

# 1 Neural Local Search Enhanced with Knowledge 2 Compilation

3 Alexandre Dubray<sup>1</sup> ✉ 

4 DTAI, KU Leuven, Belgium

5 H el ene Verhaeghe ✉ 

6 ICTEAM, UCLouvain, Belgium

7 Quentin Cappart ✉ 

8 ICTEAM, UCLouvain, Belgium

9 Polytechnique Montr al, Canada

10 Mila - Qu ebec AI Institute, Montr al, Canada

11 Siegfried Nijssen ✉ 

12 DTAI, KU Leuven, Belgium

## 13 — Abstract —

14 Constraint satisfaction, the task of finding an assignment of variables that respects a set of constraints,  
15 is a key area in artificial intelligence, with many industrial applications. In general, many constraint  
16 satisfaction problems are NP-hard; hence, exhaustive searches are intractable. Local search is a  
17 popular paradigm for solving such problems: instead of exploring the entire search space, local  
18 search solvers transition, for a limited number of steps, from one assignment to another, guided  
19 by hand-crafted, problem-dependent heuristics. Recently, neural networks have been studied as  
20 a means of automatically learning heuristics for transitioning between assignments. A particular  
21 class of neural local search solvers uses neural networks as generative models to produce joint  
22 probability distributions over the problem’s variables, conditioned on the current assignment. Then,  
23 transitioning to the next assignment is done by sampling from the generated probability distribution.  
24 However, these solvers do not integrate constraint reasoning; hence, any assignment can be generated  
25 during the local search. In this work, we propose to integrate knowledge compilation into neural  
26 local search solvers. By first compiling constraint satisfaction problems into multi-valued decision  
27 diagrams (MDDs) and sampling assignments in them, constraints can be applied to reduce the  
28 number of unsatisfying assignments in the sampling space. Our experiments show that using MDDs  
29 yields a sampling space that can be orders of magnitude smaller while incurring negligible overhead.  
30 Moreover, our MDD-based sampling procedure significantly improves the number of instances solved  
31 on Sudoku and graph colouring problems.

32 **2012 ACM Subject Classification** [Replace ccsdesc macro with valid one](#)

33 **Keywords and phrases** Dummy keyword

34 **Digital Object Identifier** 10.4230/LIPIcs.CVIT.2016.23

## 35 **1** Introduction

36 Constraint satisfaction, the task of finding an assignment to discrete variables that respects a  
37 set of constraints, and its optimisation version, are fundamental fields in artificial intelligence  
38 with numerous industrial applications, including scheduling, planning, resource management,  
39 and vehicle routing. Solving such problems often requires exploring an exponentially large  
40 solution space; hence, exhaustive searches may not finish within a reasonable time. Local  
41 search (LS) solvers alleviate this problem by avoiding classical search and instead transitioning  
42 from one assignment to another over a finite number of steps. Modifying an assignment

---

<sup>1</sup> Corresponding author



43 can be viewed as two processes: a destruction process that determines which variables  
44 of the current assignment to modify, and a repair process that assigns new values to the  
45 destroyed variables. The performance of LS solvers heavily depends on the quality of these  
46 two processes, which are often based on hand-crafted, problem-dependent heuristics.

47 In recent years, there has been a growing interest in leveraging machine learning, par-  
48 ticularly neural networks, to solve combinatorial optimisation problems [4]. Neural local  
49 search (NLS) solvers are LS solvers that rely on a heuristic learned by a neural network  
50 to transition between assignments [20, 22]. They operate in two stages. First, during the  
51 training stage, the neural network parameters are learned on small or random instances  
52 using classical machine learning techniques. Then, at inference time (i.e., when solving a  
53 new instance), the neural network is used to transition from one assignment to another,  
54 starting from a randomly chosen one. NLS solvers considered in this work generate, from an  
55 assignment, a joint probability distribution on the problem’s variables, and the local move  
56 consists of sampling new values from this distribution [20, 22]. To make sampling efficient,  
57 these solvers assume independence between the variables; hence, the neural network produces  
58 a distribution over their domain, and they are sampled independently.

59 These neural solvers present various advantages. First, they automatically learn heuristics  
60 during the training phase to explore the search space. This eases modelling optimisation  
61 problems, as there is no need to design the heuristic by hand. Moreover, it has been shown that  
62 heuristics learned by neural networks can outperform classical heuristics (i.e., they reduce the  
63 search space more in classical depth-first search) [5]. Moreover, by refining the assignment in a  
64 local-search-style manner, they will produce, with enough iterations, a satisfiable solution, in  
65 contrast to one-step neural predictors. Finally, modern methods employ neural architectures  
66 that, in theory, can operate on problems with arbitrary constraints [20, 22].

67 However, the main limitation of purely NLS solvers is that they do not integrate constraint  
68 reasoning and instead rely solely on the neural network to produce consistent distributions.  
69 For example, let us consider a problem with two binary variables linked by an inequality  
70 constraint, and the network produces uniform distributions on their domain. Then, sampling  
71 these distributions independently yields unsatisfiable assignments with probability 0.5. On  
72 the contrary, sampling one variable after another and integrating constraint reasoning always  
73 yields a satisfiable assignment. Moreover, although modern NLS solvers are designed to handle  
74 arbitrary constraints, they have only been evaluated on problems with simple constraints.  
75 For example, the **ConsFormer** solver has been benchmarked on problems containing only  
76 inequality constraints [22]. Learning to find solutions to problems with semantically more  
77 complex constraints (e.g., cardinality constraints, grammar constraints) seems currently out  
78 of reach for purely neural methods.

