Thinking with causal models: A visual formalism for collaboratively crafting assumptions

Ben Hicks
University of Technology Sydney
Sydney, New South Wales, Australia
Charles Sturt University
Albury, New South Wales, Australia
bhicks@csu.edu.au

Kirsty Kitto
University of Technology Sydney
Sydney, New South Wales, Australia
Kirsty.Kitto@uts.edu.au

Leonie Payne
University of Technology Sydney
Sydney, New South Wales, Australia
Leonie.E.Payne@student.uts.edu.au

Simon Buckingham Shum
University of Technology Sydney
Sydney, New South Wales, Australia
Simon.BuckinghamShum@uts.edu.au

ABSTRACT

Learning Analytics (LA) is a bricolage field that requires a concerted effort to ensure that all stakeholders it affects are able to contribute to its development in a meaningful manner. We need mechanisms that support collaborative sense-making. This paper argues that graphical causal models can help us to span the disciplinary divide, providing a new apparatus to help educators understand, and potentially challenge, the technical models developed by LA practitioners as they form. We briefly introduce causal modelling, highlighting its potential benefits in helping the field to move from associations to causal claims, and illustrate how graphical causal models can help us to reason about complex statistical models. The approach is illustrated by applying it to the well known problem of at-risk modelling.

CCS CONCEPTS

• **Computing methodologies → Modeling and simulation**; • **Human-centered computing → Collaborative and social computing theory, concepts and paradigms**.

KEYWORDS

causal models, directed acyclic graphs, transdisciplinary collaboration, diagrammatic reasoning

ACM Reference Format:
Ben Hicks, Kirsty Kitto, Leonie Payne, and Simon Buckingham Shum. 2022. Thinking with causal models: A visual formalism for collaboratively crafting assumptions. In *LAK22: 12th International Learning Analytics and Knowledge Conference (LAK22), March 21–25, 2022, Online, USA.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3506860.3506899

LAK22, March 21–25, 2022, Online, USA

ACM ISBN 978-1-4503-9573-1/22/03…$15.00
https://doi.org/10.1145/3506860.3506899

1 INTRODUCTION

> *To think at all is to forget. To know at all is to abstract. When we reason scientifically, we go further and deeper with thought. We become deliberate with what to forget and mindful of how to abstract.* Essa [17, p36]

Learning Analytics (LA) is sometimes referred to as a bricolage field [19] that requires people to work in a 'middle space' between the learning and analytical sciences [30, 58] to bridge epistemic boundaries [20]. As such, it is a field that can suffer from problems of communication, where people from very different cultural backgrounds talk past one another. Often, we see educational experts excluded from the conversation, unable to judge or evaluate highly technical approaches that rely upon advanced statistics, machine learning, and other methods that quickly come to resemble a black box [40]. This difficulty in traversing disciplinary boundaries leads to a number of critical challenges for the field around who gets to *Participate* in defining the questions that the field explores [64], and how educational *Theory* can be used to inform results. Furthermore, the field is often accused of having a lack of *Transparency*, which can make it difficult to *Intervene* in the system even with strong statistical results.

In this paper we will argue that many of these problems can be addressed, at least partially, by *thinking with causal models.* We will demonstrate that the visual component of causal modelling, Directed Acyclic Graphs (DAGs), can be used with little mathematical expertise to provide a highly interpretable artefact that supports genuine transdisciplinary collaboration among technical and non-technical stakeholders. After elaborating upon the problems introduced above (Section 2), and then providing a brief introduction to the framework first introduced by Pearl [41] (Section 3), we will demonstrate its utility in thinking about complex educational processes by reasoning through a real world scenario encountered by one of the authors in his position as an LA professional (in Section 4).

1.1 Challenges facing the field of Learning Analytics

Beyond the technical challenges of gathering useful data from multiple disjoint LA systems [28], building new methods of analysis

[31], and implementing them at scale [9, 61], LA faces a number of more *social* challenges that arise from its bricolage nature. Four of these will form the focus of this paper.

1.1.1 Problems of Participation. LA can be an intimidatingly technical field due to its reliance upon advanced statistics and machine learning methods in analysing educational data. This can make it hard for educators to have a say in what gets studied, as it is possible for a data analyst to collect a dataset, analyse it for patterns and report findings without any consultation with educational professionals or researchers. All too often, we see data intensive fields fall prey to 'data dredging', where a complex analysis is performed which does little to improve the actual problems faced by stakeholders [4], and LA is not immune to this problem [27]. We believe that much of this issue springs from a lack of genuine collaboration between analysts (who know and understand the methods to apply to the data) and educators (who understand the domain, and so should have influence over what critical questions are asked by LA researchers). In short: *Who gets to define the problems that LA seeks to address?* [64]

While stakeholder engagement and participatory approaches are common in the field, we must take care to ensure that a genuine collaboration emerges, where educational experts are not sidelined due to an inability to penetrate the fog often created by complex analytical methods. But how can we invite people in, to genuinely collaborate when the methodologies used can take significant expertise to master? One potential solution makes use of Participatory Design (PD) to create accessible design representations at the point of deciding what type of LA to implement [1], an example of the broader set of Human-Centered Design (HCD) methods [6]. These representations can be manipulated by participants as they think through an educational problem. Similarly, a recent paper, by Wise et al. [64] has proposed that missing data can be highlighted to help stakeholders understand what is *not* there in a LA report as much as what *is*, so helping educators to reach informed decisions from a position that is aware of potential blind spots. Such methods move us closer to what Callon et al. [10] term *technical democracy*, but they do so by attending to the need for informal representations, typically of the sociotechnical system in which a tool must be embedded, or user interface designs. Data scientists must ultimately adopt, adapt or invent algorithms informed by the qualitative insights that non-technical stakeholders typically bring, and there are to our knowledge no representational techniques that assist diverse stakeholders to formalize the domain insights gained from PD/HCD.

1.1.2 Problems of Theory. Education is a vast field, that has a long history of its own data, methods, and theories about how people learn [18]. It is essential that we work hard to link LA to well established principles and results. However, it is a well accepted challenge to map the low level activity data from online platforms/sensors, to higher level and theoretically informed educational constructs [63]. Educational data is complex, with many variables over few observations, which can make it difficult to extract insights from the chaff of clickstream data that is often collected. Despite claims to the contrary by well known researchers in Artificial Intelligence [36], big data does not always lead to better models, at least, not in education. Rather, a data analysis that is not informed by theory can fall prey to problems like Simpson's paradox, or link variables that we know cannot be correlated, and fail to spot patterns in difficult to identify subgroups [37]. Indeed, "if theory-driven models can be wrong, data-driven models can be *fragile*" [21, p4].

Unfortunately, a literature review performed by Wong et al. [65] demonstrates that while educational theories such as self-regulated learning, motivation, and social constructivism have been utilised in a number of LA studies, the results were largely correlational. As a field we are yet to move to a position where educational theories commonly drive our analysis. Moving towards more theoretically informed data analyses would help to improve participation by ensuring that LA is well linked to concepts that educational researchers recognise and understand. It would also help us to ensure that surprising data driven results could be interpreted, queried, and challenged if they were poorly founded.

