

Hybrid Learning with New Value Function for the Maximum Common Induced Subgraph Problem

Yanli Liu¹, Jiming Zhao¹, Chu-Min Li^{2*}, Hua Jiang³, Kun He⁴

¹School of Science, Wuhan University of Science and Technology, China

²MIS, Université de Picardie Jules Verne, France

³Engineering Research Center of Cyberspace & School of Software, Yunnan University, China

⁴School of Computer Science and Technology, Huazhong University of Science and Technology, China

{yanli2008, jmzhao}@wust.edu.cn, chu-min.li@u-picardie.fr
huajiang@ynu.edu.cn, brooklet60@hust.edu.cn

Abstract

Maximum Common Induced Subgraph (MCIS) is an important NP-hard problem with wide real-world applications. An efficient class of MCIS algorithms uses Branch-and-Bound (BnB), consisting in successively selecting vertices to match and pruning when it is discovered that a solution better than the best solution found so far does not exist. The method of selecting the vertices to match is essential for the performance of BnB. In this paper, we propose a new value function and a hybrid selection strategy used in reinforcement learning to define a new vertex selection method, and propose a new BnB algorithm, called McSplitDAL, for MCIS. Extensive experiments show that McSplitDAL significantly improves the current best BnB algorithms, McSplit+LL and McSplit+RL. An empirical analysis is also performed to illustrate why the new value function and the hybrid selection strategy are effective.

Introduction

Graphs have gained increasing attention in recent decades due to their natural expression in representing numerous real-world problems. Given two graphs $G_p(V_p, E_p)$ and $G_t(V_t, E_t)$, the *Maximum Common Induced Subgraph* (MCIS) problem is to find an induced subgraph G'_p of G_p and an induced subgraph G'_t of G_t , such that G'_p and G'_t are isomorphic and have the maximum number of vertices. MCIS allows to evaluate the similarity of two graphs and has broad applications in many domains, such as graph database systems (Yan, Yu, and Han 2005), biochemistry (Bonnici et al. 2013; Larsen and Baumbach 2017), malware detection (Park, Reeves, and Stamp 2013; Sun et al. 2021), cheminformatics (Raymond and Willett 2002; Antelo-Collado et al. 2020; Schmidt et al. 2021), computer vision (Solnon et al. 2015), communication networks (Nirmala, Sulochana, and Rethnasamy 2016), *etc.* There are also many MCIS variant problems, such as the Maximum Common Connected Induced Subgraph (MCCIS) problem and the Subgraph Isomorphic (SI) problem.

MCIS is NP-hard and thus computationally challenging. Despite its NP-hardness, many methods have been developed to solve MCIS, including exact and inexact algorithms.

An important class of exact methods exploits the powerful Branch-and-Bound (BnB) framework (Raymond and Willett 2002; McCreesh, Prosser, and Trimble 2017; Liu et al. 2020; Zhou et al. 2022) to travel the whole search tree and try to match each vertex of G_p with each vertex of G_t (*i.e.*, branching) to find the best matches. The key to design an efficient BnB algorithm is to reduce the search space, using techniques such as effective branching heuristic (Englert and Kovács 2015; Zhou et al. 2022) and powerful constraint filtering (Solnon et al. 2015; McCreesh et al. 2016; McCreesh, Prosser, and Trimble 2017; Schmidt et al. 2021). When exact solutions are not required, one can use inexact algorithms to find approximate solutions of acceptable quality without exhausting the search space. These inexact methods include meta-heuristics (Rutgers et al. 2010; Choi, Yoon, and Moon 2012) and spectra methods (Bai, Hancock, and Wilson 2009). Recently, classification technologies of machine learning (Zanfir and Sminchisescu 2018; Bai et al. 2021) are also adopted to solve MCIS.

Some state-of-the-art BnB algorithms combine the advantages of both search and reinforcement learning techniques to improve their branching methods to efficiently reduce the search space. Based on the BnB algorithm framework, McSplit (McCreesh, Prosser, and Trimble 2017) uses a partition method to satisfy the isomorphic constraint and a branching heuristic based on the obtained partition and vertex degree to minimize the search tree size. McSplit+RL (Liu et al. 2020) explores the vertex pair selection policy based on reinforcement learning with a value function for each vertex so that it can reach a leaf node of search tree as early as possible. McSplit+LL (Zhou et al. 2022) further proposes a long-short memory and leaf vertex union match to improve the performance of a BnB MCIS algorithm.

We observe that the learning policies in McSplit+RL and McSplit+LL only concentrate on the reduction of upper bound due to a branching. They only reward branching vertices of input graphs with upper bound reduction and select a new branching vertices in the decreasing order of their accumulative rewards, that make the algorithm to branch only on a small set of vertices so as to waste time on local optima.

To remedy these limitations, we propose a new value function, namely Domain Action Learning (DAL), for evaluating each branching, that considers both upper bound re-

*Corresponding author.

duction and real graph simplification due to a branching action. We further propose a hybrid vertex selection strategy based on different value functions to guide the search. In fact, the best branching selected by different value functions is usually different. Branching alternatively according to two value functions allows to effectively diversify the search.

Based on the above new value function and the hybrid vertex selection strategy, we propose a new BnB algorithm, termed McSplitDAL, on top of McSplit (McCreesh, Prosser, and Trimble 2017). Experiments are conducted to evaluate McSplitDAL on 24,761 instances derived from diverse applications. The experimental results show that McSplitDAL significantly outperforms McSplit+RL and McSplit+LL that are already highly efficient. We also conduct empirical analysis and provide insight on why the proposed algorithm is effective.

