# Development of a Directed Acyclic Graph through expert elicitation for Causal Inference in Built Environment Research.

Minnie Ashdown[1], Giorgos Petrou[1]
[1]UCL Institute of Environmental Design and Engineering, London, United Kingdom

## Abstract

System complexity within buildings research is high. Data are increasingly high dimensional with multiple physical and socio-technical processes driving their generation. Popular statistical and building simulation methods lack a structured procedure for incorporating intuition, limiting their potential for effect attribution in the study of buildings. This can be addressed through the use of a Directed Acyclic Graph (DAG); a graphical representation of expert knowledge and assumptions, through which interventions can be modelled, and causal assumptions are made explicit. This paper introduces the process of DAG development through expert elicitation, offering guidance to built environment researchers interested in causal modelling.

## Key Innovations

- The value of a DAG to a Building Simulation practitioner is discussed.
- Guidance in developing graphical models through expert elicitation is provided.
- A set of learnings from a case study application are discussed.

## Practical Implications

A set of recommendations are provided that can support researchers to develop DAGs that can be used to strengthen modelling decisions in Built Environment research, increasing model transparency and the strength of conclusions.

## Introduction

Improving the energy efficiency of buildings is a crucial step in global efforts to reduce Greenhouse Gas (GHG) emissions (IEA, 2021). In the United Kingdom (UK), according to the government's Heat and Building Strategy, a 'fabric-first' approach will prioritise building envelope improvements in existing homes that fall below the government's energy standard (HMG, 2021).

The improvement in home energy efficiency (HEE) offers many co-benefits in addition to the reduction in GHG emissions, such as warmer homes during the winter, especially for the fuel poor. However, concerns exist regarding the potential for unintended consequences of HEE on indoor air quality and the risk of indoor overheating (Petrou et al., 2022).

With ambitious plans to improve the fabric thermal efficiency of homes, research efforts aimed at understanding the effect of HEE on indoor overheating have intensified over the last decade. Modelling studies have suggested that HEE measures can both increase and decrease overheating risk, depending on the type of measure and its interaction with other factors, such as ventilation and solar gains (Fosas et al., 2018; Mavrogianni et al., 2012). Empirical measurements of indoor temperature from a pair of identical houses with synthetic occupancy revealed that the installation of internal wall insulation can increase summer indoor temperature, and this can be partially mitigated by the use of night ventilation and shading (Tink et al., 2018).

The real world performance of the English housing stock was studied by Petrou et al. (2019) using measurements from a sample of 800 homes monitored during the summer of 2011, and later by Lomas et al. (2021) based on data collected in 750 homes during the summer of 2018. The studies differed in their choice of dependent variable and statistical tests, but took a similar approach of sequentially assessing whether a statistically significant relationship exists between the dependent variable and a set of dwelling and household characteristics. While both studies resulted in several important findings, it is possible that un-studied factors associated with both dependent and independent variables, or *confounders*, were present. Confounders may distort or mask the effect of interest, potentially misrepresenting the underlying process and leading to unwarranted conclusions.

## Effect attribution in Building Science

Applied building scientists are often concerned with understanding the effect of an *intervention*, such as energy efficiency retrofit, or the widespread adoption of heat pumps, on desired *outcomes*, such as reduced energy demand and greenhouse gas (GHG) emissions, better internal air quality, or improvements to the well-being and productivity of occupants. To understand the mechanisms that connect two or more variables is to characterise a causal relationship, and albeit not necessarily explicit in their causal attribution, there are several means to address such questions in building science.


https://doi.org/10.26868/25222708.2023.1588

Popular approaches that rely on the use of building performance simulation tools to evaluate such mechanisms and their outcomes require detailed modelling of heat and mass transfer processes. However, computer simulations are not always able to capture system complexity, and can be time-consuming and expensive, in terms of research time, necessary skill level, and computational resources. Where parametric or optimisation methods are used, parameter dependencies are often not adequately considered, and the design space may be ambiguously defined (Chen et al., 2022). Furthermore, the knowledge and expertise needed to develop a convincing simulation often imposes strong assumptions on the model, with associated restrictions on the outputs. These factors contribute to, but are not uniquely responsible for, the performance gap between the predicted and observed effect of an intervention, a common phenomenon in building science research (de Wilde, 2014).

