Exploiting Within-Clique Factorizations in Junction-Tree Algorithms

Julian McAuley, Tibério Caetano
julian.mcauley@nicta.com.au tiberio.caetano@nicta.com.au

Abstract

We show that the expected computational complexity of the **Junction-Tree Algorithm for MAP-inference** in graphical models **can be improved**. Our results apply whenever the potentials over maximal cliques of the triangulated graph are factored over subcliques. This enlarges the class of models for which exact inference is efficient.

MAP-estimation

Passing messages in graphical models requires that we compute 'max-marginals', one step of which requires choosing the maximum product amongst two (or more) lists:

$$\hat{i} = \underset{i \in \{1 \dots N\}}{\operatorname{argmax}} \{\mathbf{v}_a[i] \times \mathbf{v}_b[i]\}.$$

Although this seems to be a **linear** time operation, it can be reduced to $O(\sqrt{N})$ (in the expected case) if we know the permutations that sort $\mathbf{v}_a$ and $\mathbf{v}_b$. Our results arise due to the fact that knowing these permutations allows us to ignore much of the search space:

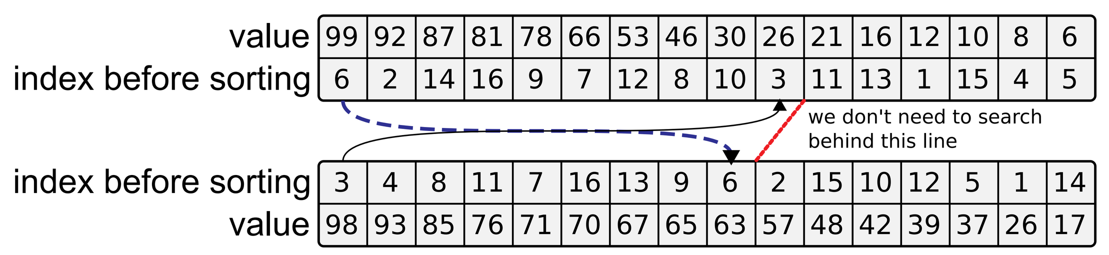

We find that these permutations can be computed efficiently **whenever the model's cliques factorize**.

Graphs whose potentials factorize

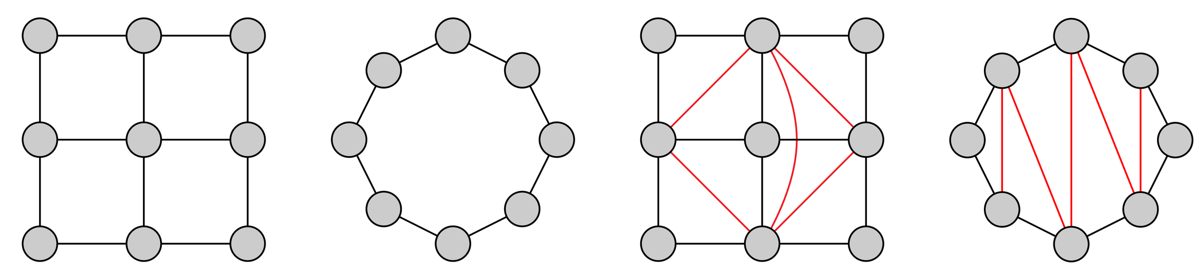

The graphical models shown above contain only pairwise factors; triangulating them increases their maximal clique size.