79 Constraint reasoning is a well-studied field in symbolic AI. *Knowledge compilation* is the  
80 task of transforming a knowledge base (e.g., a set of constraints) into a target language that  
81 supports various forms of reasoning. In this work, we consider the compilation of constraint  
82 satisfaction problems into *multi-valued decision diagrams* (MDDs). In this context, an MDD  
83 compactly encodes assignments to the variables of a CSP as a directed, layered, acyclic  
84 graph. During the compilation of such a data structure, reasoning is already applied, thereby  
85 removing infeasible solutions. Compiling an exact MDD yields the exact set of solutions;  
86 however, for NP-hard problems, this is often intractable. *Relaxed* MDDs encode a superset  
87 of the solutions; hence, they are a relaxation of exact MDDs. For instance, they can be  
88 obtained by bounding the width of the layers in an exact MDD and merging nodes that  
89 would otherwise be distinct, introducing non-solutions in the MDD [1, 6, 10].

90 The contribution of this work is the integration of constraint reasoning, via knowledge

91 compilation and MDDs, into neural local search solvers. We focus on such an integration  
 92 at inference time, when solving a new, unseen instance. We propose to integrate constraint  
 93 reasoning as a pre-processing step in the inference process of NLS solvers. First, the CSP  
 94 is compiled into a relaxed MDD using knowledge compilation techniques. During this  
 95 compilation process, constraints are applied, reducing the number of unsatisfying assignments  
 96 stored in the MDD. During the neural local search, instead of sampling independently from  
 97 the generated distributions, we sample an assignment in the MDD using these distributions.  
 98 By sampling in a relaxed MDD with fewer infeasible assignments, we expect to produce better  
 99 solutions, ultimately leading to finding feasible solutions in fewer iterations. We implemented  
 100 this framework as an extension of **ConsFormer**, a recently proposed transformer-based NLS  
 101 solver [22]. Our experiments show that integrating knowledge compilation significantly  
 102 increases the number of solved instances on the Sudoku and the graph colouring problems  
 103 while incurring a reasonable overhead.


## 104 2 Problem Definition and Formalisation

105 In this section, we give the necessary information regarding constraint satisfaction problems,  
 106 multi-valued decision diagram compilation, and neural local search solvers.

### 107 2.1 Constraint Satisfaction with Multi-valued Decision Diagrams

108 Let  $\mathcal{P} = \langle \mathbf{X}, \mathbf{D}, \mathbf{C} \rangle$  be a constraint satisfaction problem over variables  $\mathbf{X} = \{X_1, \dots, X_n\}$  with  
 109 domain  $\mathbf{D} = \{D_1, \dots, D_n\}$  and constraints  $\mathbf{C} = \{C_1, \dots, C_m\}$ . We denote  $\mathbf{x} = \langle x_1, \dots, x_n \rangle$   
 110 an assignment to the variables, with  $x_i \in D_i$ . A constraint  $C_j \in \mathbf{C}$  can be seen as a  
 111 function over its scope  $\mathbf{Y} \subseteq \mathbf{X}$ , indicating which assignment is valid or not. For example,  
 112 the **AllDifferent** ( $\mathbf{Y}$ ) returns true if and only if all variables in  $\mathbf{Y}$  take different values. A  
 113 solution to  $\mathcal{P}$  is an assignment respecting all constraints.


114 ► **Example 1** (Graph Colouring). As a running example, let us consider the graph colouring  
 115 problem. This problem consists of assigning colours from a limited set to the nodes of a  
 116 graph such that no two adjacent nodes share the same colour. Below is an example of a  
 117 graph with a valid colouring of its nodes, assuming the maximum number of colours is 3.



118  
 119 More formally, let  $G(\mathbf{V}, \mathbf{E})$  be an undirected graph with  $n = |\mathbf{V}|$  nodes and edges  $\mathbf{E}$ . To  
 120 encode the graph colouring problem as a CSP, we model each node with a variable whose  
 121 domain is the set of colours. That is, we have  $\mathbf{X} = \{X_1, \dots, X_n\}$  with  $D_1 = \dots = D_n =$   
 122  $\{1, \dots, k\}$ . There is one constraint  $C$  per edge  $(V_i, V_j) \in \mathbf{E}$  defined as  $C \equiv X_i \neq X_j$ .

123 In the context of constraint satisfaction problems, a multi-valued decision diagram is a  
 124 rooted, directed, acyclic graph that encodes the solutions to  $\mathcal{P}$ . In such a representation,  
 125 each path from the root to the sink encodes a solution; nodes correspond to variables, and  
 126 outgoing edges correspond to domains' values. In this work, we consider layered-organised  
 127 MDDs in which all nodes of a given layer correspond to the same variable.

128 ► **Example 2** (Graph Colouring (cont.)). Let us consider again the graph colouring problem  
 129 on the graph shown in Example 1. In the previous example, we showed a possible valid  
 130 assignment, but many more exist. The MDD below encodes all feasible colour assignments.  
 131 The decision variable associated with each layer is shown on the left, and labels on the edges  
 132 are values from its domain.



133


■ **Figure 1** An exact MDD encoding all solutions to the graph colouring problem in Example 1

134 Each path from the root to the sink (the white dot) is a full assignment to the CSP's variables.  
 135 It can be seen that constraints have been applied; for example, in the second layer, each  
 136 node has two outgoing edges rather than three, ensuring that the assignments are consistent  
 137 with the constraint  $X_2 \neq X_3$ .

138 Creating an MDD that represents the exact set of solutions is at least as hard as solving  
 139 the initial problem; hence, it is not always possible to construct one. Relaxed MDDs, i.e.,  
 140 MDDs representing a superset of the solutions, are designed to be easier to construct, at the  
 141 cost of not always providing a feasible solution. Typically, to construct such relaxed diagrams,  
 142 a constraint is imposed on the width of the MDD (i.e., the maximum number of nodes in a  
 143 layer). This implies merging nodes, which makes some assignments indistinguishable; hence,  
 144 unsatisfying assignments are included in relaxed MDDs.

