THE ISBA BULLETIN

Vol. 24 No. 1 | March 2017

The official bulletin of the International Society for Bayesian Analysis

A MESSAGE FROM THE PRESIDENT

- Kerrie Mengersen -
ISBA President, 2017
k.mengersen@qut.edu.au

G'day! It's a pleasure and honour to serve as ISBA President for 2017. I am following in the very big shoes left by the 25 ISBA Presidents that have preceded me, most recently Steve MacEachern who has done a great job guiding the Society through the World Meeting in Sardinia last year, the upgrading of our website and a raft of other challenges and opportunities. Many thanks to Steve and his team of supporters, including Past President Alexandra Schmidt, Secretary Amy Herring and Treasurer Murali Haran. We are very fortunate that Amy is serving as Secretary again this year and that Robert (Bobby) Gramacy has stepped up as Treasurer.

There is no doubt that the World Meeting in Sardinia was a highlight last year, so thank you very much to all those who put so much hard work into making it such a success. A particular 'thank you' goes to Michele Guindani and Chris Hans. We are now in the throes of organising the 2018 Meeting in Edinburgh, under the guidance of Chris Hans and Clair Alston. In response to feedback from members, we will be trialling a more distributed model whereby the core conference and a strong network of satellite meetings will be collectively come under the banner of the World Meeting. This will provide Society members with the opportunity to attend more focused gatherings as well as liaise with the broader ISBA community. Many of the satellite meetings, as well as the sessions at the core conference, will be led by our Society's Sections which include Bayesian Computation, Bayesian Nonparametrics, Biostatistics and Pharmaceutical Statistics, Economics Finance and Business, Environmental Science, Industrial Science, Objective Bayes and j-ISBA. Visit the website for more information!

Speaking of the website, medals for excellence and endurance must be awarded to the team who has been negotiating the new face of our Society. We are close, so please be patient and be prepared to be impressed: it's great!

Stepping outside our own Society, one of the ambitions for this year is that ISBA establishes stronger links with other cognate Societies globally. To achieve this, I need your help: please let me know if there are opportunities to link with your national professional societies, either through a formal exchange of letters or more opportunistically through co-sponsorship or co-badging of events or speakers, combined courses, etc. We already have strong connections with ASA and we are growing connections with NIPS, IMS, the Bernoulli Society and others, so let's keep up the momentum! (Continued p. 2)

In this issue

Looking further again, we recognise the vital role that students and early career professionals play in our Society in particular, and more generally in the future of Bayesian Statistics. If you align to this category of member, then we want to hear from you! What would you like from the Society and how do you think that could be effected? Join j-ISBA to have a voice and be heard.

A great way to be heard is through the Society's mouthpieces, the ISBA Bulletin, edited by Beatrix Jones and our journal, Bayesian Analysis, led by Editor-in-Chief Bruno Sansó. These are essential threads in the fabric of our profession, so please support, promote and contribute to them.

Looking forward, our future as a Society is bright. We are young at heart but we have established a solid foundation. We celebrate diversity of interests and geography, bound by a common affiliation to Bayesian analysis. Together we are much more than the sum of the parts. We are ISBA!

– Kerrie Mengersen

FROM THE EDITOR

- Beatrix jones -
m.b.jones@massey.ac.nz

This issue I would like to highlight the review of introductory Bayesian MOOC's by Richard Arnold of Victoria University, Wellington. If you don't know, MOOCs are Massive Open Online Courses. The two MOOCs featured offer free access to introductory lectures by some well known Bayesians; for a modest fee you can also be assessed on what you've learned. How did we come to feature such a review? It was suggested by a reader! If you have something you would like to see included in the *Bulletin*, particularly if you are willing to write it, send an email to me or the relevant associate editor–all their details are on the final page of this bulletin. This issue we are welcoming a new Associate Editor into the fold: Leanna House of Virginia Tech is taking over the interview section. We look forward to some compelling interviews of prominent Bayesians later this year.

As well as News of the World, and our usual communications from the BA editor and Program Council, in this issue you will also find a reflection by Alicia Carriquiry on Stephen Feinberg's contributions to our profession, and a Software Highlight by Alberto Caimo and Nial Friel featuring the BERGM R package for working with Bayesian exponential random graph models. Finally, another reader has suggested the *Bulletin* would be better off with a single column format we can scroll through on our laptops or tablets. We trial such a format over the next few pages–let me know which you prefer!

UPDATE FROM BA

From the BA Editor

- Bruno Sansó -
bruno@soe.ucsc.edu

The March issue of the journal is available online at https://projecteuclid.org/current/euclid.ba. This issue does not include a paper with discussion as there is a need to reduce the long list of papers in advance publication status.

In this issue of the *Bulletin* I would like to comment on some of the numbers that our Managing Editor, Tony Pourmohamad, has produced regarding the handling of papers submitted to Bayesian Analysis. The journal receives a number of submissions per year that, in the last four years, has been between 150 and 200 papers. The acceptance rate of papers in the journal has been decreasing steadily since 2007, when it was about 50%, to a number that is now hovering 20%. The distribution of the time to first review is somewhat irregular. It has a high concentration around the origin, that corresponds to papers that are rejected by the Editor in Chief, either directly, or after consultation with one of the Editors, the so called desk rejections. The mode and the median of the distribution are around 60 days. The data for 2016 indicate that the median time to first review may be decreasing, as we obtained a value of 43 days. But this estimate is biased by the fact that many of the papers received in the last two months of 2016 have not been counted. While we focus on providing quality feedback to the authors, the editorial board of the journal is working hard to reduce the time that papers take to be reviewed. We are particularly interested on shortening the tails of the distribution. Fortunately the journal relies on the good will of a very active Bayesian community. Keep sending your best papers to BA and make your best efforts to respond to our refereeing requests with thorough and timely reviews!

