Bayesian Artificial Intelligence for Decision Making under Uncertainty

Anthony C. Constantinou

June, 2018

Risk and Information Management Research Group, School of Electronic Engineering and Computer Science, Queen Mary University of London, London, UK, E1 4NS.

E-mail: a.constantinou@qmul.ac.uk Web: www.constantinou.info

Funder: EPSRC
Project: EP/S001646/1
Period: Jun 2018 – Jun 2021
EPSRC contribution: £475,818

Scheme: UKRI Innovation Fellowship
Priority area: New Approaches to Data Science
Topic classification: Artificial Intelligence
Full economic cost: £594,773

Project partner:
AGENARISK
www.agenarisk.com

Cite as: Constantinou, A. C. (2018). Bayesian Artificial Intelligence for Decision making under Uncertainty. *Engineering and Physical Sciences Research Council* (EPSRC), EP/S001646/1

PROJECT SUMMARY: Scientific research is heavily driven by interest in discovering, assessing, and modelling cause-and-effect relationships as guides for action. Much of the research in discovering relationships between information is based on methods which focus on maximising the predictive accuracy of a target factor of interest from a set of other related factors. However, the best predictors of the target factor are often not its causes and hence, the motto "association does not imply causation". Although the distinction between association and causation is nowadays better understood, what has changed over the past few decades is mostly the way by which the results are stated rather than the way they are generated.

Bayesian Networks (BNs) offer a framework for modelling relationships between information under causal or influential assumptions, which makes them suitable for modelling real-world situations where we seek to simulate the impact of various interventions. BNs are also widely recognised as the most appropriate method to model uncertainty in situations where data are limited but where human domain experts have a good understanding of the underlying causal mechanisms and/or real-world facts. Despite these benefits, a BN model alone is incapable of determining the optimal decision path for a given problem. To achieve this, a BN needs to be extended to a Bayesian Decision Network (BDN), also known as an Influence Diagram (ID). In brief, BDNs are BNs augmented with additional functionality and knowledge-based assumptions to support the representation of decisions and associated utilities that a decision maker would like to minimise or maximise [1]. As a result, BDNs are suitable for modelling real-world situations where we seek to discover the optimal decision path to maximise utilities of interest and minimise undesirable risk.

Because BNs come from statistical and computing sciences, and whereas BDNs come mainly from decision theory introduced in economics, research works between these two fields only occasionally extend from one field to another. As a result, it is fair to say that the landscape of these approaches has matured rather incoherently between these two fields of research. It is possible to develop a new generation of algorithms and methods to improve the way we 'construct' BDNs.

The overall goal of the project is to develop an open-source software that will enable end-users, who may be domain experts and not statisticians, mathematicians, or computer scientists, to quickly and efficiently generate BDNs for optimal real-world decision-making. The proposed system will allow users to incorporate their prior knowledge for information fusion with data, along with relevant decision support requirements for intervention and risk management, but will avoid the levels of manual construction currently required when building BDNs. The system will be evaluated with diverse real-world decision problems including, but not limited to, sports, medicine, forensics, the UK housing market, and the UK financial market.

BACKGROUND

Model A below illustrates the associations between three hypothetical factors, as typically discovered by statistical regression or classical machine learning techniques that ignore assumptions of causation or the direction of influence. Model A tells us that each factor is predictive of each other. On the other hand, model B is a directed acyclic graph and captures the same features under causal or influential assumptions. Unlike association, causal assumptions make claims about the effect of interventions. Specifically, unlike model A, model B tells us that an intervention on *Yellow teeth* will have **no** effect on *Smoking* nor on *Lung cancer*, whereas an intervention on *Smoking* will have an effect on both *Yellow teeth* and *Lung cancer*. So, in model A, the association between *Yellow teeth* and *Lung cancer* comes via the common cause *Smoking*, which can be discovered and eliminated in models of type B.