145 The framework we use in this work to construct (relaxed) MDDs is the *node splitting*  
 146 *and refinement* process [1, 6, 9]. In this framework, the compilation process starts with a  
 147 fully relaxed MDD of width 1 that encodes all assignments. In such an MDD, all paths go  
 148 through the same nodes, and assignments are indistinguishable. Then, a node is selected  
 149 for *splitting*: it is duplicated, and its incoming edges are redirected to the new nodes, while  
 150 its outgoing edges are copied. Doing so increases the number of paths that are distinct in  
 151 the succession of nodes they traverse, thereby enabling the distinction between assignments.  
 152 Hence, constraint propagation can be applied to remove some unfeasible paths; this is the  
 153 *refinement* step. Constraint propagation removes edges in the MDD and nodes that become  
 154 unreachable from the root or cannot reach the sink. After this refinement step, the number  
 155 of nodes in a layer can exceed the maximum authorised width, in which case some nodes are  
 156 merged together.

157 ▶ **Example 3** (Graph Colouring (cont.)). Let us illustrate the compilation process on our  
 158 graph colouring example. The figure below shows a relaxed MDD at different stages of  
 159 compilation, with a maximum width of 2. For clarity, we represent multiple edges with the  
 160 same source and target as a single edge with multiple assignments on its label.



164 ■ **Figure 2** Example of compilation process, by splitting and refinement, of a relaxed MDD for the  
 165 graph colouring problem of Example 1

166 In (a), the MDD is completely relaxed with one node per layer; this is the initial MDD that  
 167 encodes all assignments. Then, in (b), the node in the second layer, associated with  $X_3$ , is  
 168 selected for splitting and refinement. It can be seen that, in that case, different assignments  
 169 emerge, reflecting the possible choices for  $X_2$ . The constraint  $X_2 \neq X_3$  can be applied to  
 170 remove invalid assignments in the second layer, which are highlighted in red. For example, if  
 171  $X_2 = 1$  (leftmost outgoing edge), then  $X_3$  can not be 1, and this edge is removed. At this  
 172 step, the number of nodes in the second layer exceeds the maximum width constraint. Let  
 173 us assume that the heuristic selects the rightmost nodes for merging, resulting in the relaxed  
 174 MDD in (c). This process can be repeated for each layer, in order, resulting, for example, in  
 the final relaxed MDD shown in (d). It can be seen that this MDD is smaller than the exact  
 MDD shown in Figure 1, but includes invalid assignments.

175 **2.2 Neural Local Search Solvers**

176 Local search solvers work by producing a sequence of assignments  $\langle \mathbf{x}^1, \dots, \mathbf{x}^t \rangle$  to a CSP  
 177  $\mathcal{P}$  until either  $\mathbf{x}^t$  satisfies all constraints  $\mathbf{C}$ , or the solvers reach a limit  $T$  (e.g., a number  
 178 of steps, or a time limit). LS solvers implement various heuristics to detect which part to  
 179 modify to go from an assignment  $\mathbf{x}^i$  to an assignment  $\mathbf{x}^{i+1}$ . NLS solvers, on the other hand,  
 180 guide exploration of the search space using neural networks.

181 NLS solvers considered in this work proceed as follows [20, 22]. Intuitively, these solvers  
 182 treat the problem’s variables as random, and the neural network acts as a generative  
 183 model that outputs a joint probability distribution over these variables, conditional on the  
 184 current assignment and the constraints. Then, a new assignment can be sampled from this  
 185 distribution. To make the sampling of this distribution tractable, independence between

186 the variables is assumed, so that each variable can be sampled independently in parallel,  
 187 leveraging modern GPU architectures.

Formally, let  $\Phi$  be a trained neural network and  $\mathbf{x}^i$  an assignment to the variables  $\mathbf{X}$  of a CSP  $\mathcal{P} = \langle \mathbf{X}, \mathbf{D}, \mathbf{C} \rangle$ . Evaluating  $\Phi$  results in  $n$  probability distributions, one over the domain of each variable. Hence, the joint probability distribution to sample from can be factorised as follows

$$P(X_1, \dots, X_n \mid \mathbf{x}^i, \mathcal{P}) = \prod_{j=1}^n P(X_j \mid \mathbf{x}^i, \mathcal{P}).$$

188 For clarity, in the rest of this paper, we denote simply by  $P(X_i)$  the conditional probability  
 189 distribution  $P(X_i \mid \mathbf{x}^{i-1}, \mathcal{P})$  when  $\mathbf{x}^{i-1}$  and  $\mathcal{P}$  are clear from the context.

### 190 3 Related Work

191 On the symbolic side, multi-valued decision diagrams have emerged as a practical approach  
 192 for efficiently solving satisfiability and optimisation problems. For example, it has been used  
 193 in standalone solvers based on depth-first search [3, 7, 14, 13], branch-and-bound [18], or in  
 194 iterative solvers, refining a relaxed MDD with conflict analysis [11, 21]. Decision diagrams  
 195 have also been integrated into constraint programming [1, 6, 9, 10]. In particular, some  
 196 *global constraints* (i.e., constraints defining complex semantics over a subset of variables),  
 197 have been adapted to work on MDDs [8, 10]: given a relaxed MDD, specialised algorithms  
 198 can be applied on it to remove inconsistent arcs.

199 On the neural side, our work aligns with two recently proposed neural local search  
 200 solvers. ANYCSP uses a graph neural network as a generative model to guide exploration of  
 201 the search space [20]. They represent the CSP as a constraint-value graph with nodes for  
 202 variables, domain values, and constraints. To guide the generation of the distribution, they  
 203 label value-to-constraint edges with a binary indicator indicating whether a value will help  
 204 satisfy the constraint. On the other hand, the solver we consider in this work, referred to  
 205 as **ConsFormer**, is based on a transformer architecture [22]. The primal graph of the CSP  
 206 is used in the relative positional encoding to guide the generative model. A key aspect of  
 207 **ConsFormer** is that it changes only a subset of the assignment at each local search step.  
 208 While in the original paper, this subset is randomly selected, the authors recently proposed  
 209 an extension of **ConsFormer** that employs more advanced techniques [23].

