 doi:10.20944/preprints202210.0084.v1

Essay

The Strengths and Weaknesses of Directed Acyclic Graphs (Dags) as Cognitive, Analytical and Educational Tools for Medical Statistics

George TH Ellison[1,2,3,*]

[1]Leeds Medical School, University of Leeds, UK;
[2]Leeds Institute for Data Analytics, University of Leeds, UK;
[3]Centre for Data Innovation, University of Central Lancashire, Preston, UK
* Correspondence: George TH Ellison PhD DSc FRSH, Faculty of Science & Technology, University of Central Lancashire, JB Firth Building, Leighton Street, Preston, Lancashire, PR1 2HE, UK; gthellison@uclan.ac.uk; ORCID: 0000-0001-8914-6812

Keywords: Directed Acyclic Graph; DAG; confounding; collider bias; causal inference; epistemology

1. Introduction

The origins of directed acyclic graphs (DAGs) date back to the emergence of 'graph theory' in the early 1700s (Biggs et al. 1986). DAGs are conceptual or literal, diagrammatic representations of causal paths between variables which are constructed – as their name suggests – on the basis of two over-riding principles: first, that all causal paths are 'directed' (i.e. for each pair of variables, only one can represent the cause, while the other must be its consequence); and second, that no direct cyclical paths, or indirect cyclical pathways (comprising sequences of consecutive paths) are allowed, such that no consequence can be considered its own direct or indirect cause (hence 'acyclic'; Law et al., 2012). As such DAGs reflect the knowledge, presumptions, assumptions and/or speculation of the analyst(s) concerned regarding the causal relationships between each of the variables included therein. Current convention dictates that variables are represented as nodes/vertices, and that any causal paths between variables are represented as directed arcs/edges/lines, often in the form of arrows (see Figure 1). Although each arc indicates the presence and direction of a known/presumed/assumed/speculative causal relationship between the two variables concerned, drawing an arc does not require the sign, magnitude, precision or shape of the relationship to be known or declared (Tennant et al., 2021). In this respect, DAGs provide a simple, uncomplicated, accessible and entirely non-parametric approach for postulating causal relationships amongst any variables of interest even when these are uncertain, unknown or entirely speculative (Ellison, 2020). Nonetheless, as a result of the parametric constraints imposed by the presence/absence of possible arcs within any given DAG, these also reflect and support a number of more sophisticated statistical applications which make it possible to use DAGs to inform the design of multivariable statistical models that reflect the causal structure(s) involved – albeit without the need to know or understand the mathematical technicalities on which these are based (Lewis and Kuerbis, 2016).

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 8 October 2022 doi:10.20944/preprints202210.0084.v1

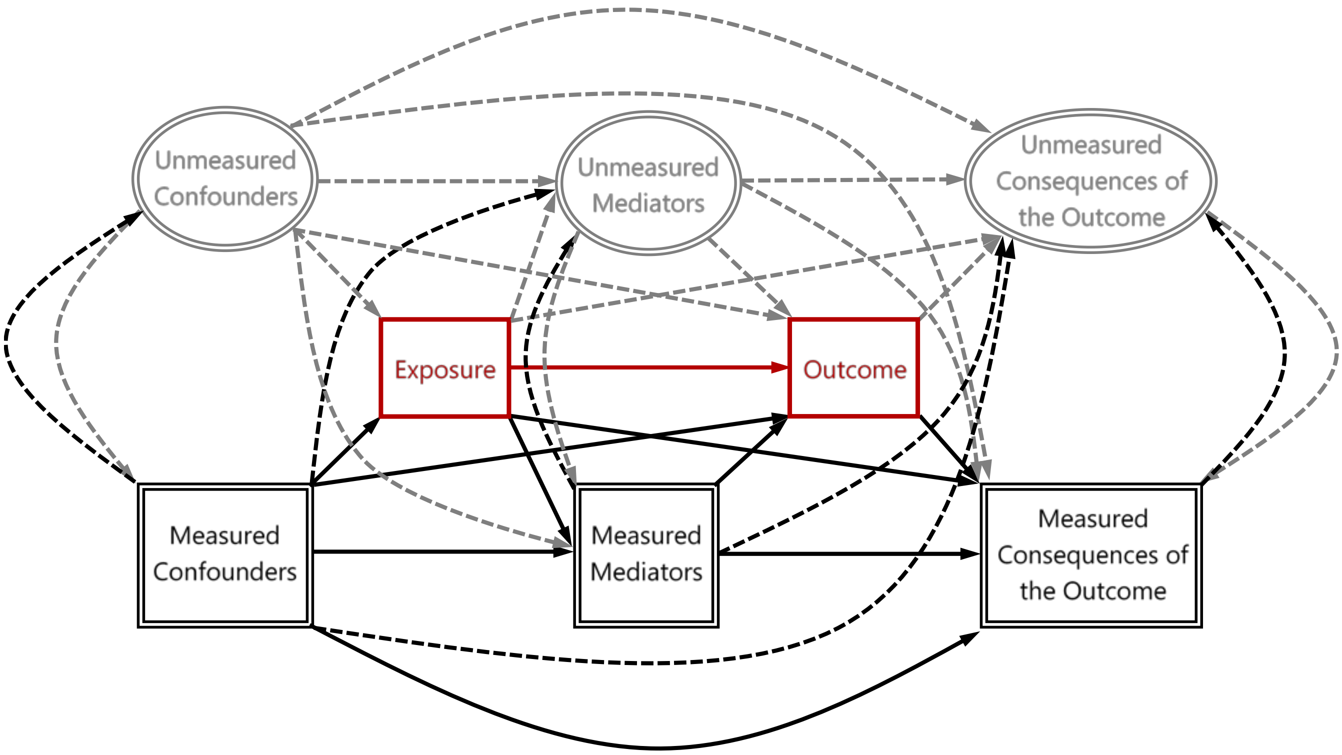

Figure 1. A comprehensive or 'universal' DAG (following Ellison, 2021: see Supplementary materials) summarising all of the conceivable variables (both known/measured and unmeasured/unknown/*latent*) that envelop the 'focal relationship' (i.e. the postulated causal relationship of interest) in what Pearl (1988) might have called a 'Markov blanket', comprising all possible variables affecting either the specified 'exposure' (or cause of interest) and the specified 'outcome' (or consequence/effect of interest). Except for the exposure and outcome variables, all covariates are represented as 'sets' (hence the double line surrounding these 'super-nodes'; as Tennant et al. 2020 described these) to indicate that more than one such covariate might co-occur during the period *before* the 'exposure', *after* the 'outcome' and the period *in between*. Because some of the covariates within each set of 'manifest' and 'latent' covariates might occur before those in another, a 'universal' DAG of this nature must include causal paths operating in both directions between each set of (measured and 'latent') 'confounders', 'mediators' and 'consequences of the outcome'.

These features make DAGs attractive cognitive, educational and analytical tools for strengthening the epistemological, theoretical and empirical basis of causal inference, and there has been a recent proliferation in the use of DAGs across a range of applied scientific disciplines (e.g. Knight and Winship, 2013), and an associated upsurge in analytical methods training (e.g. Elwert, 2011; Gilthorpe, 2017; Hernán 2018; Roy, 2021; Hünermund, 2021). This Chapter reflects on a decade of delivering medical statistics training to undergraduate medical students at the University of Leeds between 2012-2021 in which the third year research, evaluation and special studies module ('RESS3') has used DAGs to support the development of applied statistical skills relevant to the extended student-selected research and evaluation projects (ESREP) students undertake in their fourth and final years (Ellison, 2021; Ellison et al., 2014a,b).

