

Look-Ahead with Mini-Bucket Heuristics for MPE

Rina Dechter, Kaley Kask, William Lam
University of California, Irvine
Irvine, California, USA

Javier Larrosa
UPC Barcelona Tech
Barcelona, Spain

Abstract

The paper investigates the potential of look-ahead in the context of AND/OR search in graphical models using the Mini-Bucket heuristic for combinatorial optimization tasks (e.g., MAP/MPE or weighted CSPs). We present and analyze the complexity of computing the residual (a.k.a. Bellman update) of the Mini-Bucket heuristic and show how this can be used to identify which parts of the search space are more likely to benefit from look-ahead and how to bound its overhead. We also rephrase the look-ahead computation as a graphical model, to facilitate structure exploiting inference schemes. We demonstrate empirically that augmenting Mini-Bucket heuristics by look-ahead is a cost-effective way of increasing the power of Branch-And-Bound search.

Introduction

Look-ahead is known to be useful and even essential in the context of online search algorithms (e.g., game playing schemes such as alpha-beta, planning under uncertainty (Geffner and Bonet 2013; Vidal 2004) where the search space is enormous and traversing it to completion is out of the question. Look-ahead can improve the heuristic function h of a node by expanding the search tree below it to a certain depth d , evaluate the static h at tip nodes and roll back this information to n (known as a Bellman update).

Here we investigate the potential of look-ahead (Bu et al. 2014) in the different context of complete search. We consider the min-sum problem over a Graphical model (which includes, among others, the *most probable explanation* task of Markov Networks (Pearl 1988)). This problem is usually solved by depth-first search *Branch-and-Bound* (DFBnB). The search is guided by a heuristic function that provides an optimistic bound on the *cost of the best extension* of any partial assignment. It is often observed that these types of algorithms spend a significant time proving the optimality of their output (where the only use of h is pruning).

Our framework is AND/OR Branch-And-Bound (AOBB) guided by the static Mini-Bucket Elimination heuristic (MBE), which, assisted by variational cost-shifting methods, is one of the best current approaches for the min-sum problem (Marinescu and Dechter 2009; L. Otten and Dechter 2012). However, this heuristic can be weak, (e.g., mainly for

problems having a high induced width). A natural approach is to improve it via dynamic look-ahead during search itself.

In the context of depth-first BnB a better heuristic may *i)* allow finding an optimal solution early with less search (and thus the best upper bound for pruning), and *ii)* once the optimal bound is reached, it will help proving optimality by better pruning of the search space.

What we show in this paper is that in the context of graphical models and when using the mini-bucket heuristic, we can bound the computation of d -level look-ahead sufficiently, making it cost effective. This is accomplished by *i)* exploiting a relationship between the residual of the heuristic and a new concept of bucket error that can be pre-compiled, to facilitate a selective look-ahead process combined with effective pruning of the d -level lookahead tree whenever it is applied, and *ii)* using an inference-based message-passing algorithm to compute the look-ahead, thus exploiting the graphical model structure at each node.

In Section 2, we provide background. Section 3 introduces the notions of residual and local bucket-error of the mini-bucket heuristic that captures the mini-bucket approximation error. Section 4 presents the experiments and discussion. Section 5 concludes the paper.

Background

Graphical Models A *graphical model* is a tuple $\mathcal{M} = (\mathbf{X}, \mathbf{D}, \mathbf{F})$, where $\mathbf{X} = \{X_i : i \in V\}$ is a set of variables indexed by a set V and $\mathbf{D} = \{D_i : i \in V\}$ is the set of finite domains of values for each X_i , \mathbf{F} is a set of discrete functions. In this paper we focus on the *min-sum problem*, $C^* = \min_x \sum_{f \in F} f(x)$ which is applicable in a wide range of reasoning problems (e.g. MPE/MAP). The *scope* of f (i.e. the subset of \mathbf{X} relevant to f) is noted $scope(f)$. We will use lower case y to denote the assignment of a set of variables $Y \subseteq \mathbf{X}$. Thus, $f(x, Y)$ is a function with scope $X \cup Y$ partially assigned by x . In an abuse of notation we will sometimes write $f(y)$ with $scope(f) \subset Y$ assuming that the irrelevant parts of the assignment will be ignored.