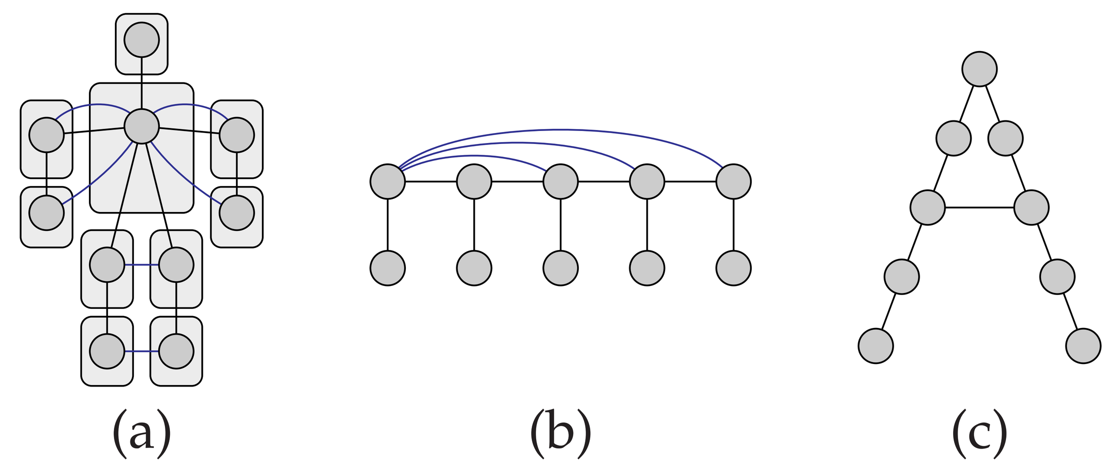

Analogous cases are common in many applications: (a) a model for pose reconstruction from [Sigal and Black, 2006]; (b) a 'skip-chain CRF' from [Galley, 2006]; (c) a model for deformable matching from [Coughlan and Ferreira, 2002]. Although the (triangulated) models have cliques of size three, they **factorize** into pairwise terms.

Computing max-marginals in cliques that factorize

Graph:

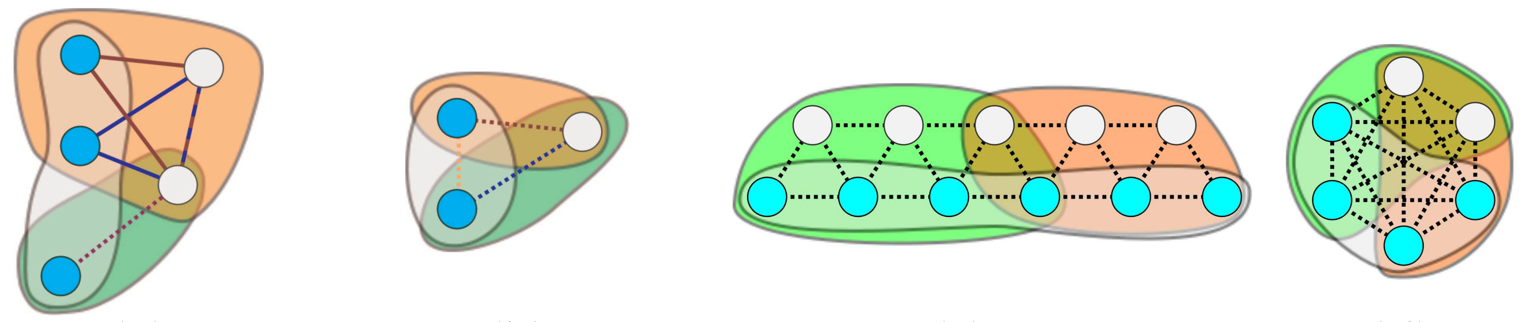

	(a)	(b)	(c)	(d)	(e)
Naïve solution:	$\Theta(N^5)$	$\Theta(N^3)$	$\Theta(N^{11})$	$\Theta(N^6)$	$\Theta(N^M)$
Our algorithm:	$O(N^3\sqrt{N})$	$O(N^2\sqrt{N})$	$O(N^6\sqrt{N})$	$O(N^5)$	$O(N^{5M/6})$
Speed-up:	$\Omega(N\sqrt{N})$	$\Omega(\sqrt{N})$	$\Omega(N^4\sqrt{N})$	$\Omega(N)$	$\Omega(N^{M/6})$

Some example cliques whose max-marginals are to be computed with respect to the coloured nodes. The factors are indicated using differently coloured edges (dotted edges indicate pairwise factors).