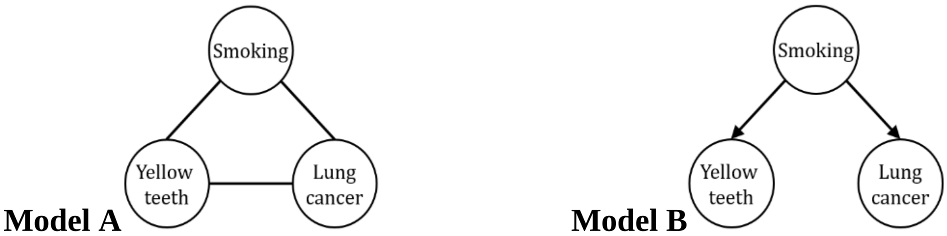

Much of the research on discovering relationships between information is based on methods which focus on maximising the predictive accuracy of a target variable of interest X from a set of observed predictors Y. However, the **best** predictors of X are often **not the causes** of X, hence the motto "*association does not imply causation*". Although the distinction between association and causation is nowadays better understood, what has changed over the last few decades is mostly the way the results are stated rather than the way they are generated. Although scientific research is heavily driven by interest in discovering, assessing, and modelling cause-and-effect relationships as guides for action and decision-making, most scientific conclusions continue to be based on results derived from outputs of models similar to A and hence, often fail to accurately answer the important questions of intervention which require models similar to B.

A Bayesian Network (BN), which is a type of a probabilistic graphical model [2], introduced by Pearl [3, 4], is a directed acyclic graph that offers a framework for modelling relationships between information under causal or influential assumptions. BNs are also widely recognised as the most appropriate method to model uncertainty in situations where the model could benefit from fusing data with knowledge; i.e., in cases where human domain experts have a good understanding of the underlying causal mechanisms and/or real-world facts of a problem where data fail to capture. As a result, BNs are suitable for modelling real-world scenarios for which we seek to simulate the impact of various interventions, available to the decision makers, for maximising outcomes of interest or managing risk. Consequently, BNs are capable of representing models similar to B above, where the nodes represent uncertain variables and the arcs represent the direction of influence. In general, there are three ways to construct BNs:

1. **Knowledge-Based**: the structure of the model and the conditional probability tables (CPTs) for each variable are determined from knowledge.

2. **Data-driven:** the structure and the CPTs are automatically discovered (i.e. learnt) from available data.

3. **Information fusion:** This involves any combination of data and knowledge and hence, combines the two above approaches. For example, a domain expert may specify the structure of the model, while the CPTs are learnt from data.

Further, there are primarily three different types of algorithms that can be used to learn the structure (i.e., the network; the directed graph) of a BN model [1, 5]. These are:

1. **Constraint-based:** These algorithms aim to establish links between variables under the assumption that the arcs represent causal relationships. They perform causal independence checks between variables in sets of triples; a process inherited from the *Inductive Causation* (IC) algorithm [6]. The *Peter and Clark* (PC) and some variants of the *Greedy Equivalent Search* (GES; also mentioned below) algorithms have had major impact in this area of research due to their simplicity and learning strategies [7, 8]. These algorithms are data-driven, but many allow the option to incorporate knowledge into the process of structure learning; e.g., variable B occurs after A and hence B cannot influence A.

2. **Score-based:** These algorithms search for different structures and score them based on one or more *scoring functions*, in terms of how well the fitting distributions agree with the empirical distributions. A large number of algorithms fall within this area of research [9, 10]. Examples include the K2 algorithm [11], Sparse Candidate Algorithm [12], the Optimal Reinsertion algorithm [13], and GES [14]. These algorithms tend to be data-driven, but some of them also allow knowledge to be incorporated into the process of structure learning, as above.

3. **Hybrid algorithms:** These simply combine the two above types of structure learning [9]. Examples include the *max-min hill-climbing* (MMHC) algorithms [15] and L1-Regularisation paths [16].

While each approach has its strengths and weaknesses, it is widely acknowledged that accurate structure learning of BN models is very difficult [5, 9, 10, 17, 18]. In general, these algorithms demonstrate promising performance when tested with simulated/synthetic data, which represent 'fake' data generated from simulation based on predetermined models that are assumed to represent the ground truth [17]. When simulated data is provided as input to these algorithms, in an effort to reverse engineer the ground truth model, the results are promising. However, simulation-based performance does not extend to real-world performance [5, 9, 17, 19-22]. This is because data observations recorded from events in the real world do not to adhere to causal representation in the same way simulated data do, which are based on well-defined causal models.

