturing* 5, 1 – 10
- 600 Quinlan, J. R. (1990). Learning logical definitions from relations. *Machine learning*
- 601 Rouveirol, C. (1992). Extensions of inversion of resolution applied to theory completion. *ILP*
- 602 Sammut, C. and Banerji, R. B. (1986). Learning concepts by asking questions. *Machine learning: An*  
603 *artificial intelligence approach*
- 604 Shavlik, J. W. and Towell, G. G. (1989). Combining explanation-based learning and artificial neural  
605 networks. In *ICML*
- 606 [Dataset] Srinivasan, A. (2007). Aleph: A learning engine for proposing hypotheses
- 607 Stroulia, E. and Goel, A. K. (1994). Learning problem-solving concepts by reflecting on problem solving.  
608 In *ECML*
- 609 Tax, D. M. J. (2001). *One-class classification: Concept learning in the absence of counter-examples*. Ph.D.  
610 thesis, TU Delft
- 611 Valiant, L. G. (1984). A theory of the learnable. *Communications of the ACM*
- 612 Vapnik, V. N. (1999). An overview of statistical learning theory. *IEEE TNN*
- 613 Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al. (2016). Matching networks for one shot learning.  
614 In *NIPS*
- 615 Vitányi, P. M. (2013). Conditional kolmogorov complexity and universal probability. *Theoretical Computer*  
616 *Science*
- 617 Vitányi, P. M., Balbach, F. J., Cilibrasi, R. L., and Li, M. (2009). Normalized information distance. In  
618 *ITSL*
- 619 Wang, Y.-X., Girshick, R., Hebert, M., and Hariharan, B. (2018). Low-shot learning from imaginary data.  
620 In *CVPR*
- 621 Wendt, K., Cortés, A., and Margalef, T. (2010). Knowledge-guided genetic algorithm for input parameter  
622 optimisation in environmental modelling. *Procedia Computer Science* 1, 1367 – 1375. doi:https:  
623 //doi.org/10.1016/j.procs.2010.04.152. ICCS 2010
- 624 Wilson, R. H. and Latombe, J.-C. (1994). Geometric reasoning about mechanical assembly. *Artificial*  
625 *Intelligence*
- 626 Wong, M. L. and Leung, K. (1997). Evolutionary program induction directed by logic grammars.  
627 *Evolutionary Computation* 5, 143–180. doi:10.1162/evco.1997.5.2.143