### 210 4 MDD-Based Sampling for Neural Local Search Solvers

211 The intuition behind our method is as follows. NLS solvers such as **ConsFormer** [22] or  
 212 **AnyCSP** [20] generate a probability distribution over all possible assignments and sample each  
 213 variable independently. In the problem described in Example 1, all  $3^5 = 243$  assignments can  
 214 be sampled when using these solvers, but there are only 24 solutions. For more challenging  
 215 CSPs, the proportion of solutions can be significantly lower. The main idea behind our  
 216 approach is to leverage relaxed MDDs to reduce the number of unfeasible assignments that  
 217 can be sampled. In Figure 2, the final relaxed MDD encodes 194 assignments, a reduction of  
 218 roughly 20%. Hence, for a set of distributions  $P(X_1), \dots, P(X_n)$ , the probability of sampling  
 219 a feasible assignment is higher when sampling a path in a relaxed MDD, since there are  
 220 fewer unfeasible assignments. Our new sampling procedure traverses the MDD from the  
 221 root to the sink and, at each layer associated with variable  $X_i$ , samples an edge according  
 222 to the neurally-generated probability distribution  $P(X_i)$ . Note that the generation of the  
 223 distribution (i.e., the neural architecture) is not changed in this work.

---

**Algorithm 1** `RelaxedMDDSampling( $P(X_1), \dots, P(X_n), \mathcal{M}$ )`


---

```

input : Probability distributions  $P(X_1), \dots, P(X_n)$ 
input : A relaxed MDD  $\mathcal{M}$  with root  $r$  and sink  $s$ 
output: An assignment  $\mathbf{x} = \langle x_1, \dots, x_n \rangle$  such that its path exists in  $\mathcal{M}$  and is
         sampled proportionally to  $\prod_{i=1}^n P(X_i)$ 
/* Empty assignment of size  $n$  */
1  $\mathbf{x} \leftarrow \langle \_ \rangle^n$ 
2  $v \leftarrow r$ 
3 while  $v \neq s$  do
4    $X_i \leftarrow$  decision variable associated with node  $v$ 
   /* Sample an edge from  $v$  to  $u$  with label  $l$  */
5    $(v, u, l) \leftarrow \text{Sample}(P(X_i), v)$ 
6    $\mathbf{x}_i \leftarrow l$ 
7    $v \leftarrow u$ 
8 end
9 return  $\mathbf{x}$ 

```

---

224 Algorithm 1 shows our new sampling procedure. It starts from an empty assignment  
 225 (line 1) and sets the current node as the root of the MDD (line 2). Then, it traverses the  
 226 MDD until the sink is reached (lines 3-8). At each step, an edge (i.e., a labelled pair or  
 227 nodes) is sampled according to the distribution associated with the current layer's decision  
 228 variable (line 5), and its label ( $l$ ) is assigned to the decision variables (line 6). Then, the  
 229 current node is updated by following the edge (line 7).

230 This sampling procedure can be integrated into a NLS solver using the loop shown  
 231 in Algorithm 2. It starts by compiling a relaxed MDD of maximum width  $W$ , encoding  
 232 assignments of the input CSP (line 1). If the MDD is empty (i.e., all nodes have been removed  
 233 during constraint propagation), then the problem is unsatisfiable (e.g., it is impossible to  
 234 colour a graph containing a clique of size 4 with 3 colours), and **UNSAT** is returned (line 2). If  
 235 the MDD contains a path passing only through non-relaxed nodes, this path is assured to be  
 236 a solution to the CSP, and the corresponding assignment is returned (line 3). For example,  
 237 in Figure 2, the leftmost path of the final MDD contains only nodes that have not been  
 238 merged; hence, it is guaranteed to be a valid colouring. If no exact path exists and the MDD  
 239 is not empty, the local search (lines 6-15) starts from a random assignment (line 4). While a  
 240 solution has not been found, or the maximum number of steps is not reached, distributions  
 241 are generated over the variables' domain (line 7) similarly to classical NLS solvers, and a  
 242 new assignment is sampled from the relaxed MDD, using Algorithm 1 (line 8).

243 Key aspects of Algorithm 2 are the **UNSAT** check as well as the extraction of a solution  
 244 directly from the relaxed MDD, when it contains an exact path. During the initial compilation,  
 245 we heuristically merge nodes linked to the most relaxed nodes; hence, this heuristic aims to  
 246 preserve as many exact paths as possible. However, if this fails, and the NLS is started, it  
 247 may be beneficial to leverage the neural network to guide a recompilation process. Hence, we  
 248 propose to recompile the MDD periodically in a neurally guided manner (lines 9-13). During  
 249 this recompilation, all solutions sampled since the last compilation are used to guide the  
 250 merging heuristic. We implemented a simple neural-based heuristic to guide the merging  
 251 process: we merge nodes whose outgoing edges are labelled with values that were least  
 252 selected in the last assignments.

---

**Algorithm 2** MDD-NLS( $\mathcal{P}, \Phi, T, W, r$ )

---

**input** : A CSP  $\mathcal{P} = \langle \mathbf{X}, \mathbf{D}, \mathbf{C} \rangle$   
**input** : A trained neural network  $\Phi$   
**input** : A maximum number of steps  $T$   
**input** : A maximum width  $W$  for the relaxed MDD and a recompilation frequency  $r$   
**output** : UNSAT, a solution to the CSP, or the last found assignment in the local search

```

1  $\mathcal{M} \leftarrow \text{CompileRelaxMDD}(\mathcal{P}, W, \langle \rangle)$ 
2 if  $\mathcal{M}$  is empty then return UNSAT
3 if  $\mathcal{M}$  has an exact path with assignment  $\mathbf{x}^s$  then return  $\mathbf{x}^s$ 
4  $\mathbf{x}^0 \leftarrow$  random assignment
5  $i \leftarrow 1$ 
6 while  $\mathbf{x}^{i-1}$  is not a solution of  $\mathcal{P}$  and  $i \leq T$  do
7    $P(X_1), \dots, P(X_n) \leftarrow \Phi(\mathbf{x}, \mathcal{P})$ 
8    $\mathbf{x}^i \leftarrow \text{RelaxedMDDSampling}(P(X_1), \dots, P(X_n), \mathcal{M})$ 
9   if  $(i \bmod r) = 0$  then
10      $\mathcal{M} \leftarrow \text{CompileRelaxMDD}(\mathcal{P}, W, \langle x^{i-r}, \dots, x^i \rangle)$ 
11     if  $\mathcal{M}$  is empty then return UNSAT
12     if  $\mathcal{M}$  has an exact path with assignment  $\mathbf{x}^s$  then return  $\mathbf{x}^s$ 
13   end
14    $i \leftarrow i + 1$ 
15 end
16 return  $\mathbf{x}^{i-1}$ 