This leads us to a second question: *How can we ensure that theory informs LA analysis?*

A number of recent papers have started to make links between educational theory and LA. Marzouk et al. [33] links the data collected in a LA system to the well regarded educational theory of Self Regulated Learning (SRL), and Epistemic Network Analysis (ENA) [54] has been used to link the coding of large data sets describing conversation to theoretically justified constructs. However, it is clear that more methods are required in this space. This raises questions about how we can wean our models from association and start to move towards stronger claims.

1.1.3 Problems of Transparency. How can non-experts be supported in critiquing LA assumptions and solutions? There is a long and distinguished line of work that has provided a critical lens on the way in which data and analytics can be misused, highlighting factors such as: the ways that worldviews underpin classification schemes [3]; risks that systematic, discriminatory biases can permeate machine learning algorithms [2]; and the dynamics by which predictive models can disconnect from reality through feedback loops [38]. Indeed, while critical data studies has started to focus upon LA [51, 53], the field has already explored many of these issues (see for example the recent special issue edited by Buckingham Shum and Luckin [7] and the references therein). The challenge to turn these academic insights into a widely accessible, practical means for *contesting* the results of an algorithm is widely recognised as a critical factor for achieving responsible LA. This problem is closely related to that of participation, as people can only challenge the results of LA if the mechanism for generating them can be inspected. If LA use models that cannot be inspected and understood, their opaqueness ensures the resulting solutions are unlikely to be trusted. There is a need to "make the assumption behind the model more explicit and transparent" [49, p2949].

In short: *Who gets to challenge the solutions delivered by LA?*

Achieving an end state where all stakeholders can critique and challenge LA solutions requires not just open, explainable or transparent algorithms; we need tools for bridging the displacement generated by sophisticated approaches and algorithms that require substantial expertise to understand and to challenge.

1.1.4 Problems of Intervention. Finally, assuming that we have passed these previous challenges, we find a new one emerging. How should we act? LA seeks to improve learning and the environments in which it occurs, but for that to occur we need to close the loop [14], by intervening in the systems that we study.

In short: *What should we do now?*

> *Regardless of how well intentioned or theoretically grounded our LA aspirations are, we must be able to implement them as running code using real data, creating a compelling user experience, with benefits for learners and/or educators.* Buckingham Shum [8, p7]

We all know the dangers associated with assuming that correlation equals causation [46], and yet, sound arguments have been presented that we have an *obligation to act* if we think that a student needs support [44]. We need ways to ensure that our LA has a strong link to reality, and makes grounded and well founded claims about what a decision maker, student, or educator could change to improve the outcomes modelled. How can we work towards ensuring that the models generated by LA are actionable [39], with results that can be mapped to definable educational constructs? As such, this problem is highly related to our problem of Theory described above.

1.2 Diagrammatic reasoning

As Wise et al. frame it, "who gets a place at the table" [64] when decisions are made to shape an analysis in LA? The previous section has illustrated that in a field as diverse as LA, it may not be enough to just bring the various stakeholders to the table, we must create tools that enable a collaborative sense-making [29].

There is a well established research literature in cognitive science and human-computer interaction, studying the properties of visual representations that (sometimes) enhance individual and collective reasoning. Their effectiveness is a function of the knowledge that a viewer brings to the diagram, they are not informative in and of themselves [11].

Designed well, a visual representation augments 'internal cognition' in the mind with a form of 'external cognition' in the environment. As analysed by Scaife and Rogers [52], this operates in multiple ways, including: (i) *Computational offloading* to reduce cognitive effort compared to other representations (ii) *Graphical constraining* to support specific kinds of reasoning; (iii) *Continuous internal/external interplay* which shapes reasoning, and in turn changes what we seek and see in the representation. Furthermore, as argued by Shum and Hammond [56], the use of computational aids to assist the design of interactive software tools brings both product and process benefits. While the *product* of modelling the problem formally is the ability for a tool to rigorously verify that key criteria have been met (e.g. statistically, or using another formal language), the *process* of modelling the problem to reach that point can itself be extremely productive, requiring the analysts to focus and structure their thinking, and explain their assumptions, drawing attention to vague or missing information.

This paper will therefore make the case that a *graphical causal structure* depicting the causal relationships between variables [41] shows promise as a form of diagrammatic reasoning, enabling a more diverse set of stakeholders to participate in a genuine and collaborative definition of a LA model, a highly abstract process that is typically left to analysts to code in opaque notation that excludes others (even if they have been consulted). We start with an introduction to causal modelling [41].

2 MODEL, OR CAUSAL MODEL?

To progress we must first understand how causal models differ from statistical models. We start by noting that all models are an abstraction, they attempt to trade a reduction in information for an increase in knowledge. As such they have little value in progressing towards 'truth', but they do help us to understand the world and avoid misconceptions [17]. The *statistical* model $Y = f(X)$ describes the effect of X on Y via a process f; it is a self-contained 'small-world' abstraction of nature. Obviously we would want our variables X and Y to closely represent reality, but how closely we require f to resemble nature's own mechanisms depends on the purpose of the model. An explanatory model requires them to match as closely as possible but predictive models do not have this restriction, as they care only for matching the patterns in X to Y in the most accurate way; the process f is less important [55]. For instance we care little about *how* a model (f) converts speech to text, only that the sound of the words we say (X) match closely with the text (Y). If instead we want to use the same techniques to understand the way in which the brain processes language we would gladly settle for a less accurate model that offers more explanatory power.

An abundance of work has, understandably, been put into improving the capacity of models and machine learning tools to *predict* an observable Y, given some input X. For example, LA has devoted much attention to the prediction of whether a student should be classified as "at risk" (the variable Y) given some set of observed behaviour (the variables X) [12]. While the utility of ever more accurate models has been questioned by some LA researchers [27], few papers have explored the distinction between prediction and explanation. However, as a field that attempts to *explain* behaviour, and moreover, to assist educators to *intervene* to improve a state of affairs, it is surely time for LA to move beyond the assumption that accurate prediction suffices [42]. It is important that we work to establish a close link between our small world statistical model f and reality, or we risk a number of pathological outcomes, such as:

- We must individually question every statistical association to determine if it is spurious or not.
- Explaining *why* something arises in the model is at best opaque, and in fact may be impossible.
- "What if" questions are off limits, as they imagine a world beyond what the data has seen.

These are the kinds of problems that causal inference frameworks [41, 50] were designed to deal with. Causal frameworks search for a causal relationship between X and Y, and in doing so work towards closing the gap between our statistical model and the real world. This enhancement requires models that know something more than just the data they were fed, because the problems above require a different level of questioning, one that can deal with interventions and counterfactuals [42]. This is commonly done by adding a theoretical structure to the model.

2.1 Graphical causal structures

Intriguingly, one way to add this theoretical knowledge consists of drawing a picture, a sketch of your assumptions about "what causes what". Augmented by this diagram (formally known as a *graphical causal structure*) the statistical model becomes a (graphical) causal model. The task that this picture performs is to describe the assumed paths of causal influence between the variables [41] in a way that a model can interpret. Because the drawing needs to be understood in the structured world of the model it must come in a particular form; a Directed Acyclic Graph (DAG), such as the one depicted in Figure 1, which could easily have been drawn in a conversation between a modeller and an educator.