REMEMBERING STEPHEN E. FIENBERG

- Alicia Carriquiry -
alicia@iastate.edu

Stephen Elliott Fienberg passed away on December 14, 2017, shortly after turning 74 years of age. He had been diagnosed with cancer about four years earlier, but kept such a demanding and productive schedule in spite of the disease, that most of us were convinced that he would prevail in the end. Steve's death was a tremendous loss for statistics and for science in general, and he will be sorely missed.

Steve was born in Toronto, Canada, on November 27, 1942. In high school, it became obvious to him that he was good at, and greatly enjoyed, the sciences, in particular the mathematical sciences. Steve liked to tell that while his mother (who passed away in Toronto less than two years ago) thought that he was a genius, he was just a good student with an aptitude for mathematics and a passion for ice hockey. Steve went on to the University of Toronto, where he obtained a degree in Mathematics in 1964. He applied to and was admitted into the doctoral program in statistics at Harvard University, and finished his PhD in 1968, under the supervision of Fred Mosteller. Meeting Fred Mosteller and working closely with him in a variety of different projects was a life-changing experience for Steve. Mosteller at the time was a rare statistician in that he was genuinely driven by interesting applied projects. The fact that statistics could be brought to bear on so many other disciplines and to such good effect was a revelation, and these early experiences had a long lasting impact on Steve's professional life. Steve had a profound respect and a deep affection for Mosteller, and often spoke of how much he had learned from his years as a graduate student working with him.

After completing his PhD, Steve was recruited by William Kruskal, then Chair of the Department of Statistics at the University of Chicago, and began his career as an Assistant Professor. Kruskal, much

Figure 1: Stephen E. Fienberg in 2012

like Mosteller, was also attracted to applications and introduced Steve to many different faculty in a wide range of disciplines with whom Steve began collaborating. In those days, political polling was becoming widespread, but polling methodology was not yet fully developed. Steve became intrigued by the political polling carried out by a local newspaper and this interest led in part to many years of research in different aspects of survey sampling.

Even though Steve enjoyed his years in Chicago, he and his wife Joyce moved to Minnesota, largely for personal reasons. In Minnesota, Steve held his first administrative position as Chair of the Department of Applied Statistics in the Saint Paul campus of the University of Minnesota. From Minnesota, Steve and Joyce moved to Carnegie Mellon University, which Steve called his academic home and where he spent the rest of his professional life. Steve joined the Department of Statistics at CMU in 1980, and with the exception of a short stint as Provost of York University in Canada, he never left. A "Conversation with Steve" by two of Steve's dearest friends Miron Straf and Judy Tanur was published in *Statistical Science* in 2013, and includes many biographical details about Steve. It also paints a wonderfully warm picure of Steve as a person.

Steve's first research contributions were largely based on his dissertation research. Mosteller introduced Steve to a National Research Council study that was known as the "National Halothane Study", and which Steve described as a "giant contingency table". For his dissertation, Steve developed loglinear model theory and methods useful for the analyses of categorical data such as those collected in the study and together with Yvonne Bishop and Paul Holland (also Mosteller students), published the well-known book entitled *Discrete Multivariate Analysis* (1975), with the green covers. Throughout his career, Steve continued to advance the theory and implementation of loglinear models, but also built world class research programs in privacy and confidentiality, machine learning and algebraic statistics.

Steve was already interested in Bayesian theory by the time he arrived in CMU, but his career as a Bayesian statistician really took off then. Steve joined Jay Kadane and Morrie DeGroot when he came to CMU, and the three of them contributed to making the department a destination for Bayesians from all over the world. In the Mosteller and Kruskal tradition, Steve developed an interest in a wide variety of problems in other disciplines, and was instrumental in the creation and editing of journals with a focus on the principled application of statistics, including the *Annals of Applied Statistics*, the *Journal of Privacy and Confidentiality*, and more recently, *The Annual Review of Statistics and Its Application*. Bayesians have much for which to be thankful to Steve; he was the second President of ISBA, and was largely responsible for attracting the funding for the ISBA 2000 World Meeting in Crete. He contributed the first article in the first issue of *Bayesian Analysis*, entitled "When did Bayesian inference become Bayesian?", a historical recount of the most important developments in Bayesian statistics between the time when Bayes' opus was published posthumously, and the end of the last century. During what the called "the Bayesian Renaissance", Steve became a tireless and effective promoter of the Bayesian paradigm worldwide.

Possibly because of Mosteller's and Kruskal's influence, Steve's passion was to advance the principled and constructive use of statistics to solve real problems in other disciplines, preferably when those problems had a public policy implication. Not long ago, Eric Lander, the renowned scientist and co-Chair of President Obama's Council of Advisors on Science and Technology (PCAST), referred to Steve as follows:

> Steve Fienberg is not just a statistician-he is a *public statistician*. He has brought his considerable statistical prowess to bear on problems of great public importance (emphasis added).