Based on successive iterations of the structure and content of the RESS3 module, together with notes made during formal and informal planning and review meetings with module leads, lecturers, tutors and students, we draw on the claims and criticisms made of DAGs in the epidemiological literature to identify a number of explicit strengths (and associated, often implicit. weaknesses) that are central to their use in prediction and causal inference modelling. While using DAGs requires (and benefits from) a clear understanding of their non-parametric nature and parametric implications, the weaknesses of DAGs seem likely to reflect both: the challenges inherent in the modelling of data generating processes when these are imperfectly understood; and troublesome cognitive and heuristic tendencies common to all analytical tools – in which the tool facilitates the task in hand by reducing the necessity (and benefits of) exploring uncertainties and identifying assumptions. These, more epistemological considerations appear particularly challenging for medical undergraduates to grasp (Ellison, 2021), but also appear poorly understood by many established analysts and clinical epidemiologists (Ellison, 2020).

2. The Strengths and Benefits of DAGs

As we have seen, DAGs offer 'principled' representations of causal pathways that are known, believed or postulated to exist – or not to exist – amongst any conceivable variables that are available for consideration; each of which might play distinct roles of specific relevance to the causal processes concerned. These include any context-relevant variables for which measurements have been made/are available (the so-called 'known knowns'; Luft and Ingham, 1955); and those for which measurements have not been made or are unavailable (i.e. the 'known unknowns' and – when unacknowledged or *in*conceivable, the 'unknown unknowns'). As such, a DAG not only represents but also reflects (and can help reveal) the premise – and its associated assumptions (both explicit and implicit) – upon which the causal model has been constructed, *provided* this has been undertaken in line with two strict and inviolable principles (or rules). These are that: any constituent causal paths are directed (i.e. uni-directional); and none of the direct or indirect paths are circular (hence acyclic – a property that reflects what is known as the topological ordering or topological sorting of uni-directional paths, and which is a definitive feature of DAGs; Zhou and Müller, 2003; Kader, 2013).

These features imbue DAGs with a number of specific properties that make them potentially useful tools to assist in the conceptualisation and operationalisation of causal processes, particularly in non-experimental (i.e. observational) contexts where the causal processes involved are incompletely understood, somewhat uncertain, or completely unknown. Indeed, even where there is empirical evidence and substantial consensus on the form such processes take, both are vulnerable to what has been dubbed "unsubstantiated certainty" (NHSE, 1999). Nonetheless, analysts have long sought to use multivariable statistical modelling to extract meaningful information from such data, based on their best estimation of the underlying data generating mechanism(s) concerned; and it is in this regard that DAGs – which offer principled representations of these mechanisms – have their potential to add value to the conceptual and statistical modelling of observational data to support prediction and causal inference. Of these, causal inference has perhaps the most widespread application and potential for impact, since improving understanding of functional mechanisms is necessary for identifying, selecting and refining interventions that aim to prevent, pre-empt, attenuate or reverse undesirable processes that might otherwise cause harm (or, indeed, enhance those processes that are likely to do good). However, causal inference is also critical to the portability of prediction models beyond the contexts, periods and datasets on which these have been developed (Piccininni, 2020). To date, robust causal inference has relied upon experimentation (i.e. the deliberate manipulation of 'exposures' using 'interventions' to evaluate their subsequent effect on 'outcomes'). Yet experimental studies are often resource intensive, have limited utility for complex exposures/interventions, and face substantial ethical constraints (Frieden, 2017). For these reasons it is worth examining, in some detail, what the potential (and realised) strengths and benefits of DAGs might be within quantitative observational analyses, focussing in particular on the contributions these might make to causal inference, but also thereby to causality-informed prediction.

2.1. Transparency

A key strength of DAGs is their ability to reveal conceptual and analytical assumptions that might otherwise remain unspecified, unclear and/or uncertain to both:

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 8 October 2022 doi:10.20944/preprints202210.0084.v1

(i) the analysts concerned (who might otherwise have: not intended to make such assumptions; been unaware of these; and/or overlooked their implications); and
(ii) third parties/others, including: peers; reviewers and end-users (who are then able to examine, and evaluate the validity and plausibility of, any such assumptions).

While transparency is, in and of itself, a tangible benefit of using DAGs (not least in terms of enhancing the reproducibility and replicability of scientific research; CRRS, 2019), it has direct methodological utility to the design and conduct of both: primary studies seeking prediction or causal inference from analyses of observational data (see 4 (i)-(v), below); and secondary studies seeking to critically appraising and/or synthesise ostensibly comparable (yet potentially incomparable) findings generated by previous primary analyses (see 4 (vi), below).

2.2. Simplicity

The ability of DAGs to improve the transparency of analysts' assumptions benefits from substantial consensus regarding the principles that govern both: what DAGs can (and cannot) represent; and how these features are represented. As theoretical and non-parametric representations of causal processes, DAGs neither reflect nor dictate the parametric features of the individual causal paths involved (i.e. the sign, magnitude, shape or form of their parametric relationships; Tennant et al. 2020) – except where the omission of a causal path imposes/reflects a specific parametric value on the relationship between the variables concerned (i.e. that the associated path coefficient = O.Ō). And while DAGs need not necessarily be operationalised as graphical diagrams (e.g. Textor et al., 2016; Fiore and Devesas Campos, 2013; Geneletti et al., 2009), all DAGs contain *directed* causal paths (i.e. unidirectional paths from preceding causes to subsequent consequences/effects) and all are acyclic (i.e. no circular paths are permitted through which a cause might directly or indirectly become a consequence/effect of itself). Ostensibly, these two simple principles appear easy to understand and apply, making DAG construction a task that is accessible even to those with little technical expertise or experience (Ellison, 2021).

2.3. Flexibility

While the twin principles of 'unidirectionality' and 'acyclicity' impose strict constraints on the form(s) that DAGs can take, the rationale applied in deciding precisely where causal paths exist (and where they do not) can: accommodate a number of very different (and potentially contradictory) considerations; and be used in both theoretical and practical applications. In applications where DAGs are used to represent entirely hypothetical causal relationships amongst variables, the presence/absence of causal paths can be determined on an entirely speculative basis. However, in applications where DAGs are intended to represent real-world 'ground truths' (i.e. the underlying 'data generating mechanism' involved), robust theoretical (i.e. contextual and functional) knowledge is used to determine where causal paths are believed/known to exist (and where they do not). And even in applications where such knowledge is equivocal, uncertain, unknown or contested, temporality alone can be used to determine where causal paths *might* exist (and where they *cannot*), simply on the basis that causes must precede any subsequent consequences or effects. In this way, the decision as to where causal paths are situated in any given DAG can be informed by speculation, knowledge, temporality or any combination thereof, making them flexible tools for a range of applications which involve the modelling of subjective, hypothetical and/or a-theoretical (and ostensibly objective) conceptualisations of the underlying data generating mechanisms involved. As will become clear below, this flexibility lies at the heart of considerable uncertainty regarding the potential utility of DAGs, and may require substantial improvements in the detail analysts provide when using and reporting DAGs.