A BN model provides an effective graphical representation of a problem, and can be used for multiple types of complex inference. Despite these benefits, a BN model alone is incapable of determining the optimal decision pathways of the problem. For example, we may want to determine the optimal treatment, or combination of treatments, to control symptoms or cure a disease, while at the same time taking care to minimise unwanted side-effects. To achieve this, a BN needs to be extended to a *Bayesian Decision Network* (BDN), also known as an *Influence Diagram* (ID). A BDN is a generalisation of a BN in which not only probabilistic inference, but also decision and utility questions can be modelled and solved. So, a BDN can be seen as a BN augmented with additional types of nodes and arcs.

Specifically, whereas in a BN all nodes are considered uncertain 'chance nodes', in a BDN if a node corresponds to a decision to be made we distinguish it as a **decision node** (drawn as rectangle). We also introduce **utility nodes (**drawn as diamonds) which are targeted for maximising or minimising a particular outcome of interest, and **information arcs (**drawn as dashed arcs) entering decision nodes, indicating that the decision is determined by information retrieved from parent nodes. In contrast to normal BN arcs (i.e., conditional arcs), information arcs only pass information forward. Specifying a BDN inevitably requires some level of knowledge since we need to specify the decision options available to the decision maker, and the utilities we seek to minimise or maximise [1]. In fact, there are certain structural rules we need to follow when transforming a BN into a BDN [23-25], such as ensuring that only informational arcs enter a decision node.

Because BNs come from statistical and computing sciences (mainly from Artificial Intelligence and Expert Systems), and whereas it can be argued that BDNs evolved from decision theory introduced in economics (mostly known as *Influence Diagrams*) as an improvement over *Decision Trees*, research works between these two fields only occasionally extend from one field to another. As a result, it is fair to say that some of the limited attempts to learn BDNs [26, 27] make limited reference, and have little relevance to the BN structure learning methods discussed above. The landscape of these approaches has matured rather incoherently between these two fields of research.

RESEARCH HYPOTHESIS

It is possible to develop a new generation of algorithms to discover the structure of BDNs for both causal inference and optimal real-world decision-making, and it is possible to incorporate these into a system that enables end users (who may be domain experts and not statisticians, mathematicians, or computer scientists) to quickly develop relevant BDN models for improved decision support. The proposed system will allow users to incorporate their prior knowledge for information fusion with data, along with relevant decision support requirements for intervention and risk management, but will avoid the levels of manual construction currently required when building BDNs.

METHODOLOGY AND OBJECTIVES:

The overall goal of the project is to develop an open-source software that will allow users to generate BDNs for optimal real-world decision-making. The system will be evaluated with diverse real-world case studies. Specifically, the programme is organised around five work packages (WPs):

1. **WP1 – Structure learning for Bayesian Decision Networks:** We will extend current algorithms that learn BNs so that they can be used to learn BDNs. The new hybrid BDN structure learning algorithm will be capable of modelling different types of information. For example, in addition to the standard BN chance nodes, the new algorithm will also account for decision and utility nodes discussed above (and possibly other types of nodes that can used in BDNs, not covered in this document, such as deterministic and function nodes). Similarly, the new algorithms will also account for different types of relationships (i.e., arcs) between the different types of nodes. For example, in addition to the standard BN conditional arcs, the new algorithm will also account for informational arcs discussed above (and possibly other types of arcs that can be used in BDNs, not covered in this document, such as functional arcs).

2. **WP2 – Knowledge engineering and information fusion:** We will work with existing [28, 29], and also introduce new, knowledge engineering and information fusion methods. The new methods will primarily focus on eliciting, incorporating, and fusing the **decision support requirements**, amongst other relevant knowledge, with data. For example, the new methods will enable end-users to quickly specify known interventions/actions or decisions available to the decision maker along with the targeted variables of interest and their desired state/s (e.g., maximising profit/minimising risk). The system will take the knowledge-based information and fuse it with data, along with other knowledge-based inputs that establish what can and cannot be 'discovered' (if known) for a particular **decision** scenario, and pass everything as different types of inputs to the hybrid BDN structure learning algorithm.

3. **WP3 – Data complexity and model evaluation:** This is based on two subtasks. First, to ensure that the hybrid algorithms can handle various data complexity issues, such as missing data values that are not missing at random, and to take into consideration regularisers to manage model overfitting. The second subtask involves devising evaluation functions that assess the 'quality' of the generated model based on factors that go beyond predictive accuracy and model complexity, and towards the usefulness of the generated model in terms of decision support/making.