Steve's first forays into public policy began shortly after he arrived in CMU; he became involved with various government agencies on matters of data collection and data sharing, and joined the Committee on National Statistics (CNSTAT) soon after it was established. Through his work with CNSTAT (which continued throughout his career), Steve had an opportunity to positively impact the work at most (if not all) federal agencies in charge of collecting, synthesizing and sharing official statistics.

After CMU, the institution in the US that most benefited from Steve's knowledge and dedication was the National Academies of Science, Engineering and Medicine (NASEM). Steve began participating in NASEM's activities in the mid-80s, but became truly involved after his election to the National Academy of Sciences in 1999 (one of Steve's proudest professional accomplishments). Not only did Steve focus much of his efforts on the NASEM; he also motivated many of us to follow in his footsteps and view the NASEM as an effective vehicle to introduce positive change in society through science-based public policy and decision-making. Steve served the Academies in a variety of roles, but possibly the most consequential of those roles was his co-chairing of the Report Review Committee, which Steve viewed as an efficient means to ensure that every report published by the Academies was based on solid science and (as appropriate) on sound statistical reasoning.

I have had the privilege of calling Steve a friend for over 25 years, and his mentoring and efforts on my behalf changed the course of my professional life. A few years ago, Steve encouraged me (and Hal Stern and Karen Kafadar) to submit a proposal to establish a NIST Center of Excellence in Forensic Statistics, which would be located at our four institutions, with "headquarters" at Iowa State. Surprisingly to me (but not to Steve!) we were successful and obtained the funds to create the center in 2015. Steve was the intellectual leader, the one with the grand vision and the far-reaching ideas, and I have great hopes that the work on which we have embarked in the center will have a positive impact on society, because Steve was instrumental in setting us off on the right path. Hal, Karen and I are tremendously thankful to Bill Eddy, who was Steve's close friend and colleague, for jumping in and picking up where Steve let off.

Steve was an affectionate and loyal friend, and he seemed to know everyone. But his world revolved around his wife Joyce and the rest of his family. Steve adored his grandchildren and loved spending time with them. He was particularly fond of having them all descend upon him and Joyce for extended summer visits. While not religious in the usual sense, Steve was proud of his Jewish heritage and culture and strongly believed in keeping the rituals and traditions, and in observing the holidays, as a means to nurture his sense of belonging and reinforce his ties to the Jewish community to which he felt so close.

Among his many other interests and activities, Steve always found time for his other "passions": ice hockey (which he continued practicing even into his 70s) and the New York Times crossword puzzle. He loved good food and fine wine (and single malt scotch) and was the instigator of the "Saturday

Figure 2: The Extravagant Dining Group at JSM 2016 in Chicago (counter-clockwise from front): Alicia Carriquiry, Steve Fienberg, Veronika Rockova, Ed George, Merlise Clyde, Jim Berger, Ann Berger

Night Extravagant Dining" group (Jim Berger, Susie Bayarri, Merlise Clyde, Ed George, Dick De Veaux, Robert Wolpert (emeritus), Veronika Rockova, myself, and anyone else reckless enough to join us) during the Joint Statistical Meetings. But he was determined to encourage good dining habits among JSM goers long before then; remember *Belizaire*, anyone?

Steve had a marshmallow core even though on occasion he could unsheathe the fangs. He was immensely patient with young faculty and students and with anyone who was really trying, but he did not suffer fools gladly. He loved a good competition but did his best to have the last word. He could be demanding, but he gave of himself generously and never ever expected anything in return. He was well respected by some, idolized by others, and ignored by no one, and sometimes he seemed invincible. His many friends will miss him dearly, for perhaps ever.

References

Fienberg, SE. 2006. When did Bayesian analysis become "Bayesian"? *Bayesian Analysis*, 1:1-40.

Straf, ML, Tanur, JM. 2013. A conversation with Stephen E. Fienberg. *Statistical Science*, 28:447-463.

FROM THE PROGRAM COUNCIL

- Clair Alston-Knox -
Chair of the Program Council c.alston-knox@griffith.edu.au

ISBA 2018 World Meeting: 24th-29th June, 2018. Edinburgh, UK After the huge success of the 2016 World Meeting which was held 13 -17 June, 2016, Forte Village Resort Convention Center, Sardinia, Italy, work is now well under way on planning our 2018 event. The 2018 World meeting will be held 24th-29th June, 2018 (please note the change of date from original proposal) at Appleton Tower, University of Edinburgh. http://www.edinburghfirst.co.uk/venues/appleton-tower.

We encourage members to start thinking about potential contributions to the program. During May, Sections will be approached to propose Special Topic Sessions of interest to their members, with contributed sessions from individuals being advertised at a later time. We will keep you updated with progress and due dates for contributions throughout the year.

Meeting Sponsorships and Co-sponsorships If you are planning a meeting in 2018 and would like to request financial sponsorship (or co-sponsorship) from ISBA, applications are due by May 30th, 2017. Information on how to submit a request can be found at https://bayesian.org/meetings/planning. Please note that ISBA is in the process of changing websites and various systems. Please email the program council at program-council@bayesian.org with requests for sponsorship (in addition, please cc this email to c.alston-knox@griffith.edu.au).

Requests for non-financial endorsement of events can be assessed by the program council as needs arise during the year (see https://bayesian.org/meetings/planning-endorsed). Inquiries can be directed to program-council@bayesian.org and c.alston-knox@griffith.edu.au.