Primal graph, pseudo-tree. The *primal graph* G of a graphical model \mathcal{M} has each variable X_i in a node and an edge (X_i, X_j) is in G iff the pair of variables appears in the scope of any $f \in \mathbf{F}$. A *pseudo-tree* $T = (V, E')$ of the




Figure 1

primal graph $G = (V, E)$ is a rooted tree over the same set of vertices V such that every edge in $E - E'$ connects a node to one of its ancestors in T .

Notation. Let T be a pseudo-tree. \bar{X}_p denotes X_p and its ancestors in T . $ch(X_p)$ denote the children of X_p in T . The subtrees rooted by each $ch(X_p)$ is denoted T_p . The same subtrees of T_p pruned below depth d is denoted $T_{p,d}$. We will often abuse notation and will consider a tree as the set of its nodes, or its variables.

Example. Figure 1a displays the primal graph of a graphical model with 10 binary cost functions, each one corresponds to an edge in the graph. Figure 1b displays a pseudo-tree. The ancestors of C are A and B . Its children are D and E . The grayed area corresponds to $T_{B,1}$.

AND/OR search. A state-of-the-art method to solve reasoning tasks on graphical models is to search their AND/OR graph (Marinescu and Dechter 2009). The AND/OR search tree, which is guided by a pseudo-tree T of the primal graph G , consists of alternating levels of OR and AND nodes. Each OR node is labeled with a variable $X_p \in \mathbf{X}$. Its children are AND nodes, each labeled with an instantiation x_p of X_p . The weight of the edge, $w(X_p, x_p)$, is the sum of costs of all the functions $f \in F$ that are completely instantiated at x_p but are not at its parent X_p . Children of AND nodes are OR nodes, labeled with the children of X_p in T . Each child represents a conditionally independent subproblems given assignments to their ancestors. Those edges are not weighted. The root

of the AND/OR search tree is the root of T . The path from an AND node labeled x_p to the root corresponds to a partial assignment to \bar{X}_p denoted \bar{x}_p .

A solution tree in the AND/OR tree is a subtree that (1) contains its root, (2) if an OR node is in the solution subtree, then exactly one of its children is in the solution tree, (3) if an AND node is in the solution tree, then all its children are. A solution tree corresponds to a complete assignment and its cost is the sum-cost of its arc weights.

The AND/OR search tree can be converted into a graph by merging nodes that root identical subproblems. It was shown that identical subproblems can be identified by their context (a partial instantiation that decomposes the subproblem from the rest of the problem graph), yielding an AND/OR search graph (Marinescu and Dechter 2009).

Example. Figure 1d displays the AND/OR search graph of the running example. A solution tree, corresponding to the assignment $(A = 0, B = 1, C = 1, D = 0, E = 0, F = 0, G = 0)$, is highlighted.

AND/OR Branch-And-Bound (AOBB) is a state-of-the-art algorithm for solving combinatorial optimization problems over graphical models (Marinescu and Dechter 2009). It explores the set of (partial) assignments by traversing the weighted AND/OR context-minimal search graph in a depth-first manner, while keeping track of the current best solution (upper bound, ub) of the subtree rooted at each node \bar{x}_p , denoted $ub(\bar{x}_p)$. Each AND node has an associated heuristic value $h(\bar{x}_p)$ which is a lower bound on the best cost

extension of \bar{x}_p to its descendants in T_p . Then whenever $h(\bar{x}_p) \geq ub(\bar{x}_p)$, search below \bar{x}_p can be safely pruned.

The heuristic is also used to guide the order in which OR nodes and AND nodes are expanded (namely, in which order independent subproblems are considered and in which order values for the current variable are considered).