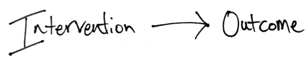

Figure 1: A simple causal Directed Acyclic Graph (DAG)

A *graph* is simply some blobs (nodes) with lines between them (edges). The graph is *directed* because the lines are arrows: the model needs to learn about which variables are *causes* and which are *effects*; does $X \rightarrow Y$ (X cause Y) or does $Y \rightarrow X$ (Y cause X)? The graph is *acyclic* (no loops) to enforce a simplicity into the causal description: no variable ends up a cause of itself [1]. Despite it's simplicity, Figure 1 encodes rich information that may be unknowable from data alone: Some sort of an *Intervention* results in (i.e. causes) a better *Outcome* for a student.

Structural Equation Models (SEM), in particular their model diagrams, are akin to a graphical causal model but over the years they have moved away from making causal claims [41, p138]. So despite their common ancestry and similarities, described in Pearl and Mackenzie [42], the differences between causal DAGs and SEMs are telling. DAGs emphasise the causal relationship between variables without initially specifying its functional form, delaying the technical requirements for engaging with the model construction.

DAGs have a number of properties that could prove highly beneficial to LA: they can elegantly communicate abstract structural information, are underpinned by a rigorous mathematical framework [41], and most importantly to the current argument, they can open conversations with stakeholders using terms as simple as "blobs" and "arrows".

2.2 How does Learning Analytics currently make causal claims?

Education has traditionally relied upon Randomised Controlled Trials (RCTs) to make causal claims, if at all. Thus, a sample is randomly split into a control group, and a group that receives an intervention, with the difference in outcomes between the two groups then used to link the intervention to a causal claim [60]. However, such approaches are both expensive to run, and fall prey to ethical quandaries. In particular, if we truly think that an intervention is likely to lead to better student outcomes then we must weigh up the importance of building up our evidence base against the duty of care we owe to our students [44]. Ethical tensions quickly mount up [26] when we attempt to run RCTs in a field with outcomes that are so likely to impact upon (and possibly entrench) long term advantages and disadvantages [51].

[1] Although we note there are ways to deal with this kind of causal feedback e.g. by indexing the variables in time to stop the loop returning to the same 'place'.

If LA groups decide not to run an RCT, then they need other tools and frameworks to make causal claims. Interestingly, explicit graphical causal models are rare in the educational domain. One good example is provided by Ramirez-Arellano et al. [45] which uses prior theory to build a causal model, in the form of a SEM. Additionally, some studies have explicitly used causal discovery techniques [13], where causal structure is discovered from a data set via an algorithm.

More common are quasi-experimental techniques which induce an implicit causal model through experimental design. One method, *matching* [35], involves finding pairs of similar (matched) data points, one who has received an educational intervention and one who has not, to make semi-causal claims. Matching methods have been used in the LA and Educational Data Mining (EDM) fields to evaluate the impact of student facing dashboards [25] and assess the effectiveness of intelligent tutoring systems [34]. The common thread with all these applications was the goal of estimating the effect of an intervention under the influence of a selection bias [5].

Throughout these examples researchers are encountering a similar problem; how do I find the effect I am interested in when there are confounding factors? This is a common enough problem to have some call for an increase in attempts to address this complex issue, but exactly how we might do this remains uncertain [48].

2.3 A hypothesis

We believe that causal models are useful *per se*, for the modelling opportunities that they provide. As such, they should become an object of increasing attention for the educational data sciences. However, in this paper we will focus upon a slightly different benefit that we believe flows from a serious adoption of causal modelling:

> *Causal Models show promise for supporting a theoretically informed and more collaborative communication mechanism between stakeholders of different backgrounds when defining LA models.*

Rather than performing user studies, this paper will remain largely theoretical. We will walk the reader through an example of how DAGs help us to reason about a LA system, illustrating the new insights that can be gained from an attempt to explicitly impose causal structure on a dataset. However, before we can apply causal modelling to educational data, we first need to introduce the reader to the approach.

3 A *VERY* BRIEF INTRODUCTION TO CAUSAL DAGS

In this section we provide a very brief introduction to causal modelling with DAGs. Our discussion will privilege concepts that support transdisciplinary communication over statistical rigour. For a deeper exploration of causal models see Pearl and Mackenzie [42] for an approachable introduction, or Pearl [41] for the full mathematical apparatus.

Let us start by supposing we want to model the effect of having an extensive home library on later academic achievement. We might start by assuming that children with access to a *Home Library* demonstrate improved *Academic Achievement*, as they are able to make use of those books to practice reading and learn about the world beyond the family home. This relationship is introduced by the simple DAG shown in Figure 2(a).

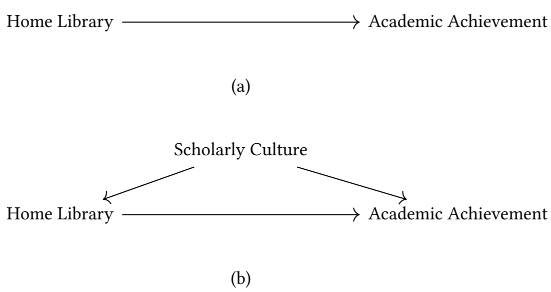

Figure 2: Two simple alternative causal models that describe the effect of a home library on a student's academic achievement.

However, on presenting such a diagram to an educator they would be able to raise an immediate (rather obvious to them) challenge to the model it represents, by pointing out that a child with an extensive home library is likely to have parents with a more scholarly background (after all, someone bought the books in the library!) But, says the educator, would such parents not have an impact upon the *Academic Achievement* of the student as well? At this point the modeller could respond that this is actually a concept that is well understood statistically: it is a possible confounding effect upon our model [57]. The modeller could construct the DAG shown in figure 2(b) to incorporate this new causal claim, and the educator would be able to interrogate and approve (or challenge) the new representation. While simple, this vignette demonstrates the utility of a simple pictorial representation which has a well defined modelling apparatus to support it. The educator is facilitated in understanding and contributing to the development of the models shown in Figure 2, without needing to understand the full complexity of the statistical model it represents.

Unpacking the statistical apparatus introduced in Figure 2 a little more, the blobs (nodes) of the DAG represent variables that we have decided to include in the model. The arrows (edges) in the DAG represent our assumptions about how those variables influence one another. A useful way of analysing a DAG is to pick two variables that you want to explore the causal relationship between, (in this case from *Home Library* to *Academic Achievement*), and classify them as (i) open or closed, and (ii) causal or non-causal.

3.1 Causal Paths

Put your finger on any node of a graph and then trace your finger along the edges to reach another node and you have described a *path*. For instance starting at *Home Library* in Figure 2(b) then moving up to *Scholarly Culture* and across to *Academic Achievement* traces the path *Home Library* ← *Scholarly Culture* → *Academic Achievement*. A path like this shows a connection between the variables *Scholarly Culture* and *Academic Achievement*, but what kind of connection?