```

---

## 5 Experimental Results

To showcase the effectiveness of our approach, we implemented our MDD-based neural local search solver as a modification of **ConsFormer** [22]. We evaluate our approach against **ConsFormer** on two satisfaction problems: Sudoku and graph colouring.

### 5.1 Experimental setup

We based our implementation on **ConsFormer**'s source code, which is freely available online<sup>2</sup>. We trained the neural networks using the hyperparameters reported by **ConsFormer**'s authors [22], and our baseline's performance matches theirs. We also use the same train and test datasets as the ones used by **ConsFormer** [22].

For the Sudoku, training instances consist of  $9 \times 9$  grids with between 39 and 50 missing cells. For conciseness, we omit benchmarks on the test instances, as they are easily solved by all methods, and we focus on out-of-distribution instances that are  $9 \times 9$  grids containing between 47 and 64 missing cells. For graph colouring, training instances consist of random graphs with 50 nodes that must be coloured with 5 colours. Out-of-distribution instances are graphs with 100 nodes that must be coloured with 5 colours. We set the maximum number of steps to 20,000 for Sudoku and 10,000 for graph colouring instances.

Regarding the MDD compilation, we used a modified version of the **MaxiCP** solver,

---

<sup>2</sup> At the time of the writing, only the source code for the initial **ConsFormer** is available; we leave experiments with the advanced subset selection strategy for further work.

implemented in Java, that integrates MDD compilation using the split and refine framework, as well as constraint propagation for `AllDifferent` and inequality constraints on MDDs [8, 19]. We limit the available memory to 16 GB for the compilation process. We use the following heuristic to order the variables in the MDDs. For Sudoku, the variables in the first layer correspond to prefilled values; for graph colouring, the first variable corresponds to the node of highest degree. Then we use the same heuristic for both problems: the variable selected at layer  $i$  is the one most constrained by variables in layers  $j < i$ . We use this heuristic because it has been shown to perform well for solving graph colouring problems with MDDs [21].

We implemented our sampling method in PyTorch [2], so that it is integrated with the current `ConsFormer` codebase and allows for solving multiple instances at the same time. Given a batch of instances, the MDDs are compiled using the `MaxiCP` Java implementation and are then stored in tensors. An MDD can be stored in two tensors, using a flattened representation. The  $N$  nodes of the MDD are implicitly stored using indexes ranging from 0 to  $N - 1$ . The first tensor is a  $(N, D)$ -shaped tensor ( $D$  being the number of values in the variables' domain) encoding the structure of the MDD. It stores, for each node  $v$  and possible label  $l$ , the index of the node reached by following the outgoing edge of  $v$  labelled with  $l$ . If an edge does not exist,  $-1$  is used, as it does not correspond to any node. The second tensor is a  $(N, |\mathbf{X}|)$ -shaped tensor, mapping each node to its decision variable. Using these two tensors, it is possible to traverse the MDD and assign values to the decision variables.

We evaluate two versions of our method. First, the complete implementation of our method, denoted `ConsFormer-MDD`. Then, to evaluate the impact of sampling paths in the MDD, compared to just detecting UNSAT and exact paths, we evaluate a version of our method in which the MDD is compiled, but the distributions are sampled independently, as in the initial `ConsFormer`, denoted `ConsFormer-MDD-mix`.


## 5.2 Sudoku

We implemented Sudoku using the classical encoding with one `AllDifferent` constraint per row, column, and block. Figure 3 shows the proportion of solved instances over the number of iterations. We show one cactus plot per configuration, where  $r = 0$  indicates no recompilation. We also show a virtual best solver (VBS) across all configurations. When the maximum width for MDD compilation is 1, we do not recompile because no splitting or merging is performed; hence, recompiling would result in the same MDD.

It can be seen that, overall, integrating knowledge compilation results in a higher number of solved instances, even when the maximum width is 1. In this configuration, only the `AllDifferent` propagator is launched (no nodes are split), but infeasible values are still removed (e.g., if a cell is prefilled with 8, this value is removed from all other cells in its row, column, and block). Due to the number of missing values and the strong propagation power of the `AllDifferent` constraint, the number of unfeasible assignments removed is high. On average, in our benchmark, `ConsFormer` samples from  $7 \times 10^{59}$  assignments while the average number of paths encoded in the relaxed MDDs is  $6 \times 10^{40}$ . Applying constraint propagation drastically reduced the number of assignments to sample from.

Increasing the maximum width for compilation does not significantly increase the overall number of solved instances, but it increases the number of instances for which an exact path is found during compilation. Moreover, periodically recompiling MDDs helps with the solving process, as small bumps are observed at recompilation iterations, especially the first.

Finally, it can be seen that our mixed approach `ConsFormer-MDD-mix` performs better than the basic `ConsFormer`, but worse than our full approach `ConsFormer-MDD`. Hence, on this problem, sampling paths in the MDD is beneficial.




■ **Figure 3** Proportion of instances solved over the number of iterations on the out-of-distribution dataset for the Sudoku (47 to 64 missing cells) for various configurations for the MDD (re-)compilation.