Mini-Bucket Elimination Heuristic. Current implementations of AOBB (Kask and Dechter 2001; Marinescu and Dechter 2009) are guided by the mini-bucket heuristic. This heuristic is based on *Mini-Bucket Elimination* and is called $MBE(i)$ where i , called *i-bound*, allows trading off pre-processing time and space for heuristic accuracy with actual search. It belongs to the class of *static heuristics*, meaning that it is pre-processed before AOBB starts. $MBE(i)$ works relative to the same pseudo-tree T which defines the AND/OR search graph. Each node X_p of T is associated with a bucket B_p that includes a set of functions. The *scope of a bucket* is the union of scopes of all its functions before it is processed as described next. First, each function f of the graphical model is placed in a bucket B_p if X_p is the deepest variable in T s.t. $X_p \in \text{scope}(f)$. Then $MBE(i)$ processes the buckets of the pseudo-tree from leaves towards the root. Processing a bucket may require partitioning the bucket's functions into *mini-buckets* $B_p = \cup_r B_p^r$, where each B_p^r includes no more than i variables. Then, processing each mini-bucket separately, all its functions are combined (by sum in our case) and the bucket's variable (X_p) is eliminated (by min in our case). Each resulting new function, also called a *message*, is placed in the closest ancestor bucket whose variable is contained in its scope. $MBE(i)$ is time and space exponential in the i -bound.

Bucket Elimination (BE). In the special case in which the i -bound is high enough so there is no partitioning into mini-buckets, the algorithm is the exact bucket-elimination (BE) scheme (Dechter 2013). The time and space complexity of BE is exponential in the size of the largest bucket scope encountered which is called the *induced width* and is denoted w^* , for that ordering.

Notation. In the following, f_p denotes an original function placed in Bucket B_p (if there is more than one, they are all summed into a single function), and $\lambda_{j \rightarrow p}$ denotes a message created at bucket B_j and placed in bucket B_p . Processing bucket B_p produces messages $\lambda_{p \rightarrow i}$ for some ancestors X_i of X_p (i.e. $X_i \in \bar{X}_p - X_p$).

Example. Figure 1c illustrates the execution of $MBE(3)$ in our running example by means of the so-called mini-bucket tree (nodes are buckets and tree-edges show message exchanges). In this example, bucket B_D is the only one that needs to be partitioned into mini-buckets. Each mini-bucket generates a message. In the figure, messages are displayed along an edge to emphasize the bucket where they are generated and the bucket where they are placed. For instance, $\lambda_{D \rightarrow C}$ is generated in bucket D (as $\lambda_{D \rightarrow C} = \min_D \{f(B, D) + f(B, D)\}$) and placed in bucket C .

Extracting the mini-bucket heuristic (Kask and Dechter

1999). The (static) mini-bucket heuristic used by AOBB requires a $MBE(i)$ execution prior to search keeping all the message functions. Let \bar{x}_p be a partial assignment. Λ_j denotes the sum of the messages sent from bucket B_j to all of the instantiated ancestor variables,

$$\Lambda_j(\bar{x}_p) = \sum_{X_q \in \bar{X}_p} \lambda_{j \rightarrow q}(\bar{x}_p) \quad (1)$$

The heuristic value at node \bar{x}_p is,

$$h(\bar{x}_p) = \sum_{X_j \in T_p} \Lambda_j(\bar{x}_p) \quad (2)$$

Example. In our example, the heuristic function of partial assignment $(A = 0, B = 1)$ is $h(A = 0, B = 1) = \lambda_{D \rightarrow A}(A = 0) + \lambda_{C \rightarrow B}(B = 1) + \lambda_{F \rightarrow B}(A = 0, B = 1)$

AND/OR Look-Ahead

OR Look-Ahead In the field of *heuristic search* the general task is to find the least cost path in a weighted directed graph from one initial node to one of the possibly many goal nodes. The search is guided by an admissible heuristic function $h(n)$ which underestimates the least cost path from n to a goal state. More accurate $h(n)$ decreases the expanded search space (Nilsson 1980). If h is not accurate enough it can be improved by means of a look-ahead, a technique useful in game playing and in planning with uncertainty, especially for anytime solvers (Geffner and Bonet 2013; Vidal 2004).