2.4. Methodological utility

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 8 October 2022 doi:10.20944/preprints202210.0084.v1

Improving the transparency of the (acknowledged, unacknowledged and unintentional) assumptions that analysts make can help to improve the methods they use at every stage in the research cycle – i.e. during problem/hypothesis identification, study design, sampling, measurement, analysis, interpretation and synthesis (i.e. in the critical appraisal and meta-analysis of primary studies):

2.4.1. Hypothesising:

Where hypotheses rely on the presence or absence of specific causal pathways, DAGs can help in exploring and evaluating the potential implications and consequences of the causal assumptions involved, and the plausibility of each of these hypotheses. In this way, and even in the absence of data or analysis thereon, DAGs are tools that can improve the key conceptual phases of the research process – phases that are likely to extend the application of DAGs from the modelling of observational data to the design of experimental studies (see for example: Tafti and Shmueli, 2020).

2.4.2. Sampling:

Where the research involves a choice of secondary datasets, or the prospective collection of data *de novo*, prior specification of a DAG can help identify the potential risk of collider bias (Griffith et al., 2020) that might be involved in selecting datasets, or generating new samples, likely to have been affected by, or vulnerable to, unrepresentative sampling procedures.

2.4.3. Data availability/collection:

Prior specification of a DAG can likewise help identify which covariates might need to be measured/available (as 'know*able* knowns') for inclusion within the covariate adjustment set(s) used in multivariable statistical models where the intended estimands are either: the total causal effect (Hudgens and Halloran, 2008); and/or the naïve direct causal effect between a specified 'exposure' and 'outcome' (Baron and Kenny, 1986; VanderWeele, 2014) – where adjustment for covariates acting as potential confounders or mediators are required, respectively (Groenwold et al., 2021).

2.4.4. Data analysis:

DAGs have particular utility in helping analysts identify measured (and unmeasured) covariates acting as colliders (including mediators and consequences of the outcome; Viswanathan et al., 2013) or potential confounders (Al-Jewair TS et al., 2017). The risk of bias due to collider adjustment or unadjusted/unobserved confounding can then be mitigated (through the exclusion of colliders, and the inclusion of measured confounders, in the covariate adjustment sets used in the study's multivariable statistical analyses; van der Zander et al., 2014). And wherever the speculative, theoretical and/or temporal rationale(s) applied when constructing DAGs involves the omission of causal paths where these might nonetheless exist (i.e. without introducing cyclical causal paths), a number of alternative yet equivalent adjustment sets may exist, each containing a different selection of covariates; Greenland et al. 1999). Under such circumstances, it is then possible to optimise the adjustment set selected on the basis that this contains covariates offering the most detailed and most accurate information available on potential confounding (i.e. from those sets of covariates that capture the most variance in confounding and have been measured with the greatest accurately; Law et al., 2012).

2.4.5. Interpretation:

DAGs also have substantial utility in interpreting findings generated by multivariable statistical analyses where either: adjustment for one or more colliders is necessary or unavoidable; or where one or more potential confounders have not or could not be measured (and therefore could not be included in the covariate adjustment sets). The former

 doi:10.20944/preprints202210.0084.v1

(collider bias) may occur when adjustment for mediators is used to generate naïve estimates of direct effects between the specified exposure and outcome (e.g. Dondo et al., 2020), or when consequences of the outcome are mistaken for competing exposures and included in the covariate adjustment set to improve the precision of the path coefficient estimated for a focal relationship (a practice that can even undermine the validity of experimental studies; Freedman, 2008). Endogenous selection bias may also make adjustment for colliders inevitable wherever it is likely that the sample of data available/generated for analysis is unrepresentative of the population to whom the analyses' findings are intended to apply (i.e. when the sampling procedures used generate a dataset that conditions on a collider; Elwert and Winship, 2014). Likewise, the latter (unadjusted/unobserved confounder bias) may occur whenever unmeasured/latent confounders are likely to be present and are therefore omitted from the covariate adjustment sets within the study's multivariable statistical analyses. Indeed, in most multivariable analyses of observational data, undeclared/unacknowledged naïvety extends beyond the inappropriate application of simplistic mediator-adjustment procedures to estimate direct causal effects, since it is implausible that: all sampling procedures generate absolutely representative samples that do not (unintentionally) condition on potential colliders; any covariate adjustment set can include all potential confounders (given these will include not only the so-called 'known knowns' and 'known unknowns' but also the 'unknown unknowns'); and that all confounders included in any given covariate adjustment set have been measured with absolute precision ('residual confounding' being that proportion of potential confounder bias remaining after the adjustment for confounders that can be attributed to random or systematic measurement error).

2.4.6. Critical appraisal and synthesis:

Although as yet unrealised (Mueller et al., 2018; Dekkers et al., 2019; Sarri et al., 2020), DAGs have substantial potential utility for strengthening the critical appraisal and synthesis of findings generated by previous studies if only by facilitating assessments of the risk of bias therein. Indeed, even where the original studies concerned have not used DAGs to inform the methodological approach(es) adopted (or have not described/reported the DAG[s] used in any/sufficient detail; Tennant et al., 2021), critical appraisal can be applied to discrete focal relationships within carefully defined contexts based on speculation, theoretical knowledge and/or temporal/probabilistic considerations concerning the underlying data generating mechanism(s) from contexts that are likely to be similar. In such instances, DAGs can be developed *de novo* simply on the basis of the covariates available to any of the primary studies concerned, augmented by consideration of any potential unmeasured covariates (particularly those positioned before the specified exposure which might thereby act as sources of unadjusted/unobserved confounder bias in the coefficient estimates reported for the focal relationship involved; Alfawaz, 2017). Such DAGs can then be used across multiple studies to assess the risk of bias in their multivariable statistical analyses that might arise from: endogenous selection bias or unrepresentative sampling (collider bias); under-adjustment for potential confounders (confounder bias), and particularly those measured by, or available to, some but not all of the studies examined; or over-adjustment (for consequences of the outcome or mediators; e.g. Dekkers et al., 2019; Sarri et al., 2020).

2.5. Consistency evaluation

In those studies where DAGs are intended to: inform multivariable statistical analyses capable of supporting causal inference; and/or accurately reflect the underlying data generating mechanism involved, it may also be possible to use these as a basis for evaluating DAG-analysis, and DAG-dataset, consistency. DAG-analysis consistency can be evaluated for any DAGs, regardless of their structure or the rationale(s) involved in constructing these. Such evaluations involve examining whether the covariate adjustment

 doi:10.20944/preprints202210.0084.v1

set(s) used in the associated multivariable statistical analyses only include an appropriate selection of: measured/available and latent/unmeasured confounders sufficient to optimally mitigate the risk of confounder bias; any relevant mediators (where the estimand(s) concerned include naïve estimation of any direct causal effects); and any covariates identified as competing exposures (where the precision of the causal estimates generated is considered sufficiently important to warrant the potential risks involved). In contrast, DAG-dataset evaluations are only possible for DAGs in which the speculative, theoretical and/or temporal rationale(s) applied when constructing these support the omission of one or more causal paths where these might nonetheless exist (i.e. without breaching the principle of acyclicity). In these instances, the non-parametric features of the DAGs concerned impose testable parametric constraints on the data these DAGs are intended to represent (Textor et al. 2016). It is therefore possible to establish whether such constraints actually apply within these datasets *and* to identify a comprehensive set of any (and all) alternative DAGs which are consistent with the datasets concerned (regardless of whether any of these DAGs reflect [m]any of the features of the DAG proposed by the analysts concerned).

Although these assessments do not represent formal 'tests' as to whether any given DAG is consistent with any given dataset, they can help to evaluate: whether DAGs constructed on theoretical, speculative and/or temporal grounds might actually reflect the real-world data (assuming, of course, that the analysts' DAGs were actually intended to accurately reflect the data generating mechanism concerned); and the range of DAGs that might be (parametrically) plausible for the dataset(s) in hand (thereby prompting subsequent consideration of the basis on which one or more of these DAGs might actually/better/best reflect the underlying data generating mechanism involved).