Another approach is based on the use of observational data, both to develop data-driven machine learning models, and to determine whether a statistically significant difference exists between the quantity of interest for groups of dwellings with and without an intervention (Zhuravchak et al., 2022). A crucial limitation with this commonly used approach is that as discussed, *confounding* can result in the inaccurate attribution of effect.

Across disciplines, a widely accepted method to determine the effect of an intervention is a Randomised Control Trial (RCT), where randomisation directly addresses the problem of confounding and causality. However, the many challenges associated with implementation of RCTs in built environment research, including ethical concerns and high implementation costs, have resulted in the scarce use of this gold standard approach within the field.

Causal Inference techniques offer an alternative; the opportunity to bring together statistical analysis with domain expertise and intuition about a particular phenomenon to address the mechanistic relationships between variables. They offer a natural means to bring together classical physics and Bayesian methods, address the problem of confounding, and go beyond the study of association to quantify the causal effect of an intervention using observational data, graphical models and do-calculus (Pearl, 2009). A Directed Acyclic Graph (DAG), developed on the basis of the understanding of multiple experts, may be able to provide more weight to the conclusions drawn in analysis of correlations in large data sets by highlighting the theoretical underpinnings of model choices, and provide a transparent representation of modelling assumptions that are straightforward to critique.

## Causal Inference

Causal Inference is a logical strategy popularised and given mathematical rigour by Judea Pearl (Pearl, 2009), along with the statisticians and mathematicians such as Granger (1969), and Goldberger (1972). Although well established in fields such as epidemiology, econometrics, and social science, causal inference is seldom utilised explicitly in building science (with few exceptions, for example the work of Chen et al. (2022)). It exists at the boundaries of intuition, statistics, and probabilistic reasoning, and attempts to make use of the combined force of these three tenets to answer questions of *why* things happen, which constitutes a much stronger epistemological claim than simply stating a correlation (association). Broadly, it is a strategy with which to reason about the effect of certain interventions or counterfactual queries, exploiting the notion of conditional probability coupled with intuition to develop a theory of causation.

Pearl (2009) articulates what is called a general theory of causation in the form of a Structural Causal Model (SCM). An SCM is a representation of a system in terms of a set of variables, each of which may influence one or more other variables. The relationships between the variables are functionally represented, and this system is depicted in the form of a DAG (see Figure 1). A useful characteristic of this framework is that it is not always necessary to commit to a particular functional form of a variable, in order to exploit the invariant characteristics of the structural equations. This means that causal effects can still be estimated in both non-parametric and non-linear models.

Despite the fact that the data available to building science are growing, and indeed there are very large, rich data sets, the data alone are unable to capture the entire system behaviour. They are still a partial depiction of a highly complex and dynamic system, whose behaviour neither statistical nor physical modelling is able to completely capture. A purely statistical model, especially in highly complex systems, may not be robust to changes in the behaviour of the underlying system, and a purely physical model imposes hard constraints on the behaviour represented, restricted by the knowledge of the modeller (Berliner, 2003). Causal Inference provides a systematic way to integrate knowledge of the data generating process into models, without imposing hard, potentially limiting constraints on the components of the model.

Within the field of built environment research, the use of causal inference techniques and DAGs is scarce despite the many potential applications, including: the development of building performance simulation models that better-incorporate diverse expert input; achieving deeper understanding of variables leading to the generation of energy and indoor environmental data; and the integration of models from domains that may not straightforwardly be linked together, for example, a model for occupant behaviour may be integrated with a building physics model. In a notable recent publication, Chen et al. (2022) proposed the use of causal inference, together with building simulations tools, to inform energy efficient building design. The authors introduced a four-step process that relies on an automated approach to causal structure finding, to enable the integration of domain knowledge and data-driven techniques, resulting in a more efficient design process.

Another application of causal inference is the study of causal effects based on large observational data sets. To


https://doi.org/10.26868/25222708.2023.1588

explore this under-researched area within the context of Built Environment research, and attempt to answer a pertinent question in the field of indoor overheating, the CIMBER project (Pilot study on the application of Causal Inference Methods in complex Built Environment Research) aims to investigate the causal effect of home energy efficiency on summer indoor temperature in the English housing stock. The starting point for this analysis is the development of a DAG. DAGs may be developed through expert knowledge of a system, or they may be learned from the data. While DAG learning is a vast and expanding field in computer science, which is beyond the scope of this paper , this conference paper specifically focuses on the former approach. In complex problem domains like building science, it is highly valuable to employ methods that effectively harness the collective knowledge of multiple experts for model development. This expertise is not available from the data alone, and this approach therefore has the potential to significantly enhance coherence and transparency, by leveraging the collective expertise available. The resulting model will provide a rigorous basis from which to develop inferential analysis, giving additional weight to future machine learning and building simulation analyses.