4. **WP4 – Case studies:** The algorithms and methods will be evaluated by applying them to a number of different and diverse case studies. We will focus on (but not limited to) sports, medicine, forensics, the UK housing market, and the UK financial market. We have relevant available data from past academic and industrial works [30-44] that can be used in this project. Because the datasets from these studies come from our own previous works, this also enables us to compare how our previously published BNs/BDNs compare to the BDNs generated when based on the software that will be developed for this project. Each of the case studies also has potential to extend to answering other questions of interest in all of these areas of application.

5. **WP5 - System development driven by end-user requirements:** We will implement the new algorithms and resulting methods into an open-source software with a Graphical User Interface (GUI) and a detailed user manual targeted for general users. Both academic and industry end-users will contribute to the specifications of the overall system. We will elicit end-user requirements through questionnaires and interviews. Agena Ltd (see below) will provide us with access to over 1,000 potential end-users from diverse areas of industry, who have strong interest in Bayesian modelling for decision analysis. We will also build and maintain a dedicated website for the project and the open-source software.

The objectives and outputs of the project are summarised in the diagram presented in Figure 1.

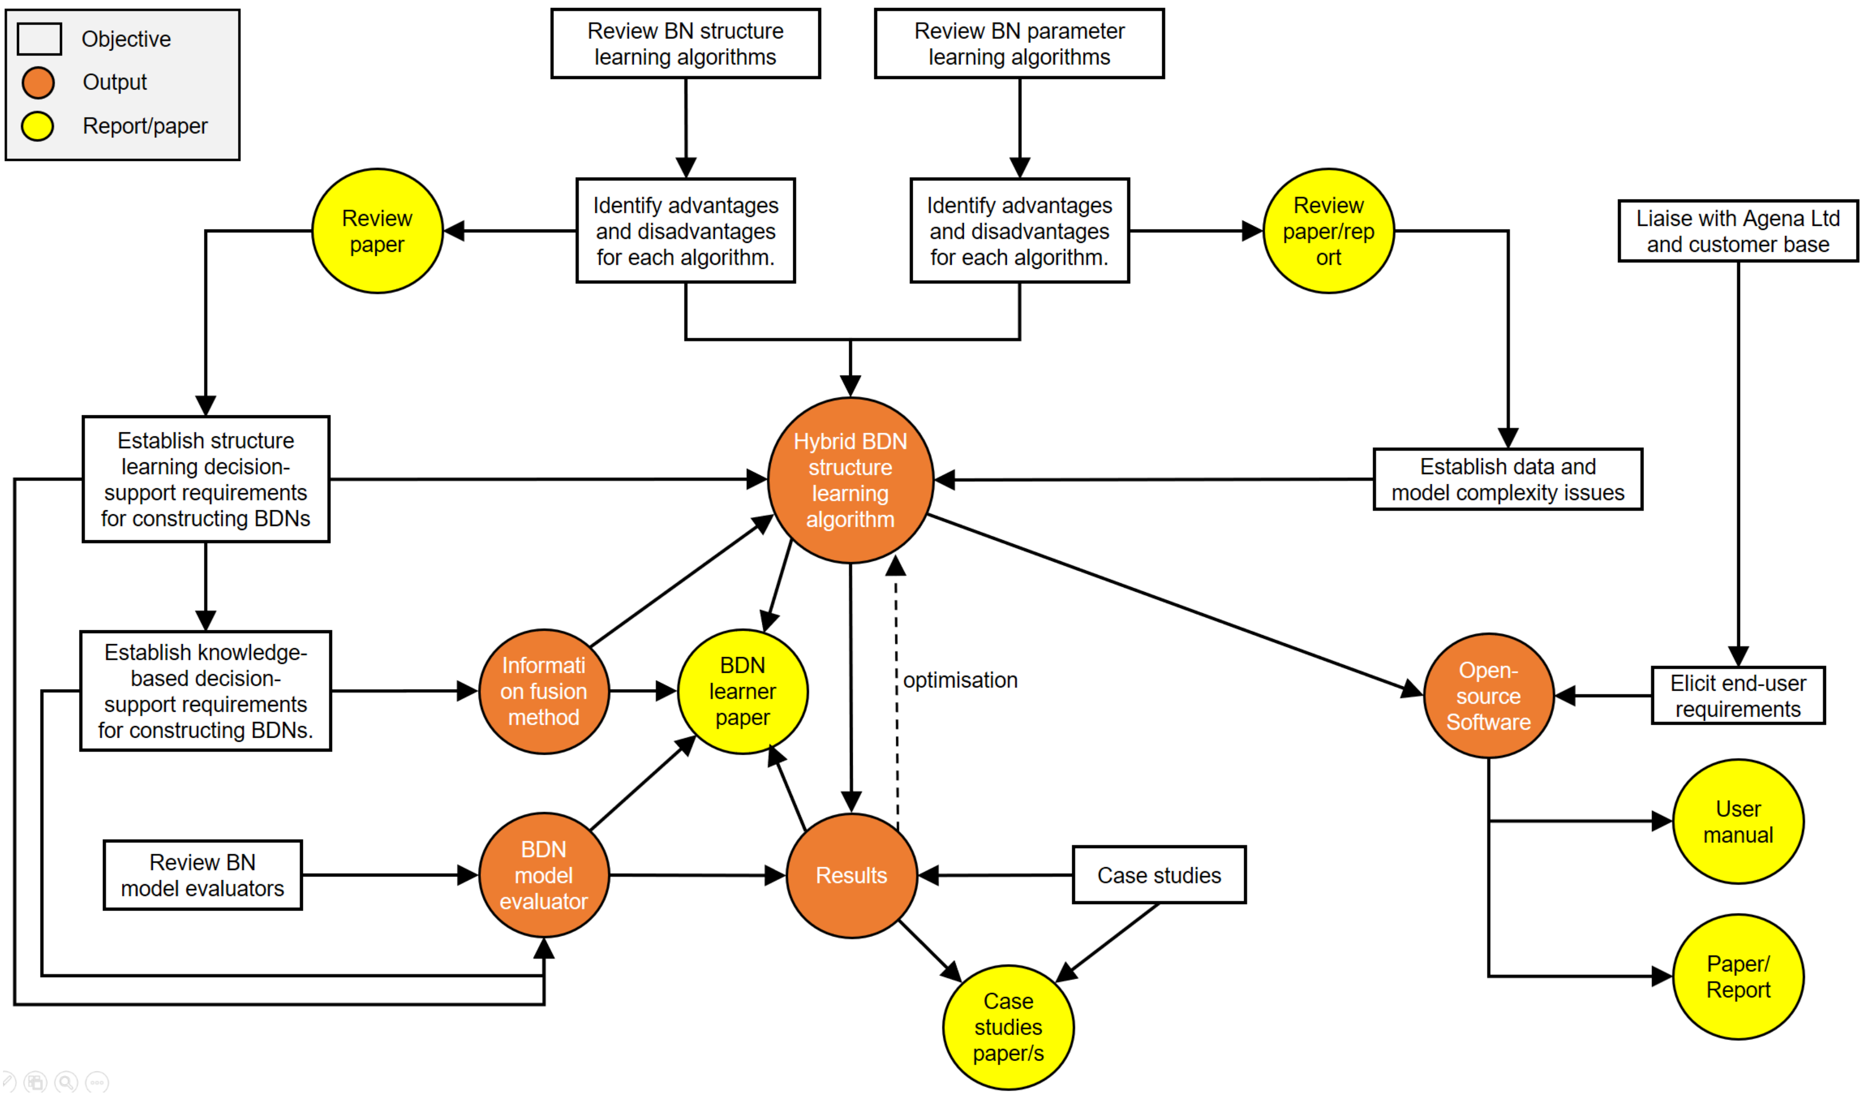

Figure 1. Diagram of project objectives, expected outputs and reports/papers.

INDUSTRIAL ENGAGEMENT

This project will be carried out in collaboration with Agena Ltd (www.agenarisk.com); a UK-based company, with strong QMUL links, that provides state-of-the-art risk analysis and decision support BN software to customers world-wide (e.g., GE Healthcare, Software Engineering Institute Carnegie Mellon, Australian Government, General Dynamics, and many more). Agena Ltd will provide us with full access to their AgenaRisk BN API engine and Desktop versions. The project will benefit from the AgenaRisk Bayesian inference engine, which is unique in enabling us to model variables under the assumption of many different types of continuous distributions, based on their breakthrough *Dynamic Discretisation* algorithm [45]. Agena Ltd will also benefit from this project since the development of the open-source software will be based on their public API programming engine, which should allow them to migrate any developments of interest into their commercial engine with relatively little effort. Agena Ltd has agreed to make the resulting BDN structure learning software open-source. Total in-kind contribution from Agena Ltd is set to £38,100 (for details please refer to resource summary from project partners or letter of support). The Business Development team at QMUL will also support this project by identifying further companies which may benefit from this research.

