# Diverse $M$-Best Solutions by Dynamic Programming

Carsten Haubold[1], Virginie Uhlmann[2], Michael Unser[2], and Fred A. Hamprecht[1(✉)]

[1] IWR/HCI, University of Heidelberg, 69115 Heidelberg, Germany
{carsten.haubold,fred.hamprecht}@iwr.uni-heidelberg.de
[2] BIG, École Polytechnique Fédérale de Lausanne (EPFL), 1015 Lausanne, Switzerland

**Abstract.** Many computer vision pipelines involve dynamic programming primitives such as finding a shortest path or the minimum energy solution in a tree-shaped probabilistic graphical model. In such cases, extracting not merely the best, but the set of $M$-best solutions is useful to generate a rich collection of candidate proposals that can be used in downstream processing. In this work, we show how $M$-best solutions of tree-shaped graphical models can be obtained by dynamic programming on a special graph with $M$ layers. The proposed multi-layer concept is optimal for searching $M$-best solutions, and so flexible that it can also approximate $M$-best diverse solutions. We illustrate the usefulness with applications to object detection, panorama stitching and centerline extraction.

## 1 Introduction

A large number of problems in image analysis and computer vision involve the search for the *shortest path* (*e.g.*, finding seams and contours) or for the *maximum-a-posteriori* (MAP) configuration in a tree structured graphical model, as in hierarchies of segmentation hypotheses or deformable part models. To compute the solution to those problems, one relies on efficient and optimal methods from dynamic programming [4] such as Dijkstra's algorithm [7]. In many of these scenarios, it is of interest to find not merely the single lowest energy (*i.e.*, MAP) solution, but the $M$ solutions of lowest energy (*$M$-best*) [3,17,19,23,26]. This can *e.g.* be useful for learning [16], tracking-by-detection methods that allow competing hypotheses [13,18,24], or for re-ranking [28] solutions based on higher order features which would be prohibitively complex for the original optimization problem. If these $M$ solutions are required to differ in more than one label, the problem is referred to as *diverse $M$-best* [2,14,22].

**Electronic supplementary material** The online version of this chapter (doi:10.1007/978-3-319-66709-6_21) contains supplementary material, which is available to authorized users.

© Springer International Publishing AG 2017
V. Roth and T. Vetter (Eds.): GCPR 2017, LNCS 10496, pp. 255–267, 2017.
DOI: 10.1007/978-3-319-66709-6_21

*Contributions:* In this work, we show how the optimal second best ($M = 2$) solution of a tree-shaped graphical model can be found through dynamic programming in a multi-layer graph by using a replica of the original graph as second layer and connecting both layers through edges with special jump potentials (Sect. 3). Using these building blocks, we extend our approach to exactly find the $M > 2$ best solutions sequentially by constructing $M$-layer graphs (Sect. 4). While the above can be seen as a special case of [29], our multi-layer approach is an intuitive interpretation that allows flexible modeling of the desired result. We thus develop two heuristics using multiple layers to find the approximate diverse $M$-best solutions for tree-shaped graphical models (Sect. 5). Lastly, we experimentally compare the different diversity approximations to prior work, and show results for a variety of applications, namely: *(i)* panorama stitching, *(ii)* nested segmentation hypotheses selection, and *(iii)* centerline extraction (Sect. 6).

## 2 Related Work

*M-best MAP:* An algorithm for sequentially finding the $M$ most probable configurations of general combinatorial problems was first presented in [17]. To find the next best solution, they branch on the state of every single variable, resolve, and finally choose the best of all resulting configurations. While this works for any optimization method and model, it is in practice prohibitively expensive. Several works have extended this to junction trees [19,26] that work in $O(|\mathcal{V}|(L^2 + M + M\log(|\mathcal{V}|M)))$, while [9,23] developed a similar bucket elimination scheme ($O(M|\mathcal{V}|L^{|\mathcal{V}|})$). A similar idea was applied in [8] to find the $M$ shortest paths jointly by building an auxiliary graph with a heap at every node that contains the $M$-best paths to reach that node. For situations where the optimal or approximative max-marginals can be computed, [29] derived an improvement on [19] such that the max-marginals have to be computed only $2M$ times, yielding the same runtime complexity ($O(M|\mathcal{V}|L^2)$) as the method we present here. A method that finds the $M$ best solutions on trees in only $O(L^2V + \log(L)|\mathcal{V}|(M-1))$ by an algebraic formulation that is similar to sending messages containing $M$ best values as in [9] was presented by [25, Chap. 8]. However, in contrast to [25,29], our approach provides a lot of modeling flexibility, allowing it to be used to approximate diverse $M$ best solutions as well. A polyhedral optimization view of the sequential $M$-best MAP problem is given in [10]. There, a linear programming (LP) relaxation is constructed by characterizing the assignment-excluding local polytope through spanning-tree-inequalities. This LP relaxation is tight for trees for $M = 2$, but not for higher $M$ or loopy graphs because the assignment-excluding inequalities could together cut away other integral vertices of the polytope. An efficient message passing algorithm for the same LP relaxation exploiting the structure of the polytope was designed by [3]. A completely orthogonal way to explore solutions around the optimum would be to sample from the modeled distribution, *e.g.* using Perturb-and-MAP [20].