This paper is organized as follows. Section 2 gives some basic graph definitions and related concepts used in this paper. Section 3 reviews existing BnB algorithms and learning methods for MCIS. Section 4 describes our new value function and hybrid vertex selection policy, followed by our McSplitDAL algorithm. Section 5 presents the empirical results and analysis. Section 6 concludes.

Preliminaries

Consider a simple, undirected and unlabelled graph $G = (V, E)$, where V is the set of vertices, $E \subseteq V \times V$ is the set of edges. Two vertices u and v are adjacent (or neighbors) if $(u, v) \in E$. The degree of a vertex v is the number of its adjacent vertices. A subgraph of G induced by a vertex subset $V' \subseteq V$ is defined by $G[V'] = (V', E')$, where $E' = \{(u, v) \in E | u, v \in V'\}$.

Given a graph $G_p = (V_p, E_p)$ (named pattern graph) and a graph $G_t = (V_t, E_t)$ (named target graph), if there exist an induced subgraph $G'_p = (V'_p, E'_p)$ of G_p , an induced subgraph $G'_t = (V'_t, E'_t)$ of G_t , and a bijection $\phi : V'_p \rightarrow V'_t$, such that any v and v' of V'_p are adjacent in G_p if and only if $\phi(v)$ and $\phi(v')$ of V'_t are adjacent in G_t , then G'_p or G'_t is called an induced common subgraph of G_p and G_t . In this case, we say that v and $\phi(v)$ are matched and $(v, \phi(v))$ is a match. The *Maximum Common Induced Subgraph (MCIS) problem* is to find a common induced subgraph of G_p and G_t with the maximum number of vertices. Let $V'_p = \{v_1, v_2, \dots, v_{|V'_p|}\}$, a feasible solution of MCIS can be represented as a set of matched pairs $\{(v_1, w_1), (v_2, w_2), \dots, (v_{|V'_p|}, w_{|V'_p|})\}$, where $w_j = \phi(v_j)$ for $j \in \{1, \dots, |V'_p|\}$.

A variant of MCIS called the Maximum Common Connected Induced Subgraph (MCCIS) problem requires that the maximum common induced subgraph is connected. Another variant called Subgraph Isomorphism (SI) requires $V'_p = V_p$.

A common induced subgraph G'_p is maximal if it cannot be extended to a larger common induced graph. If a feasible solution $\{(v_1, w_1), (v_2, w_2), \dots, (v_{|V'_p|}, w_{|V'_p|})\}$ is not maximal, it induces a nonempty set of s vertex subset pairs $Ev = \{(V_{1p}, V_{1t}), \dots, (V_{sp}, V_{st})\}$, called an environment of MCIS, where for $1 \leq i \leq s$, $V_{ip}(V_{it}) \subseteq V_p(V_t)$ and for




Figure 1: An illustration of an environment of MCIS and its related concepts. A non-maximal common induced subgraph $\{(0, a), (1, b)\}$ induces a set of vertex subset pairs $Ev = \{(\langle 3 \rangle, \langle d \rangle), (\langle 4, 2 \rangle, \langle f \rangle), (\langle 5 \rangle, \langle c, e \rangle)\}$, in which the vertex subset pair $(\langle 4, 2 \rangle, \langle f \rangle)$ is labelled ‘10’, because vertices 4 and 2 are all adjacent to 0 and f is adjacent to a , while 4 and 2 are all non-adjacent to 1 and f is non-adjacent to b . Other labels in the graphs are interpreted similarly. Vertices with the same color (or the same label) are in the same domain. Note that there is no vertex labelled with ‘00’ in G_p , so that vertices labelled with ‘00’ in G_t cannot be matched with any vertex in G_p . In fact, this Ev can provide at most 3 additional matches.

any $i' \neq i$, $V_{i'p}(V_{i't}) \cap V_{ip}(V_{it}) = \emptyset$, with the following property (McCreesh, Prosser, and Trimble 2017):

- For any $1 \leq i \leq s$ and any $1 \leq j \leq |V'_p|$, either all vertices in V_{ip} are adjacent to v_j and all vertices in V_{it} are adjacent to w_j , or all vertices in V_{ip} are non-adjacent to v_j and all vertices in V_{it} are non-adjacent to w_j .

As illustrated in Figure 1, if V_{ip} and V_{it} are not empty for some $1 \leq i \leq s$, choosing any vertex v in V_{ip} and w in V_{it} allows to extend the induced common subgraph by this vertex pair. Clearly, Ev can be defined to be $\{(V_p, V_t)\}$ when the common induced subgraph is empty.

For any $1 \leq i \leq s$ and any $1 \leq j \leq |V'_p|$, since all vertices in V_{ip} and in V_{it} have the same (non-)adjacency to v_j and w_j , we can use a bit 1 (0) to say that all vertices in V_{ip} or V_{it} are (non-)adjacent to v_j or w_j , respectively. So, a vertex subset pair (V_{ip}, V_{it}) can be labelled using a $|V'_p|$ -bit string, in which the j^{th} bit indicates whether vertices in V_{ip} and V_{it} are adjacent to v_j and w_j (McCreesh, Prosser, and Trimble 2017).

