

1 **Review #1:** Thanks for your comments and suggestions.

2 We are aware of the weighted model counters by Chavira and Darwiche which are based on either search or knowledge  
3 compilation ideas. However, these algorithms count **\*all\*** solutions whereas our proposed algorithms focus on counting  
4 **the optimal solutions** only. It is important to emphasize that counting optimal solutions in a compiled structure such as  
5 a decision diagram (e.g., arithmetic circuit, multi-valued AND/OR decision diagram, OBDD) require two passes over  
6 the compiled structure and therefore will yield the count of optimal solutions by enumeration. In contrast, our algorithms  
7 are not dependent on the number of optimal solutions and therefore are much more efficient as we demonstrate in the  
8 empirical section.

9 Regarding the sum-product networks, we believe that they fall within the same category of algorithms that count all  
10 solutions. In principle, we think that they can also be extended to count optimal solutions but they will compute the  
11 number of optimal solutions by enumeration as well.

12 Finally, extending the current weighted model counters and sum-product networks to counting optimal solutions is not  
13 yet clear and therefore will be considered as part of our future work.

14 **Review #2:** Thanks for your comments and suggestions.

15 The semiring is a general framework that allows us to specify a wide range of reasoning tasks for deterministic and/or  
16 probabilistic graphical models in a unified manner. In this paper we show for the first time that counting the number of  
17 optimal solutions can be formulated within this general framework as well and solved exactly with variable elimination  
18 and search algorithms. Therefore, our algorithms can be extended to more general semiring based graphical models.  
19 For example to finding the number of optimal policies in an influence diagram.

20 We also show that, unlike many other reasoning tasks that are formulated within the semiring framework and admit  
21 simple partitioning based bounding schemes (e.g., using for instance the mini-bucket approach), the #opt task cannot be  
22 approximated using these kinds of schemes to produces bounds (upper or lower) on the number of optimal solutions.

23 We will fix all typos and presentation issues. Regarding the small fonts used for the algorithms, we will use an extra  
24 page to avoid all the space issues. This will also allow us to expand on the motivation behind our work and discuss  
25 additional examples where we believe that #opt is important.

26 **Review #3:** Thanks for your comments and suggestions.

27 We are aware of the Yanover and Weiss work on finding the M most probable configurations and we will definitely  
28 cite it. In principle, this work together with the more recent work by Flerova, Marinescu and Dechter (which we cite)  
29 on finding the M best solutions using search can be extended to compute the number of optimal solutions but these  
30 specialized m-best algorithms will compute the count of optimal solutions by enumeration which is much less effective  
31 (especially on problems with many optimal solutions) than our proposed algorithms which do not depend on the actual  
32 number of optimal solutions.

33 In the extreme case where there is only 1 optimal solution a search based m-best algorithm should be as good as our  
34 AOBB approach for #opt. In all other cases, the m-best algorithms must iterate for several values of m in order to find  
35 the number of optimal solutions. Therefore, we emphasize that our AND/OR branch and bound for #opt is always  
36 better than a specialized m-best approach. We will extend the discussion in the paper to clarify this point.