An empirical study of Gaussian belief propagation and application in the detection of F-formations

Francois Kamper

Stellenbosch University

October 23, 2017

What is Gaussian belief propagation (GaBP)?

- Belief propagation applied on a Markov graph (MG) constructed from a multivariate Gaussian distribution in canonical form.
- Can be viewed as an iterative message-passing algorithm.
- When constructing a message from node i to a neighbour node j, node i collects all incoming messages from neighboring nodes (except from j).
- These messages are used by node i to form a belief and this belief is then propagated to node j.

GaBP can be used for ...

- performing approximate marginalization on a Gaussian MG in the sense that (assuming convergence) it provides the correct marginal means and (potentially loose) approximations for the marginal precisions.
- solving linear systems (variational inference) and approximating inverse diagonal blocks without direct matrix inversion.
- other novel applications (an example is given later).

Shortcomings of GaBP include ...

- convergence is not guaranteed in loopy graphs.
- convergence can be slow when the precision matrix is ill-conditioned.
- even if convergence occurs the approximations for the marginal precisions can be poor.

What do we propose to do about these shortcomings?

- we contract the beliefs formed in the current round of message-passing to beliefs formed in the previous round using a L_2 penalty through a tuning parameter λ.
- after a round of updates, damping is performed on the means and the damping factors is automatically computed from λ (adaptive damping).
- this preserves the exactness of the means and the penalized BP will converge for sufficiently large λ.
- the marginal precision approximations provided by the penalized BP can be more accurate compared to those from normal GaBP.
- empirical evidence suggest that the λ yielding the best marginal precision approximations is close to the λ yielding the fastest convergence.

Purpose and Construction of the empirical study.

- Purpose of the empirical study is to investigate some of the claims made in the previous section.
- The convergence speed of GaBP is heavily influenced by the spectral radius of $\mathbf{I} - \mathbf{S}$, where $\mathbf{S}$ is the precision matrix scaled to have all diagonal entries equal to one.
- We generate random precision matrices and potential vectors where the spectral radius of $\mathbf{I} - \mathbf{S}$ is set to a specific value.
- For each generated pair of precision matrix and potential vector we compare the output generated by different GaBP variants.
- For the regularized GaBP-variants we used a heuristic measure to determine λ.

Table of number of iterations required for convergence.

Radius	GaBP	hGaBP	GaBP-m	hGaBP-m
0.9	137.77	18.72	60.31	19.07
0.905	146.47	18.68	60.02	18.99
0.91	152.75	18.63	63.52	19.11
0.915	161.61	18.89	65.55	19.37
0.92	170.53	18.79	65.10	19.20
0.925	183.23	18.98	68.32	19.42
0.93	194.41	19.10	68.61	19.65
0.935	210.74	19.03	70.48	19.60
0.94	230.64	19.26	72.93	19.74
0.945	247.37	19.16	75.94	19.87
0.95	272.07	19.21	78.94	19.74
0.955	304.92	19.30	80.69	19.87
0.96	342.12	19.49	80.81	20.06
0.965	391.43	19.43	86.39	19.96
0.97	455.64	19.59	87.84	20.09
0.975	547.96	19.60	90.07	20.14
0.98	689.92	19.85	93.72	20.32
0.985	NA	19.60	96.30	20.25
0.99	NA	19.83	98.69	20.56
0.995	NA	20.09	105.84	20.33

Table 1: Summary of mean number of iterations required for convergence as a function of the zero-diagonal spectral radius.

Table of mean KL-distances.

Radius	GaBP	hGaBP	GaBP-m	hGaBP-m
0.9	0.38	0.04	2.71	0.85
0.905	0.38	0.04	2.69	0.93
0.91	0.40	0.03	2.89	0.90
0.915	0.44	0.06	3.12	0.91
0.92	0.47	0.05	3.27	1.01
0.925	0.49	0.05	3.47	1.12
0.93	0.52	0.05	3.61	1.40
0.935	0.52	0.07	3.67	1.18
0.94	0.54	0.05	3.81	1.62
0.945	0.57	0.06	3.94	1.03
0.95	0.59	0.05	4.14	1.51
0.955	0.59	0.04	4.22	1.36
0.96	0.65	0.07	4.55	1.34
0.965	0.66	0.07	4.63	1.27
0.97	0.66	0.08	4.64	1.37
0.975	0.72	0.07	4.99	1.44
0.98	0.72	0.09	5.03	1.38
0.985	0.72	0.07	5.01	1.30
0.99	0.80	0.09	5.51	1.56
0.995	0.83	0.06	5.76	1.45