2.6. Epistemological

For those studies engaged in generating causal hypotheses, analyses and inferences from observational data, DAGs have benefits that extend beyond their impact on the coherence and consistency of multivariable statistical modelling. Indeed, the cognitive and conceptual impact of DAGs on collective understanding of data generating mechanisms (and on how these might be modelled using statistical techniques to generate insight and foresight) may prove to be just as important (if not more so) for identifying and elucidating hitherto poorly understood, under-acknowledged or completely hidden/unknown analytical opportunities and sources of bias. These benefits are evident in: the role that the concept of 'colliders' have played in understanding the impact of unrepresentative sampling and selection procedures on biased causal inference (hence the name 'collider bias'); and the recent identification of 'M-bias' and 'Butterfly-bias' – two forms of bias whose nomenclature stems from the shapes these take when elucidated within topologically arrayed DAGs (Ding and Miratrix, 2015). Ongoing applications of DAGs within causally-infused prediction models (Piccininni et al., 2020) are likewise capitalising on the cognitive and conceptual value that these bring to understanding the data generating mechanisms on which interpolative and extrapolative predictive modelling rely.

3. The Limitations and weaknesses of DAGs

Despite the many potential (and realised) benefits of using DAGs – to enhance the transparency, reproducibility and analytical integrity of efforts to generate causality-enhanced prediction, and robust causal inference, from observational data – variation in their uptake and application (Tennant et al., 2021) suggest there are multiple limitations and weaknesses (in both the conceptualisation and operationalisation of DAGs) that warrant further consideration. This is important not least because: such limitations/weaknesses might undermine the willingness of analysts to consider using DAGs; and the misapplication of DAGs reduces not only their evident utility but also the perceived improvements in analytical practice they might otherwise enable. Indeed, at present, the use of

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 8 October 2022 doi:10.20944/preprints202210.0084.v1

DAGs in analytical studies offers little in the way of reassurance that the analyses concerned have been any less inappropriately designed (or that they have better conducted or more robustly interpreted) than traditional/established practices (in which numerous biases and errors remain commonplace, routine and widely accepted/overlooked – as described by: Pocock et al. 2004; von Elm and Egger 2004; Blair et al. 2007; Detweiler et al. 2016; Ellison, 2021; Griffith et al., 2020). As such, were analysts, reviewers or end-users to naïvely accept/interpret references to, or the use/inclusion of, DAGs as evidence of sophisticated, advanced or robust analytical practice, these might simply serve to further distract attention away from improvements in technique that are long overdue.

These concerns affect any advances in technique that require specialist understanding, knowledge and/or skill; and it seems likely that much of the poor quality analytical practice that pervades contemporary studies involving prediction or causal inference from observational data might simply reflect the paucity of expertise (amongst analysts and reviewers) required to recognise and prevent this. Since the use of causal path diagrams (and particularly DAGs) constitutes a substantial departure from established analytical practice, the potential for misuse and mis-application is likely to introduce limitations and weaknesses across all of the strengths and benefits summarised above (see Section 2). It is therefore worth considering each of these (putative/potential benefits and strengths) in turn to identify those where: variation (in understanding and/or practice) might benefit from greater consensus or standardisation; and this simply requires greater clarity regarding:

(i) the inherent flexibility and adaptability of DAGs;
(ii) their potential utility across a range of applications (i.e. not simply to strengthen the modelling of observational data to support prediction and causal inference); and
(iii) substantive limitations (both conceptual and operational) that might either require and benefit from further development, or might pose enduring constraints on the contributions DAGs can make to future analytical practice.

3.1. Transparency

Revealing analytical assumptions that might otherwise remain unspecified, unclear and/or uncertain is undoubtedly a key potential benefit of using DAGs to support the modelling of observational data. Nonetheless, their utility in this regard is constrained not only by the understanding of the analysts concerned (and of peers, reviewers and end-users), but also by: the size and complexity of the DAGs themselves (which can be challenging to represent in diagrammatic form); and the accessibility (readability and interrogability) of the form(s) in which they are presented. Physical constraints will inevitably place limits on the number of variables and causal paths that can be presented in any finite space, while there will be similar constraints on the ability of the human eye to interpret cluttered and fine-grained images of complex diagrams. Indeed, in Tennant et al.'s (2021) review of 144 published DAGs – all of which had been published/presented as fixed/static two-dimensional figures – the co-authors/reviewers involved made more errors recording the numbers of variables and paths in those DAGs containing larger numbers of variables and/or larger numbers of paths; and such errors occurred in well over a third (39%) of the DAGs examined. However, data extraction errors were lower amongst DAGs drawn using specialist DAG-specification software (www.daggity.net; Textor et al., 2011; Textor, 2020), and amongst those that were topologically arrayed (though only when their causal paths had been aligned vertically [i.e. top ↔ bottom] or horizontally [i.e. left ↔ right], and not those arranged diagonally). It is tempting to conclude from these findings that the benefits of DAGs in supporting greater transparency will be limited to leaner, simpler DAGs, and those amenable to specification using specialist software. Yet Tennant et al.'s (2021) review did not include DAGs presented in alternative formats (such as the list-wise representation developed by Stacey et al., 2019 (see Figure S1 therein) or those summarised using specialist technical notation which has the added benefit of being machine-

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 8 October 2022 doi:10.20944/preprints202210.0084.v1

readable (hence its utility within the *R* package 'daggity'; Textor et al. 2016). While it seems unlikely that the former can address the inherent space constraints and risk of cluttered and indecipherable images (a risk that undermines the utility of large and complex DAGs presented as static, two-dimensional images); the latter clearly offers greater scope for readability/interrogability, albeit at the cost of simplicity and accessibility to non-specialists and those unfamiliar with any software required.

3.2. Simplicity

The apparent simplicity of generating diagrammatic representations of causal path diagrams based on two ostensibly accessible principles (that the paths therein are directed and acyclic) masks the more complex conceptual challenges this entails (Ellison, 2021). Indeed, regardless of the application involved (and notwithstanding their apparent flexibility; see 2.3, above and 3.3, below), the use of DAGs to support the modelling of observational data relies upon an understanding of what these diagrams aim to represent – namely the underlying 'data generating mechanism(s)' responsible for the relationships observed between all conceived and *conceivable* variables. These variables include not only those for which measurements are/are not available (the *known* knowns and the *unknown* knowns) but also those for which measurements are not available because the analyst is (as yet) unaware of their (possible) existence (the so-called *unknown unknowns*). Since all three sets of variables are required to comprehensively characterise the underlying data generating mechanism(s) involved, analysts need to be able to apply this as perhaps a third (hitherto undeclared) principle, namely that: DAGs which seek to represent real-world causal processes need to include **all** potential variables necessary to inform prediction and/or generate robust causal inference (with particular emphasis on '**all**'). In applying this principle, analysts require not only: a substantial degree of humility (given our limited and incomplete understanding of the functional mechanisms involved in most real-world systems – except, perhaps, those where the systems concerned are discrete artefacts of deliberate human designs, or are based on established physical laws); but also an awareness of the critical role that context can play, and how contexts themselves vary over time and space.