Upcoming ISBA Events We would like to highlight the following upcoming meetings that are being co-sponsored by ISBA:

- BNP 11: 11th Conference on Bayesian Nonparametrics, June 26 -30, 2017. Paris, France. https://www.ceremade.dauphine.fr/~salomond/BNP11/index.html

- BISP 10: Bayesian Inference in Stochastic Processes, June 13-15, 2017. Milano, Italy. https://www.unibocconi.eu/wps/wcm/connect/Site/BISP10/Home/

- Summer school on advanced Bayesian methods September 11- 15, 2017. Leuven, Belgium. https://ibiostat.be/seminar/summerschool2017/Summer2017Bayesian

- BAYES 2017: May 22 ? 25, 2017. Albacete, Spain. http://www.bayes-pharma.org/

- School of Statistics for Astrophysics 2017: Bayesian Methodology 9 -13 October 2017, Autrans, France. https://stat4astro2017.sciencesconf.org

- COBAL V: 7 - 10 June, 2017. CIMAT, Guanajuato, Mexico. http://cobal2017.eventos.cimat.mx/

- O'Bayes 2017: 10-13th December 2017. Austin, Texas. Website will go live soon.

REVIEW: INTRODUCTORY BAYESIAN MOOCS

- Richard Arnold -
Victoria University, Wellington richard.arnold@vuw.ac.nz

For large classes doing standard introductory material, Massive Open Online Courses (MOOCs) provide unparalleled efficiency for both students and teachers. Each works in his or her own time. The lecturer only needs to deliver lectures once, but can also hone, polish and rework material. The student can view and re-view material as slowly or as quickly as they like. Students don't hold each other up, or feel as if they are left behind.

Whether MOOCs end up replacing a substantial proportion of conventional University teaching remains to be seen: they suffer from the same drawbacks of conventional distance learning (reduced student engagement, the lack of the motivating experience of being in a room with peers, lack of direct access to the lecturer, as well as the problem of verifying who has actually completed an assignment or quiz). However the new and developing technology associated with MOOCs (new video presentation techniques, integration of written materials with video and online content, peer discussion forums) do mean that MOOCs are a significant step forward for distance learners. For Universities they provide a showcase for teaching staff, attracting new students to enroll, and also a convenient way to allow students to access introductory pre-requisite material prior to enrolment in a regular taught programme.

So what MOOCs are available for students of Bayesian statistics? Here I review two introductory Bayesian statistics courses which are available on the Coursera platform (www.coursera.com). They are 'Bayesian Statistics: From Concept to Data Analysis', taught by Herbie Lee from University of California at Santa Cruz, and the 'Bayesian Statistics' module of the 'Master Statistics with R' programme taught by Mine Çetinkaya-Rundel, David Banks, Colin Rundel and Merlise Clyde of Duke University.

The Coursera platform has some very nice features for its students. Courses are bundled into week-long blocks, with short videos (3-10 minutes typically), full video transcripts, supporting documents, and quizzes (mostly multi-choice answers, though some allowing numerical or text answers to questions). Videos can be played at variable speed (happily without altering the instructor's voice pitch), and instructors can arrange for the video to pause and ask the student a question at key points.

Students pay around US$50 per month for access to a course – which entitles them to submit quizzes for credit, and participate in peer discussion forums. For the Santa Cruz course, this covers the full 4 week course. The Duke Bayesian course could also be completed in 4 weeks, though it is one of 5 courses in a 'Specialization' in introductory statistics with with R – which at the recommended pacing would take 5-6 months to complete. Students can of course go faster if they wish. At the end of the course students receive a certificate. It is also possible to gain access to the course content without payment, but this removes access to assessments, feedback and peer discussion forums – and there is no recognition of learning available at the end.

The two courses contain broadly similar content – Bayes Theorem in the familiar and simple discrete and continuous settings seen in any Bayesian course (Beta-Binomial, Gamma-Poisson, Normal-Normal). Both spend time talking about conjugacy, Bayesian hypothesis testing, and the Bayesian approach to linear regression. Both spend time comparing the frequentist and Bayesian approaches, with the Duke course going further and spending time on Lindley's Paradox.

Courses on Bayesian statistics have two very particular demands which distinguish them from other courses, even from frequentist courses: namely requirements for mathematics and computational methods. The former presents complications for statistics programs of all kinds worldwide: we have an increasing demand for statistical training for students who have a low level of mathematical knowledge. The latter is more easily addressed by teaching students the computational methods they need as and when they need them.

The two courses differ in their pre-requisite levels of mathematics. The Duke course assumes only very basic mathematics, whereas the Santa Cruz course explicitly states that first year University calculus is required – although more at the level of familiarity than mastery. Both courses nevertheless spend time developing expressions for posteriors, involving the necessary concepts of finite and infinite sums and integrals. Consistent with its higher mathematical pre-requisite, the Santa Cruz course had more of the feel of a traditional University course with instructor Herbie Lee writing out extensive expressions and derivations using a Lightboard. In contrast, the Duke course instructors stood in front of their content slides, where fully formed mathematical expressions appeared – for the most part as additional information for those students who were interested in the detail.