DEFINITION 1 (look-ahead, residual). Let $ch(n)$ be the children of node n in the search graph and h be a heuristic function. The d look-ahead of n is,

$$h^d(n) = \min_{n_i \in ch(n)} \{w(n, n_i) + h^{d-1}(n_i)\}$$

where $w(n, n_i)$ is the weight of the arc, and $h^0 = h$. The look-ahead residual, defined by,

$$res(n) = h^1(n) - h(n)$$

measures the error of h with respect to the next level.

AND/OR Look-Ahead. Focusing now on AND/OR search for graphical models, let n be an AND node labeled by x_p terminating a partial assignment \bar{x}_p . Extending the definition of look-ahead to the next level of AND nodes (i.e. two levels in the AND/OR graph) we get:

$$h^d(\bar{x}_p) = \sum_{X_q \in ch(\bar{x}_p)} \min_{x_q} \{w(X_q, x_q) + h^{d-1}(\bar{x}_q)\} \quad (3)$$

Specializing to the MBE heuristic we can show

PROPOSITION 1. If h denotes the mini-bucket heuristic, \bar{x}_p is a partial assignment and $x_{p,d}^\downarrow$ is an extension of the assignment to all the variables in $T_{p,d}$

$$h^d(\bar{x}_p) = h(\bar{x}_p) + L^d(\bar{x}_p) - \sum_{X_k \in T_{p,d}} \Lambda_k(\bar{x}_p) \quad (4)$$

where,

$$L^d(\bar{x}_p) = \min_{x_{p,d}^\downarrow} \left\{ \sum_{X_k \in T_{p,d}} [f_k(\bar{x}_p, x_{p,d}^\downarrow) + \sum_{X_j \in T_p - T_{p,d}} \lambda_{j \rightarrow k}(\bar{x}_p, x_{p,d}^\downarrow)] \right\} \quad (5)$$

Proof. (sketch) The reformulation is obtained after expanding the recursive definition of lookahead (eq. 3) and replacing the generic heuristic by the MBE heuristic (eq. 2 and 1) \square

$L^d(\bar{x}_p)$ is a min-sum problem over a graphical model, as defined next.

DEFINITION 2 (look-ahead graphical model). Given a graphical model $\mathcal{M} = (\mathbf{X}, \mathbf{D}, \mathbf{F})$ and messages generated by $MBE(i)$ along pseudo-tree T , $\mathcal{M}(\bar{x}_p, d)$ is the depth- d look-ahead graphical model at \bar{x}_p . It has as variables the set $\{X_k \mid X_k \in T_{p,d}\}$ and has functions $\{f_k(X, \bar{x}_p) \mid X_k \in T_{p,d}\} \cup \{\lambda_{j \rightarrow k}(X, \bar{x}_p) \mid X_j \in T_p - T_{p,d}, X_k \in T_{p,d}\}$.

Clearly, $T_{p,d}$ is a valid pseudo-tree of $\mathcal{M}(\bar{x}_p, d)$, that we call the look-ahead tree. The induced width of $\mathcal{M}(\bar{x}_p, d)$ is noted $lw_{p,d}$.

Example. In Figure 1, the depth 1 look-ahead graphical model relative to assignment $(A = 0, B = 1)$ contains variables $\{C, F\}$, and the functions, $\{f(B = 1, C), f(B = 1, F), \lambda_{D \rightarrow C}(B = 1, C), \lambda_{E \rightarrow C}(B = 1, C), \lambda_{G \rightarrow F}(A = 0, F)\}$, whose induced width is 1. Figure 1e displays the depth 2 look-ahead tree at variable B .

Clearly, $lw_d \leq \min\{d, w^*\}$ and the bound is tight. In practice, only shallow look-aheads may be feasible because the complexity of solving the look-ahead graphical model is exponential in lw_d . A main source for improving the efficiency in our look-ahead is using effective variable elimination algorithms to compute the look-ahead at each node. We next focus on the main method for bounding the look-ahead overhead using the notion of bucket error.

Bounding the Look-Ahead Overhead

Bucket Error

The residual measures the error of one level look-ahead. For the mini-buckets heuristic the error is originated by the mini-bucket partitioning. In the following we will compare the message a bucket would have computed without partitioning (called exact bucket message and noted μ_k^*) against the messages computed by mini-buckets of the bucket (called combined mini-bucket message and noted μ_k). The difference between the two captures the error introduced by partitioning the bucket.