If for some $1 \leq i_1 < i_2 \leq s$, the pairs (V_{i_1p}, V_{i_1t}) and (V_{i_2p}, V_{i_2t}) have the same label, then they should be combined into one pair. Any pair (V_{ip}, V_{it}) in which V_{ip} or V_{it} is empty is removed from Ev . So, we assume the pairs in Ev cannot have the same labels and that there exists no pair such that one of its sets of vertices is empty.

Consequently, for any $1 \leq i_1 < i_2 \leq s$, any vertex v in $V_{i_1p}(V_{i_2p})$ cannot be matched with any vertex w in $V_{i_2t}(V_{i_1t})$ such that $G'_p = G_p[\{v_1, \dots, v_{|V'_p|}\}]$ extended with v and $G'_t = G_t[\{w_1, \dots, w_{|V'_p|}\}]$ extended with w remain isomorphic, because there is a j ($1 \leq j \leq |V'_p|$) such that v

is adjacent to v_j , but w is not adjacent to w_j , or vice versa. Thus, the sum $\sum_{(V_{ip}, V_{it}) \in Ev} \min(|V_{ip}|, |V_{it}|)$ provides an upper bound of the number of vertices that can be added into the common induced subgraph G'_p and G'_t .

In this paper, each labelled pair (V_{ip}, V_{it}) is called a domain, because it specifies two sets of vertices that can be matched, and Ev is called an environment in which a BnB MCIS algorithm works.

Search for MCIS and Learning Policy

This section first presents a state-of-the-art BnB search framework as shown in Algorithm 1, which allows an exploration of search space and enforces the isomorphism constraint. Thus, it serves as the backbone of McSplit+RL (Liu et al. 2020), McSplit+LL (Zhou et al. 2022) and our McSplitDAL. Then, we review representative learning policies that tell a BnB algorithm how to select a branching pair.

Branch and Bound for MCIS

To simplify the description, we suppose that two input graphs are undirected and unlabelled, and search methods can be easily extended to other kinds of graphs (McCreesh et al. 2016).

Given a pattern graph G_p and a target graph G_t , the BnB algorithm MCIS depicted in Algorithm 1 works with an environment Ev (i.e., a set of domains), a policy π_v to select a vertex in G_p , a policy π_w to select a vertex w in G_t to match with v , a current growing solution $curSol$, and the best solution $MaxSol$ found so far. At the beginning, $curSol$ and $MaxSol$ are both empty, and $Ev = \{(V_p, V_t)\}$ contains only one domain, meaning that every vertex in V_p is a candidate to match every vertex in V_t .

MCIS first estimates an upper bound on the number of matches that can be found with the current Ev . If the UB is not larger than the size of the best solution $MaxSol$ found so far, the algorithm prunes this branch and backtracks (Line 1–4). Otherwise, the algorithm selects a new vertex pair (v, w) such that $v \in V_p$ and $w \in V_t$ to match using policy π_v and π_w respectively. As a consequence of matching v with w , (v, w) is added into $curSol$, and each domain in Ev is split into two domains D_1 (D_2): domain in which each vertex of V_p is (non-)adjacent to v and each vertex of V_t is (non-)adjacent to w (Line 10–14). Domains with at least one empty vertex subset are removed. Afterwards, the algorithm runs recursively on the new domains (Line 15). After finishing the search of the subtree rooted at (v, w) , the algorithm tries to match v with other vertices in V_t (Line 16) selected using policy π_w . Then, it removes v from Ev and runs recursively. At last, the optimal solution is returned (Line 18–20).

The policies π_v and π_w are both based on a $selectD(\cdot)$ function that returns a domain from Ev , in which v and w are selected. (McCreesh, Prosser, and Trimble 2017) provide a $selectD(\cdot)$ function, by defining the size of a domain (V_{ip}, V_{it}) to be $\max(|V_{ip}|, |V_{it}|)$ and returning the domain with the smallest size from Ev , with ties broken by the largest vertex degree in V_{ip} . This function is used in Algorithm 1.

Algorithm 1: MCIS($Ev, \pi_v, \pi_w, curSol, MaxSol$)

Input: a domain set Ev ; policies π_v and π_w for selecting the matching pair (v, w) ; the current solution $curSol$ and the best solution found so far $MaxSol$

Output: $MaxSol$

```

1:  $UB \leftarrow |curSol| + \sum_{(V_{ip}, V_{it}) \in Ev} \min(|V_{ip}|, |V_{it}|)$ 
2: if  $UB \leq |MaxSol|$  then
3:   return  $MaxSol$ 
4: end if
5:  $(V_{ip}, V_{it}) \leftarrow selectD(Ev)$ 
6:  $v \leftarrow selectV(V_{ip}, \pi_v)$ 
7: for  $k$  in range( $|V_{it}|$ ) do
8:    $w \leftarrow selectW(V_{it}, \pi_w)$ 
9:    $V_{it} \leftarrow V_{it} \setminus \{w\}$ 
10:   $curSol \leftarrow curSol \cup \{(v, w)\}$ 
11:  if  $|curSol| > |MaxSol|$  then
12:     $MaxSol \leftarrow curSol$ 
13:  end if
14:   $Ev' \leftarrow$  a new domain set obtained by splitting domains in  $Ev$ 
15:   $MaxSol \leftarrow MCIS(Ev', \pi_v, \pi_w, curSol, MaxSol)$ 
16:   $curSol \leftarrow curSol \setminus \{(v, w)\}$ 
17: end for
18:  $Ev' \leftarrow$  a new domain set by removing  $v$  from  $Ev$ 
19:  $MaxSol \leftarrow MCIS(Ev', \pi_v, \pi_w, curSol, MaxSol)$ 
20: return  $MaxSol$ 