### 317 5.3 Graph Colouring

318 Now, we consider the graph colouring problem, as described in Example 1. We encode this  
 319 problem with one variable per node and inequality constraints between adjacent nodes. To  
 320 be efficient, we must implement two optimisations. First, we detect symmetrical nodes using  
 321 `Nauty` [12]. Such nodes can be swapped without changing the graph’s structure; hence,  
 322 given variables  $X_1, \dots, X_m$  representing  $m$  symmetrical nodes, we impose  $X_1 \leq \dots, \leq X_m$ .  
 323 Moreover, we detect cliques in the graph to strengthen constraint propagation. Indeed,  
 324 nodes in a clique must take different values, which can be encoded using an `AllDifferent`  
 325 constraint, whose propagation is stronger than the pairwise inequality constraint [16, 17].  
 326 However, finding all cliques is computationally intractable; hence, given a graph colouring  
 327 problem with  $k$  colours, we find all cliques of size at most  $k + 1$  using `Cliquer` [15]. Finding  
 328 larger cliques is unnecessary, since finding a clique of size  $k + 1$  makes the instance unsatisfiable.

329 Figures 4 and 5 show the proportion of solved instances over the number of iterations  
 330 for, respectively, the in-distributions and out-of-distribution datasets. It can be seen that,  
 331 while MDD-based sampling helps solve more instances, the improvement is less marked  
 332 than on the Sudoku. This can be explained by considering how much the search space is



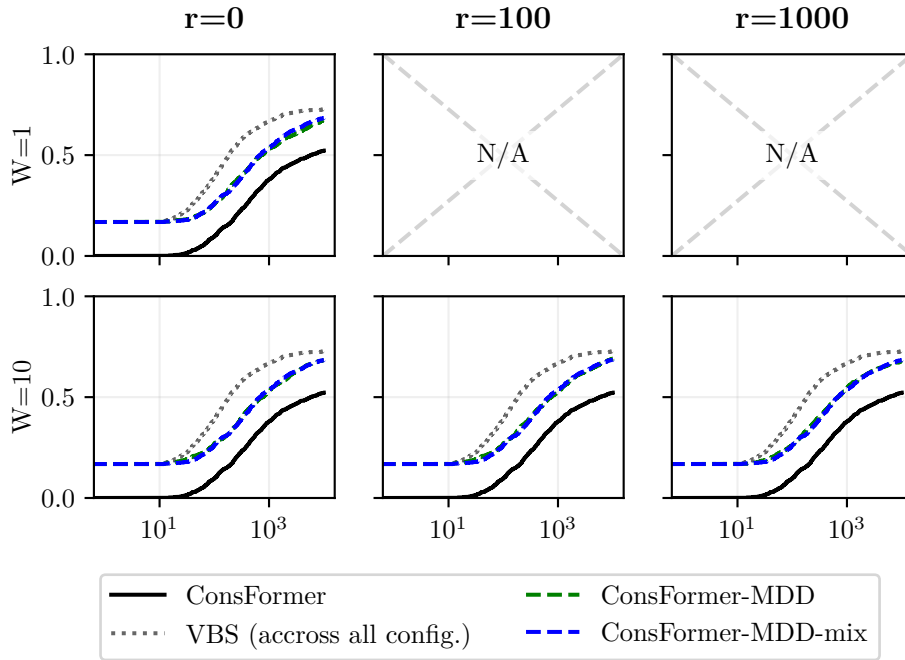
■ **Figure 4** Proportion of instances solved over the number of iterations on the in-distribution dataset for the graph colouring (50 vertices and 5 colours) for various configurations for the MDD (re-)compilation.

333 reduced when applying constraint propagation (i.e., MDD of width 1). While for Sudoku,  
 334 the number of encoded paths was  $10^{19}$  times smaller, for graph colouring, a 1-width relaxed  
 335 MDD encodes between 1% and 2% of the initial number of assignments. Due to this weaker  
 336 pruning capability, the maximum width allowed during compilation is also smaller. In our  
 337 experiments, instances with 100 nodes reach an out-of-memory state when compiled with a  
 338 maximum width  $W = 100$ .

339 However, a key advantage of our method is its ability to detect UNSAT instances. While  
 340 ConsFormer does not check for unsatisfiable instances, our method handles that seamlessly.  
 341 AllDifferent constraints on cliques detect roughly 10% of UNSAT instances in both in-  
 342 distribution and out-of-distribution datasets.

## 343 5.4 Runtime Overhead

344 Finally, we analyse the overhead induced by Algorithm 2 compared to ConsFormer. There  
 345 are two sources of overhead, one from compiling relaxed MDDs and one from the sequential  
 346 sampling of an assignment, compared to the parallel sampling of ConsFormer. Table 1 shows



■ **Figure 5** Proportion of instances solved over the number of iterations on the out-of-distribution dataset for the graph colouring (100 vertices and 5 colours) for various configurations for the MDD (re-)compilation.

347 the average overhead of our method. For each problem, we show the average time to compile  
 348 the relaxed MDDs for various maximum widths and the slowing factor of the sequential  
 349 sampling. We also report the average time required by **ConsFormer** to solve a problem  
 350 instance.

351 The cost of compiling a relaxed MDD can be controlled with the maximum width, allowing  
 352 a balance between search space reduction and runtime. For example, for the graph colouring  
 353 instances we evaluated, many instances can be detected as infeasible by compiling a relaxed  
 354 MDD of width 1, which takes one second for the larger instances.

355 The overhead induced by our sequential sampling is also limited. While the number of  
 356 sampling steps increases linearly with the number of variables, the slowing factor on the  
 357 iteration’s runtime increases more slowly. For example, for the out-of-distribution graph  
 358 colouring instances, our sampling strategy traverses a MDD with 100 layers, but the iterations’  
 359 runtime is less than twice the **ConsFormer**’s iteration runtime. For the smaller graphs, the  
 360 overhead is more significant, but our method is still less than an order of magnitude slower  
 361 than the initial **ConsFormer**. Further development is needed to optimise our MDD tensor  
 362 implementation to reduce this overhead.