*Diverse M-Best:* For general graphical models, the first formulation of the *diverse M-best* problem can be found in [2]. Even though their Lagrangian

relaxation of the diversity constraint can work with any choice of metric, it is not even tight for Hamming distances of $k > 1$. Different diversity metrics are explored in [22], where a greedy method to find good instances from the (exponentially big) set of possible solutions is designed by setting up a factor graph with higher order potentials, assuming that the diversity metric is submodular. In [14], the authors construct a factor graph that *jointly* finds the diverse $M$-best solutions by replicating the original model $M$ times and inserting factors for the diversity penalty depending on the structure of the chosen distance. It is shown that maximizing the diversity of the $M$-best solutions jointly, not only sequentially as in [2], can yield better results. They propose a reformulation that preserves solvability with $\alpha$-$\beta$-*swap*-like methods. Still, when applied to trees, the factors introduced for diversity unfortunately turn the problem into a loopy graph and prevent the application of dynamic programming. All those approaches incorporate the diversity constraint into the original optimization problem by Lagrangian relaxation. In contrast, the constructions we propose in this work yield a solution with the desired diversity in a single shot. A different line of work [5,6] focused on extracting the $M$-best *modes*, but those methods' computational complexity renders them intractable for large graphs.

## 3 Optimal Second Best Tree Solutions

We now present how the second best solution of a tree-shaped graphical model can be obtained using dynamic programming on a special graph construction. By second best, we mean a solution that differs from the best configuration in at least one node, *i.e.*, that has a Hamming distance of $k \geq 1$ to the best solution. We begin with an informal motivation based on the search for the second best shortest path.

*Motivation:* The Dual Dijkstra method from [11] allows finding not only the best, but a collection of $M$ shortest paths from a source $\sigma$ to a target $\tau$ in a graph. To do so, two shortest path trees are constructed, one starting at the source and one at the target. Thus, for every node $v$, the shortest path from the source $\mathcal{P}_{\sigma,v}$ and to the target $\mathcal{P}_{v,\tau}$ is known. Summing the distances to source and target gives the length of the shortest path from $\sigma$ to $\tau$ via $v$. An important property is that, for all vertices along the shortest path from $\sigma$ to $\tau$, this sum is equal to the length of the shortest path.

Now imagine these shortest path trees as two copies of the initial graph stacked as two layers, as seen in Fig. 1(a). The lower layer indicates the lowest cost to reach every vertex $v$ from $\sigma$, and the upper layer the cost of the shortest path to reach $\tau$ from $v$. By selecting any vertex $v$ and connecting the paths at $v$ in the lower and upper layer, one can again find the shortest path from $\sigma$ to $\tau$ via $v$, this time by introducing an auxiliary *jump* edge between the two layers.

The benefit of this two layer setup is that to find the *second* best solution, we simply have to search for the vertex that *does not* lie on the best path, at which jumping between the layers leads to the minimal cost path.

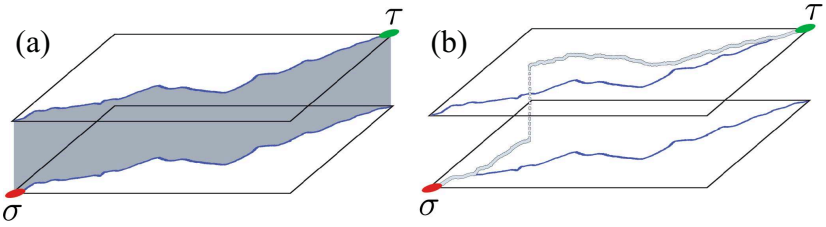


**Fig. 1.** A schematized two-layer grid-graph construction to find the second best shortest path from the source $\sigma$ in the lower, to the target $\tau$ in the upper layer. A valid path is required to jump between layers, which is allowed everywhere for the best path. (a) Shown in blue is the best solution, which could have jumped to the upper layer at every node along the path with the lowest cost. (b) To find the second best path, layer jumps are forbidden at the nodes used by the best solution. Thus the second path diverges to the jump location leading to the next minimal cost path. (Color figure online)

*Dynamic Programming:* Let us briefly review the dynamic programming (DP) paradigm on an undirected tree-shaped graph $\bar{G} = (\mathcal{V}, \bar{\mathcal{E}})$. We denote the state of a node $v \in \mathcal{V}$ as $\mathbf{x}_v$, and the full state vector as $\mathbf{x} = \{\mathbf{x}_v : v \in \mathcal{V}\}$. The potentials of node $v$ (*unary* potential) and of the edge connecting nodes $u$ and $v$ (*pairwise* potential) are represented by $\theta_v(\mathbf{x}_v)$ and $\theta_{uv}(\mathbf{x}_u, \mathbf{x}_v)$, respectively. From this, we define the inference problem as an energy minimization task [15] with objective

$$\min_{\mathbf{x}} \sum_{v \in \mathcal{V}} \theta_v(\mathbf{x}_v) + \sum_{(u,v) \in \bar{\mathcal{E}}} \theta_{uv}(\mathbf{x}_u, \mathbf{x}_v). \tag{1}$$