DEFINITION 3 (bucket and mini-bucket messages at B_k). Given a mini-bucket partition $B_k = \cup_k B_k^r$, we define the combined mini-bucket message at B_k ,

$$\mu_k = \sum_r (\min_{X_k} \sum_{f \in B_k^r} f) \quad (6)$$

In contrast, the exact message that would have been generated with no partitioning at B_k is,

$$\mu_k^* = \min_{X_k} \sum_{f \in B_k} f \quad (7)$$

Notice that while μ_k^* is exact for B_k it may contain partitioning errors introduced in earlier buckets.

DEFINITION 4 (local bucket-error at B_k). Given a run of MBE, the local bucket error function at B_k denoted Err_k is,

$$Err_k = \mu_k^* - \mu_k \quad (8)$$

The scope of Err_k is a subset of $\bar{X}_k - X_k$

Example. In an MBE(3) execution as in Figure 1c the bucket error of D is $Err_D = \min_D [f(A, D) + f(B, D) + f(C, D)] - (\lambda_{D \rightarrow A}(A) + \lambda_{D \rightarrow C}(B, C))$.

Algorithm 1 (BEE) computes the bucket errors. Following MBE(i), it executes a second pass from leaves to root along the pseudo-tree. When considering bucket B_k , it computes μ_k , μ_k^* and Err_k . The complexity is exponential in the scope of the bucket after the $MBE(i)$ execution. The total cost is dominated by the largest scope. This number is captured by a new parameter that we call pseudo-width.

Algorithm 1: Bucket Error Evaluation (BEE)

Input: A Graphical model $\mathcal{M} = (\mathbf{X}, \mathbf{D}, \mathbf{F})$, a pseudo-tree T , i -bound

Output: The error function Err_k for each bucket

Initialization: Run $MBE(i)$ w.r.t. T .

for each $B_k, X_k \in \mathbf{X}$ **do**

 Let $B_k = \cup_k B_k^r$ be the partition used by $MBE(i)$

$\mu_k = \sum_r (\min_{X_k} \sum_{f \in B_k^r} f)$

$\mu_k^* = \min_{X_k} \sum_{f \in B_k} f$

$Err_k \leftarrow \mu_k^* - \mu_k$

return Err functions

DEFINITION 5 (pseudo-width(i)). Given a run of $MBE(i)$ along pseudo-tree T , the pseudo-width of B_j , $psw_j^{(i)}$ is the number of variables in the bucket at the moment of being processed. The pseudo-width of T relative to $MBE(i)$ is denoted $psw(i) = \max_j psw_j^{(i)}$.

THEOREM 1 (complexity of BEE). The complexity of BEE is $O(n \cdot k^{psw(i)})$, where n is the number of variables, k bounds the domain size and $psw(i)$ is the pseudo-width along T relative to $MBE(i)$.

The pseudo-width lies between the width and the induced width w^* of the ordering, and it grows with the i -bound. When the i -bound of MBE is large, computing the error exactly may not be possible.

We next show (Theorem 2) the relationship between the residual of the mini-bucket heuristic and the bucket errors. To prove it we need the following lemmas which relates μ_k (Eq. 6) to the MBE heuristic of its parent X_p , and relates μ^* (Eq. 7) to the combinatorial part of the look-ahead computation (Eq. 5).

Lemma 1. If X_k is a child of a variable X_p . Then, $\Lambda_k(\bar{X}_p) = \mu_k(\bar{X}_p)$

Proof. $\Lambda_k(\bar{X}_p)$ (see Eq. 1) is the sum of messages that $MBE(i)$ sends from B_k to the buckets of variables in \bar{X}_p .