```

Related Learning Policy

In a BnB algorithm for MCIS, the branching heuristic to select v and w to match is crucial to reduce the size of the search tree. Early heuristics mainly rely on the properties of input graphs (Solnon et al. 2015; Englert and Kovács 2015; Bonnici and Giugno 2017; McCreesh, Prosser, and Trimble 2017) and focus on selecting v , while w is selected in turn to be matched with v in their natural order or the decreasing degree order. For instance, the degree heuristic first matches the vertex with the highest degree (Solnon et al. 2015). The degree-weighted-domains heuristic is to select a vertex with the greatest degree in the smallest domain (Boussemart et al. 2004). The neighbourhood heuristic selects a vertex that is a neighbor of the current partial order of matched pairs (Cibej and Mihelic 2014). McSplit (McCreesh, Prosser, and Trimble 2017) first selects a domain in Ev with the smallest $\max(|V_{ip}|, |V_{it}|)$ value, and then the vertex v (w) with the greatest degree in V_{ip} (V_{it}).

Recent heuristics use reinforcement learning to improve the branching heuristic of McSplit. They regard the BnB algorithm as an agent having a goal of reaching a search tree leaf as soon as possible. An action of the agent is to match a vertex v in V_p with a vertex w in V_t . A value function is defined based on a reward given to each performed action, then reinforcement learning is used to recognize the best action to choose at each step based on the accumulative rewards of each action received in the past. So, the key issue here is how to define a reward and a value function, and how to exploit them to select an action.

McSplit+RL (Liu et al. 2020) defines the reward of match-

ing (v, w) to be the upper bound reduction produced by the matching. Then, both v and w receive this reward. The policy π_v selects v with the highest accumulative rewards in the smallest domain (*i.e.*, a domain with the smallest size as defined in McSplit) and w is selected in the same domain in the decreasing order of their accumulative rewards, to be matched with v in turn.

As can be seen in Algorithm 1, the algorithm reaches a leaf when $UB \leq |MaxSol|$. Therefore, picking a vertex with the greatest accumulated reductions of upper bound can help McSplit+RL reach a leaf quickly. We refer the policies of McSplit+RL to select v and w to match by *RL*.

McSplit+LL (Zhou et al. 2022) further reduces the size of the search tree with Long-Short Memory (LSM) and Leaf vertex Union Match (LUM) techniques, which employs the same matching reward as McSplit+RL. But McSplit+LL manages the vertex value differently from McSplit+RL. Specifically, LSM records the accumulative rewards of each vertex in G_p , and the accumulative rewards of each vertex pair (v_i, w_j) matched in the past. At each step, it picks the vertex v in G_p with the greatest accumulative rewards in the smallest domain, as McSplit+RL does, then picks a vertex w in G_t in the same domain such that the vertex pair (v, w) has the greatest accumulative rewards among $\{(v, w_1), (v, w_2), \dots, (v, w_{|V_t|})\}$. LUM is to simultaneously match the leaf neighbors of v to the leaf neighbors of w after matching (v, w) . The leaf neighbor of a vertex is its neighbor with degree 1.

We observe that there are two limitations in the above learning policies. First, the reward for a matching action is only defined by its effect on upper bound. However, consider two possible matches (v, w) and (v', w') . The graph simplification due to these two matches can be very different, even if they give the same upper bound reduction. Second, these policies tend to produce a kind of ‘‘Matthew effect’’¹: the vertices with high accumulated rewards will be chosen again and again upon backtracking, and get their accumulated rewards higher, while the vertices with low accumulated rewards have little chance to be chosen and their accumulated rewards stay low. The Matthew effect can make the algorithm mainly branch on a small subset of vertices, so that the search wastes more time and is less efficient.


In the next section, we will propose a new value function based on a new reward definition, and a new hybrid vertex selection strategy to overcome these two limitations.

Proposed Method


In this section, we define a new reward to an action of matching a vertex v in G_p and a vertex w in G_t , in order to reflect more accurately the consequence of the action, further obtain a new vertex selection policy. We then propose a hybrid strategy combining the new vertex selection policy with *RL*, the policy of McSplit+RL, allowing to overcome the Matthew effect of a single policy.

Note that the tabular method in reinforcement learning is not directly applicable to BnB MCIS algorithms, because


¹See https://en.wikipedia.org/wiki/Matthew_effect



(I) $Ev = \{(\langle 2, 3 \rangle, \langle b, c, d \rangle), (\langle 1, 4, 5, 6, 7 \rangle, \langle e, g, f \rangle)\}$, induced by $\{(0, a)\}$



(II) $Ev' = \{(\langle 2 \rangle, \langle d \rangle), (\langle 3 \rangle, \langle b, c \rangle), (\langle 4, 6, 7 \rangle, \langle f \rangle), (\langle 5 \rangle, \langle g \rangle)\}$, induced by $\{(0, a), (1, e)\}$



(III) $Ev'' = \{(\langle 2 \rangle, \langle c \rangle), (\langle 1, 4, 5, 6, 7 \rangle, \langle e, g, f \rangle)\}$, induced by $\{(0, a), (3, b)\}$

Figure 2: An example of modifying environment by a match, where $Ev = \{(\langle 2, 3 \rangle, \langle b, c, d \rangle), (\langle 1, 4, 5, 6, 7 \rangle, \langle e, g, f \rangle)\}$ is induced by matching vertex 0 with vertex a . Then it transforms into a new domain set containing simpler subgraphs by matching vertex 1 with vertex e than by matching vertex 3 with vertex b . The vertices with the same color are in the same domain.

there are too many states and actions to store, and approximate functions needs lots of computation to fit state value or action value for MCIS.