363 Overall, there is a tradeoff between compilation time and epoch overhead. For instance, in  
 364 Sudoku, there is no overhead to running MDD-based sampling, which improves the number  
 365 of solved instances. On the other hand, for graph colouring, our current implementation  
 366 incurs overhead and does not scale well with MDD width, without increasing the solver’s  
 367 capabilities. In this case, a mix strategy such as **ConsFormer-MDD-mix** is preferable.

Problem	Baseline avg. runtime (s)	Metric	W=1	W=10	W=100	W=1000
Sudoku	275.23	Compilation time (s)	0.22 ± 0.01	-	0.68 ± 0.33	2.20 ± 2.41
		Iteration runtime increase factor	1.00	-	1.02	1.19
GCOL ( V  = 50)	24.60	Compilation time (s)	0.49 ± 0.14	1.71 ± 1.00	7.61 ± 5.55	-
		Iteration runtime increase factor	3.40	3.37	3.34	-
GCOL ( V  = 100)	92.88	Compilation time (s)	1.17 ± 0.42	4.82 ± 1.76	-	-
		Iteration runtime increase factor	1.70	1.67	-	-

**Table 1** Runtime overhead of MDD-based sampling for different maximum widths. For each problem, the average runtime needed by ConsFormer to solve one instance is indicated (second column). Dashed entries indicate configurations that have not been run in our experiments. The top line of each problem shows the average runtime (in seconds) for compiling the relaxed MDD, along with the standard deviation. The bottom line for each problem shows the slowing factor of MDD-based sampling.

## 6 Conclusion

In this work, we proposed a new methodology for solving constraint satisfaction problems using neural local search solvers. While the neural search solvers we studied in this work do not integrate constraint reasoning into their local search, we propose integrating knowledge compilation to guide the search towards higher-quality assignments. In a first phase, our method compiles the CSP into a relaxed multi-valued decision diagram, enabling the application of constraints and thereby removing infeasible assignments. Then, during the local search, generated assignments are restricted to those that appear in the relaxed MDD. We implemented our framework as an extension of the ConsFormer solver, using PyTorch. Our experiment shows that integrating constraint reasoning into neural local search solvers significantly improves their performance while incurring a reasonable overhead.

In this paper, we implemented our framework using existing solvers, that is ConsFormer and MaxiCP. MaxiCP has not been designed to be integrated with neural local search solvers. Moreover, when recompiling relaxed MDDs, the compiler starts from scratch, considering only the solutions generated during the previous steps of local search. We leave as future work the design of a compiler aimed towards neural solvers. In this work, we focused on modifying the inference stage of NLS solvers. How to integrate knowledge compilation into the training pipeline for NLS solvers remains an open question. Another interesting research direction is applying our framework to optimisation problems, as they are encoded differently than CSP's when solved using MDDs. Finally, problems with semantically more challenging constraints are also left for future work.

## References

- 1 H. R. Andersen, T. Hadzic, J. N. Hooker, and P. Tiedemann. A Constraint Store Based on Multivalued Decision Diagrams. In Christian Bessière, editor, *Principles and Practice of Constraint Programming – CP 2007*, volume 4741, pages 118–132. Springer Berlin Heidelberg, Berlin, Heidelberg, 2007. doi:10.1007/978-3-540-74970-7\_11.
- 2 Jason Ansel, Edward Yang, Horace He, Natalia Gimelshein, Animesh Jain, Michael Voznesensky, Bin Bao, Peter Bell, David Berard, Evgeni Burovski, Geeta Chauhan, Anjali Chourdia, Will Constable, Alban Desmaison, Zachary DeVito, Elias Ellison, Will Feng, Jiong Gong, Michael Gschwind, Brian Hirsh, Sherlock Huang, Kshiteej Kalambarkar, Laurent Kirsch, Michael Lazos, Mario Lezcano, Yanbo Liang, Jason Liang, Yinghai Lu, C. K. Luk, Bert Maher, Yunjie Pan, Christian Puhersch, Matthias Reso, Mark Saroufim, Marcos Yukio Siraichi, Helen Suk, Shunting Zhang, Michael Suo, Phil Tillet, Xu Zhao, Eikan Wang, Keren Zhou, Richard