Table 2: Summary of mean KL-distance ($\times 10^3$) of the converged posteriors to the true marginals as a function of the zero-diagonal spectral radius. In general the posterior precisions have better convergence behavior than the posterior means, hence the availability of values in the last 3 entries of the first column.

F-formations.

- A F-formation arises whenever two or more people sustain a spatial and orientational relationship in which the space between them is one to which they have equal, direct, and exclusive access (Kendon, 1990).
- We are interested in detecting F-formations from data obtained from cameras during the SALSA poster session.
- For each of the 18 individuals taking part in the poster session we have their xy-coordinates as well as their head- and body-poses. For our analysis we used the ground-truth data.

The basic Idea

- We use the data to obtain association scores between the 18 individuals. These association scores are such that individuals closer to each other (xy-coordinates) and with aligning poses will have a higher score. All scores are positive.
- Among 5 individuals we might have the following scores:

$$\begin{bmatrix} 1.00 & 0.26 & 0.06 & 0.13 & 0.19 \\ 0.26 & 1.00 & 0.18 & 0.06 & 0.26 \\ 0.06 & 0.18 & 1.00 & 0.17 & 0.27 \\ 0.13 & 0.06 & 0.17 & 1.00 & 0.11 \\ 0.19 & 0.26 & 0.27 & 0.11 & 1.00 \end{bmatrix}$$

The basic Idea (Continued)

- Instead of using this matrix directly we perform regularized GaBP and replace the off-diagonal entries with the precision-component of the message between two individuals.
- The matrix (rounded to 2 decimals) changes to :

$$\begin{bmatrix} 1.00 & -0.07 & 0.00 & -0.02 & -0.04 \\ -0.08 & 1.00 & -0.04 & 0.00 & -0.08 \\ 0.00 & -0.04 & 1.00 & -0.03 & -0.08 \\ -0.02 & 0.00 & -0.03 & 1.00 & -0.01 \\ -0.04 & -0.08 & -0.08 & -0.02 & 1.00 \end{bmatrix}$$

The basic Idea (Continued)

- If we rescale the matrix from the previous slide such that the off-diagonal entries are all positive with mean equal to the mean of the off-diagonal entries of the original matrix we see the following
$$\begin{bmatrix} 1.00 & 0.32 & 0.02 & 0.09 & 0.17 \\ 0.34 & 1.00 & 0.18 & 0.02 & 0.34 \\ 0.02 & 0.17 & 1.00 & 0.15 & 0.33 \\ 0.08 & 0.02 & 0.13 & 1.00 & 0.06 \\ 0.18 & 0.35 & 0.35 & 0.07 & 1.00 \end{bmatrix}$$
- Note the changes in the magnitude of the off-diagonal entries.
- We can perform thresholding on this matrix (instead of the original) to detect F-formations. Two individuals i, j are defined to be in a F-formation if entry i, j or entry j, i of the thresholded matrix is non-zero.

Computing the association matrix

- Individual i receives a positive definite 2×2 score matrix, $\mathbf{S}_{ii}$.
- The association between individual i and j is described through a 2×2 matrix $\mathbf{S}_{ij}$.
- The 2×2 matrices are used to incorporate both types of poses.
- The weight between individual i and j is $\exp[-\tau_1||\mathbf{x}_i - \mathbf{x}_j||_2^2]$ where $\mathbf{x}_k$ represents the xy-coordinates of individual k.
- Each individual receives two coordinates, one based on the head-pose and the other on the body-pose.
- Suppose the body-pose of individual i is θ_i, then the body-pose coordinate for individual i is $\mathbf{z}_i(R_i) = \mathbf{x}_i + R_i\mathbf{v}_i$ with $\mathbf{v}_i = (\cos(\theta_i), \sin(\theta_i))'$ and $R_i > 0$.