These considerations aside, wherever 'robust causal inference' involves consideration of a finite number of causal paths between variables (i.e. one or more 'focal relationships' between specified exposures and specified outcomes), then it is usually unnecessary to generate extensive or comprehensive DAGs (such as those detailing all possible interactions amongst variables acting as potential confounders) since all that may be necessary to mitigate the most important biases (the 'tigers' as opposed to the 'mice', as the statistician George Box described them; Box, 1976) will be a focus on the 'Markov blanket' of key influences that precede and surround the focal relationships concerned (Pearl, 1988). Arguably, such considerations detract from the perceived simplicity of DAGs, or at least require far greater thoughtfulness and attention when constructing, reporting and reviewing these – an issue that is also necessary to ensure that the flexibility of DAGs (i.e. their utility for a range of applications relevant to the modelling of observational data) does not undermine their transparency (see 3.3, below).

3.3. Flexibility

The implicit conceptual considerations that belie the apparent simplicity (and much vaunted transparency) of DAGs extend to their flexibility; and, therein, to those applications whose specific intention is to strengthen the modelling of real-world observational data. This is because:

(i) DAGs can be developed on the basis of speculation, theoretical knowledge, temporal/probabilistic considerations, or a combination of all three; and
(ii) the rationale involved in DAG development sets constraints on their likely (internal) validity and (external) applicability.

Thus, where DAGs are constructed on the basis of entirely speculative causal relationships between each of the constituent variables, these can still be conceptually valid even when they bear little relation to the observational data available. Likewise, where DAGs are constructed on the basis of theoretical (be that mechanistic and/or empirical) knowledge of the causal relationships believed to be present (or absent) amongst each of the variables involved, then these DAGs will offer valid representations of the theoretical causal structure concerned even when these are at odds with the real-world data the DAG is intended to represent. Indeed, those analysts who apply temporal/probabilistic considerations (in which the time-point at which each of the constituent variables occurs or crystallises determines their potential causal relationships with every other variable) – to generate ostensibly a-theoretical and objective DAGs that reflect the probabilistic causal processes involved – may still find that their DAGs deviate from the data for which they were generated. Yet, in each of these three examples, assessing whether the analysts concerned have generated DAGs that fit their intended (speculative, theoretical or temporal/probabilistic) application(s) requires that these intentions are clearly reported/declared, since any given DAG – regardless of how this is represented (whether as a static, two-dimensional diagram or in machine-readable notation) – does not, in and of itself, reveal the rationale involved when deciding what variables to include (see 3.2, above) and which causal paths do/do not exist between and amongst these.

For this reason, knowing the rationale involved is critical when interpreting the likely value, insight and inference that might be drawn from the modelling of observational data based thereon. Indeed, encouraging analysts to declare the intended application(s) of their DAG(s), and the associated rationale involved, when presenting these for dissemination/publication might represent another, fourth, principle: Analysts using DAGs which seek to represent real-world causal processes need to report the applications for which these were designed **and** the rationale(s) involved when constructing these. Like each of the three earlier principles, adopting such an approach when reporting the development and application of DAGs would: not only help others (peers, reviewers and end-users) better understand and critically appraise the decisions made, and the likely internal and external validity of any modelling based thereon; but might also prompt analysts to more carefully reflect on the (explicit and implicit) rationale(s) applied and any assumptions and potential inconsistencies involved therein. The latter may be critical given that speculation, theoretical understanding (whether formal, informal or empirically derived), and temporal/probabilistic considerations *all* rely on cognitive reasoning (involving conscious and unconscious heuristics that are prone to error and bias; Hume, 1738). All of the decisions made regarding the presence/absence of causal paths informed by any of these approaches will therefore be vulnerable to confusion and conflation, as well as to error and bias. These include those decisions based on temporality, wherever there is uncertainty as to the precise point in time that a variable occurs or its value crystallises relative to the exposure and outcome (see Figure 1); and particularly when the variables involved are time-variant features of distinct entities/processes rather than discrete (time-invariant) phenomena that can be conceptualised (or operationalised) as 'time-stamped' events.

3.4. Methodological utility

To a large extent, the methodological utility of DAGs relies upon the competency, thoughtfulness, diligence and open-ness of the analysts concerned, which together determine their ability to generate principled representations of data generating mechanisms offering valid summaries of the speculative, theoretical or temporal/probabilistic rationale(s) involved, and the subsequent applications for which these were intended. Beyond the constraints these considerations also place on the transparency, simplicity and flexibility of DAGs (and how improved reporting might be required to secure and enhance each of the related potential benefits; see 3.1-3.3, above), the methodological utility of DAGs extends beyond their *internal* validity (i.e. whether, as specified, these accurately reflect the assumptions involved, rationale[s] used and application[s] intended) to their

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 8 October 2022 doi:10.20944/preprints202210.0084.v1

external validity. Following George Box's (1976) adage that "*all models are wrong, but some are useful*", the potential methodological limitations of DAGs principally stem from the challenges involved in generating and applying DAGs as (imperfect/'wrong' but nonetheless 'useful') representations of one (or more) unknown and uncertain, underlying data generating mechanism(s). Indeed, assessing whether any such models are nonetheless 'useful' should clearly involve evaluating whether these are capable of supporting improved estimates through prediction (i.e. interpolation and extrapolation) and/or causal inference. Put simply, incorrectly specified DAGs that do not *closely* (or, at the very least, usefully) represent the underlying data generating mechanism(s) involved are unlikely to provide a useful basis on which multivariable statistical models can be designed that generate causality-enhanced predictions or valid causal inference.

However, unlike the considerations brought to bear on transparency, simplicity and flexibility (see 3.1-3.3, above), methodological concerns are primarily relevant only to those applications where DAGs are intended to strengthen statistical estimation of variables or focal relationships (through analyses of data to generate predictions or causal inferences); rather than more conceptual, theoretical and potentially nebulous considerations undertaken in the absence of data (such as those necessary to speculate the existence of M-bias and 'butterfly-bias'; Ding and Miratrix, 2015). As such, the methodological strengths of using DAGs relies upon the careful application of the most realistic (and pragmatic) assumptions and rationale available to: minimise the extent to which any modelling based thereon might be wrong; and strengthen the extent to which the model's findings might be useful (Box, 1976). In most contexts, speculation, theoretical understanding (based on empirical and experiential knowledge), and temporal/probabilistic considerations may all be appropriate and necessary in this regard – not least because (and as discussed earlier, and despite the apparent benefits of a temporal/probabilistic rationale), operationalising variables as phenomena/events requires both theoretical understanding and speculation to specify precisely where (with respect to all other variables) each individual variable is believed to have occurred/crystallised. Importantly, incorporating a temporal/probabilistic rationale when generating DAGs (either exclusively or in combination with theoretical and speculative considerations) imposes two substantive consequences on the methodological utility of all such DAGs:

(i) first, this tends to require that all DAGs intended to represent uncertain, real-world data generating mechanisms are 'saturated' (i.e. contain all possible arcs, such that each variable is assumed to cause all subsequent variables, while ensuring that all paths are directed and directly/indirectly acyclic; Foraita et al., 2014); and
(ii) second, this eliminates the possibility that any variables might operate independently of (all) previous variables – except those variables situated at the very beginning of any causal process (any preceding causes for which may therefore not have been measured, considered, or even imagined/conceived).