Both courses use R for analysis, providing detailed code and demonstration examples of Bayesian analyses. The Santa Cruz course also provides equivalent instruction for each example in Excel. These examples are an excellent start for students who want to go off and do something practical after the course – the Duke course in particular giving an extensive worked example of model selection and averaging in linear regression using the BAS package in R.

These courses both provide a learner a motivating start for learning more about Bayesian statistics– the different styles of presentation are engaging and the amount of thought and preparation that has gone into the two courses is very apparent. The Duke course is conceptually ambitious – with its coverage of frequentist paradoxes and model selection – and a nice feature is a set of four short interviewers with Bayesian academics and practitioners. The Santa Cruz course covers less ground, but does so with a well-paced higher technical intensity, building on an assumption of greater mathematical knowledge.

I'd recommend both courses to students who want to know what Bayesian statistics is about – whether or not they intend to take it further. Those intending to carry on with further study will be well prepared for higher level courses in Bayesian statistics.

NEWS FROM THE WORLD

Meetings and conferences

Workshop on Statistical Social Network Analysis with R, Birkbeck, University of London. June 21-23, 2017.

Statistical network analysis plays an integral role in data science across multiple disciplines such as social sciences, business, and management helping make decisions based on the interpretation of complex relational data structures. This hands-on workshop will provide an overall understanding of statistical models for the analysis of relational data with application to real-world problems in social sciences. It is open to postgraduates, early career researchers and scientists from academia, industry and government agencies.

While no prior experience with social network analysis will be assumed, participants will be expected to understand fundamental social network concepts and terminology, and to have some knowledge of basic concepts in statistical inference. Participants are encouraged to bring their own data to work on. Live demonstrations and social interactions between participants will be an important part of this workshop. Participants will be introduced to several R packages.

The workshop will be limited to 30 participants. Registration will open soon. For full details, please visit the conference website http://bida.bbk.ac.uk/SSNAR17 or following the twitter account @SSNAR17.

11th Annual RCEA Bayesian Econometric Workshop, Melbourne, Australia. July 3 - 4, 2017.

The 11th Annual Bayesian Econometric workshop of the Rimini Centre for Economic Analysis (RCEA) will take place at the University of Melbourne, Australia, on July 3rd and 4th, 2017. Keynote speakers are Sylvia FrÃ¼hwirth-Schnatter from the Vienna University of Economics and Business, and Sylvia Kaufmann from the Study Center Gerzensee.

This is the first time the RCEA Bayesian workshop will run in Australia and is the result of a partnership with the Bayesian Analysis and Modeling Research Group at the University of Melbourne. In addition to the keynote speakers, there will be a full set of contributed sessions. Papers in all areas of Bayesian econometrics are welcome.

Registration will start on April 1, 2017. Detailed information will be available at the websites http://fbe.unimelb.edu.au/economics/bam and www.rcfea.org.

IMSM Graduate Student Modeling Workshop, North Carolina State University, NC. July 16-26, 2017.

The 23rd Industrial Mathematical & Statistical Modeling (IMSM) Workshop for Graduate Students will take place at North Carolina State University, between 16-26 July 2017. The workshop is sponsored by the Statistical and Applied Mathematical Sciences Institute (SAMSI) together with the Center for Research in Scientific Computation (CRSC) and the Department of Mathematics at North Carolina State University.

The IMSM workshop exposes graduate students in mathematics, engineering, and statistics to exciting real-world problems from industry and government. The workshop provides students with experience in a research team environment and exposure to possible career opportunities. On the first day, a Software Carpentry bootcamp will bring students up-to-date on their programming skills in Python/Matlab and R, and introduce them to version control systems and software repositories.

Local expenses and travel expenses will be covered for students at US institutions. The application deadline is April 15, 2017. Information is available at http://www.samsi.info/IMSM17 and questions can be directed to grad@samsi.info.

SOFTWARE HIGHLIGHT

BERGM: BAYESIAN EXPONENTIAL RANDOM GRAPH MODELS IN R

Alberto Caimo, Nial Friel
alberto.caimo@dit.ie, nial.friel@ucd.ie

Networks are relational data that can be defined as a collection of nodes interacting with each other and connected in a pairwise fashion. From a statistical point of view, networks are relational data represented as mathematical graphs. A graph consists of a set of n nodes and a set of m edges which define some sort of relationships between pair of nodes called dyads. The connectivity pattern of a graph can be described by an $n \times n$ adjacency matrix y encoding the presence or absence of an edge between node i and j:

$$y_{ij} = \begin{cases} 1, & \text{if } (i,j) \text{ are connected,} \\ 0, & \text{otherwise.} \end{cases}$$

Two nodes are adjacent or neighbours if there is an edge between them. If $y_{ij} = y_{ji}, \forall i, j$ then the adjacency matrix is symmetric and the graph is undirected, otherwise the graph is directed and it is often called a digraph. Edges connecting a node to itself (self-loops) are generally not allowed in many applications and will not be considered in this context.

Exponential random graph models

Introduced by [9] to model individual heterogeneity of nodes and reciprocity of their edges, the family of exponential random graph models (ERGMs) was generalised by [6], [18] and [16]. ERGMs constitute a broad class of network models (see [15] for an introduction) that assume that the observed network y can be explained in terms of the relative prevalence of a set of network statistics $s(y)$:

$$p(y|\theta) = \frac{\exp\{\theta^t s(y)\}}{z(\theta)} \tag{1}$$

where the normalising constant $z(\theta)$ is intractable for non trivially-small networks.