NATIONAL IMPORTANCE

The proposal falls within the priority area of *New approaches to data science*, and is a perfect fit for the *Decision making under uncertainty* initiative, which has been identified as an area of significant interest by multiple research councils in the UK. New approaches and tools that enable improved decision-making have become particularly important for policy makers, industrial organisations and academic research. This is because such tools are becoming increasingly useful for optimal decision making to maximise or minimise a particular outcome of interest, whether this involves the most cost-effective route, maximum impact, minimum risk, or an equilibrium between them. Policy makers and industrial organisations who invest in big-data solutions also tend to be particularly interested in these alternative advancements because they can provide them with potential to reduce at source much unnecessary data collection, and at the same time improve decision making performance. This proposal is timely in the sense that it combines three emerging technologies; data science, causal modelling, and decision making under uncertainty.

ACADEMIC IMPACT

This project is expected to extend the state-of-the-art in disciplines related to data science, causal modelling, decision making under uncertainty, and knowledge-based systems. Specifically, the project will contribute to a) **Data science and Machine Learning**, with new data-driven hybrid BDN structure learning algorithms, b) **Information fusion**, with new methods that will allow us to construct machine learnt models that can take into consideration knowledge-based information about what must, can and cannot be discovered for decision support purposes, c) **Causal modelling**, with extended techniques to establish whether a relationship between two factors can be assumed to be causal or some other type of decision-modelling relationship, and d) **Decision Science**, with an improved way of constructing models for optimal decision making under uncertainty. Additionally, the project will contribute to the areas of sports, medicine, forensics, the UK housing market, and the UK financial market through the various case studies. As a result, it is realistic to expect that the project will have positive repercussions in a broad range of disciplines.

The project will benefit, as well as complement, another two projects. First, the EPSRC PAMBAYESIAN project, which runs from June 2017 to May 2020 (www.pambayesian.org), which is a collaboration between researchers from the School of Electronic Engineering and Computer Science (same School with the PI) and the Barts and The London School of Medicine and Dentistry. PAMBAYESIAN is supported by numerous digital health firms and hence, offers potential for further industrial engagement for this project. We will also benefit from access to further medical data already cleared as part of the PAMBAYESIAN project, which focuses on medical decision support based on home-based and wearable real-time monitoring systems for chronic conditions. Second, the MRC UKRI Innovation project on population genomics and health data science, that will support four postdoctoral fellows on Health Data Research for three years, starting mid-2018. This project is collectively led by the medical (SMD), computing (EECS), biology (SBCS), and materials (SEMS) schools at QMUL. As a Co-Investigator on this project under the theme "*Machine Learning and Artificial Intelligence*", I will assist with postdoctoral secondary/tertiary supervision within the School of EECS.

References

[1] Constantinou, A. C., & Fenton, N. (2018). Things to know about Bayesian Networks. *Significance*, 15(2): 19-23. [Open access DOI, PDF]

[2] Koller, D., & Friedman, N. (2009) Probabilistic graphical models: principles and techniques The MIT Press.

[3] Pearl, J. (1982). Reverend Bayes on Inference Engines: A Distributed Hierarchical Approach. *AAAI - 82 Proc.*, 133-136.

[4] Pearl, J. (1985). A model of activated memory for evidential reasoning. *In Proc. of the Cognitive Science Society*, 329-334.

[5] Koski, T. J. T., & Noble, J. M. (2012). A review of Bayesian Networks and Structure Learning. *Mathematica Applicanda*, Vol. 40(1), 53-103.

[6] Verma, T. S., & Pearl, J. (1990). Equivalence and synthesis of causal models. *In Proceedings of the Sixth Conference on Uncertainty in Artificial Intelligence*, 220–227.

[7] Glymour, C., & Cooper, G. F. (1999). Computation, Causation & Discovery. AAAI Press/MIT Press, Menlo Park, California, Cambridge Massachusetts, London, England.