3.1.1 Causal or non-causal paths. A *causal* path from X to Y means that if X is forced to change then Y will change also. This is not an easy phenomena to directly observe in the data alone. Data excels at showing how the variables are associated, not how influence flows from one variable to the next — this is why the direction of the arrows is an essential ingredient of a causal model. However, it is simple in a DAG to discriminate between causal and non-causal paths: a causal path is one where all the arrows go in the same direction.

3.1.2 Open or closed paths. An *open* path allows information to flow along it. The way this manifests in the data is that the variables along an open path will be associated with one another — they will *not* change independently. A *closed* path will have a feature blocking the flow of association so that the two variables at each end of the path are independent of each other. To decide if a path is open or closed it is enough to examine a DAG to find one of the three possible patterns that can arise in them; the chain, the fork and the collider, depicted in Figure 3.

Chain	X ⟶ C ⟶ Y	*open*
Fork	X ⟵ C ⟶ Y	*open*
Collider	X ⟶ C ⟵ Y	*closed*

Figure 3: The (unadjusted) three elemental DAG patterns; the Chain, Fork and Collider.

Why is the fork open? C is a common cause of X and Y, so when C changes we see changes in X and Y *together.* As they move together they will appear associated in the data. Why is the collider different? X and Y are common causes of C, so whilst C depends on the values of X and Y there is no reason that X and Y share information with each other. Plants need water and sunlight to grow, $Water \rightarrow Grow \leftarrow Sunlight$. But knowing how much *Water* a plant is getting does not tell us anything about much *Sunlight* it is getting; information does not flow along this path automatically, it is blocked by the collider *Grow*. This all changes if we learn something about the collider, and for this we will need to discuss *adjustment.*

3.1.3 Adjustment, opening and closing paths. *Adjusting*[2] for a variable means that we include our knowledge about the variable in the analysis. Take the plant growth example, which begins with the variables *Water* and *Sunlight* independent of one another. Let us say we separate our study into groups; we look at plants that are growing well compared to plants that are not, that is we add knowledge of the variable *Grow* to our model. In this case, knowing something about the *Water* variable within the context of, for instance, dead plants, does say something about the status of *Sunlight*. Indeed, if you are looking a plant that is not growing, knowing that it has plenty of *Water* tells you something about the likelihood it is getting enough *Sunlight* — information now flows between *Water* and *Sunlight*, but only if we know the value of *Grow* as well. To express this again for any DAG, adjusting for the common cause in a collider pattern *opens* the path and allows the flow of association between variables. Visually, an adjusted $\boxed{\text{variable}}$ is placed in a box.

[2]Adjustment is also called *controlling*, *conditioning on* or *stratifying*.

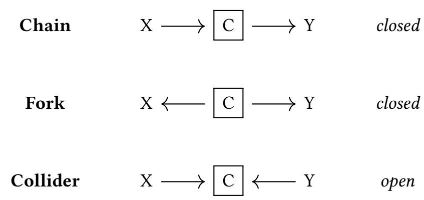

Figure 4: The three elemental DAG patterns; with central variable adjustment.

Why do the other patterns in Figure 4, the chain and the fork, close when adjusting for the central variable? Remember that adjusting for a variable can be conceptualised as *learning about* that variable. To illustrate this for a chain let us extend the plant growth model by adding the path *Grow* → *Flower*:

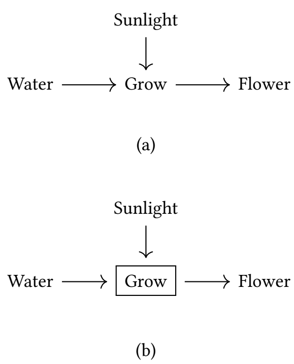

Figure 5: DAGs with a chains (e.g. *Water* → *Grow* → *Flower*) and a collider (*Water* → *Grow* ← *Sunlight*). In (a) association flows between *Water* and *Flower* but not between *Water* and *Sunlight*. When we adjust for *Grow* in (b) this opens the path between *Water* and *Sunlight* but closes the path between *Water* and *Flower*.

Knowing if a plant has had enough water certainly tells us about the chance of it flowering (*Water* and *Flower* are associated). But if we then learn that the plant is growing well, our knowledge of *Water* (and *Sunlight*) loses relevance. To see why, imagine we are looking at a verdant specimen brimming with life; its high level of *Grow* tells us all we need to know about the plants prospect of flowering — the importance of knowing if it is getting enough *Water* is diminished because the effect of *Water* on *Flower* is mediated through *Grow*, which we now have knowledge of. As such, within groups of plants with a similar level of growth (adjusting for *Grow*) there is no association between *Water* and *Flower*, the path is now blocked.

3.2 Confounding

Confounding occurs when we have association flowing along non-causal paths between the two variables we are trying to compare, which is what is happening when we try to measure the effect of *Home Library* on *Academic Achievement*. Returning to our original example, the path *Home Library* → *Academic Achievement* is an open causal path. It is a causal path because all the arrows flow in the direction we are thinking about, and it is an open path because nothing is 'blocking' the flow of information. As the path is *open* we would expect to see association between the variables *Home Library* and *Academic Achievement* in the data. As the path is *causal* we would expect that manipulating *Home Library* would change *Academic Achievement*. This is the path we are interested in measuring the strength of.

In contrast, the path *Home Library* ← *Scholarly Culture* → *Academic Achievement* is an open *non-causal* path. As the path is non-causal, manipulating *Home Library* will not influence *Academic Achievement* along this path. It represents the causal claim that adding books to the home library *does not* change the scholarly culture and in turn change academic achievement. As the path between the variables *Home Library* and *Academic Achievement* is *open* we would expect to see an association emerge in a dataset with these variables — even if there was no direct effect between *Home Library* and *Academic Achievement*! This is due to the common cause *Scholarly Culture*; as it increases we expect an (on average) increase in size of the *Home Library* and also in a student's *Academic Achievement*. While this results in an association between *Home Library* and *Academic Achievement* the cause of the association is not the *Home Library*. It is the harder to measure, or latent, variable *Scholarly Culture*. This is why *Scholarly Culture* is called a confounder, because the total association we see in the data between *Home Library* and *Academic Achievement* comes from two sources, *Home Library* → *Academic Achievement* (causal) and *Home Library* ← *Scholarly Culture* → *Academic Achievement* (non-causal); the direct causal effect of *Home Library* on *Academic Achievement* is confounded by the existence of other open non-causal paths. If we want to measure the direct effect of *Home Library* on *Academic Achievement* we need to block the flow of association along the non-causal path, isolating the causal path that we are interested in. There are two main methods that can achieve this goal; *randomisation* or *back-door adjustment*.

3.3 Removing confounding through randomisation

Randomisation involves manipulating how the influencing variable (*Home Library*) is generated, and forcing its value (the number of books) to be randomly assigned. This is what a Randomised Control Trial does, and in terms of the DAG it effectively removes all arrows *into* the influencing variable. This becomes governed purely by a random process and is now no longer influenced by any other variable in our model (See Figure 6). However randomisation is not generally available as a tool in observational studies such as this, so another way to block the path is needed.