This conference paper describes the process of developing a DAG appropriate for this pilot study through expert elicitation. Reflections from this process can inform and accelerate future DAG development within the field.

## Causal modelling with graphical heuristics

A DAG is a type of causal diagram, which is a qualitative, parsimonious summary of the process that generates the data (Cinelli et al., 2022). DAGs are composed of *nodes* (random variables), and *directed edges* (arrows) that denote causal paths between variables. An interpretation of these edges offered by Pearl (2022), when for example directed from variable $X$ to variable $Y$, is that they represent the answer to the question: what are the *sources* of variation we observe in variable $Y$? These paths are necessarily directed, because they denote the direction of causation, and are acyclic, to denote the topological ordering of the variables in the graph. The acyclic property thus means that there can be no feedback loops between variables.

As discussed above, a *confounding* variable is one that has an effect on both the independent and dependent variables, and if one wishes to isolate the effect of the independent on the dependent variable, their effect must be controlled for. A graphical criterion that allows the researcher to eliminate the effect of confounding is known as the backdoor criterion (Pearl, 2009), which supports the blocking of spurious (non causal) paths between treatment and outcome. Judicious application of this criterion further means that it is possible, assuming the DAG is correct, to assert the causal effect of an intervention on the outcome, or investigate counterfactual queries. A simple DAG is illustrated as follows (see Figure 1). A direct path exists between $A$ and $C$, represented in text as $A \rightarrow C$. The direction of the arrow denotes the direction of causality. From the figure, there exists an additional open (backdoor) path between $A$ and $C$, via $B$, which is a common cause of both $A$ and $B$. Any effect on $C$ due to $A$ will also contain information about $B$, which is confounding the relationship between $A$ and $C$. Controlling for $B$ is this situation is necessary to obtain an unbiased estimate of the Average Causal Effect (ACE) of $A$ on $C$ (Cinelli et al., 2022). An overview of the implications of different model structures, and the biases that may be implied is available in Cinelli et al. (2022).

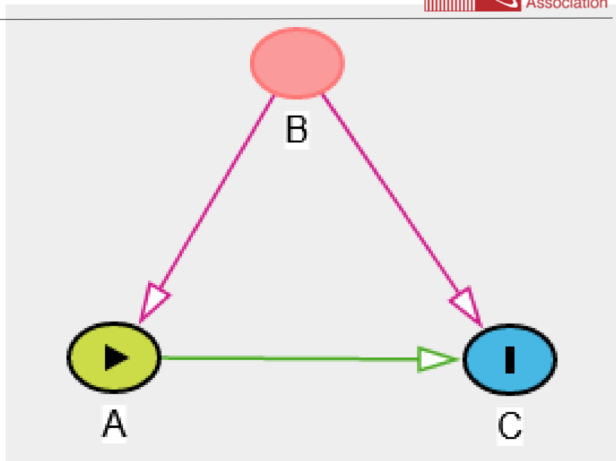

*Figure 1: A simple Directed Acyclic Graph, with two paths between Exposure A and Outcome C, produced using DAGitty.*

In the context of building science, there may be a direct path between indoor temperature ($C$) and glazing fraction ($A$), however this relationship may be confounded by building orientation ($B$). As such, a DAG provides a static visual summary of the model assumptions and the variables that need to be measured in order to be controlled (Greenland et al., 1999).

This provides a flexible counterpart to regression models, which rely on strong assumptions about the relationships between variables and their distributions, and may not represent the underlying causal relationships in a system, in particular their direction. Through the use of DAGs, Causal Inference allows researchers to explicitly model their beliefs about underlying causal relationships, and conditional dependence structure within models without necessarily requiring knowledge of the functional form of relationships between variables or their distributions.

## Eliciting a causal model for Residential Indoor Temperatures

To develop a DAG for the causal effect of home energy efficiency on summer indoor temperature, a process of expert elicitation was used as shown in Figure 2. This process is discussed in the following section.