New Value Function

A BnB algorithm as specified in Algorithm 1 works with a set $Ev = \{(V_{1p}, V_{1t}), \dots, (V_{sp}, V_{st})\}$ of domains, where each vertex subset pair (V_{ip}, V_{it}) ($1 \leq i \leq s$) is a domain. We

say that Ev is the environment that the learning agent stays. Each environment induces an upper bound of the number of matches that can be added into the current growing solution. An action matching a vertex v and a vertex w changes the current environment. The changed environment induces a new upper bound. The difference between the old upper bound and the new one is used as the reward to the action (v, w) in McSplitRL and McSplitLL. We argue that this reward is not accurate enough for an action, because two actions inducing the same upper bound reduction can result in different domains. Figure 2 gives an illustrative example.

Example 1 Figure 2 shows two graphs G_p and G_t , $V_p = \{0, 1, 2, 3, 4, 5, 6, 7\}$, $V_t = \{a, b, c, d, e, f, g\}$. The vertex pair $(0, a)$ has been first matched, so that the current $Ev = \{(\langle 2, 3 \rangle, \langle b, c, d \rangle), (\langle 1, 4, 5, 6, 7 \rangle, \langle e, g, f \rangle)\}$, that can at most provide $2 + 3 = 5$ vertex pairs to extend current solution $\{(0, a)\}$. Thus, the current upper bound is 5.

If the second matching is $(1, e)$, Ev will be modified into $Ev' = \{(\langle 2 \rangle, \langle d \rangle), (\langle 3 \rangle, \langle b, c \rangle), (\langle 4, 6, 7 \rangle, \langle f \rangle), (\langle 5 \rangle, \langle g \rangle)\}$. The upper bound induced by Ev' is 4 and $|Ev'| = 4$.

If the branching strategy picks $(3, b)$ instead of $(1, e)$ as the second matching, Ev will be modified into $Ev'' = \{(\langle 2 \rangle, \langle c \rangle), (\langle 1, 4, 5, 6, 7 \rangle, \langle e, g, f \rangle)\}$. The upper bound induced by Ev'' is 4 and $|Ev''| = 2$.

The two matches $(1, e)$ and $(3, b)$ induce the same upper bound reduction. However, the environment change from Ev to Ev' induced by the matching $(1, e)$ is clearly more important than the change from Ev to Ev'' induced by the matching $(3, b)$, because the domains in Ev are split into more new domains, meaning that the problem is more simplified by the branching on $(1, e)$. We use the number of domains in Ev' or Ev'' to measure the environment change. Intuitively, the more the environment changes, the more the subproblem is easier to solve.

Based on the above observation, we propose a new reward defined in Equation 1, where Ev' is the environment modified from Ev by the match (v, w) .

$$R(v, w) = \sum_{(V_{ip}, V_{it}) \in Ev} \min(|V_{ip}|, |V_{it}|) - \sum_{(V'_{ip}, V'_{it}) \in Ev'} \min(|V'_{ip}|, |V'_{it}|) + |Ev'| \quad (1)$$

Equation 1 uses both the upper bound reduction and the number of domains contained in the new environment Ev' to reward an action (v, w) . A new value function called Domain and Action Learning (DAL) is defined by Equation 2 and Equation 3.

$$DAL(v) \leftarrow DAL(v) + R(v, w) \quad (2)$$

$$DAL(v, w) \leftarrow DAL(v, w) + R(v, w) \quad (3)$$

The DAL value function considers both the upper bound reduction and the number of domains contained in the new environment Ev' . A greater upper bound reduction presumably allows to prune the search earlier. A greater number

of domains in Ev' presumably implies a subproblem easier to solve. Our purpose is to combine the two advantages to speed up the search.

A new vertex selection policy is thus defined using the DAL value function, which gives the score $DAL(v)$ for each vertex v in V_p and the score $DAL(v, w)$ for each match (v, w) , all initialized to 0. Then at each step, after selecting the smallest domain (V_{ip}, V_{it}) from the current environment Ev in Algorithm 1 (Line 5), Equation 2 is used to select a vertex v in V_{ip} with the highest $DAL(v)$, and Equation 3 is used to select a vertex w in V_{it} such that $DAL(v, w)$ is maximum. After matching v with w , $R(v, w)$ defined in Equation 1 is added into $DAL(v)$ and $DAL(v, w)$.

As in McSplitLL, if $DAL(v)$ and $DAL(v, w)$ reach T_v and T_{vw} , respectively, where T_v and T_{vw} are two parameters as in McSplitLL, all vertex values in $DAL(v)$ and $DAL(v, w)$ decay to a half.

Hybrid Branching Policy

As is explained in the previous section, the current vertex selection policies based on reinforcement learning can suffer from the Matthew effect, so does the new vertex selection policy based on the DAL value function.