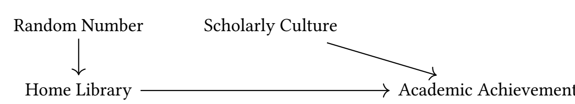

Figure 6: A randomised control trial DAG.

3.4 Removing confounding with the back-door adjustment

The non-causal path *Home Library* ← *Scholarly Culture* → *Academic Achievement* is known as a *back-door path* since it begins flowing against the direction of the arrow. The back-door adjustment provides a method of closing back-door paths by *adjusting* for key variables to block the flow of association along the path.

In terms of our example adjusting for *Scholarly Culture* blocks the flow of information along the path *Home Library* ← *Scholarly Culture* → *Academic Achievement*. This is because knowing about the level of *Scholarly Culture* tells us everything we want to know about the strength of the association between *Home Library* and *Academic Achievement* due to this common cause, information no longer flows along this path.

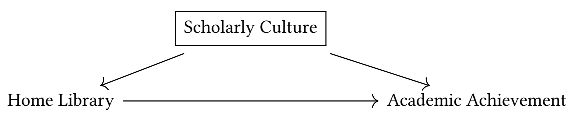

Figure 7: Blocking the non-causal path by adjusting for Scholarly Culture.

So if we can measure the association between *Home Library* and *Academic Achievement* within levels of *Scholarly Culture* and then pool the results we obtain an unbiased estimate of the causal effect of *Home Library* on *Academic Achievement*, assuming the DAG is sufficient. We could do this by surveying all families in our system, to measure their scholarly culture, and then segmenting our sample accordingly.

4 THINKING WITH A CAUSAL MODEL

To think *with* something is to inhabit its world in a meaningful way [16], with the hope for exploring the diversity of sense making and so cultivating a more collaborative and inclusive process [62]. Models are often used to explain and predict, but the traditional focus has been upon the final model, not in the process of constructing the model itself. We posit here that this *process* of model construction, aided by causal modelling tools, may be as fruitful to LA as the final *product*. This is because modelling forces us to make explicit what we consider important in a system, and what can be safely ignored. We propose that deeply engaging with the world of causal modelling has the potential to catalyse rich collaborative sense making. Specifically, the development of a graphical causal model makes explicit an often obscured process; the abstraction of a model from the context in which it resides.

The DAG serves as an artefact to help facilitate collaboration across disciplinary boundaries and so leverage a wider variety of expert knowledge, potentially inviting a greater breadth of expertise to the LA development table. DAGs are a way of reducing artificial complexity in understanding relationships between variables. For instance $A \rightarrow B \rightarrow C$ is rich in data about conditional independence between the variables A, B and C; you could state that C and B are not independent, B and A are not independent, and C is independent of A conditional on B (with notation $C \perp A|B$), and those statements do not capture the postulated *direction* of the effect. These are also more foreign mathematical concepts and symbols for a non-technical audience to interpret and challenge. The causal model and DAG seems to keep the natural complexity of the model intact, while minimising its artificial complexity.

5 THINKING WITH A STUDENT RETENTION MODEL

To make the above rather abstract concepts more explicit, we will now consider a real world modelling scenario, to demonstrate the utility of causal modelling to support thinking and collaboration in transdisciplinary teams. In what follows we will make use of the widely studied problem of student retention and support [22]. In the spirit of "thinking with diagrams", we present this in a stylised form of analyst's narrative, as though we were thinking aloud with colleagues, using the causal diagrams to clarify our thinking.

5.1 Causal models in the retention space

Not everyone succeeds in their academic studies, but a better understanding of what helps a student persist is important. In LA, student retention normally centres around the prediction of "at risk" students who need to be supported in some way via an intervention. While some early attempts used student dashboards [43], it has by now become common to contact students identified as at risk via a phone call that offers them support, and potentially discuss changing their study load. Finding the best way to intervene in this complex scenario is a challenge [15], it remains hard to know which lever to pull. The effectiveness of an intervention requires a model that can offer explanation as well as prediction. Since this area of LA has been extensively studied, we need to demonstrate what more could be gained through the use of a causal model.

Interestingly this space has already been approached from a causal perspective, years before the formalisation of methods involving DAGs [41] emerged. Kember proposed a causal model of student persistence [23, 24], that operationalises earlier work by Tinto on a conceptual model of student dropout [59]. Woodley et al. [66] has since pointed out that dropout is complex and comes in many forms, possibly requiring a different model for each.

Our example will centre around a particular case of a student support system at Charles Sturt University (CSU) aimed at early intervention, primarily via a phone call [32]. This example was chosen as it is a central part to the work done by the first author at CSU (BH). Students are identified in week 3 or 4 of the session and put on a list to be called by a student outreach team. Students on the "at-risk" list are all contacted in some form (sms, email) and phoned multiple times, but a conversation with an outreach team member can only happen if the student answers the phone. The overall aim is to reduce failure rates in subjects, particularly due to the non-submission of assessments. As a result, 'success' may be due to one of two possible outcomes: guiding the student to existing support structures, or towards withdrawing from the subject. In the scenario to be modelled here, CSU was trying to ascertain what effect the intervention has on the academic outcome of a student identified as at risk of failure.

On ethical grounds, a randomised controlled trial was not possible, since the goal of the project was primarily to support as many students as possible; understanding the projects' effectiveness was a secondary level goal, but not the principle *raison d'être* for the project. Nonetheless, we see the potential for using two subgroups of students to make claims about the effectiveness of the project; students who received the intervention (by answering a phone call) and those who do not (they did not answer any of the calls). Membership of either group is not randomly assigned, however, so any attempt to assign a *causal* effect of the intervention on the outcome would need to also understand and account for how students are assigned to the two groups. In what follows we will illustrate how causal modelling can help us to reason about this scenario.

5.2 The Causal Model

5.2.1 Starting somewhere. The model construction begins simply with a hypothesis that: $Intervention \rightarrow Outcome$. In this case, $Intervention$ denotes whether or not the student had a conversation about their options with an outreach caller, while *Outcome* denotes if the student passed all their subjects they remained enrolled in or not. We immediately encounter a problem in our modelling.

Problem 1: The intervention does not *directly* change the outcome, rather, it changes some unmeasured attributes of the student that drive their academic decision making.

5.2.2 Where lies what we do not know? To address this problem, we can add a variable in between *Intervention* and *Outcome* that represents the changeable academic attributes of the student that affect their chance of success. It is these changeable attributes that we are trying to influence with our phone intervention. Let us call this variable *Mutable*, to represent that these student attributes are malleable. Our resulting enhanced model is depicted in Figure 8.

$$\begin{tikzcd}
\text{Intervention} \arrow[r] & \text{Mutable} \arrow[r] \arrow[l, bend left, dashed] & \text{Outcome}
\end{tikzcd}$$

Figure 8: A DAG relating an *Intervention* to an *Outcome*, via a mediator of *Mutable* student academic attributes.