When applying dynamic programming, one successively computes the energy $E$ of optimal solutions of subproblems of increasing size. One node of the graph $\bar{G}$ is arbitrarily selected as the root node $r$. This results in a directed graph $G = (\mathcal{V}, \mathcal{E})$ where edges point towards the root. Let $\overleftarrow{N}(v)$ denote the neighboring nodes along incoming edges of $v$ in $G$. Using the tree-imposed ordering of edges, one starts processing at the leaves and sends messages embodying the respective subproblem solutions towards the root. Whenever a node $v$ has received a message from all incoming edges, it can – disregarding its successors in $G$ – compute the lowest energies $E_v(\mathbf{x}_v)$ of the subtree rooted at $v$ for every state $\mathbf{x}_v$, and send a message to its parent [21]. Because leaves have no incoming edges, their energy is equal to their unary potentials. All subsequent nodes combine the incoming messages with their unary potentials to obtain the energy of the subtree rooted at them by

$$E_v(\mathbf{x}_v) := \theta_v(\mathbf{x}_v) + \sum_{u \in \overleftarrow{N}(v)} \min_{\mathbf{x}_u} \left[\theta_{uv}(\mathbf{x}_u, \mathbf{x}_v) + E_u(\mathbf{x}_u)\right]. \tag{2}$$

While sending these messages, each node $v$ stores which state ($\arg\min_{\mathbf{x}_u}$) of the previous node ($u \in \overleftarrow{N}(v), v$) along each incoming edge led to the minimal energy of every state $\mathbf{x}_v$. When the root has been processed, the state that led to the minimal energy is selected and, by backtracking all the recorded arg min, the best global configuration $\mathbf{x}^\star$ can be found. Figure 2 shows a minimal tree example.

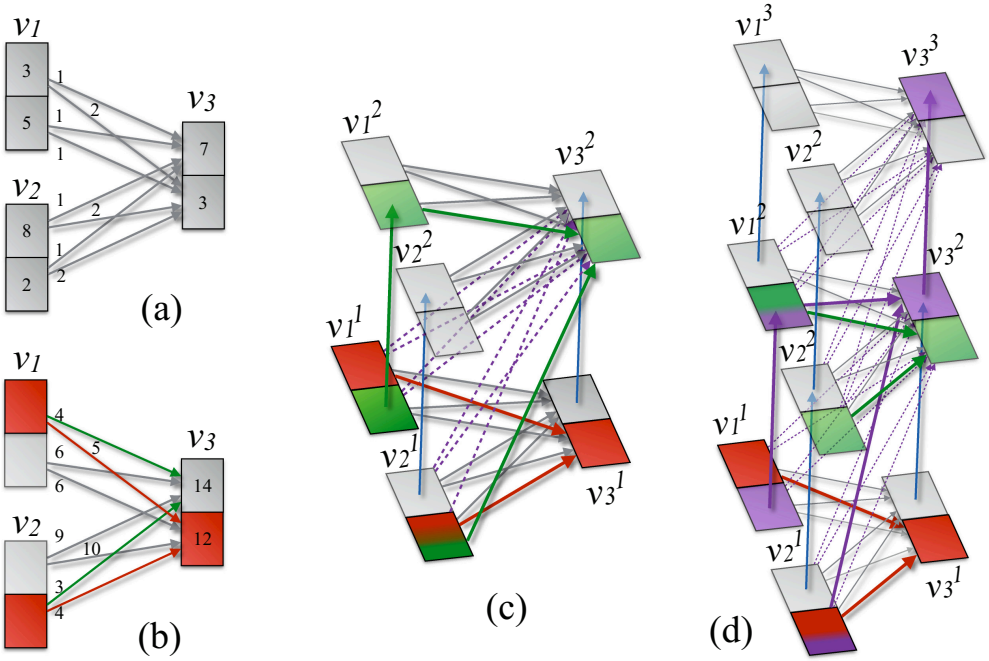


**Fig. 2.** (a) Minimal tree structured graph of nodes $v_1$, $v_2$, and $v_3$, with two states each, visualized as stacked boxes. $v_3$ is arbitrarily designated as the root, or target. Unary and pairwise costs are shown as numbers in boxes and along edges, respectively. Its optimal solution is highlighted in red in (b). Green edges correspond to the arg min incoming configurations for each state, and green numbers depict the accumulated min-sum messages. (c) Two-layer tree used to find the second best solution. Blue arcs represent layer-jump-edges with finite potential, which are available at states not occupied by the best solution. Purple dashed edges need to be considered if, at a branching point (such as $v_3^2$), not all incoming messages are coming from the upper layer. The second best solution is represented in green. (d) Searching for the $3^{rd}$-best solution (blue) with a Hamming distance of $k = 1$ to the best (red) and second best solution. The new solution must jump twice to reach the upper layer, by taking a state that was not used in the configuration represented by layers 1 and 2. (Color figure online)

Regarding DP runtime complexity, consider that (2) needs to be evaluated for every state of every node exactly once. In addition, in (2), we consider all states of every incoming edge, of which there are $|\mathcal{E}| = |\mathcal{V}| - 1$ in a tree. If $L$ denotes the maximum number of states, one obtains $O(|\mathcal{V}|L^2)$.