Although neither of these consequences (and the assumptions they entail) may accurately reflect the data generating mechanisms operating within (m)any given real-world contexts, some (though not all) of their impacts on multivariable statistical models designed to support prediction or causal inference may prove to be trivial, while some (if only a few) may require careful consideration:

3.4.1. Prediction modelling:

In prediction modelling of observational data, the value of DAGs lies primarily in identifying covariates that might potentially contribute or reflect (either directly or indirectly) information relevant to those causal pathways associated with the target variable for which the model aims to generate accurate predictions. And while covariates with strong (direct or indirect) causal links to the target variable often provide substantial information of value to the prediction model (information that warrants their consideration

 doi:10.20944/preprints202210.0084.v1

as 'candidate predictors'), they are routinely excluded from models where their *net* contribution comes at the cost of simplicity, accuracy or precision (Rothman and Lash, 2020). However, where 'transportability' is considered a critical feature of the prediction models concerned – such that optimising accuracy over time, or across a range of contexts, is considered more important than optimising accuracy at any single point in time or within any discrete context – the utility of DAGs lies in prioritising information from candidate predictors whose contribution to the model stems, in some/large part, from the direct and indirect causal role(s) they play within the underlying data generating mechanism (Piccininni et al., 2020). Under these circumstances, using a temporal rationale that assumes that preceding variables are probabilistic causes of all subsequent variables risks prioritising consideration of more covariates as *candidate* predictors in DAG-informed (and transportability-focused) prediction modelling than is likely to have been the case had the DAG concerned been formulated using a purely theoretical and/or speculative rationale. However, since prediction modelling routinely examines multiple combinations of alternative candidate predictors (or groups thereof), preferencing covariates for consideration therein on the mistaken assumption of (indirect/direct, probabilistic) causal involvement seems unlikely to affect the performance of the optimum model(s) available or selected – though it might perhaps complicate/extend the process required to identify these; and this issue warrants further investigation, not least within machine-generated prediction techniques.

3.4.2. Causal inference modelling:

In contrast, the principal benefit of using DAGs to generate causal inference from observational data stems from the way these facilitate the identification of covariates acting as potential confounders (see 2.4 (iv), above). Facilitating the identification of such variables ensures that all can then be considered for inclusion in the covariate adjustment sets of multivariable statistical models to mitigate bias in the estimation of total causal effects between any given exposure and any given outcome. In this regards, an *a priori* assumption that all preceding variables should be viewed as possible (if not likely probabilistic) causes of all subsequent variables – at least in the absence of unequivocal evidence to the contrary – might seem unlikely to compromise the ability of DAGs to identify potential confounders and to assist in mitigating the impact of confounder bias through their inclusion in covariate adjustment sets of any multivariable statistical models. This is because all variables interpreted as having occurred/crystallised before the specified exposure will commonly be viewed as likely to be probabilistic causes of *both* the exposure and *any* subsequent outcome.

Likewise, adjustment for covariates that do not actually cause the exposure, or the outcome (or either) might seem unlikely to affect the strength or direction of the path coefficient between exposure and outcome (the 'focal relationship'). This is often considered to be the case for covariates acting as 'competing exposures' (see Figure 1 in Tennant et al., 2021) which are defined as having a causal effect on the outcome but no direct/indirect causal relationship with the exposure – features which have led some analysts to include such covariates within the adjustment sets of multivariable statistical models on the basis that they should not affect the strength or direction of the path estimate between exposure and outcome, but can help to improve its precision. Setting aside the inappropriate conflation of confidence intervals and precision with hypothesis testing that such practices reveal (Gardner and Altman, 1986), these also risk overlooking the possibility that competing exposures are actually: caused by the exposure; or caused by one or more mediators; or, on occasion, may actually represent 'consequences of the outcome' – such that any improvement in precision comes at an increased (and some might argue, unnecessary) risk of bias. In this instance, assuming that all variables should be considered probabilistic causes of all subsequent variables would help to militate against the risk of bias associated with the inclusion of competing exposures within the covariate adjustment sets of multivariable statistical models. This is because no such variables would exist within

DAGs drawn using a primarily or exclusively temporal/probabilistic rationale (except in the highly unlikely/improbable scenario where there is unequivocal evidence that the competing exposures concerned had *no* causal relationship with *any* [measured, unmeasured or unknown] preceding variables).

Nonetheless, beyond the benefit of discouraging (unnecessary or risky) adjustment for competing exposures, assuming that all possible (directed and acyclic) causal paths exist between preceding and subsequent variables might nonetheless risk introducing additional/alternative sources of bias. In particular, in the case of "mediator-outcome confounders" (MOCs; Tennant et al., 2021; see Figure 1 therein) – which are covariates that have no direct causal relationship with the specified exposure, but have an indirect causal relationship with the outcome through a mediator that is itself a consequence of the exposure – adjustment for such covariates would introduce the risk of biases associated with mediator adjustment (i.e. the reversal paradox and collider bias; Tu et al., 2008; Schisterman et al., 2009; Richiardi et al., 2013). Whether such risks are common or have substantive impact on the estimated path coefficient between exposure and outcome will depend not only upon the strength and direction of each of the constituent (causal) paths involved, but also on whether the apparent MOC actually occurred/crystallised prior to the exposure (in which case it would represent a misclassified confounder) or after the exposure (in which case it would represent a misclassified mediator). As such, the issue that will inevitably prove most critical to the benefits and risks of assuming that all variables (are assumed to) act as probabilistic causes of all subsequent variables will be accurately identifying when each variable occurred/crystallised relative to all of the other variables considered. In most (but not all) current applications of DAGs within causal inference modelling, this issue relies less on temporality/probabilistic considerations than on speculation and theoretical knowledge. Exploring how the misspecification of 'when' and 'where' each variable sits within a DAG's temporally dependent pathways affects the risk of bias (from under-, over- and inappropriate covariate adjustment) remains an area worthy of further exploration.

3.5. Consistency evaluation

As mentioned previously (see 3.4, above), a further consequence of the assumption that variables be considered probabilistic causes of all subsequent variables is that the saturated DAGs these create are not amenable to DAG-data consistency assessment using the *R* package 'dagitty' (Textor et al., 2016). For these reasons, the rationale(s) used when generating DAGs (be this on the basis of speculation, theoretical knowledge or temporal/probabilistic considerations) dictate both DAG-theory and DAG-data consistency assessment. Greater clarity and precision regarding the intended application *for* which, and/or the rationale(s) *on* which, analysts have generated their DAG(s) will ensure these can support DAG-theory consistency assessment. However, DAG-data assessment will not be available for any DAGs in which temporal/probabilistic considerations constitute the only (or pre-eminent) rationale involved and thereby impose saturation thereon. DAG-data assessment of these DAGs will only be possible where analysts are prepared to speculate (i.e. consider the possibility) that one or more of the causal paths are missing, or are confident that definitive theoretical knowledge exists to support such an assumption. Whether DAG-data consistency evaluation might retain the capacity to address any uncertainty regarding precisely when (relative to other variables) each of the included covariates occurred/crystallised (relative to one another and to the specified exposure and outcome) remains to be seen.

3.6. Epistemological

Finally, while it is true that using DAGs has helped analysts to identify potential sources of bias (particularly those relevant to colliders, as mentioned under 2.4, above) that had proved challenging to conceptualise or operationalise, it is also possible that

 doi:10.20944/preprints202210.0084.v1

DAGs might lead to levels of abstraction that – though theoretically and methodologically insightful – bear little relation to the forms that observational datasets might plausibly, or most commonly take. In this regard, it seems likely that many of the possible roles that variables might play within DAGs – such as competing exposures and mediator-outcome confounders (MOCs; Tennant et al. 2021) – might be illusory or very unlikely to exist in real-world contexts. Certainly, from a temporal/probabilistic perspective neither of these roles could exist within a saturated DAG; and provided this perspective is not itself an abstraction of reality (which would be ironic given it makes assumptions that are generally intended to be plausible, objective and likely), then such roles constitute an unnecessary, and perhaps unhelpful, distraction to models that aim to reflect the underlying, real-world data generating mechanism(s) involved. Further research may be warranted to map all of the potential *additional* roles that covariates might play within an otherwise simplistic and unsaturated DAG (i.e. one that simply includes: an exposure; an outcome; and one or more confounders, mediators and consequences of the outcome), and evaluate both: the potential risk of bias each of these additional roles might pose to robust modelling of causal inference; and their likely occurrence in real-world contexts (based on understanding informed by theoretical knowledge, speculation and/or temporality/probabilistic considerations).