In order to overcome the Matthew effect, we propose to hybrid the RL policy of McSplitRL and the DAL policy defined in this paper. Concretely, let $\Pi \in \{RL, DAL\}$ denote the current policy, and be initialized to be RL. Every time v or w is selected using Π , a counter $NbApp$ is incremented by 1. When $NbApp$ reaches a fixed threshold $MaxNbApp$, it is reset to 0, and Π is changed to another policy in $\{RL, DAL\}$. An exception happens when a better solution is found. In this case, $NbApp$ is reset to 0, and the same policy continues to be used.

Note that the vertices with the highest value DAL or RL are usually different. The hybrid branching policy allows to branch on different vertices, thus diversifying the search while keeping good quality branchings.

Experiments

The proposed algorithm McSplitDAL is implemented in C++ on top of McSplit and compiled using g++. We conduct experiments to evaluate the new algorithm and the proposed strategies. All experiments were performed on Intel Xeon CPUs E5-2680 v4@2.40 G under Linux with 4G memory.

The three parameters T_v , T_{vw} and $MaxNbApp$ are set to 10^5 , 10^9 , $2 \times \min(|V_p|, |V_t|)$, respectively.

Benchmarks

The benchmark datasets ² include 24,761 instances, which are divided into two sets.

- Biochemical reactions (Gay et al. 2014): including 136 directed unlabelled bipartite graphs. The number of vertices varies from 9 to 386. All graphs describe the biochemical reaction networks. This dataset provides 9316 instances by pairing any two graphs (including 136 self-match pairs).

²Available at <http://iris.cnrs.fr/csolnon/SIP.html>

- Large SI instances (Damiand et al. 2011; Solnon et al. 2015; Hoffmann, McCreesh, and Reilly 2017; McCreesh, Prosser, and Trimble 2017; Liu et al. 2020; Zhou et al. 2022): including 15,445 instances generated from the real-world problems or random models, such as segmented images, modelling 3D objects, and scale-free networks. Specifically, this instance set contains: 6,278 *Images-CVIU11*, 1225 *LV*, 3,430 *LargerLV*, 24 *Image-PR15*, 1170 *SI*, 100 *Scalefree*, 3018 *Meshes-CVIU11* and 200 *phase*. The number of vertices varies from 22 to 6,671.

The time limit for each instance in the experiments is 1800 seconds.

Solvers

We compare the new algorithm McSplitDAL with two state-of-the-art BnB algorithms: McSplit+LL (Zhou et al. 2022) and McSplit+RL (Liu et al. 2020). To better understand the proposed policies, four variants of these algorithms are also included in the experiments.

- McSplitDAL: our implementation of Algorithm 1 on top of McSplit (McCreesh, Prosser, and Trimble 2017) with the new value function *DAL* and the hybrid vertex selection policy $\Pi \in \{RL, DAL\}$.

- McSplit+RL (Liu et al. 2020): An implementation of the Algorithm 1 on top of McSplit with the value function *RL*, which significantly improves McSplit.

- McSplit+LL (Zhou et al. 2022): An implementation of the Algorithm 1 on top of McSplit with the LSM and LUM techniques.

- McSplitRLD: a variant of McSplitRL using the new *DAL* value function instead of the *RL* policy.

- McSplitLLD: a variant of McSplit+LL using the new *DAL* value function instead of its own policy.

- McSplitDAL+rand: A variant of McSplitDAL, which applies one of two branching policies $\{RL, DAL\}$ in random at each branch node, instead of applying each policy *MaxNbApp* times alternatively.

- McSplitDAL+depth: A variant of McSplitDAL, which changes the policy according to the depth of the search tree, instead of applying each policy *MaxNbApp* times alternatively. Concretely, let $Maxdep = \min(|V_p|, |V_t|)$. When the tree depth is in range of $[1, \frac{1}{4}Maxdep]$ and $[\frac{1}{2}Maxdep, \frac{3}{4}Maxdep]$, McSplitDAL+depth uses *RL* policy. Otherwise, it uses *DAL* policy.

Comparison of Performance

The first experiment compares the general performance of McSplit+RL, McSplit+LL and McSplitDAL on the benchmarks, excluding the too easy instances that can be solved by all the compared solvers within 10 seconds and the too hard instances that cannot be solved by any compared solver within the time limit to make the comparison clearer. The average runtimes of McSplit+RL, McSplit+LL and McSplitDAL on the excluded easy instances are 0.59s, 0.58s and 0.54s, respectively.

Figure 3 shows the cumulative numbers of solved instances by the compared three solvers over the remaining 2,229 instances. McSplitDAL solves 292 (437)




Figure 3: Cumulative numbers of instances solved by McSplit+RL, McSplit+LL and McSplitDAL on the 2,229 MCIS instances.




Figure 4: Cumulative plots of McSplit+RL, McSplit+LL and McSplitDAL on recursive calls for proving the optimality of the found solution.

more instances than McSplit+LL (McSplit+RL). In other words, McSplitDAL solves 16.3% more instances than McSplit+LL. Note that McSplitDAL, McSplit+LL and McSplit+RL all share the same implementation of Algorithm 1 and the unique difference between McSplitDAL and McSplit+LL is the branching heuristic, while the difference between McSplit+LL and McSplit+RL includes the branching heuristic and the LUM technique. However, the performance improvement of McSplitDAL w.r.t. McSplit+LL is greater than the performance improvement of McSplit+LL w.r.t. McSplit+RL. Considering the high performance of baseline algorithms and the NP-hardness of MCIS, the results show that new value function *DAL* and hybrid branching strategy are very effective for BnB MCIS algorithms.