*Two-Layer Model:* Once the optimal solution is found, we might be interested in the second best solution $\mathbf{x}$, which assigns a different state to at least one node $\exists v \in \mathcal{V} : \mathbf{x}_v \neq \mathbf{x}_v^\star$. Because messages in DP only convey the optimal subtree energies, we cannot immediately extract this second best solution. Hence we are looking for a way to enforce that a different state is attained at least once, but we do not know at which node(s) this should happen to yield the optimal energy. Fortunately, we can apply the same idea as in the second shortest path example: We duplicate the graph to get a second layer and insert edges connecting the two layers such that jumping is only permitted at states not used in the optimal

solution $\mathbf{x}^\star$. After propagating messages through both layers, the second best solution can be obtained by backtracking from the minimum energy state of the root in the second layer to leaves in the first layer. This means that messages must have jumped to the second layer at least once at some node $v$ with a state different to $\mathbf{x}^\star_v$, fulfilling our requirement for the second best solution.

We here state the layer setup conceptually and provide the formal construction in the Supplementary. To create the two layers, we duplicate graph $G$ (Fig. 2a) such that we get a layer 1, and a layer 2 replica. We address the instances of every node $v \in \mathcal{V}$ by $v^1$ and $v^2$ for layer 1 and layer 2, respectively. When duplicating the graph, the unary and pairwise potentials of nodes and edges are copied to layer 2. At every node $v \in \mathcal{V}$, we insert a *layer-jump-edge* from $v^1$ to $v^2$ (blue edges in Fig. 2c) with a pairwise potential $\theta_{v^1v^2}$ that is only zero if both variables take the same state $\mathbf{x}_{v^1} = \mathbf{x}_{v^2}$ different from $v$'s state in $\mathbf{x}^\star$, and infinity (forbidden) otherwise. This way, finite valued messages in layer two represent configurations that did differ from $\mathbf{x}^\star$ at least once. These jump edges would suffice for a chain graph, but the branching points in a tree need special consideration. When a layer 2 branching point is not reached by a layer jump, the current construction only allows considering incoming messages from layer 2. However, since we only require *one* variable to take a new state, only one branch is necessary to reach layer 2 on a path with finite cost. To cope with this situation, we insert *layer-crossing* edges from $u^1$ to $v^2$ for all edges $(u, v) \in \mathcal{E}$ (dashed purple edges in Fig. 2c) with the same pairwise potential as in the original graph $\theta_{u^1v^2} = \theta_{uv}$, and alter the DP update equation for nodes in layer 2 to

$$E_{v^2}(\mathbf{x}_{v^2}) := \min \Bigg( \theta_{v^1v^2}(\mathbf{x}_{v^1}, \mathbf{x}_{v^2}) + E_{v^1}(\mathbf{x}_{v^1}), \tag{3}$$

$$\theta_{v^2}(\mathbf{x}_{v^2}) + \min_{\substack{L_2 \subseteq \overleftarrow{N}(v) \\ |L_2| \geq 1}} \sum_{u \in L_2} \min_{\mathbf{x}_{u^2}} \left[ \theta_{u^2v^2}(\mathbf{x}_{u^2}, \mathbf{x}_{v^2}) + E_{u^2}(\mathbf{x}_{u^2}) \right]$$

$$+ \sum_{u \in \overleftarrow{N}(v) \setminus L_2} \min_{\mathbf{x}_{u^1}} \left[ \theta_{u^1v^2}(\mathbf{x}_{u^1}, \mathbf{x}_{v^2}) + E_{u^1}(\mathbf{x}_{u^1}) \right] \Bigg). \tag{4}$$

Compared to (2), we now have two options instead of one at every node $v$ in layer 2. Firstly, we can reach $v^2$ by a layer jump. Note that, in case of a jump (3), we do not account for the unary $\theta_{v^2}(\mathbf{x}_{v^2})$ as $E_{v^1}(\mathbf{x}_{v^1})$ contains the same term already. Alternatively, at least one of the incoming messages must come from a nonenpty set $L_2$ of predecessors in layer 2 (4), while the remaining messages could *cross* layers. These options are visualized in Fig. 2c.

*Optimality and Runtime:* By duplicating the directed graph and inserting two sets of new edges which are oriented towards the root in layer two, the topology of the graph remains a directed acyclic graph, and DP hence yields the optimal configuration. As long as the solution has finite energy, no forbidden *layer-jump-edge* is used, giving us the second best solution. In terms of runtime

complexity, we have duplicated the number of vertices and have four times as many edges, which are small constant factors that disappear in $O(|\mathcal{V}|L^2)$. For optimal performance, one can reuse the messages in layer 1 because these do not change.