Since X_p is the parent of X_k , $\Lambda_k(\bar{X}_p)$ is the sum of all the messages departing from X_k , which is the definition of $\mu_k(\bar{X}_p)$ \square

Lemma 2. *If X_k is a child of a variable $X_p \in T$. Then, $L^1(\bar{x}_p) = \sum_{X_k \in ch(X_p)} \mu_k^*(\bar{x}_p)$*

Proof. (sketch) In the expression of $L^1(\bar{x}_p)$ (see Eq. 5) it is possible to push the minimization into the summation. Thus,

$$L^1(\bar{x}_p) = \sum_{X_k \in ch(X_p)} \min_{x_k} \{f_k(\bar{x}_p, x_k) + \sum_{X_j \in T_p - ch(X_p)} \lambda_{j \rightarrow k}(\bar{x}_p, x_k)\}$$

The set of functions inside of each \min_{x_k} are, by definition, the set of functions in B_k placed there either originally or are messages received from its descendent, yielding (Eq. 7)

$$= \sum_{X_k \in ch(X_p)} \min_{x_k} \sum_{f \in B_k} f = \sum_{X_k \in ch(X_p)} \mu_k^*$$

\square

THEOREM 2 (residual and bucket-error). *Assume an execution of $MBE(i)$ along T yielding heuristic h . Then, for every \bar{x}_p*

$$res(\bar{x}_p) = \sum_{X_k \in ch(X_p)} Err_k(\bar{x}_p) \quad (9)$$

Proof. (sketch) From the definition of residual and Prop 1,

$$res(\bar{x}_p) = L^1(\bar{x}_p) - \left(\sum_{X_k \in ch(X_p)} \Lambda_k(\bar{x}_p) \right)$$

From the Lemma 1 and Lemma 2,

$$res(\bar{x}_p) = \sum_{X_k \in ch(X_p)} \mu_k^*(\bar{x}_p) - \sum_{X_k \in ch(X_p)} \mu_k(\bar{x}_p)$$

Grouping together expressions that refer to the same child proves the theorem. \square

Corollary 1. *When a bucket is not partitioned into mini-buckets, its bucket error is 0, and therefore it contributes zero to the residual and the look-ahead of its parent.*

Pruned Look-Ahead Trees

Given the error functions, we can now identify variables that are *irrelevant* for lookahead (that is, their look-ahead does not produce any improved heuristic). Once those are identified, for each variable and look-ahead depth d , we will prune its look-ahead subtree to include only paths that touch variables that are relevant. If the bucket error function is zero for a variable, the variable is clearly irrelevant for 1-level look-ahead. However to allow more flexibility we will identify those buckets whose *average relative bucket error* is below a given threshold ϵ . Below, $dom(B_j)$ denotes all the assignments to variables in the scope of bucket B_j .

DEFINITION 6 (average relative bucket error). *The average relative bucket error of X_j given a run of $MBE(i)$ is*

$$\tilde{E}_j = \frac{1}{|dom(B_j)|} \sum_{\bar{x}_j} \frac{Err_j(\bar{x}_j)}{\mu^*(\bar{x}_j)}$$

DEFINITION 7 (irrelevant variable). *A variable X_j is ϵ -irrelevant iff $\tilde{E}_j \leq \epsilon$.*

DEFINITION 8 (minimal pruned d -level subtree). *Given a threshold ϵ , a minimal pruned look-ahead subtree $T_{p,d}^*$ for variable X_p is obtained from $T_{p,d}$ by removing ϵ -irrelevant leaves X_j (recursively until quiescence).*

Example. Corollary 1 identifies trivial 0-irrelevant variables. For instance, in the MBE execution displayed in Figure 1c, all variables but D are 0-irrelevant. It may happen that variable D is also 0-irrelevant or ϵ -irrelevant for some small ϵ , but we cannot tell from the symbolic execution displayed in the figure. It can be detected using bucket errors.

Figure 1e shows a depth 2 look-ahead tree for B and (grayed) the minimal pruned subtree assuming $\epsilon = 0$ and only irrelevant variables as for Corollary 1.