4. Conclusion

DAGs – like all analytical tools – benefit from doubt, circumspection and careful deliberation to ensure their thoughtful specification can harness the opportunities they provide for 'discovery' (alongside the self-evident contribution their careful application *should* make to 'translation' through greater consistency, competency and transparency). Although many analysts, educationalists and students may be drawn to DAGs as accessible tools for conceptualising and operationalising data generating mechanisms, the two ostensibly simple *principles* involved (of unidirectionality and acyclicity) still require thoughtful consideration and careful application – not least when DAGs are used for very different purposes, and are specified on the basis of very different rationales (i.e. on the basis of speculation, theoretical knowledge, and/or temporal/probabilistic considerations).

For this reason we conclude that to ensure the use of DAGs optimises the benefits in competency, consistency and transparency they should provide, all analysts need to provide greater detail of the rationale(s) used when specifying their DAGs and the application(s) for which their DAGs have been designed. Where these applications involve the need to represent real-world (rather than entirely speculative) causal processes so as to mitigate the risk of bias in the modelling and estimation of causal relationships (or when optimising the portability of causality-enhanced prediction models), we recommend that – regardless of the role that speculative, theoretical and/or probabilistic/temporal considerations might play therein – such DAGs need to include all possible/conceivable variables necessary to generate robust causal inference and/or inform prediction – where 'conceivable' will include both 'known/know-able knowns' and 'unknown/unknow-able (un)knowns'. Specifying *all* such variables will not only help analysts and educators to acknowledge the inherent and irreducible uncertainties that bedevil our understanding of most real-world data generating mechanism(s); but will also prompt both analysts and students to more fully explore and address the potential biases that these unacknowledged uncertainties might otherwise conceal.

As Mark Twain is often quoted as saying (though he may have been unfairly credited with ideas promulgated by many others; including Josh Billings/Henry Wheeler Shaw in 1874; Artemus Ward/Charles Farrar Browne before his death in 1867; Kin Hubbard/ Frank McKinney Hubbard in 1958; and Will Rogers in 1978; O'Toole, 2018):

"It ain't what you don't know that gets you into trouble. It's what you know for sure that just ain't so."

5. References

Alfawaz RA. Exploring the Relationship between Metabolic Syndrome and Sleep amongst Adults in the UK. Doctoral dissertation. Leeds; University of Leeds: 2017.

Al-Jewair TS, Pandis N, Tu YK. Directed acyclic graphs: A tool to identify confounders in orthodontic research, Part II. American Journal of Orthodontics and Dentofacial Orthopedics 2017; 151: 619-21.

Baron RM, Kenny DA. The moderator–mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. Journal of Personality and Social Psychology 1986; 51: 1173-1182.

Biggs NL, Lloyd EK, Wilson RJ. Graph Theory, 1736-1936. Oxford, Oxford University Press: 1976; 239pp.

Blair A, Stewart P, Lubin JH, Forastiere F. Methodological issues regarding confounding and exposure misclassification in epidemiological studies of occupational exposures. American Journal of Industrial Medicine 2007; 50: 199-207.

Box GEP. Science and statistics. Journal of the American Statistical Association 1976; 71: 791-9.

Committee on Reproducibility and Replicability in Science (CRRS). Reproducibility and Replicability in Science. Washington DC; National Academies Press; 2019.

Dekkers OM, Vandenbroucke JP, Cevallos M, Renehan AG, Altman DG, Egger M. COSMOS-E: guidance on conducting systematic reviews and meta-analyses of observational studies of etiology. PLoS Medicine 2019; 16: p.e1002742.

Detweiler BN, Kollmorgen LE, Umberham BA, Hedin RJ, Vassar BM. Risk of bias and methodological appraisal practices in systematic reviews published in anaesthetic journals: a meta-epidemiological study. Anaesthesia 2016; 71: 955-68.

Ding P, Miratrix LW. To adjust or not to adjust? Sensitivity analysis of M-bias and Butterfly-bias. Journal of Causal Inference 2015; 3: 41-57.

Dondo TB, Hall M, Munyombwe T, Wilkinson C, Yadegarfar ME, Timmis A, Batin PD, Jernberg T, Fox KA, Gale CP. A nationwide causal mediation analysis of survival following ST-elevation myocardial infarction. Heart 2020; 106: 765-71.

Ellison GTH. COVID-19 and the epistemology of epidemiological models at the dawn of AI. Annals of Human Biology 2020; 47: 506-13.

Ellison GTH. Might temporal logic improve the specification of directed acyclic graphs (DAGs)? Journal of Statistics and Data Science Education 2021; 29: 202-213.

Ellison, G.T.H., Harrison, W., Law, G.R., Textor, J. (2014a), "Graphical, cross-tabulatory and relational: which is best for drawing DAGs?" Proceedings of VI European Congress of Methodology, Jul 23, 6, 11.

Ellison, G.T.H., Harrison, W., Law, G.R., Textor, J. (2014b), "Teaching DAGs to support MBChB students design, analyze and critically appraise clinical research," Proceedings of the 35th Annual Burwalls Meeting for Teachers of Medical Statistics, Jul 3, 35, 3.

Elwert F. Causal Inference with DAGs. Population Health Sciences, University of Wisconsin-Madison, WI; 2011. Available at: https://www.ssc.wisc.edu/~felwert/causality/?page_id=140 [accessed 30 November 2021].

Elwert F, Winship C. Endogenous selection bias: The problem of conditioning on a collider variable. Annual Review of Sociology 2014; 40: 31-53.

Fiore M, Devesas Campos M. The algebra of directed acyclic graphs. In: Coecke B, Ong L, Panangaden P. (Eds) Computation, Logic, Games, and Quantum Foundations. The Many Facets of Samson Abramsky. Lecture Notes in Computer Science 2013; 7860: 37-51.

Foraita R, Spallek J, Zeeb H. Directed Acyclic Graphs. In: Ahrens W, Pigeot I (Eds). Handbook of Epidemiology. New York: Springer; 2014.

Freedman DA. On regression adjustments to experimental data. Advances in Applied Mathematics 2008; 40: 180-93.

Frieden TR. Evidence for health decision making—beyond randomized, controlled trials. New England Journal of Medicine 2017; 377: 465-75.

Gardner MJ, Altman DG. Confidence intervals rather than P values: estimation rather than hypothesis testing. British Medical Journal 1986; 292: 746-50.

Geneletti S, Richardson S, Best N. Adjusting for selection bias in retrospective, case–control studies. Biostatistics 2009; 10: 17-31.

Gilthorpe MS. Advanced Modelling Strategies: Challenges and Pitfalls in Robust Causal Inference with Observational Data. Society for Social Medicine Summer School, Leeds Institute for Data Analytics, UK; 2017. Available at: http://lida.leeds.ac.uk/event/advanced-modelling-strategies-challenges-pitfalls-robust-causal-inference-observational-data/ [accessed 30 November 2021].