Analysis

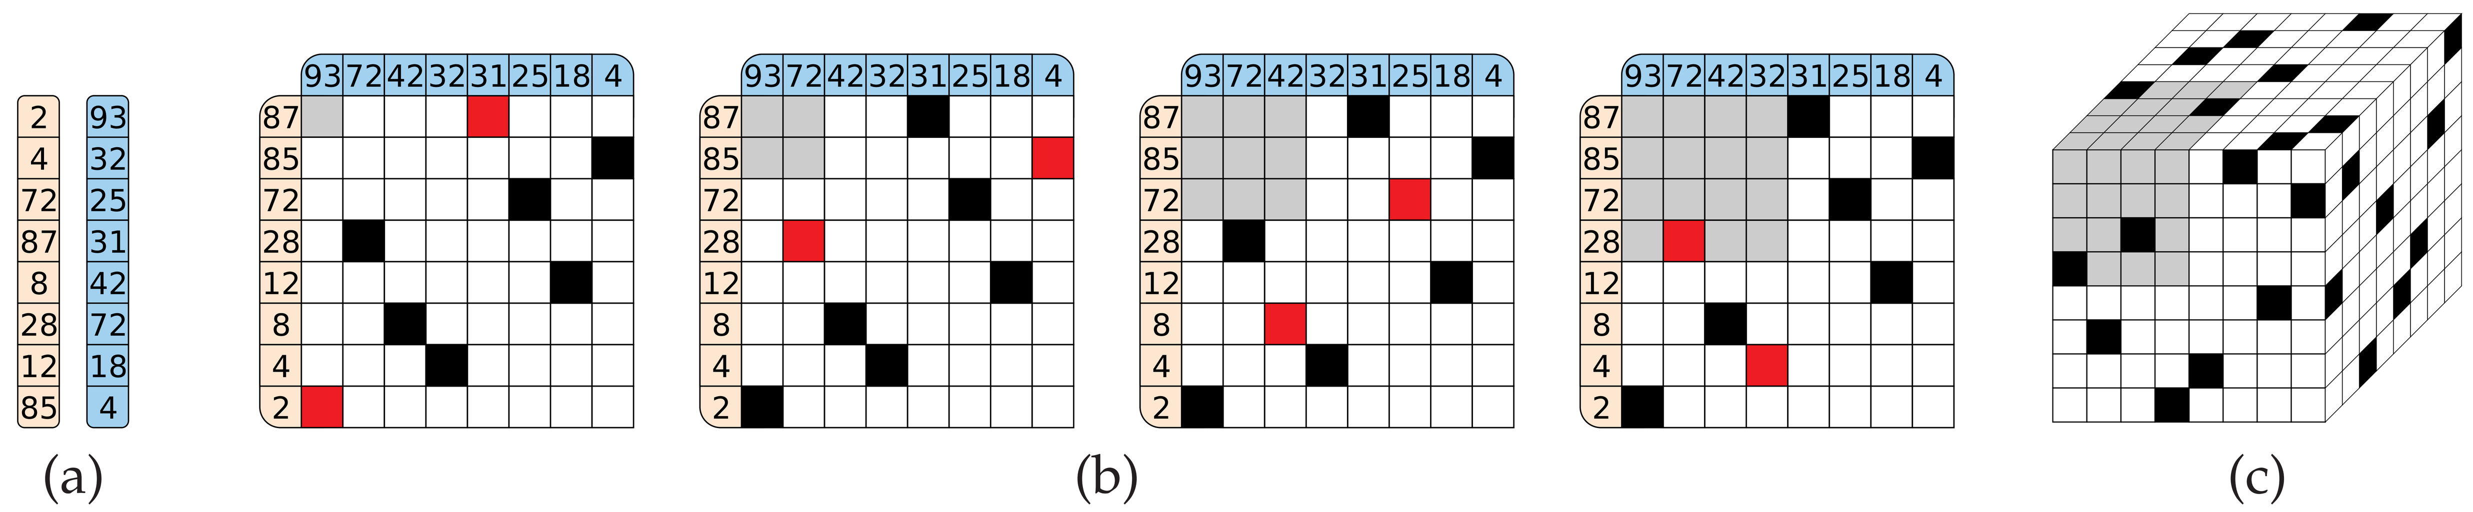

(a) Two lists for which we want to compute $\operatorname{argmax}_{i \in \{1 \dots N\}} \{\mathbf{v}_a[i] \times \mathbf{v}_b[i]\}$. (b) The black squares show the permutation from $\mathbf{v}_a$ to $\mathbf{v}_b$ after sorting; the red squares show the products being computed at each step; the algorithm terminates once the grey box contains an entry. (c) Our results generalize to several lists.