## 4 Optimal $M$-Best Tree Solutions

The two-layer setup can easily be extended to multiple layers, which allows us to search for the $M$-best solutions with a Hamming distance of $k \geq 1$. We use one additional layer per previous best configuration; that is, $M$ layers. Each layer is responsible for one of the previous solutions, hence its *layer-jump-edges* are restricted according to the respective solution. Solutions must be ordered by increasing cost; such that the first layer constrains jumps with respect to the best configuration, the second layer for the second best, and so on. The new update rules from in Sect. 3 can then be applied to every consecutive pair of layers. Figure 2d shows an example.

*Optimality and Runtime:* When considering more than one previous solution using the multi-layer setup, the jump restrictions encoded in the *layer-jump-edge* potentials are independent at each layer. For any given node and state in a layer, the cost and path to reach it are optimal with respect to all layers below. This straightforwardly holds for the one and two layer cases, and is the reason why layers must be ordered by increasing cost of the represented previous solutions. For the sake of argument, layers could be flattened as they are getting processed, bringing back the problem to a series of $M - 1$ optimal two-layer cases, which yields a computational complexity of $O(M|\mathcal{V}|L^2)$.

## 5 Approximate Diverse $M$-Best Solutions

In the classical diverse-$M$-best setting [2], additional solutions are required to have *e.g.* a Hamming distance of $k > 1$. Here, we look at the straightforward multi-layer extension of Sect. 3 to handle $k > 1$. We argue that this approach is suboptimal, and present a two-layer approximation that trades quality for efficiency. Lastly, we discuss how this could be used to find $M$ diverse solutions.

*Multi-layer Model:* To ensure that the next solution differs by at least $k$ from the best one, we could construct a $k+1$-layer graph using the same jumping criteria between all layers. To reach the top layer, a solution must hence jump $k$ times. This raises two challenges: *(a)* a branching point at layer $N$ can be reached by a combination of edges from different layers such that the predecessors *in total* account for a Hamming distance of $N$, and *(b)* a solution should never jump more than once at a single node, otherwise it will not have the desired diversity. Both can be achieved by adjusting the DP update equation to consider a set of admissible incoming edge combinations. We provide the precise expression in the Supplementary.

Unfortunately, this simple setup does not yield optimal solutions. To forbid two jumps in a row, one needs to introduce a dependence on a previously made decision. These dependencies invalidate the subproblem optimality criterion for DP to yield the correct result. It is thus possible that DP does not reach the root on layer $k$ with finite cost, as shown in the Supplementary. Using the same reasoning, even if a valid solution is found, it is not necessarily optimal. Additionally, the set of admissible combinations of incoming edges grows combinatorially, making this approach unsuited for large $k$.

*Diversity Accumulation:* Instead of using $k$ layers, one can also formulate a heuristic on a two-layer graph that ensures that any found solution contains the desired amount of diversity. To do so, we reformulate the Hamming distance constraint (that the new solution must differ from the previous one at $k$ nodes) as a constraint on accumulated diversity, *i.e.*, that $\sum_{v \in \mathcal{V}} \alpha_v(\mathbf{x}_v) > T$, where $\alpha$ is a measure of diversity per node and state, and $T$ a threshold. We change the DP update rules as follows. First, while propagating messages from the leaves of the tree to the root in layer 1, one must also propagate the amount of diversity accumulated by the corresponding configuration of the subtree. Let us denote nodes and edges of the subtree rooted at node $v$ in layer 1 by $\overleftarrow{\mathcal{V}_{v^1}}$ and $\overleftarrow{\mathcal{E}_{v^1}}$, respectively. The accumulated diversity $\mathcal{A}$ is given by

$$\mathcal{A}_{v^1}(\mathbf{x}_{v^1}) := \sum_{i \in \overleftarrow{\mathcal{V}_{i^1}}} \alpha_i(\mathbf{x}_i) + \sum_{(i,j) \in \overleftarrow{\mathcal{E}_{i^1}}} \alpha_{ij}(\mathbf{x}_i, \mathbf{x}_j). \tag{5}$$

Then, we define the layer jump potential $\tilde{\theta}_{v^1 v^2}(\mathbf{x}_{v^1}, \mathbf{x}_{v^2})$ to be infinity as long as the accumulated diversity is below the desired threshold $\mathcal{A}_{v^1}(\mathbf{x}_{v^1}) < k$. The limitation of this heuristic is that, at each node and state, we find the optimal subtree configuration by minimizing the energy without considering diversity. This can prevent us from finding solutions with large diversity. Yet, as we will see in the experiments, this approach has an attractive runtime because it only requires two layers to find a solution with any Hamming distance $k$, and thus has the constant runtime complexity of $O(|\mathcal{V}|L^2)$ per solution.

*Extension to M Diverse Solutions:* Finding $M$ solutions with a Hamming distance of $k$ could be achieved by stacking $M \times (k+1)$ layers, but then the long range dependency problems depicted above are even more prominent. With diversity accumulation on the other hand, $M$ diverse solutions can be obtained heuristically by using one diversity map $\alpha$ and one accumulator $\mathcal{A}$ per previous solution. The jump criterion must then ensure that enough diversity has been accumulated with respect to each previous solution.