Due to the complexity of networks, it is necessary to reduce the information to describe essential properties of the network. Usually this is done via network statistics, a series of counts of subgraph configurations (e.g., the number of edges, stars, triangles, functions of degree distributions, edgewise shared parters, etc.), catching the relevant information [16].

Bayesian inference

In the ERGM context (see [18] and [15]), the posterior distribution of model parameters θ given an observed network y on n nodes maybe written as:

$$p(\theta|y) = \frac{p(y|\theta)\, p(\theta)}{p(y)} = \frac{\exp\{\theta^t s(y)\}}{z(\theta)}\, \frac{p(\theta)}{p(y)}, \tag{2}$$

where $s(y)$ is a known vector of sufficient network statistics [11], $p(\theta)$ is a prior distribution placed on θ, $z(\theta)$ is the intractable likelihood normalising constant, and $p(y)$ is the model evidence. The presence of the intractable ERGM likelihood implies that the usual suite of standard Bayesian inferential methods, especially standard MCMC tools are not possible in this context. However recent work has shown that the ERGM can be given the full Bayesian treatment as we now outline.

The `Bergm` package for R

The `Bergm` package [4] for `R` [12] implements Bayesian analysis for ERGMs [2, 3, 5, 17, 1]. The package provides a comprehensive framework for Bayesian inference using Markov chain Monte Carlo (MCMC) algorithms. It can also supply graphical Bayesian goodness-of-fit procedures that address the issue of model adequacy.

The package is simple to use and represents an attractive way of analysing network data as it offers the advantage of a complete probabilistic treatment of uncertainty. `Bergm` is based on the `ergm` package [10] which is part of the `statnet` suite of packages [8] and therefore it makes use of the same model set-up and network simulation algorithms. The `ergm` and `Bergm` packages complement each other in the sense that `ergm` implements maximum likelihood-based inference whereas `Bergm` implements Bayesian inference. The `Bergm` package has been continually improved in terms of speed performance over the last years and we feel that this package now offers the end-user a feasible option for carrying out Bayesian inference for networks with several thousands of nodes.

Approximate exchange algorithm

In order to approximate the posterior distribution $p(\theta|y)$, the `Bergm` package uses the exchange algorithm described in Section 4.1 of [2] to sample from the following distribution:

$$p(\theta', y', \theta|y) \propto p(y|\theta)p(\theta)\epsilon(\theta'|\theta)p(y'|\theta')$$

where $p(y'|\theta')$ is the likelihood on which the simulated data y' are defined and belongs to the same exponential family of densities as $p(y|\theta)$, $\epsilon(\theta'|\theta)$ is any arbitrary proposal distribution for the augmented variable θ'. As we will see in the next section, this proposal distribution is set to be a normal centred at θ.

At each MCMC iteration, the exchange algorithm consists of a Gibbs update of θ' followed by a Gibbs update of y', which is drawn from $p(\cdot|\theta')$ via an MCMC algorithm [10]. Then a deterministic exchange or swap from the current state θ to the proposed new parameter θ'. This deterministic proposal is accepted with probability:

$$\min\left(1, \frac{q_\theta(y')p(\theta')\epsilon(\theta|\theta')q_{\theta'}(y)}{q_\theta(y)p(\theta)\epsilon(\theta'|\theta)q_{\theta'}(y')} \times \frac{z(\theta)z(\theta')}{z(\theta)z(\theta')}\right),$$

where q_θ and $q_{\theta'}$ indicate the unnormalised likelihoods with parameter θ and θ', respectively. Notice that all the normalising constants cancel above and below in the fraction above, in this way avoiding the need to calculate the intractable normalising constant.

The approximate exchange algorithm is implemented by the `bergm` function in the following way:

`for` $i = 1, \ldots, N$

1. `generate` θ' `from` $\epsilon(\cdot|\theta)$
2. `simulate` y' `from` $p(\cdot|\theta')$
3. `update` $\theta \to \theta'$ `(log) probability:`

$$\min\left(0,\ [\theta - \theta']^t\,[s(y') - s(y)] + \log\left[\frac{p(\theta')}{p(\theta)}\right]\right)$$

`end for`

where $s(y)$ is the observed vector of network statistics and $s(y')$ is the simulated vector of network statistics. Step 2. above requires a draw from the ERGM likelihood and perfect sampling in principle is a possibility, however practically this is out of reach as no such sampler has yet been developed. Therefore the pragmatic approach we take is to run a Gibbs sampler for `aux.iters` iterations targetting $p(\cdot|\theta')$. In order to improve mixing a parallel adaptive direction sampler (ADS) approach [7, 14] is considered as the default procedure.