Thinking: In Figure 8 we have started with a near minimal DAG. Although far from ideal, once drawn it can be interrogated by others, and the building of the causal model becomes a collaborative process of thinking with the model. However, in making explicit our assumptions, we quickly hit a new problem.

Problem 2: What is it we are not seeing? This modelling exercise began because an educational expert had voiced a suspicion that the selection of students into the *Intervention* group (i.e. those who answer the call) is not random. But if this the case, and we have wrapped these student attributes into *Mutable*, then *Mutable* influences *Intervention* as well as the other way around. We have a loop. This is not allowed in DAG (remember the A is for Acyclic).

5.2.3 Making time acyclic. We can avoid loops in a DAG by carefully splitting variables into 'before' and 'after' some epoch, in this case we can split *Mutable* into before the phone call, $Mutable_{t=0}$, and after the phone call, $Mutable_{t=1}$, an update that is now reflected in Figure 9.

Thinking: The maxim that causes should precede effects is a useful device to help sharpen thinking about a causal model. However, working out exactly *where* to divide the model of data into before and after can require some finesse.

$$\begin{tikzcd}
 & \text{Intervention} \arrow[d] & \\
\text{Mutable}_{t=0} \arrow[r] \arrow[ur] & \text{Mutable}_{t=1} \arrow[r] & \text{Outcome}
\end{tikzcd}$$

Figure 9: Adding an epoch, before and after the phone call.

Problem 3: At the moment, because the *Mutable* variables are unmeasured, there is no way to adjust for the back-door path $Intervention \leftarrow Mutable_{t=0} \rightarrow Mutable_{t=1} \rightarrow Outcome$. This means there is no approach, with the current model, to find the causal effect of *Intervention* on *Outcome*.

5.2.4 Leveraging educational theory. We have placed unknown (i.e. latent) student attributes into our model, now we add further structure to our model using an insight raised by our educational expert; there are student attributes that we cannot change. We denote these *Fixed*, and in Figure 10 we use them to represent any student qualities that we can measure prior to the intervention that: (i) remain fixed throughout the study session, (ii) are measurable, (iii) influence the chance of the student answering the phone, (iv) influence student outcomes through the *Mutable* attributes. The most important thing to note when augmenting a graphical causal model like this is the *lines that are not drawn.*

$$\begin{tikzcd}
\text{Fixed} \arrow[r] \arrow[d] & \text{Intervention} \arrow[d] & \\
\text{Mutable}_{t=0} \arrow[r] \arrow[ur, dashed] & \text{Mutable}_{t=1} \arrow[r] & \text{Outcome}
\end{tikzcd}$$

Figure 10: Adding *Fixed*, a variable of known (relatively fixed) student attributes.

In the case of Figure 10, we should note that there is no direct path between $Fixed$ and $Outcome$. This amounts to using the model to claim that the effect of $Fixed$ on $Outcome$ is always mediated by $Mutable$. The dashed path from $Mutable_{t=0}$ to $Intervention$ indicates some doubt over including that path. Indeed, this dashed path indicates our concern that there may be some student characteristics in *Mutable* that impacts directly upon the likelihood of a student receiving an intervention (i.e. in this case answering the phone call) that is not explained through the common cause, *Fixed*.

Thinking: We have split the variables describing our students into two parts: mutable ones that cannot be measured, but which we are trying to influence with our intervention, and a separate set of fixed, measurable variables that are likely to be useful in segmenting our student population. At this point it is not necessary to determine exactly what *Fixed* or *Mutable* are comprised of; these are place holders for everything we know about the student that we think affects the other variables in a particular way. This level of abstraction helps move forward in discussing the causal structure quickly, without getting stuck in details that someone unfamiliar with the system may not (care to) know about. However, an educational expert could work with the modeller to decide what belongs in them. *Fixed* might include variables relating to: socio-economic status, gender, and first language, whereas *Mutable* could include variables relating to: employment status, study load, grit, and work ethic.

Having worked through this exercise, we can now see how to measure the causal effect of interest, but how we do this will depend on our assumptions around the path $Mutable_{t=0} \dashrightarrow Intervention$. If this path is weak enough to be ignored, meaning we believe the association between *Mutable* and *Intervention* arises mostly from their common cause *Fixed*, then there is only one non-causal path to block: $Intervention \leftarrow Fixed \rightarrow Mutable_{t=0} \rightarrow Mutable_{t=1} \rightarrow Outcome$. This can be blocked by adjusting for *Fixed*.

Problem 4: On the other hand, if we believe the path $Mutable_{t=0} \dashrightarrow Intervention$ should be included in the model then there is no way to use adjustment to block the non-causal path $Intervention \leftarrow Mutable_{t=0} \rightarrow Mutable_{t=1} \rightarrow Outcome$ as we cannot adjust for the unmeasured *Mutable* variables. In this case there is no way to get an unbiased estimate of the effect of the *Intervention* on the academic *Outcome* of the student. This would be reflected in the data, and so we now have a testable set of claims that can be used to investigate this scenario further.

5.3 What was gained by this approach?

Despite the ambiguous end result of the retention causal model, the final model did provide two clear pieces of knowledge: (i) it helped us to clearly scope the assumptions required for estimating a causal effect, and (ii) it provided us with a method for estimating this effect if the assumptions are met. The *process* of constructing the model through dialogue and diagram also provided us with insights that helped clarify our thinking about this scenario. The first was the introduction of the *Mutable* variable and thinking about what exactly is it that the intervention was trying to change. The second was a clearer understanding of *why* the system was complex — the many confounding factors are challenging to place in the model structurally and temporally. These insights were gained *prior* to looking at the data, through the careful incorporation of concepts from educational theory, and in dialogue with educational experts.

6 CONCLUSION

> *Perhaps the greatest benefit of an explicit causal inference framework is that it requires us to be more precise about the causal questions we are asking, thus enforcing conceptual consistency.* Rohrer et al. [47, p4]

Established participatory design tools empower non-technical stakeholders by using "low tech, high touch" materials or software. However, intentionally sacrificing formality in the process excludes those stakeholders from downstream modelling. Diagrammatic reasoning with causal models complements such techniques, seeking to sustain that participatory engagement by making transparent, and contestable, the causal assumptions that the computational/statistical modellers would otherwise be left to code alone. The DAG as a diagrammatic tool demystifies an abstract process by stripping away much of the notation and language of statistical modelling, and in doing so enables a greater diversity of researchers and practitioners to decide upon what LA gets created, and how the models are developed. It is important that no mention of regression or more advanced statistical techniques is required for the graphical model construction; no one involved in the discussion needs to know which variables to adjust for. Nonetheless, the graphical structure of the model provides a collaborative mechanism for people with strong theoretical foundations, but less technical expertise, that helps them to contribute their knowledge to the LA modelling process. This knowledge, translated by the DAG, provides very concrete implications for the analyst to use as they implement the resulting model. The framework represented by DAGs also brings with it an extended apparatus for reasoning causally about systems, and so delivers new actionable insights.

In summary, this paper has contributed a new perspective on the question of "Who decides what learning analytics get created and implemented?" We hope that this rich and rigorous technique appeals to the field, as its widescale adoption would help LA to start moving up the "ladder of causation" [42] from association and towards models that can make stronger causal claims about learning and the environments in which it occurs.