## 6 Applications and Experiments

We now evaluate the performance of our heuristics to obtain diverse solutions with prior work, and demonstrate its applicability to several problems in Computer Vision[1].

[1] See the Supplementary for an application to depth estimation from stereo.

*Comparison with Existing Works:* [2,14,28] search for the diverse-$M$-best solutions by incorporating the diversity constraint via Lagrangian relaxation. Our heuristics follow a different approach and turn the constraint into a lower bound instead of relaxing it. The resulting advantage is that we guarantee the set of solutions to be as diverse as required, at the possible expense of a higher cost or the inability to find a solution at all. On 50 random trees, with 100 nodes each, all nodes having 3 states with unary and pairwise potentials drawn uniformly from the range $[0, 1]$, we evaluate different Hamming distances in Fig. 3. We let the method of [2] run for 100 iterations with a step size of $1/n$ in iteration $n$ or stop at convergence. In terms of runtime, diversity accumulation stands out as it constantly requires only two layers. Because the distance to the best configuration is not enforced by hard constraints, solutions found by [2] often contain too little diversity, yielding a too low mean Hamming distance. Diversity accumulation gives solutions with more diversity than required, and hence also deviates more from the optimal energy. In terms of returned diversity and energy, the multi-layer dynamic programming solution yields favorable results compared to the other two methods, but is unfortunately slower – it suffers from the combinatorial explosion of admissible edge sets to consider – and fails to find a valid solution on several trees due to the limitations described in Sect. 5.

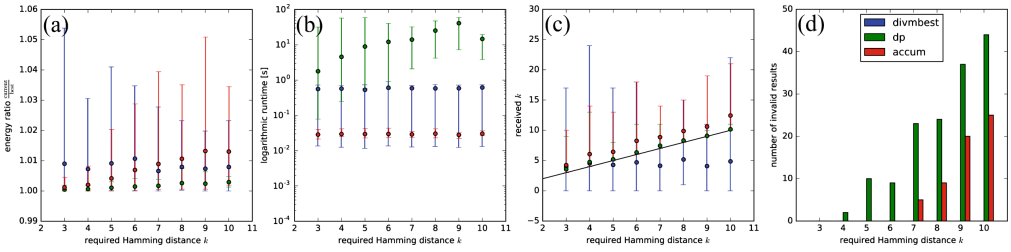


**Fig. 3.** Comparison of the $k+1$-layer `dp` and diversity `accumulation` heuristic for obtaining diverse solutions (Sect. 5) against `divmbest` [2]. All results show mean, minimum and maximum over the valid solutions obtained for every setting on 50 random trees, where (d) shows the number of experiments that did not find a valid solution. (a) Energy ratio between the optimal unconstrained solution and the one with Hamming distance $k$. (b) Runtime. (c) Hamming distance of the resulting solution. Lower is better in all plots but (c), where the returned Hamming distance should be close to, or preferably above the drawn diagonal.

*Medial Axis Identification in Biological Objects:* Identifying the medial axis of biological objects is a common problem in bioimage analysis, as it serves as a basis for length or growth estimation and tracking-by-assignment. Simple dynamic programming can achieve this task given the end points, although, as biological images tend to get noisy or crowded, designing a robust cost function is difficult. In Fig. 4, we illustrate the usefulness of searching for a collection of possible best solutions instead of only one shortest path in brightfield microscopy images of *C. elegans* nematodes.

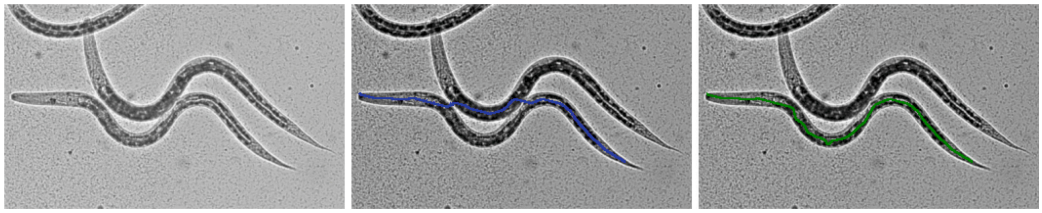


**Fig. 4.** Diverse shortest path finding in noisy bioimages featuring objects in close contact. (a) Raw brightfield microscope images of *C. elegans*, (b) first best path, and (c) 5th best path between auto-selected end points using an exclusion corridor of 30 pixels and a required accumulated diversity of $k \geq 25$.

*Selection of Segmentation Hypotheses:* In datasets with cell clumps, it is often hard to select the correct detections from a set of segmentation hypotheses. We illustrate this problem in images from the Mitocheck project dataset[2] [12] using the tree model proposed in [1]. There, the task is to assign a class label to each element of a set of nested maximally stable extremal regions. The labels indicate the number of objects that each particular region represents. In the tree, nodes correspond to regions, and edges between parent and child node model the nestedness properties. In Fig. 5, we show results obtained when constraining dynamic programming with our $M$-best approach. This is useful to generate segmentation or pose candidates as needed by joint segmentation and tracking procedures, *e.g.* [13,24].