[8] Spirtes, P., Glymour, C., & Scheines, R. (2000). Causation, Prediction, and Search : 2nd Edition. The MIT Press, Cambridge, Massachusetts, London, England.

[9] Korb, K., & Nicholson, A. (2011). *Bayesian Artificial Intelligence* (Second Edition). CRC Press, London, UK.

[10] Pearl, J. (2013). Causality: Models, Reasoning, and inference (2nd Edition). Cambridge University Press, University of California, Los Angeles.

[11] Cooper, G. F., and E. Herskovits. 1992. A Bayesian method for the induction of probabilistic networks from data. *Machine Learning*, 9: 309–347.

[12] Friedman, N., Nachman, I., & Pe'er, D. (1999). Learning Bayesian network structure from massive datasets: the 'sparse candidate' algorithm. *In Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence (UAI '99)*, pp. 196-205.

[13] Moore, A., & Wong, W. K. (2003). Optimal Reinsertion: A new search operator for accelerated and more accurate Bayesian network structure learning. *In Proceedings. of the 20th International Conference on Machine Learning (ICML - 2003)*, Washington DC.

[14] Chickering, D. M. (2002). Optimal structure identification with greedy search. *Journal of Machine Learning Research*, 507–554.

[15] Tsamardinos, I., Brown, L. E., & Aliferis, C. F. (2006). The Max-Min Hill-Climbing Bayesian Network Structure Learning Algorithm. *Machine Learning*, Vol. 65, pp. 31-78.

[16] Schmidt, M., Niculescu-Mizil, A., & Murphy, K. (2007). Learning graphical model structure using L1-regularization paths. *In Proceedings of the National Conference on Artificial Intelligence*, Vol 22(2), pp. 12-78.

[17] Constantinou, A., Fenton, N., & Neil, M. (2018). How do some Bayesian Network machine learned graphs compare to causal knowledge? *Under review*, 2018.

[18] Scutari, M., & Brogini, A. (2012). Bayesian Networks Structure Learning with Permutation Tests. *Communications in Statistics – Theory and Methods*, 41: 3233-3243.

[19] Freedman, D., & Humphreys, P. (1999). Are there algorithms that discover causal structure? *Synthese*, Vol 121, pp. 29-54.

[20] Zhang, J. (2008). On the completeness of orientation rules for causal discovery in the presence of latent confounders and selection bias. *Artificial Intelligence*, 172(16-17): 1873-1896.

[21] Dawid, A. P., Musio, M., & Stephen, F. (2015). From Statistical Evidence to Evidence of Causality. *Bayesian Analysis*, 11(3): 725-752.

[22] Triantafillou, S., & Tsamardinos, I. (2016). Score based vs constraint based causal learning in the presence of confounders. *In Proceedings of the Causation: Foundation to Application Workshop, 32nd Conference on Uncertainty in Artificial Intelligence (UAI 2016)*, New York City, USA, June 29, 2016.

[23] Constantinou, A. C., Fenton, N., Marsh, W., & Radlinski, L. (2016). From complex questionnaire and interviewing data to intelligent Bayesian Network models for medical decision support. *Artificial Intelligence in Medicine*, 67: 75-93. [draft, DOI]

[24] Hagmayer, Y., Sloman, S. A., Lagnado, D. A., & Waldmann, M. R. (2007). Causal reasoning through intervention. In Causal learning: Psychology, philosophy

and computation, ed. A. Gopnik & L. Schulz. Oxford University Press, Oxford, England.

[25] Yet, B., Neil, M., Fenton, N., Constantinou, A., & Dementiev, E. (2018). An Improved Method for Solving Hybrid Influence Diagrams. *International Journal of Approximate Reasoning*, 95: 93-112. [draft, DOI]

[26] Suryadim D. & Gmytrasiewicz, P. (1999). Learning models of other agents using influence diagrams. *User Modeling*, 407:223–232.

[27] Conroy, R., Zeng, Y., Cavazza, M., Chen, Y. (2015). Learning Behaviors in Agents Systems with Interactive Dynamic Influence Diagrams. *International Joint Conference on Artificial Intelligence*, 2015: 39-45.