To illustrate the inferential procedure, we fit a 4-dimensional ERGM to the Faux Mesa High School network data [13] including uniform homophily between students with the same 'grade' (`nodematch('Grade')`), and statistics capturing the degree distribution (`gwdegree`) and transitivity effect (`gwesp`):

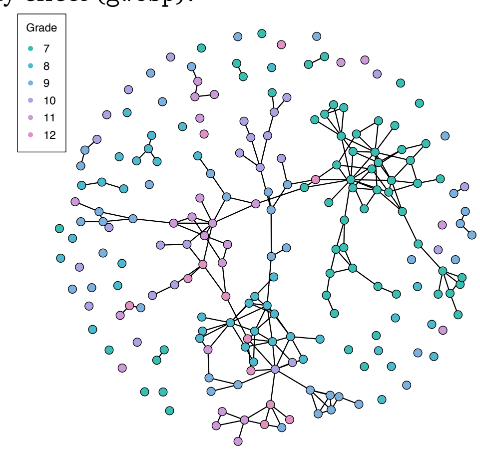


```
> model <- y ~ edges +
+              nodematch('Grade') +
+              gwdegree(0.2, fixed = TRUE) +
+              gwesp(0.2, fixed = TRUE)
```

and we use the `bergm` function with $20,000$ auxiliary iterations for network simulation and 6 MCMC chains for the ADS procedure consisting of $2,000$ main iterations each:

```
> bergm.post <- bergm(model,
+                     burn.in = 300,
+                     main.iters = 2000,
+                     aux.iters = 20000,
+                     nchains = 6,
+                     gamma = 0.6)
```

The estimation took about 200 seconds. A summary of the MCMC results is available via the `bergm.output` command:

```
> bergm.output(post)
```

```
                                Mean        SD    Naive SE
theta1 (edges)            -6.4539945 0.2269798 0.002072032
theta2 (nodematch.Grade)   2.0653066 0.1562896 0.001426723
theta3 (gwdegree)          0.1555102 0.2156994 0.001969057
theta4 (gwesp.fixed.0.2)   1.6045295 0.1624254 0.001482734

                          Time-series SE
theta1 (edges)               0.013682987
theta2 (nodematch.Grade)     0.009303603
theta3 (gwdegree)            0.013919679
theta4 (gwesp.fixed.0.2)     0.009959682

                                2.5%        25%        50%
theta1 (edges)            -6.8988256 -6.6003492 -6.4532048
theta2 (nodematch.Grade)   1.7879473  1.9557575  2.0537140
theta3 (gwdegree)         -0.2825274  0.0158655  0.1618899
theta4 (gwesp.fixed.0.2)   1.2841624  1.4983524  1.6031113
                                 75%      97.5%
theta1 (edges)            -6.3043630 -6.0255210
theta2 (nodematch.Grade)   2.1665645  2.3952768
theta3 (gwdegree)          0.3118872  0.5568012
theta4 (gwesp.fixed.0.2)   1.7099524  1.9425482

Acceptance rate: 0.1965833
```

Density and trace plots are produced automatically by the `bergm.output` function:

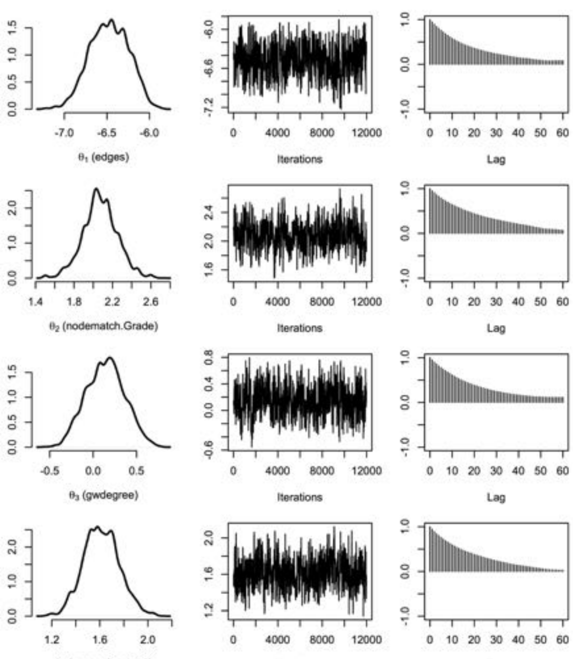

Posterior predictive goodness-of-fit diagnostics plots are available via the `bgof` command, as shown in the figure below.

```
> bgof(bergm.post,
+      aux.iters = 20000,
+      n.deg = 14,
+      n.dist = 15,
+      n.esp = 10)
```

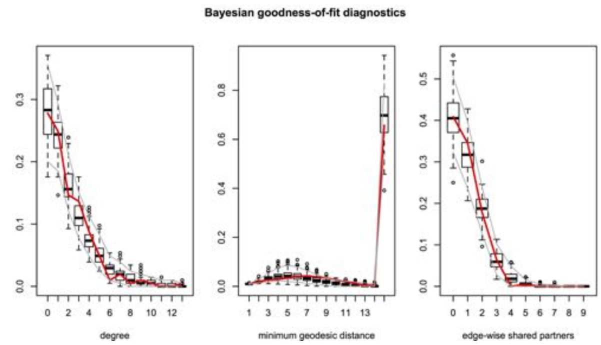

The plots in the figure indicate a very good fit of the model in terms of a higher-level network statistics in the data.