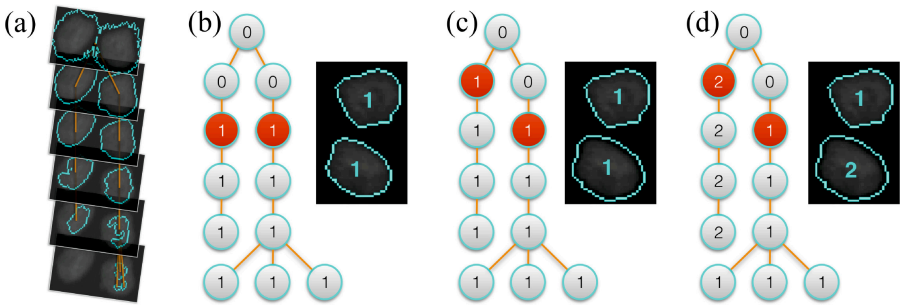


**Fig. 5.** Finding the $M$-best configurations of a tree (a) of MSER segmentation hypotheses as in [1]. The best (b), second (c) and third (d) best configuration found by blocking previous solutions in the respective *layer-jump-edge* potentials. The selected label at each node denotes the predicted object count of the first nonzero ancestor in the tree.

*Panorama Stitching:* In our motivation in Sect. 3, we mentioned that the proposed multi-layer setup can also be used for shortest paths. Here, we apply that in the context of boundary seam computation for panorama stitching [27]. We

[2] http://www.mitocheck.org/.

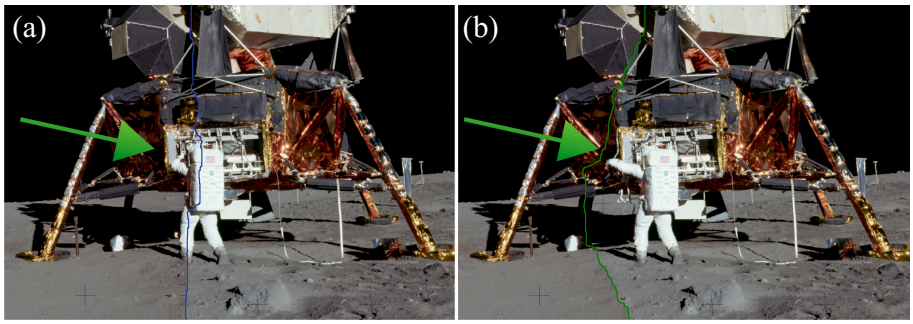


**Fig. 6.** Finding diverse best paths (seams) for panorama stitching. Once the best solution (a) has been found, layer-jump-edges were blocked in a corridor around it to obtain the diverse second best solution (b).

stitch images taken during the Apollo 11 moon landing (Apollo-Armstrong: 2 images of $2349 \times 2366$, courtesy of NASA). As observed in Fig. 6, the second diverse shortest path also corresponds to a visually correct stitching, although the resulting path significantly differs from the globally optimal one.

## 7 Conclusion

We have presented a multi-layer graph construction that allows formulating the $M$-best problem for tree-shaped graphical models efficiently through dynamic programming. This flexible framework can be used to find $M$-best solutions for a Hamming distance of $k = 1$ optimally. For $k > 1$, we present two heuristics, one using a multi-layer graph, and one using two-layers where each new configuration must accumulate diversity before it can reach the upper layer. We evaluated both heuristics against diverse-$M$-best [2], revealing that both perform favourably with certain strengths over the baseline. We demonstrated for several practical applications that the presented methods can reveal interesting alternative solutions.

**Acknowledgements.** This work was partially supported by the HGS MathComp Graduate School, DFG grant HA 4364/9-1, SFB 1129 for integrative analysis of pathogen replication and spread, and the Swiss National Science Foundation under Grant 200020_162343/1.

## References

1. Arteta, C., Lempitsky, V., Noble, J.A., Zisserman, A.: Learning to detect partially overlapping instances. In: Proceedings of the IEEE Conference on Computer Vision And Pattern Recognition (CVPR 2013), Portland, OR, USA, 25–27 June 2013, pp. 3230–3237 (2013)
2. Batra, D., Yadollahpour, P., Guzman-Rivera, A., Shakhnarovich, G.: Diverse M-best solutions in Markov random fields. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Sato, Y., Schmid, C. (eds.) ECCV 2012. LNCS, vol. 7576, pp. 1–16. Springer, Heidelberg (2012). doi:10.1007/978-3-642-33715-4_1