- 401 Zou, Xiaodong Wang, Ajit Mathews, William Wen, Gregory Chanan, Peng Wu, and Soumith  
 402 Chintala. PyTorch 2: Faster Machine Learning Through Dynamic Python Bytecode Trans-  
 403 formation and Graph Compilation. In *Proceedings of the 29th ACM International Conference*  
 404 *on Architectural Support for Programming Languages and Operating Systems, Volume 2*, pages  
 405 929–947, La Jolla CA USA, April 2024. ACM. doi:10.1145/3620665.3640366.
- 406 3 David Bergman, Andre A. Cire, Willem-Jan Van Hove, and John Hooker. *Decision Diagrams*  
 407 *for Optimization*. Artificial Intelligence: Foundations, Theory, and Algorithms. Springer  
 408 International Publishing, Cham, 2016. doi:10.1007/978-3-319-42849-9.
- 409 4 Quentin Cappart, Tias Guns, Michele Lombardi, Gilles Pesant, and Dimos Tsouros. Com-  
 410 bining Constraint Programming and Machine Learning: From Current Progress to Future  
 411 Opportunities. *Journal of Artificial Intelligence Research*, 84, 2025.
- 412 5 Félix Chalumeau, Ilan Coulon, Quentin Cappart, and Louis-Martin Rousseau. SeaPearl: A  
 413 Constraint Programming Solver Guided by Reinforcement Learning. In Peter J. Stuckey,  
 414 editor, *Integration of Constraint Programming, Artificial Intelligence, and Operations Research*,  
 415 volume 12735, pages 392–409. Springer International Publishing, Cham, 2021. doi:10.1007/  
 416 978-3-030-78230-6\_25.
- 417 6 Rebecca Gentzel, Laurent Michel, and W.-J. Van Hove. HADDOCK: A Language and  
 418 Architecture for Decision Diagram Compilation. In Helmut Simonis, editor, *Principles and*  
 419 *Practice of Constraint Programming*, volume 12333, pages 531–547. Springer International  
 420 Publishing, Cham, 2020. doi:10.1007/978-3-030-58475-7\_31.
- 421 7 Xavier Gillard, Vianney Coppé, and Pierre Schaus. Ddo, a generic and efficient framework  
 422 for mdd-based optimization. In *International Joint Conference on Artificial Intelligence*  
 423 *(IJCAI20)*, 2020.
- 424 8 Amaury Guichard, Laurent Michel, H el ene Verhaeghe, and Pierre Schaus. Towards Bound  
 425 Consistency for the No-Overlap Constraint Using MDDs. *arXiv preprint arXiv:2601.14784*,  
 426 2026. arXiv:2601.14784.
- 427 9 Tarik Hadzic, John N. Hooker, Barry O’Sullivan, and Peter Tiedemann. Approximate  
 428 Compilation of Constraints into Multivalued Decision Diagrams. In David Hutchison, Takeo  
 429 Kanade, Josef Kittler, Jon M. Kleinberg, Friedemann Mattern, John C. Mitchell, Moni Naor,  
 430 Oscar Nierstrasz, C. Pandu Rangan, Bernhard Steffen, Madhu Sudan, Demetri Terzopoulos,  
 431 Doug Tygar, Moshe Y. Vardi, Gerhard Weikum, and Peter J. Stuckey, editors, *Principles and*  
 432 *Practice of Constraint Programming*, volume 5202, pages 448–462. Springer Berlin Heidelberg,  
 433 Berlin, Heidelberg, 2008. doi:10.1007/978-3-540-85958-1\_30.
- 434 10 Samid Hoda, Willem-Jan Van Hove, and J. N. Hooker. A Systematic Approach to MDD-Based  
 435 Constraint Programming. In David Hutchison, Takeo Kanade, Josef Kittler, Jon M. Kleinberg,  
 436 Friedemann Mattern, John C. Mitchell, Moni Naor, Oscar Nierstrasz, C. Pandu Rangan,  
 437 Bernhard Steffen, Madhu Sudan, Demetri Terzopoulos, Doug Tygar, Moshe Y. Vardi, Gerhard  
 438 Weikum, and David Cohen, editors, *Principles and Practice of Constraint Programming –*  
 439 *CP 2010*, volume 6308, pages 266–280. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010.  
 440 doi:10.1007/978-3-642-15396-9\_23.
- 441 11 Anthony Karahalios and Willem-Jan van Hove. Column Elimination: An Iterative Approach  
 442 to Solving Arc-Flow Formulations.
- 443 12 Brendan D. McKay and Adolfo Piperno. Practical graph isomorphism, II. *Journal of symbolic*  
 444 *computation*, 60:94–112, 2014.
- 445 13 Quentin Meurisse, Pierre Schaus, Renaud De Landtsheer, Fabian Germeau, Roger Kameugne,  
 446 and Xavier Gillard. DDOLib: Librairie pour les Diagrammes de D ecision en Optimisation. In  
 447 *ROADEF 2026*, 2026.
- 448 14 Laurent Michel and Willem-Jan Van Hove. CODD: A Decision Diagram-Based Solver for  
 449 Combinatorial Optimization. In Ulle Endriss, Francisco S. Melo, Kerstin Bach, Alberto  
 450 Bugar ın-Diz, Jos e M. Alonso-Moral, Sen en Barro, and Fredrik Heintz, editors, *Frontiers in*  
 451 *Artificial Intelligence and Applications*. IOS Press, October 2024. doi:10.3233/FAIA240997.
- 452 15 Sampo Niskanen and Patric  osterg ard. Cliquer user’s guide, version 1.0. 2003.

- 453 16 Jean-Charles Régin. A filtering algorithm for constraints of difference in CSPs. In *AAAI*,  
454 volume 94, pages 362–367, 1994.
- 455 17 Francesca Rossi, Peter Van Beek, and Toby Walsh. *Handbook of Constraint Programming*.  
456 Elsevier, 2006.
- 457 18 Isaac Rudich, Quentin Cappart, and Louis-Martin Rousseau. Improved peel-and-bound:  
458 Methods for generating dual bounds with multivalued decision diagrams. *Journal of Artificial*  
459 *Intelligence Research*, 77:1489–1538, 2023.
- 460 19 Pierre Schaus, G. Derval, A. Delecluse, L. Michel, and P. V. Hentenryck. Maxicp: A  
461 constraint programming solver for scheduling and vehicle routing. URL [https://github.com/aia-](https://github.com/aia-uclouvain/maxicp)  
462 [uclouvain/maxicp](https://github.com/aia-uclouvain/maxicp), 2024.
- 463 20 Jan Tönshoff, Berke Kisin, Jakob Lindner, and Martin Grohe. One Model, Any CSP: Graph  
464 Neural Networks as Fast Global Search Heuristics for Constraint Satisfaction, August 2022.  
465 [arXiv:2208.10227](https://arxiv.org/abs/2208.10227), [doi:10.48550/arXiv.2208.10227](https://doi.org/10.48550/arXiv.2208.10227).
- 466 21 Willem-Jan van Hoeve. Graph coloring with decision diagrams. *Mathematical Programming*,  
467 192(1):631–674, 2022.
- 468 22 Yudong W. Xu, Wenhao Li, Scott Sanner, and Elias B. Khalil. Self-Supervised Transformers  
469 as Iterative Solution Improvers for Constraint Satisfaction, June 2025. [arXiv:2502.15794](https://arxiv.org/abs/2502.15794),  
470 [doi:10.48550/arXiv.2502.15794](https://doi.org/10.48550/arXiv.2502.15794).
- 471 23 Yudong W. Xu, Wenhao Li, Scott Sanner, and Elias B. Khalil. Large Neighborhood Search  
472 meets Iterative Neural Constraint Heuristics. *arXiv preprint arXiv:2603.20801*, 2026. [arXiv:](https://arxiv.org/abs/2603.20801)  
473 [2603.20801](https://arxiv.org/abs/2603.20801).