Greenland S, Pearl J, Robins JM. Causal diagrams for epidemiologic research. Epidemiology 1999; 10:37-48.

Griffith GJ, Morris TT, Tudball MJ, Herbert A, Mancano G, Pike L, Sharp GC, Sterne J, Palmer TM, Smith GD, Tilling K. Collider bias undermines our understanding of COVID-19 disease risk and severity. Nature Communications 2020; 11: 1-12.

Groenwold RH, Palmer TM, Tilling K. To adjust or not to adjust? When a "confounder" is only measured after exposure. Epidemiology 2021; 32: 194-201.

Hernán MA. Causal Diagrams: Draw Your Assumptions Before Your Conclusions. TH Chan School of Public Health, Harvard University, MA; 2018. Available at: https://www.edx.org/course/causal-diagrams-draw-your-assumptions-before-your accessed 25th June 2021 [accessed 30 November 2021].

Hernán MA, Robins JM. Causal Inference: What If. Boca Raton: Chapman & Hall/CRC: 2020. Pp 312.

Hudgens MG, Halloran ME. Toward causal inference with interference. Journal of the American Statistical Association 2008; 103: 832-842.
Hume D. A Treatise of Human Nature. London: John Noon; 1738.
Hünermund P. Copenhagen Business School, University of Copenhagen, DK; 2021. Available at: https://www.udemy.com/course/causal-data-science/ [accessed 30 November 2021].
Kader IA. Path partition in directed graph–modeling and optimization. New Trends in Mathematical Sciences 2013; 1: 74-84.
Knight CR, Winship C. The causal implications of mechanistic thinking: Identification using directed acyclic graphs (DAGs). In: Handbook of Causal Analysis for Social Research 2013 (pp. 275-299). Springer, Dordrecht.
Law GR, Green R, Ellison GTH. Confounding and causal path diagrams. In: Tu Y-K, Greenwood DC. Modern Methods for Epidemiology. Pp 1-13. Springer: Dordrecht; 2012.
Lewis M, Kuerbis A. An overview of causal directed acyclic graphs for substance abuse researchers. Journal of Drug and Alcohol Research 2016; 5: 1-8. DOI: 10.4303/jdar/235992
Luft J, Ingham H. The Johari window: a graphic model of interpersonal awareness. Proceedings of the Western Training Laboratory in Group Development. Los Angeles, CA: UCLA; 1955.
Mueller M, D'Addario M, Egger M, Cevallos M, Dekkers O, Mugglin C, Scott P. Methods to systematically review and meta-analyse observational studies: a systematic scoping review of recommendations. BMC Medical Research Methodology 2018; 18:1-8.
NHS Executive (NHSE). Asking the question: Finding the evidence. Evidence-Based Health Care. Luton: NHS Executive, Unit 2-CASP/HCLU: 1999.
O'Toole G. "It ain't what you don't know that gets you into trouble. It's what you know for sure that just ain't so": Mark Twain? Josh Billings? Artemus Ward? Kin Hubbard? Will Rogers? Anonymous? Available from: https://quoteinvestigator.com/2018/11/18/know-trouble/ [accessed 30 November 2021].
Pearl J. Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. San Mateo, CA: Representation and Reasoning Series, Morgan Kaufmann; 1988.
Piccininni M, Konigorski S, Rohmann JL, Kurth T. Directed acyclic graphs and causal thinking in clinical risk prediction modeling. BMC Medical Research Methodology 2020; 20: 179. DOI: 10.1186/s12874-020-01058-z
Pocock SJ, Collier TJ, Dandreo KJ, de Stavola BL, Goldman MB, Kalish LA, Kasten LE, McCormack VA. Issues in the reporting of epidemiological studies: a survey of recent practice. British Medical Journal 2004; 329: 883.
Richiardi L, Bellocco R, Zugna D. Mediation analysis in epidemiology: methods, interpretation and bias. International Journal of Epidemiology 2013; 42: 1511-9.
Rothman KJ, Lash TL. Epidemiologic study design with validity and efficiency considerations. Chapter 6 in: Lash TL, VanderWeele TJ, Haneuse S, Rothman KJ (Eds). Modern Epidemiology. Baltimore, MD: Lippincott, Williams and Wilkins; 2021.
Roy JA. A Crash Course in Causality: Inferring Causal Effects from Observational Data. Department of Biostatistics and Epidemiology, Rutgers University, NJ; 2021. Available at: https://www.coursera.org/learn/crash-course-in-causality [accessed 30 November 2021].
Sarri G, Patorno E, Yuan H, Guo JJ, Bennett D, Wen X, Zullo AR, Largent J, Panaccio M, Gokhale M, Moga DC. Framework for the synthesis of non-randomised studies and randomised controlled trials: a guidance on conducting a systematic review and meta-analysis for healthcare decision making. BMJ Evidence-Based Medicine 2020; Dec 9. DOI: 10.1136/bmjebm-2020-111493
Schisterman EF, Cole SR, Platt RW. Overadjustment bias and unnecessary adjustment in epidemiologic studies. Epidemiology 2009; 20: 488-95.
Stacey T, Tennant PWG, McCowan LME, Mitchell EA, Budd J, Li M, Thompson JMD, Martin B, Roberts D, Heazell AEP. Gestational diabetes and the risk of late stillbirth: a case–control study from England, UK. British Journal of Obstetrics and Gynaecology 2019; 126: 973– 982.
Tafti A, Shmueli G. Beyond overall treatment effects: Leveraging covariates in randomized experiments guided by causal structure. Information Systems Research 2020; 31: 1183-99.
Tennant PWG, Murray EJ, Arnold KF, Berrie L, Fox MP, Gadd SC, Harrison WJ, Keeble C, Ranker LR, Textor J, Tomova GD, Gilthorpe MS, Ellison GTH. Use of directed acyclic graphs (DAGs) to identify confounders in applied health research: review and recommendations. International Journal of Epidemiology. 2020; 1-13. doi: 10.1093/ije/dyaa213
Textor J. DAGbase: A database of human-drawn causal diagrams. Proceedings of the European Causal Inference Meeting 2020; Apr 23.
Textor J, Hardt J, Knüppel S. DAGitty: a graphical tool for analyzing causal diagrams. Epidemiology 2011; 22: 745.
Textor J, van der Zander B, Gilthorpe MS, Liśkiewicz M, Ellison GTH. Robust causal inference using directed acyclic graphs: the R package 'dagitty'. International Journal of Epidemiology 2016; 45: 1887-94.
Tu YK, Gunnell D, Gilthorpe MS. Simpson's Paradox, Lord's Paradox, and Suppression Effects are the same phenomenon–the reversal paradox. Emerging Themes in Epidemiology 2008; 5:1-9.
VanderWeele TJ. A unification of mediation and interaction: a four-way decomposition. Epidemiology 2014; 25:749.
van der Zander B, Liśkiewicz M, Textor J. Constructing separators and adjustment sets in ancestral graphs. Proceedings of the UAI 2014 Conference on Causal Inference: Learning and Prediction, Aachen 2014; 1274: 11–24.
Viswanathan M, Berkman ND, Dryden DM, Hartling L. Assessing Risk of Bias and Confounding in Observational Studies of Interventions or Exposures: Further Development of the RTI Item Bank. Rockville MD: Agency for Healthcare Research and Quality; 2013

von Elm E, Egger M. The scandal of poor epidemiological research. British Medical Journal 2004; 329: 868-9.
Zhou J, Müller M. Depth-first discovery algorithm for incremental topological sorting of directed acyclic graphs. Information Processing Letters 2003; 88: 195-200.