Algorithm 2: Minimal Pruned Look-Ahead Subtree

Input: A Graphical model $\mathcal{M} = (\mathbf{X}, \mathbf{D}, \mathbf{F})$, a pseudo-tree T , i -bound, threshold ϵ , depth d
Output: Minimal pruned look-ahead subtrees
Initialization: Run $MBE(i)$ w.r.t. T .
 $X' \leftarrow \mathbf{X}$
Run $BEE(\mathcal{M}, T, i)$
for each $B_j, X_j \in \mathbf{X}$ **do**
 $\tilde{E}_j \leftarrow \frac{1}{|dom(B_j)|} \sum_{\bar{x}_j} \frac{Err_j(\bar{x}_j)}{\mu^*(\bar{x}_j)}$
 if $\tilde{E}_j < \epsilon$ **then** $X' \leftarrow X' - \{X_j\}$;
for each X_p **in** \mathbf{X} **do**
 Initialize $T_{p,d}^*$ to $T_{p,d}$
 while $T_{p,d}^*$ has leaves in X' **do**
 Remove leaves $X_j \notin X'$ from $T_{p,d}^*$
return $T_{p,d}^*$ for each X_p

Algorithm 2 describes the generation of pruned look-ahead trees for each variable.

Experimental Evaluation

Algorithm Setup. We experimented with running AND/OR branch-and-bound (AOBB) guided by $MBE(i)$ with moment-matching (Ihler et al. 2012). To ensure a fixed search space for all algorithms we used pre-computed variable orderings. We tried 2-3 different i -bounds for each problem instance, aiming at a diversity of heuristic accuracy.

We run Algorithm 2 for pre-processing, yielding a minimal pruned look-ahead subtree for each variable. When the BEE computation (and its table) gets too large (e.g. over 10^6) we sample (e.g. 10^5 assignments). Within each i -bound setting, we varied the look-ahead depth from 0 to 6.

Benchmarks. Includes instances from genetic linkage analysis (pedigree, largeFam) and medical diagnosis (promedas) (see Table 2). In total, we evaluated 59 problem instances with induced widths ranging from 19 to 120.