Further Analysis

The search process of an exact MCIS algorithm can be divided into two phases: find an optimal solution (Finding time) and prove it is optimal (Proving time). Total time of solving an instance is the sum of finding time and proving time. The experimental results in Figure 3 explain partially why the hybrid learning policy based on the new value function improves the McSplit+LL and McSplit+RL for MCIS.




Figure 5: Cumulative numbers of instances solved by McSplit+RL, McSplit+LL, McSplitRLD, McSplitLLD and McSplitDAL on 2,361 MCIS instances.




Figure 6: Cumulative numbers of solved instances by McSplitDAL, McSplitDAL+rand and McSplitDAL+depth on 2,277 MCIS instances.

For a BnB MCIS algorithm, it is easier to reach the pruning condition if the optimal solution is found earlier. Figure 3 shows that McSplitDAL generally finds optimal solutions earlier than McSplit+LL and McSplit+RL, due to the effectiveness of the new value function DAL and the hybrid vertex selection strategy.

Figure 4 shows the number of recursive calls of McSplitDAL, McSplit+LL and McSplit+RL for proving the optimality of the found solution on the same instances as in Figure 3. The number of recursive calls of McSplitDAL is clearly the smallest, suggesting that the new value function DAL and the hybrid vertex selection policy in McSplitDAL allows better branching and is also efficient to overcome Matthew effect of a single policy, so that McSplitDAL diversifies better search than McSplit+RL and McSplit+LL, and significantly reduces the number of recursive calls.

Ablation Study

To further access the effectiveness of the proposed value function, we compare the performance of McSplit+RL with McSplitRLD, the performance of McSplit+LL with McSplitLLD, and with the performance of McSplitDAL. The results are showed in Figure 5 (after excluding the too easy instances solved by all the 5 solvers within 10s).

Recall that the only difference of McSplit+RL (McSplit+LL) and McSplitRLD (McSplitLLD) is the value function, for McSplitRLD and McSplitLLD do not employ the hybrid vertex selection strategy. As Figure 5 shows, McSplitLLD (McSplitRLD) solves 154 (138) more instances than McSplit+LL (McSplit+RL). In other words, McSplitLLD (McSplitRLD) solves 8.4% (8.2%) more instances than McSplit+LL (McSplit+RL). So, the results in Figure 5 show that the new value function DAL is indeed more effective for MCIS, because DAL considers both upper bound reduction and environment changes, while the policies in McSplit+LL and McSplit+RL only consider upper bound reduction.

Note that McSplitDAL solves 138 (299) more instances than McSplitLLD (McSplitRLD), thanks to the hybrid vertex selection strategy (cf. Figure 3). For 2087 medium instances solved by McSplitDAL in Figure 3, there are 1847 instances that use DAL policy more times than RL policy in the search process.

The hyper-parameter $MaxNbApp$ in the switching policy conditions is important to McSplitDAL. We leverage McSplitDAL, McSplitDAL+rand and McSplitDAL+depth to evaluate the hybrid policy. Figure 6 shows the comparison of performance of the three solvers over 2,277 instances (after excluding the too easy instances and the too hard instances w.r.t. the three compared algorithms as before). The comparison shows that McSplitDAL has the best performance, solving 121 and 286 more instances than McSplitDAL+rand and McSplitDAL+depth, respectively. The goal of MCIS algorithm is to find an optimal solution which size is at most $\min(|V_p|, |V_t|)$. Experimental results show the switch condition related to the optimal solution size have a better performance than random choice and fixed tree depth.

Conclusion

In this paper, we propose a new value function and a hybrid branching strategy in a Branch-and-Bound (BnB) algorithm based on reinforcement learning for the Maximum Common Induced Subgraph (MCIS) problem. The new value function considers both upper bound reduction and environment change to reward an action of matching two vertices. It allows to select vertices to better simplify the graphs. The hybrid branching strategy guides the search by employing alternatively two different branching heuristics to diversify the search and find optimal solutions earlier. We implement the new approaches into a BnB algorithm called McSplitDAL. Extensive experimental results show that the proposed methods significantly improve the efficiency of the BnB MCIS algorithm.

In the future, we will continue to study the interplay of the search and the learning, and try machine learning to design a scheme for policy selecting. These methods solving Mathew effect can be applied to solve other NP-hard problems.

Acknowledgements

This work is partially funded by the French Agence Nationale de la Recherche, reference ANR-19-CHIA-0013-01, the National Natural Science Foundation of China (Grants

U22B2017, 62162066), the Scientific Research Program of Hubei Provincial Department of Education for Youth (Grant Q20211111), the Open Research Fund of Hubei Society for Industrial and Applied Mathematics (Grant 2022003), and supported by the Matrics platform of Université de Picardie Jules Verne.