Results

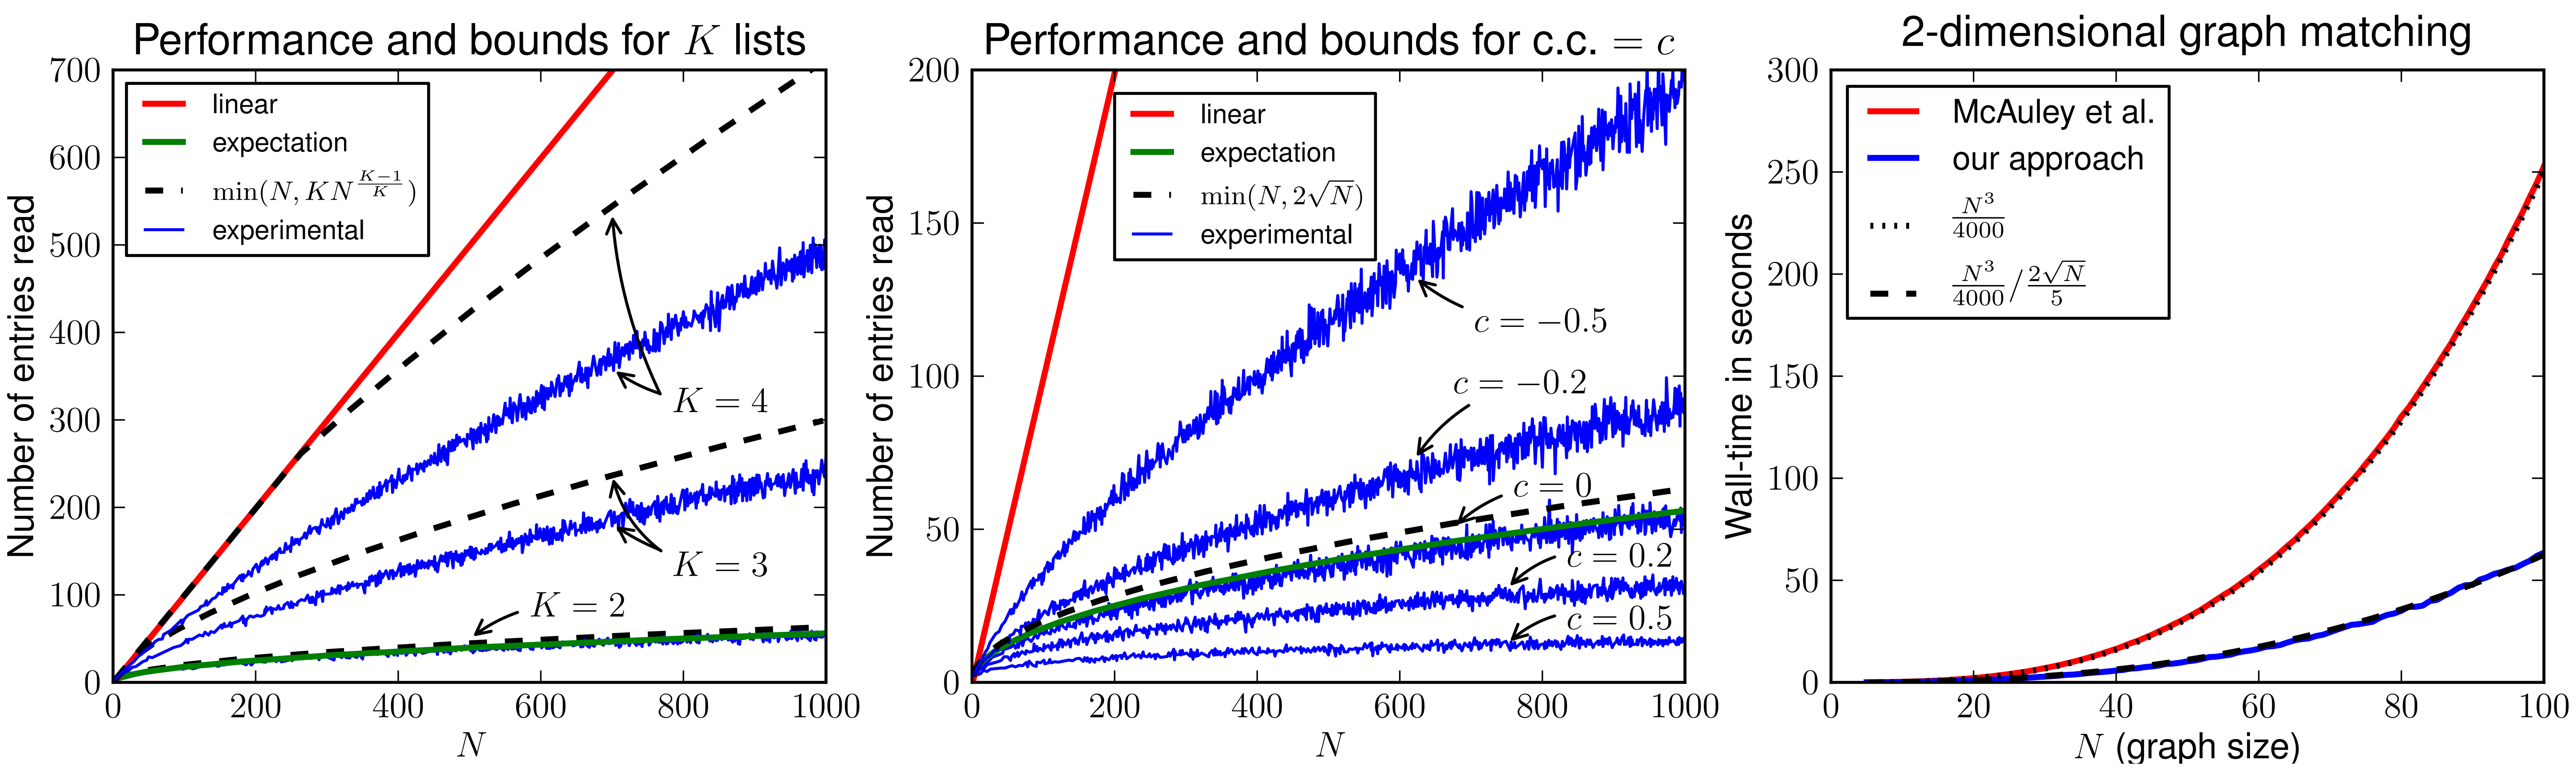

Left: Performance of our algorithm over 100 trials; the dotted lines show the bounds. Centre: Performance of our algorithm for different correlation coefficients. Right: The running time of our method on a graph matching experiment over 10 trials [McAuley et al., 2008].

Bibliography

[Coughlan and Ferreira, 2002] Coughlan, J. M. and Ferreira, S. J. (2002). Finding deformable shapes using loopy belief propagation. In *ECCV*.

[Galley, 2006] Galley, M. (2006). A skip-chain conditional random field for ranking meeting utterances by importance. In *EMNLP*.

[McAuley et al., 2008] McAuley, J. J., Caetano, T. S., and Barbosa, M. S. (2008). Graph rigidity, cyclic belief propagation and point pattern matching. *IEEE Transanctions on Pattern Analysis and Machine Intelligence*, 30(11):2047–2054.

[Sigal and Black, 2006] Sigal, L. and Black, M. J. (2006). Predicting 3D people from 2D pictures. In *AMDO*.