3. Batra, D.: An efficient message-passing algorithm for the M-best map problem. In: Proceedings of the 28th Conference on Uncertainty in Artificial Intelligence (UAI 2012) (2012)
4. Bellman, R.: On the theory of dynamic programming. Proc. Nat. Acad. Sci. **38**(8), 716–719 (1952)
5. Chen, C., Liu, H., Metaxas, D., Zhao, T.: Mode estimation for high dimensional discrete tree graphical models. In: Advances in Neural Information Processing Systems (NIPS 2014), Montréal, Canada, 8–13 December 2014, pp. 1323–1331 (2014)
6. Chen, C., Kolmogorov, V., Zhu, Y., Metaxas, D.N., Lampert, C.H.: Computing the M most probable modes of a graphical model. In: AISTATS, pp. 161–169 (2013)
7. Dijkstra, E.: A note on two problems in connexion with graphs. Numer. Math. **1**(1), 269–271 (1959)
8. Eppstein, D.: Finding the k shortest paths. SIAM J. Comput. **28**(2), 652–673 (1998)
9. Flerova, N., Rollon, E., Dechter, R.: Bucket and mini-bucket schemes for M Best solutions over graphical models. In: Croitoru, M., Rudolph, S., Wilson, N., Howse, J., Corby, O. (eds.) GKR 2011. LNCS, vol. 7205, pp. 91–118. Springer, Heidelberg (2012). doi:10.1007/978-3-642-29449-5_4
10. Fromer, M., Globerson, A.: An LP view of the M-best MAP problem. In: Advances in Neural Information Processing Systems, pp. 567–575 (2009)
11. Fujita, Y., Nakamura, Y., Shiller, Z.: Dual Dijkstra search for paths with different topologies. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA 2003), vol. 3, Taipei, Taiwan, 14–19 September 2003, pp. 3359–3364 (2003)
12. Held, M., Schmitz, M., Fischer, B., Walter, T., Neumann, B., Olma, M., Peter, M., Ellenberg, J., Gerlich, D.: Cellcognition: time-resolved phenotype annotation in high-throughput live cell imaging. Nat. Methods **7**(9), 747–754 (2010)
13. Jug, F., Pietzsch, T., Kainmüller, D., Funke, J., Kaiser, M., van Nimwegen, E., Rother, C., Myers, G.: Optimal joint segmentation and tracking of Escherichia coli in the mother machine. In: Proceedings of the First International Workshop on Bayesian and grAphical Models for Biomedical Imaging (BAMBI 2014), Cambridge, MA, USA, 18 September 2014, pp. 25–36 (2014)
14. Kirillov, A., Savchynskyy, B., Schlesinger, D., Vetrov, D., Rother, C.: Inferring M-best diverse labelings in a single one. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV 2015), Santiago, Chile, 13–16 December 2015, pp. 1814–1822 (2015)
15. Koller, D., Friedman, N.: Probabilistic Graphical Models: Principles and Techniques. MIT Press, Cambridge (2009)
16. Lampert, C.H.: Maximum margin multi-label structured prediction. In: Advances in Neural Information Processing Systems, pp. 289–297 (2011)
17. Lawler, E.: A procedure for computing the k best solutions to discrete optimization problems and its application to the shortest path problem. Manag. Sci. **18**(7), 401–405 (1972)
18. Milan, A., Schindler, K., Roth, S.: Detection-and trajectory-level exclusion in multiple object tracking. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 3682–3689 (2013)
19. Nilsson, D.: An efficient algorithm for finding the m most probable configurationsin probabilistic expert systems. Stat. Comput. **8**(2), 159–173 (1998)
20. Papandreou, G., Yuille, A.L.: Perturb-and-map random fields: using discrete optimization to learn and sample from energy models. In: 2011 International Conference on Computer Vision, pp. 193–200. IEEE (2011)

21. Pearl, J.: Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. Morgan Kaufmann, Burlington (1988)
22. Prasad, A., Jegelka, S., Batra, D.: Submodular meets structured: finding diverse subsets in exponentially-large structured item sets. In: Advances in Neural Information Processing Systems (NIPS 2014), Montréal, Canada, 8–13 December 2014, pp. 2645–2653 (2014)
23. Rollon, E., Flerova, N., Dechter, R.: Inference schemes for M best solutions for soft CSPs. In: Proceedings of the Seventh International Workshop on Preferences and Soft Constraints, vol. 2. Sitges, Spain, 1 October 2011
24. Schiegg, M., Hanslovsky, P., Haubold, C., Koethe, U., Hufnagel, L., Hamprecht, F.: Graphical model for joint segmentation and tracking of multiple dividing cells. Bioinformatics **31**(6), 948–956 (2015)
25. Schlesinger, M.I., Hlavác, V.: Ten Lectures on Statistical and Structural Pattern Recognition, vol. 24. Springer Science & Business Media, New York (2013)
26. Seroussi, B., Golmard, J.L.: An algorithm directly finding the k most probable configurations in Bayesian networks. Int. J. Approx. Reason. **11**(3), 205–233 (1994)
27. Summa, B., Tierny, J., Pascucci, V.: Panorama weaving: fast and flexible seam processing. ACM Trans. Graph. **31**(4), 83:1–83:11 (2012)
28. Yadollahpour, P., Batra, D., Shakhnarovich, G.: Discriminative re-ranking of diverse segmentations. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), June 2013
29. Yanover, C., Weiss, Y.: Finding the m most probable configurations using loopy belief propagation. In: Advances in Neural Information Processing Systems, vol. 16, p. 289 (2004)