References

- Antelo-Collado, A.; Carrasco-Velar, R.; García-Pedrajas, N.; and García, G. C. 2020. Maximum common property: a new approach for molecular similarity. *J. Cheminformatics*, 12(1): 61.
- Bai, X.; Hancock, E. R.; and Wilson, R. C. 2009. A generative model for graph matching and embedding. *Comput. Vis. Image Underst.*, 113(7): 777–789.
- Bai, Y.; Xu, D.; Sun, Y.; and Wang, W. 2021. GLSearch: Maximum Common Subgraph Detection via Learning to Search. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, Virtual Event*, volume 139, 588–598. PMLR.
- Bonnici, V.; and Giugno, R. 2017. On the Variable Ordering in Subgraph Isomorphism Algorithms. *IEEE ACM Trans. Comput. Biol. Bioinform.*, 14(1): 193–203.
- Bonnici, V.; Giugno, R.; Pulvirenti, A.; Shasha, D. E.; and Ferro, A. 2013. A subgraph isomorphism algorithm and its application to biochemical data. *BMC Bioinform.*, 14(S-7): S13.
- Boussemart, F.; Hemery, F.; Lecoutre, C.; and Sais, L. 2004. Boosting Systematic Search by Weighting Constraints. In *Proceedings of the 16th European Conference on Artificial Intelligence, ECAI 2004, Valencia*, 146–150. IOS Press.
- Choi, J.; Yoon, Y.; and Moon, B. R. 2012. An efficient genetic algorithm for subgraph isomorphism. In *Genetic and Evolutionary Computation Conference, GECCO 2012, Philadelphia, PA*, 361–368. ACM.
- Cibej, U.; and Mihelic, J. 2014. Search Strategies for Subgraph Isomorphism Algorithms. In *Applied Algorithms - First International Conference, ICAA 2014, Kolkata*, volume 8321, 77–88. Springer.
- Damiand, G.; Solnon, C.; de la Higuera, C.; Janodet, J.; and Samuel, É. 2011. Polynomial algorithms for subisomorphism of nD open combinatorial maps. *Comput. Vis. Image Underst.*, 115(7): 996–1010.
- Englert, P.; and Kovács, P. 2015. Efficient Heuristics for Maximum Common Substructure Search. *J. Chem. Inf. Model.*, 55(5): 941–955.
- Gay, S.; Fages, F.; Martinez, T.; Soliman, S.; and Solnon, C. 2014. On the subgraph epimorphism problem. *Discret. Appl. Math.*, 162: 214–228.
- Hoffmann, R.; McCreesh, C.; and Reilly, C. 2017. Between Subgraph Isomorphism and Maximum Common Subgraph. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, AAAI 2017, California*, 3907–3914. AAAI Press.
- Larsen, S. J.; and Baumbach, J. 2017. CytoMCS: A Multiple Maximum Common Subgraph Detection Tool for Cytoscape. *J. Integr. Bioinform.*, 14(2).
- Liu, Y.; Li, C.; Jiang, H.; and He, K. 2020. A Learning Based Branch and Bound for Maximum Common Subgraph Related Problems. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, New York*, 2392–2399. AAAI Press.
- McCreesh, C.; Ndiaye, S. N.; Prosser, P.; and Solnon, C. 2016. Clique and Constraint Models for Maximum Common (Connected) Subgraph Problems. In *Principles and Practice of Constraint Programming the 22nd International Conference, CP 2016, Toulouse*, volume 9892, 350–368. Springer.
- McCreesh, C.; Prosser, P.; and Trimble, J. 2017. A Partitioning Algorithm for Maximum Common Subgraph Problems. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia*, 712–719. ijcai.org.
- Nirmala, P.; Sulochana, L. R.; and Rethnasamy, N. 2016. Vertex cover-based binary tree algorithm to detect all maximum common induced subgraphs in large communication networks. *Knowl. Inf. Syst.*, 48(1): 229–252.
- Park, Y. H.; Reeves, D. S.; and Stamp, M. 2013. Deriving common malware behavior through graph clustering. *Comput. Secur.*, 39: 419–430.
- Raymond, J. W.; and Willett, P. 2002. Maximum common subgraph isomorphism algorithms for the matching of chemical structures. *J. Comput. Aided Mol. Des.*, 16(7): 521–533.
- Rutgers, J. H.; Wolkotte, P. T.; Hölzenspies, P. K. F.; Kuper, J.; and Smit, G. J. M. 2010. An Approximate Maximum Common Subgraph Algorithm for Large Digital Circuits. In *13th Euromicro Conference on Digital System Design, Architectures, Methods and Tools, DSD 2010, Lille*, 699–705. IEEE Computer Society.
- Schmidt, R.; Krull, F.; Heinzke, A. L.; and Rarey, M. 2021. Disconnected Maximum Common Substructures under Constraints. *J. Chem. Inf. Model.*, 61(1): 167–178.
- Solnon, C.; Damiand, G.; de la Higuera, C.; and Janodet, J. 2015. On the complexity of submap isomorphism and maximum common submap problems. *Pattern Recognit.*, 48(2): 302–316.
- Sun, Y.; Bashir, A. K.; Tariq, U.; and Xiao, F. 2021. Effective malware detection scheme based on classified behavior graph in IIoT. *Ad Hoc Networks*, 120: 102558.
- Yan, X.; Yu, P. S.; and Han, J. 2005. Substructure Similarity Search in Graph Databases. In *Proceedings of the ACM SIGMOD International Conference on Management of Data, Baltimore, SIGMOD 2005, Maryland*, 766–777. ACM.
- Zanfir, A.; and Sminchisescu, C. 2018. Deep Learning of Graph Matching. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City*, 2684–2693. IEEE Computer Society.
- Zhou, J.; He, K.; Zheng, J.; Li, C.; and Liu, Y. 2022. A Strengthened Branch and Bound Algorithm for the Maximum Common (Connected) Subgraph Problem. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna*, 1908–1914. ijcai.org.