REFERENCES

[1] Carlos Prieto Alvarez, Roberto Martinez-Maldonado, and Buckingham Shum. 2020. LA-DECK: A card-based learning analytics co-design tool. In *Proceedings of the tenth international conference on learning analytics & knowledge*. 63–72.
[2] Solon Barocas and Andrew D Selbst. 2016. Big data's disparate impact. *California Law Review* 104 (2016), 671.
[3] Geoffrey C Bowker and Susan Leigh Star. 2000. *Sorting things out: Classification and its consequences*. MIT press.
[4] danah boyd and Kate Crawford. 2012. Critical questions for big data: Provocations for a cultural, technological, and scholarly phenomenon. *Information, communication & society* 15, 5 (2012), 662–679.
[5] Christopher Brooks, Omar Chavez, Jared Tritz, and Stephanie Teasley. 2015. Reducing selection bias in quasi-experimental educational studies. In *Proceedings of the fifth international conference on learning analytics and knowledge*. 295–299.
[6] Simon Buckingham Shum, Rebecca Ferguson, and Roberto Martinez-Maldonado. 2019. Human-centred learning analytics. *Journal of Learning Analytics* 6, 2 (2019), 1–9.
[7] Simon Buckingham Shum and Rosemary Luckin. 2019. Learning analytics and AI: Politics, pedagogy and practices. *British journal of educational technology* 50, 6 (2019), 2785–2793.
[8] Simon J Buckingham Shum. 2019. Critical data studies, abstraction and learning analytics: Editorial to Selwyn's LAK keynote and invited commentaries. *Journal of Learning Analytics* 6, 3 (2019), 5–10.
[9] Simon J. Buckingham Shum and Timothy A. McKay. 2018. Architecting for Learning Analytics: Innovating for Sustainable Impact. *EDUCAUSE Review* 53, 2 (2018), 25–37.
[10] Michel Callon, P Lascoumes, and Y Barthe. 2009. *Acting in an uncertain world. An essay on technological democracy*. Cambridge: MIT Press.
[11] Peter C-H Cheng, Ric K Lowe, and Mike Scaife. 2001. Cognitive science approaches to understanding diagrammatic representations. *Thinking with diagrams* (2001), 79–94.
[12] Kwok Tai Chui, Dennis Chun Lok Fung, Miltiadis D Lytras, and Tin Miu Lam. 2020. Predicting at-risk university students in a virtual learning environment via a machine learning algorithm. *Computers in Human Behavior* 107 (2020), 105584.
[13] Clayton Clemens, Vivekanandan Kumar, David Boulanger, Jérémie Seanosky, et al. 2018. Learning Traces, Competence Assessment, and Causal Inference for English Composition. In *Frontiers of cyberlearning*. Springer, 49–67.
[14] Doug Clow. 2012. The learning analytics cycle: closing the loop effectively. In *Proceedings of the 2nd international conference on learning analytics and knowledge*. 134–138.
[15] Shane Dawson, Jelena Jovanovic, Dragan Gašević, and Abelardo Pardo. 2017. From prediction to impact: Evaluation of a learning analytics retention program. In *Proceedings of the seventh international learning analytics & knowledge conference*. 474–478.
[16] Maria Puig De La Bellacasa. 2012. 'Nothing comes without its world': thinking with care. *The Sociological Review* 60, 2 (2012), 197–216.
[17] Alfred Essa. 2019. Is Data Dark? Lessons from Borges's "Funes the Memorius". *Journal of Learning Analytics* 6, 3 (2019), 35–42.
[18] Frank Fischer, Cindy E Hmelo-Silver, Susan R Goldman, and Peter Reimann. 2018. *International handbook of the learning sciences*. Routledge.
[19] Dragan Gašević, Vitomir Kovanović, and Srećko Joksimović. 2017. Piecing the learning analytics puzzle: A consolidated model of a field of research and practice. *Learning: Research and Practice* 3, 1 (2017), 63–78.