[28] Constantinou, A. C., Fenton, N., & Neil, M. (2016). Integrating expert knowledge with data in Bayesian networks: Preserving data-driven expectations when the expert variables remain unobserved. *Expert Systems with Applications*, 56: 197-208. [draft, DOI]

[29] Fenton, N., Neil, M., Lagnado, D., Marsh, W., Yet, B., & Constantinou, A. (2016). How to model mutually exclusive events based on independent causal pathways in Bayesian network models. *Knowledge-Based Systems*, 113: 39-50. [Open access DOI, PDF]

[30] Constantinou, A. C., Fenton, N. E. & Neil, M. (2012). pi-football: A Bayesian network model for forecasting Association Football match outcomes. *Knowledge-Based Systems, 36: 322, 339*. [DOI, draft]

[31] Constantinou, A. C. & Fenton, N. E. (2012). Solving the Problem of Inadequate Scoring Rules for Assessing Probabilistic Football Forecast Models. *Journal of Quantitative Analysis in Sports: Vol. 8: Iss. 1, Article 1*. [DOI, draft]

[32] Constantinou, A. C., Fenton, N. E. & Neil, M. (2013). Profiting from an inefficient Association Football gambling market: Prediction, Risk and Uncertainty using Bayesian networks. *Knowledge-Based Systems, 50: 60-86*. [Open access DOI, PDF].

[33] Constantinou, A. C. & Fenton, N. E. (2013). Determining the level of ability of football teams by dynamic ratings based on the relative discrepancies in scores between adversaries. *Journal of Quantitative Analysis in Sports. Vol. 9, Iss. 1, 37–50*. [DOI, draft]

[34] Constantinou, A. C. & Fenton, N. E. (2013). Profiting from arbitrage and odds biases of the European football gambling market. *The Journal of Gambling Business and Economics, Vol. 7, 2: 41-70*. [PDF]

[35] Constantinou, A. C., Fenton, N. E., & Pollock, L. J. H. (2014). Bayesian networks for unbiased assessment of referee bias in Association Football. *Psychology of Sport and Exercise, Vol. 15, 5: 538-547*. [DOI, draft]

[36] Constantinou, A. C., Freestone, M., Marsh, W., & Coid, J. (2015). Causal inference for violence risk management and decision support in Forensic Psychiatry. *Decision Support Systems*, 80: 42-55. [DOI, draft]

[37] Constantinou, A. C., Freestone, M., Marsh W., Fenton, N., & Coid, J.W. (2015). Risk assessment and risk management of violent reoffending among prisoners. *Expert Systems with Applications*, 42(21): 7511-7529. [DOI, draft]

[38] Coid, J. W., Ullrich S., Kallis, C., Freestone, M., Gonzalez, R., Bui, L., et al. (2016). Improving risk management for violence in mental health services: a multimethods approach. *Programme Grants for Applied Research*, Vol. 4, Iss. 16. [DOI, PDF]

[39] Constantinou, A. C., Yet, B., Fenton, N., Neil, M., & Marsh, W. (2016). Value of Information Analysis for Interventional and Counterfactual Bayesian Networks in Forensic Medical Sciences. *Artificial Intelligence in Medicine*, 66: 41-52. [DOI, draft]

[40] Yet, B., Constantinou, A. C., Fenton, N., Neil, M., Luedeling, E., & Shepherd, K. (2016). A Bayesian Network Framework for Project Cost, Benefit and Risk Analysis with an Agricultural Development Case Study. *Expert Systems with Applications*, 60: 141-155. [draft, DOI]

[41] Constantinou, A. C., & Fenton, N. (2017). The future of the London Buy-To-Let property market: Simulation with Temporal Bayesian Networks. PLoS ONE, 12(6): e0179297. [Open access DOI, PDF]

[42] Constantinou, A. C., & Fenton, N. (2017). Towards Smart-Data: Improving predictive accuracy in long-term football team performance. *Knowledge-Based Systems*, 124: 93-104. [draft, DOI]

[43] Yet, B., Constantinou, A., Fenton, N., & Neil, M. (2018). Expected Value of Partial Perfect Information in Hybrid Models using Dynamic Discretization. *IEEE Access*, 6: 7802-7817. [draft, DOI]

[44] Constantinou, A. (2018). Dolores: A model that predicts football match outcomes from all over the world. *Machine Learning*, 1-27. [Free view, DOI, draft]

[45] Neil, M., Tailor, M., & Marquez, D. (2007). Inference in hybrid Bayesian networks using dynamic discretization. *Statistics and Computing*, 17.