Pseudo-posterior calibration

An alternative approach to Bayesian inference for ERGMs has been proposed by [1] based on replacing the intractable ERGM likelihood with a tractable pseudo-likelihood approximation. This results in a so-called pseudo-posterior distribution for which it is straightforward to sample from using the usual MCMC toolbox, for example. However it is well understood that Bayesian inference based on the pseudolikelihood can yield poor estimation and this motivated [1] to develop an approach which allows one to correct or calibrate a sample from such a pseudo-posterior distribution so that it is approximately distributed from the target posterior distribution. This is achieved by estimating the maximum a posteriori (MAP) of the posterior distribution and also estimating the Hessian of the posterior distribution at the MAP. Both of these quantites can then be used to define an affine transformation of the pseudo-posterior distribution to one that is approximately distributed as the posterior distribution. The pseudo-posterior calibration approach can be carried out using the `calibrate.bergm` function:

```
> cbergm.post <- calibrate.bergm(model,
+                  iters = 1000,
+                  aux.iters = 20000,
+                  noisy.nsim = 100,
+                  noisy.thin = 1000,
+                  mcmc = 10000)
```

The estimation took about 80 seconds and the MCMC output can be analysed by using the `bergm.output` function.

In the plots below, we see that the posterior estimates from the `calibrate.bergm` function is in good agreement with that corresponding to the `bergm` function:

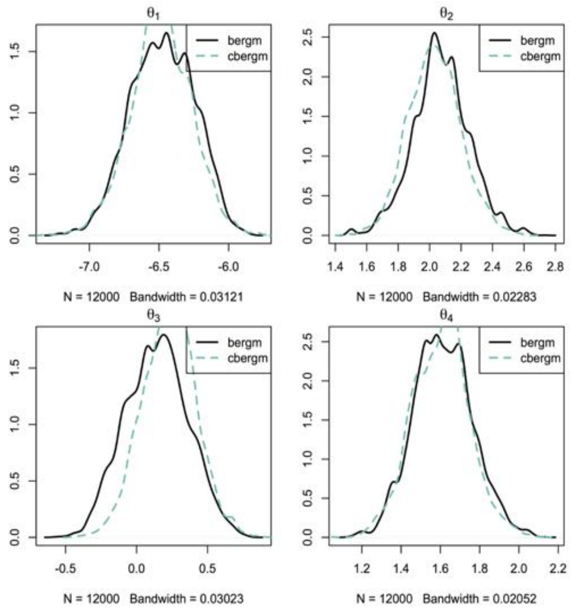

More information

The `Bergm` package is available on the CRAN at: https://CRAN.R-project.org/package=Bergm, and also on GitHub at: https://github.com/acaimo/Bergm.

References

[1] L. Bouranis, N. Friel, and F. Maire. Bayesian inference for misspecified exponential random graph models. *arXiv preprint arXiv:1510.00934*, 2015.

[2] A. Caimo and N. Friel. Bayesian inference for exponential random graph models. *Social Networks*, 33(1):41 – 55, 2011.

[3] A. Caimo and N. Friel. Bayesian model selection for exponential random graph models. *Social Networks*, 35(1):11 – 24, 2013.

[4] A. Caimo and N. Friel. Bergm: Bayesian exponential random graphs in R. *Journal of Statistical Software*, 61(2):1–25, 2014.

[5] A. Caimo and A. Mira. Efficient computational strategies for doubly intractable problems with applications to bayesian social networks. *Statistics and Computing*, 25(1):113–125, 2015.

[6] O. Frank and D. Strauss. Markov graphs. *Journal of the American Statistical Association*, 81:832–842, 1986.

[7] W. R. Gilks, G. O. Roberts, and E. I. George. Adaptive direction sampling. *Statistician*, 43(1):179–189, 1994.

[8] M. S. Handcock, D. R. Hunter, C. T. Butts, S. M. Goodreau, and M. Morris. statnet: Software tools for the representation, visualization, analysis and simulation of network data. *Journal of Statistical Software*, 24(1):1–11, 2007.

[9] P. W. Holland and S. Leinhardt. An exponential family of probability distributions for directed graphs (with discussion). *Journal of the American Statistical Association*, 76:33–65, 1981.

[10] D. R. Hunter, M. S. Handcock, C. T. Butts, S. M. Goodreau, and M. Morris. ergm: A package to fit, simulate and diagnose exponential-family models for networks. *Journal of Statistical Software*, 24(3):1–29, 2008.

[11] M. Morris, M. S. Handcock, and D. R. Hunter. Specification of exponential-family random graph models: terms and computational aspects. *Journal of statistical software*, 24(4):15–48, 2008.

[12] R Development Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2011.

[13] M. D. Resnick, P. S. Bearman, R. W. Blum, K. E. Bauman, K. M. Harris, J. Jones, J. Tabor, T. Beuhring, R. E. Sieving, and M. Shew. Protecting adolescents from harm: findings from the national longitudinal study on adolescent health. *Jama*, 278(10):823–832, 1997.

[14] G. O. Roberts and W. R. Gilks. Convergence of adaptive direction sampling. *Journal of Multivariate Analysis*, 49(2):287–298, 1994.

[15] G. Robins, P. Pattison, Y. Kalish, and D. Lusher. An introduction to exponential random graph models for social networks. *Social Networks*, 29(2):169–348, 2007.

[16] T. A. B. Snijders, P. E. Pattison, G. L. Robins, and Handcock M. S. New specifications for exponential random graph models. *Sociological Methodology*, 36:99–153, 2006.

[17] Stephanie Thiemichen, Nial Friel, Alberto Caimo, and Göran Kauermann. Bayesian exponential random graph models with nodal random effects. *Social Networks*, 46:11–28, 2016.

[18] S. Wasserman and P. Pattison. Logit models and logistic regression for social networks: I. An introduction to Markov graphs and p^*. *Psycometrica*, 61:401–425, 1996.