[20] Andrew P Gibson. 2017. *Reflective writing analytics and transepistemic abduction.* Ph.D. Dissertation. Queensland University of Technology. https://eprints.qut.edu.au/106952/
[21] David Hand. 2019. What is the Purpose of Statistical Modelling? *Harvard Data Science Review* 1, 1 (1 7 2019). https://hdsr.mitpress.mit.edu/pub/9qsbf3hz
[22] Dirk Ifenthaler and Jane Yin-Kim Yau. 2020. Utilising learning analytics to support study success in higher education: a systematic review. *Educational Technology Research and Development* 68, 4 (2020), 1961–1990.
[23] David Kember. 1995. *Open learning courses for adults: A model of student progress.* Educational Technology.
[24] David Kember, David Murphy, Irene Siaw, and KS Yuen. 1991. Towards a causal model of student progress in distance education: Research in Hong Kong. *American Journal of Distance Education* 5, 2 (1991), 3–15.
[25] Hassan Khosravi, Kirsty Kitto, and Joseph Jay Williams. 2019. RiPPLE: a crowdsourced adaptive platform for recommendation of learning activities. *Journal of Learning Analytics* 6, 3 (2019), 91–105.
[26] Kirsty Kitto and Simon Knight. 2019. Practical ethics for building learning analytics. *British Journal of Educational Technology* 50, 6 (2019), 2855–2870. Publisher: Wiley Online Library.
[27] Kirsty Kitto, Simon Buckingham Shum, and Andrew Gibson. 2018. Embracing imperfection in learning analytics. In *Proceedings of the 8th international conference on learning analytics and knowledge.* 451–460.
[28] Kirsty Kitto, John Whitmer, Aaron E. Silvers, and Michael Webb. 2000. *Creating Data for Learning Analytics Ecosystems.* Technical Report. Society for Learning Analytics Research (SoLAR). https://www.solaresearch.org/wp-content/uploads/2020/09/SoLAR_Position-Paper_2020_09.pdf
[29] Simon Knight, Simon Buckingham Shum, and Karen Littleton. 2013. Collaborative sensemaking in learning analytics. In *CSCW and Education Workshop (2013): Viewing education as a site of work practice, co-located with the 16th ACM Conference on Computer Support Cooperative Work and Social Computing (CSCW 2013).*
[30] Simon Knight, Simon Buckingham Shum, and Karen Littleton. 2014. Epistemology, assessment, pedagogy: where learning meets analytics in the middle space. *Journal of Learning Analytics* 1, 2 (2014), 23–47.
[31] Charles Lang, George Siemens, Alyssa Wise, and Dragan Gašević (Eds.). 2017. *The Handbook of Learning Analytics.* The society for learning analytics research. Available at: https://solaresearch.org/hla-17/.
[32] Kelly Linden. 2021. Improving Student Retention by Providing Targeted Support to University Students who do not Submit an Early Assessment Item. *Student Success* 12, 3 (Nov. 2021). https://doi.org/10.5204/ssj.2152
[33] Zahia Marzouk, Mladen Rakovic, Amna Liaqat, Jovita Vytasek, Donya Samadi, Jason Stewart-Alonso, Ilana Ram, Sonya Woloshen, Philip H Winne, and John C Nesbit. 2016. What if learning analytics were based on learning science? *Australasian Journal of Educational Technology* 32, 6 (2016).
[34] Shirin Mojarad, Alfred Essa, Shahin Mojarad, and Ryan S Baker. 2018. Studying adaptive learning efficacy using propensity score matching. In *Companion Proceedings of the 8th International Conference on Learning Analytics and Knowledge (LAK'18).* 5–9.
[35] Amin Mousavi, Matthew Schmidt, Vicki Squires, and Ken Wilson. [n.d.]. Assessing the Effectiveness of Student Advice Recommender Agent (SARA): the Case of Automated Personalized Feedback. *International Journal of Artificial Intelligence in Education* ([n. d.]).
[36] Peter Norvig. 2012. Colorless green ideas learn furiously: Chomsky and the two cultures of statistical learning. *Significance* 9, 4 (2012), 30–33.
[37] Alexandra Olteanu, Carlos Castillo, Fernando Diaz, and Emre Kıcıman. 2019. Social data: Biases, methodological pitfalls, and ethical boundaries. *Frontiers in Big Data* 2 (2019), 13.
[38] Cathy O'Neil. 2017. *Weapons of math destruction: How big data increases inequality and threatens democracy.* Broadway Books.
[39] Abelardo Pardo, Negin Mirriahi, Roberto Martinez-Maldonado, Jelena Jovanovic, Shane Dawson, and Dragan Gašević. 2016. Generating actionable predictive models of academic performance. In *Proceedings of the sixth international conference on learning analytics & knowledge.* 474–478.
[40] Frank Pasquale. 2015. *The black box society.* Harvard University Press.
[41] Judea Pearl. 2009. *Causality.* Cambridge University press.
[42] Judea Pearl and Dana Mackenzie. 2018. *The book of why: the new science of cause and effect.* Basic books.
[43] Matthew D Pistilli, Kimberly Arnold, and Matt Bethune. 2012. Signals: Using academic analytics to promote student success. *EDUCAUSE Review Online* 10, 8 (2012).
[44] Paul Prinsloo and Sharon Slade. 2017. An elephant in the learning analytics room: The obligation to act. In *Proceedings of the seventh international learning analytics & knowledge conference.* 46–55.
[45] Aldo Ramirez-Arellano, Elizabeth Acosta-Gonzaga, Juan Bory-Reyes, and Luis Manuel Hernandez-Simon. 2018. Factors affecting student learning performance: A causal model in higher blended education. *Journal of Computer Assisted Learning* 34, 6 (DEC 2018), 807–815.
[46] Julia M Rohrer and Ruben C Arslan. 2021. Precise answers to vague questions: Issues with interactions. *Advances in Methods and Practices in Psychological Science* 4, 2 (2021), 25152459211007368.
[47] Julia M Rohrer, Stefan Schmukle, and Richard McElreath. 2021. The Only Thing That Can Stop Bad Causal Inference Is Good Causal Inference. (2021).
[48] Carolyn Penstein Rosé. 2019. Monolith, Multiplicity, or Multivocality: What Do We Stand for and Where Do We Go from Here?. *Journal of Learning Analytics* 6, 3 (2019), 31–34.
[49] Carolyn P Rosé, Elizabeth A McLaughlin, Ran Liu, and Kenneth R Koedinger. 2019. Explanatory learner models: Why machine learning (alone) is not the answer. *British Journal of Educational Technology* 50, 6 (2019), 2943–2958.
[50] Donald B Rubin. 2005. Causal inference using potential outcomes: Design, modeling, decisions. *J. Amer. Statist. Assoc.* 100, 469 (2005), 322–331.
[51] Otto Sahlgren. 2021. The politics and reciprocal (re)configuration of accountability and fairness in data-driven education. *Learning, Media and Technology* 0, 0 (2021), 1–14. https://doi.org/10.1080/17439884.2021.1986065
[52] Mike Scaife and Yvonne Rogers. 1996. External cognition: how do graphical representations work? *International journal of human-computer studies* 45, 2 (1996), 185–213.
[53] Neil Selwyn. 2020. Re-imagining 'Learning Analytics'... a case for starting again? *The Internet and Higher Education* 46 (2020), 100745.
[54] David Williamson Shaffer, Wesley Collier, and Andrew R Ruis. 2016. A tutorial on epistemic network analysis: Analyzing the structure of connections in cognitive, social, and interaction data. *Journal of Learning Analytics* 3, 3 (2016), 9–45.
[55] Galit Shmueli. 2010. To explain or to predict? *Statistical science* 25, 3 (2010), 289–310.
[56] Simon Buckingham Shum and Nick Hammond. 1994. Delivering HCI modelling to designers: A framework and case study of cognitive modelling. *Interacting With Computers* 6, 3 (1994), 314–341.
[57] Joanna Sikora, MDR Evans, and Jonathan Kelley. 2019. Scholarly culture: How books in adolescence enhance adult literacy, numeracy and technology skills in 31 societies. *Social science research* 77 (2019), 1–15.
[58] Dan Suthers and Katrien Verbert. 2013. Learning analytics as a" middle space". In *Proceedings of the Third International Conference on Learning Analytics and Knowledge.* 1–4.
[59] Vincent Tinto. 1975. Dropout from higher education: A theoretical synthesis of recent research. *Review of educational research* 45, 1 (1975), 89–125.
[60] Carole J Torgerson and David J Torgerson. 2001. The need for randomised controlled trials in educational research. *British journal of educational studies* 49, 3 (2001), 316–328.
[61] Yi-Shan Tsai, Pedro Manuel Moreno-Marcos, Ioana Jivet, Maren Scheffel, Kairit Tammets, Kaire Kollom, and Dragan Gašević. 2018. The SHEILA framework: Informing institutional strategies and policy processes of learning analytics. *Journal of Learning Analytics* 5, 3 (2018), 5–20.
[62] Thom van Dooren. 2018. Thinking with Crows:(Re) doing philosophy in the field. *Parallax* 24, 4 (2018), 439–448.
[63] Alyssa Wise, Simon Knight, and Simon Buckingham Shum. 2021. Collaborative Learning Analytics. *International Handbook of Computer-Supported Collaborative Learning* (2021). In Press.
[64] Alyssa Friend Wise, Juan Pablo Sarmiento, and Maurice Boothe Jr. 2021. *Subversive Learning Analytics.* Association for Computing Machinery, New York, NY, USA, 639–645. https://doi.org/10.1145/3448139.3448210
[65] Jacqueline Wong, Martine Baars, Björn B de Koning, Tim van der Zee, Dan Davis, Mohammad Khalil, Geert-Jan Houben, and Fred Paas. 2019. Educational theories and learning analytics: From data to knowledge. In *Utilizing learning analytics to support study success.* Springer, 3–25.
[66] Alan Woodley, Paul de Lange, and George Tanewski. 2001. Student progress in distance education: Kember's model re-visited. *Open Learning: The Journal of Open, Distance and e-Learning* 16, 2 (2001), 113–131.