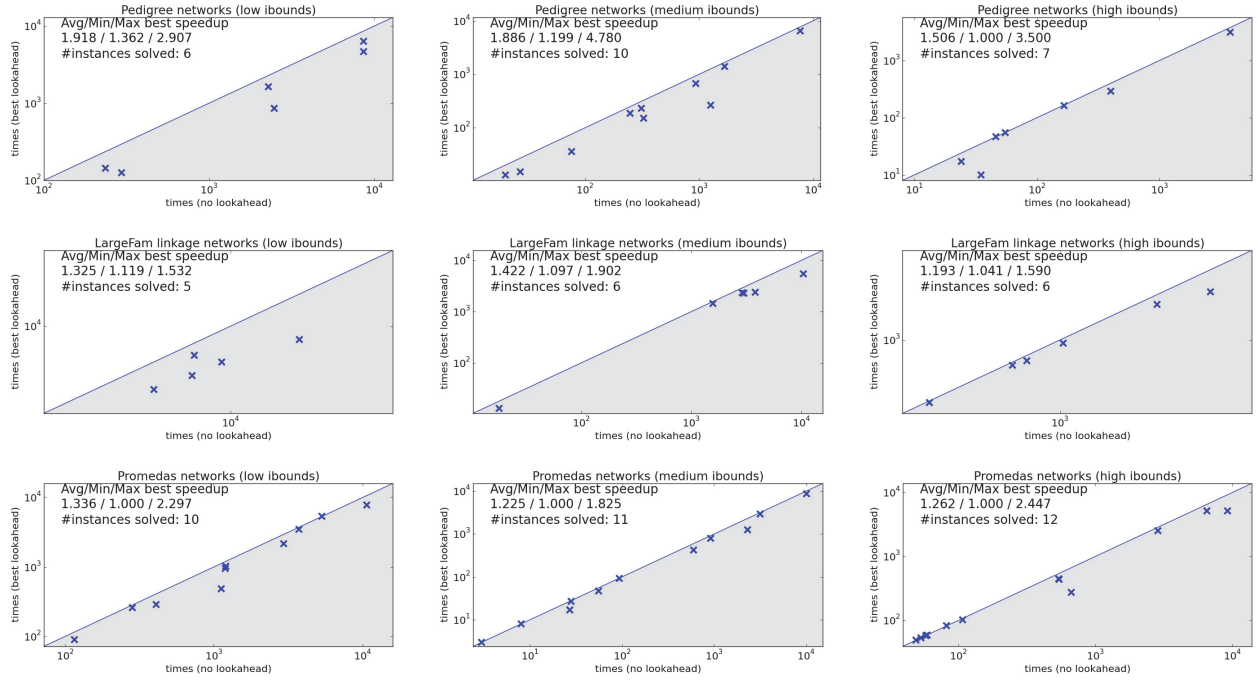


Figure 2: Total CPU times for using no look-ahead plotted against best total CPU times for any level of look-ahead (in log scale). Each point in the gray region represents an instance where the look-ahead heuristic has better time performance for some setting of the depth. We also report the number of instances which were actually solved and the average speedup across these instances.




Figure 3: Average speedups relative to using no look-ahead per depth for the minimal and full look-ahead subtrees.

lookahead naively. We observe here the advantage of using minimal pruned trees. For example, on the pedigree instances, there is almost no speedup on average when using the full look-ahead subtree by a depth of 3, but using the minimal pruned tree allows look-ahead to remain cost-effective.

Conclusion

This paper opens up an investigation on the use of look-ahead in graphical models. We address the challenge of finding an effective balance between the look-ahead overhead and its punning power, by exploiting the graphical structure. We have showed a general relationship between the residual of the mini-bucket heuristic and a local bucket-error due to the mini-bucket algorithm. This relationship shows that the residual (and consequently depth-1 look-ahead) can be pre-compiled using a message-passing like computation. As such, we showed that the bucket-errors can assist in controlling the

look-ahead at any depth. We showed empirically that even this simple mechanism of identifying irrelevant variables can be instrumental in making deep-level look-head cost-effective. In addition, for the mini-bucket heuristic, we have characterized look-ahead as an inference task, which allows a full inference solving approach. In the future, we plan to automate parameter selection for optimizing look-ahead, and explore the potential for any other combinatorial optimization.

Acknowledgements

We thank the reviewers for their valuable feedback. This work was supported in part by NSF grants IIS-1065618, IIS-1526842, and IIS-1254071, the US Air Force under Contract No. FA8750-14-C-0011 under the DARPA PPAML program, and MINECO under projects TIN2013-45732-C4-3-P and TIN2015-69175-C4-3-R.

References

- Bu, Z.; Stern, R.; Felner, A.; and Holte, R. C. 2014. A* with lookahead re-evaluated. In Edelkamp, S., and Barták, R., eds., *Proceedings of the Seventh Annual Symposium on Combinatorial Search, SOCS 2014, Prague, Czech Republic, 15-17 August 2014*. AAAI Press.
- Dechter, R. 2013. *Reasoning with Probabilistic and Deterministic Graphical Models: Exact Algorithms*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers.
- Geffner, H., and Bonet, B. 2013. *A Concise Introduction to Models and Methods for Automated Planning*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers.
- Ihler, A.; Flerova, N.; Dechter, R.; and Otten, L. 2012. Join-graph based cost-shifting schemes. In *Uncertainty in Artificial Intelligence (UAI)*. Corvallis, Oregon: AUAI Press. 397–406.
- Kask, K., and Dechter, R. 1999. Branch and bound with mini-bucket heuristics. *Proc. IJCAI-99*.
- Kask, K., and Dechter, R. 2001. A general scheme for automatic search heuristics from specification dependencies. *Artificial Intelligence* 91–131.
- L. Otten, A. Ihler, K. K., and Dechter, R. 2012. Winning the pascal 2011 map challenge with enhanced and/or branch-and-bound. In *Workshop on DISCML 2012 (a workshop of NIPS 2012)*.
- Marinescu, R., and Dechter, R. 2009. Memory intensive and/or search for combinatorial optimization in graphical models. *Artif. Intell.* 173(16-17):1492–1524.
- Nilsson, N. J. 1980. *Principles of Artificial Intelligence*. Tioga, Palo Alto, CA.
- Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann.
- Vidal, V. 2004. A lookahead strategy for heuristic search planning. In *Proceedings of the Fourteenth International Conference on Automated Planning and Scheduling (ICAPS 2004), June 3-7 2004, Whistler, British Columbia, Canada*, 150–160.