Computing the association matrix (Continued)

- To find R_i for all individuals i we propose finding

$$\underset{R_k,\,\forall k}{\operatorname{argmin}}\{\sum_{i \neq j} W_{ij}||\mathbf{x}_i - \mathbf{x}_j + (R_i\mathbf{v}_i - R_j\mathbf{v}_j)||_2^2\},$$

  under the restriction that $R_k > 0 \ \forall \ k$.
- Individuals with coinciding poses will be moved towards each other, but the effect is dampened depending on the distance between their xy-coordinates. The same can be done for the head-poses.
- $\mathbf{S}_{ii}$ is obtained by computing the radial kernel, with scale parameter τ_2, based on the Euclidean distance between the head- and body-pose coordinates of individual i.

Computing the association matrix (Continued)

- $\mathbf{S}_{ij}$ is obtained by computing the radial kernel, with scale parameter τ_2, based on the Euclidean distance between the head- and body-pose coordinates of individual i and individual j.

Constructing the MG

- The MG is constructed by assigning a 2-dimensional node to each individual.
- We associate with node i the precision matrix $\mathbf{S}_{ii}$.
- The link between node i and node j is $\mathbf{S}_{ij}$.
- After GaBP (only considering the precision-components) we have a 2×2 matrix $\mathbf{Q}_{ij}$ for all $i \neq j$.
- We define a new matrix $\mathbf{L}$, where all diagonal entries equal one and entry i, j is the spectral radius of $\mathbf{Q}_{ij}$.
- We apply our thresholding technique on $\mathbf{L}$ to detect F-formations.

SALSA: Poster session

- SALSA recorded social interactions among 18 participants for over 60 minutes in a poster presentation and cocktail party.
- The ground-truth data (for the poster session) consists of head- and body-poses along with the *xy*-coordinates for all participants over 645 frames taken in 3 second intervals.
- We define a grid of values $0.1 \leq \tau_1, \tau_2 \leq 0.3$. For each pair in the grid we apply the model as discussed and determine the threshold giving the highest F_1-score.
- By best possible estimation we mean that we select the value of τ_1, τ_2 giving the highest F_1-score for each frame individually.
- By best total estimation we select the value of τ_1, τ_2 (used for all frames) giving the highest F_1-score over all the frames.

Summary of F_1-scores.

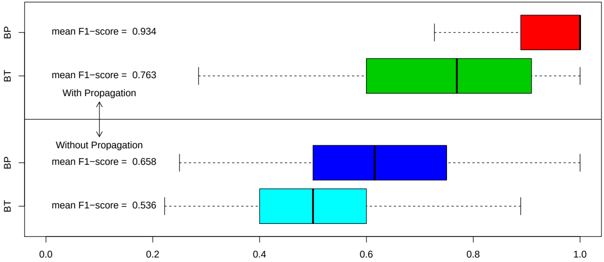

Figure 1: Empirical Results for the poster session. Figure consists of box-plots of the F_1-score associated with each method. On the vertical axis BP = Best Possible and BT = Best Total. Belief propagation substantially improves performance.

SALSA: Poster session (Continued)

- The values used for the best total estimation are $\tau_1 = 0.3$ and $\tau_2 = 0.24$. These values give a precision, recall and F_1 -measure of 0.728, 0.78 and 0.753, respectively.
- Dynamic modeling can be used to improve on the precision, recall and F_1 -measure up to 0.86, 0.9 and 0.88 respectively. Some state of the art scores.

	HVFF lin	HVFF ms	GC
GT-B	0.66 / 0.65 / 0.65	0.72 / 0.69 / 0.71	0.74 / 0.72 / 0.73
GT-H	0.65 / 0.63 / 0.64	0.70 / 0.67 / 0.69	0.72 / 0.70 / 0.71

- We showed empirically that the use of belief propagation over basic distance measures can provide better F-formation predictions.
- The use of the mean-components of messages for detecting F-formations was ignored.
- More refined ways of employing belief propagation should be considered (for instance introducing a L_1 penalty for inference).
- Novel message-passing procedures (not necessarily belief propagation) attempting to emulate how people communicate in groups.