Expert elicitation is a commonly used research tool, especially in contexts where there exists significant epistemic uncertainty due to a lack of data, or the data are challenging or expensive to collect. There are thus precedents in elicitation techniques across multiple disciplines. Kuhnert et al. (2010) provide guidance in the context of


https://doi.org/10.26868/25222708.2023.1588

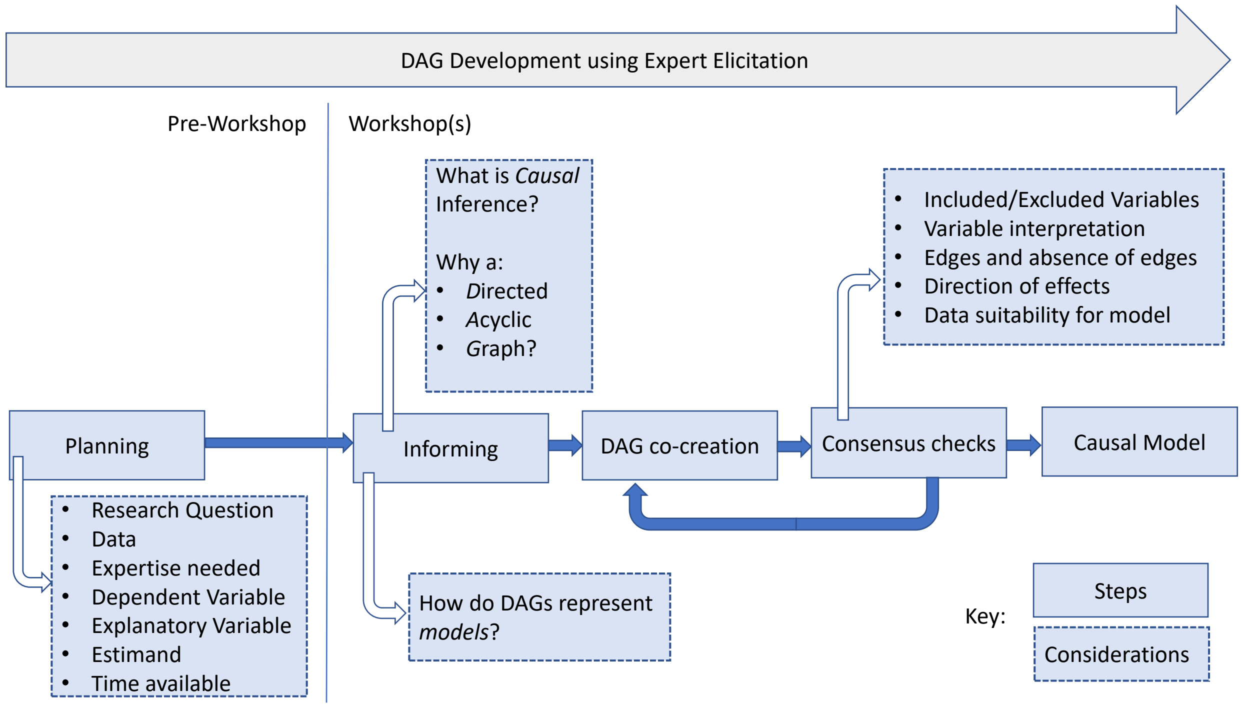


*Figure 2: Schematic view of the iterative process of DAG development using Expert Elicitation.*

Bayesian modelling of ecological models, arguing that expert knowledge should be utilised in models with the same degree of rigour as the gathering and subsequent use as empirical data. For this reason, this paper pays close attention to this process, aiming to aid researchers in developing this rigour for the complex field of building simulation.

Rodrigues et al. (2022) describe the process of DAG development in health services using a workshop composed of representatives of multiple stakeholder groups. The workshop objective was to construct a DAG for an intervention, made by a healthcare professional in the National Health Service in England, and the journal article offers guidance to practitioners interested in using a workshop to elicit information from domain experts. Noting that the problem space is substantially different, nevertheless, a similar workshop format was followed for our DAG development.

**Planning**

As highlighted in Figure 2, the first step in the process is planning, with particular focus on the following elements: developing the Research Question and estimand, along with the explanatory and dependent variables, as well as determining the data requirements, time, and necessary expertise to develop the model.

The datasets intended to be used for this causal analysis are the linked 2011 English Housing Survey (EHS) and Energy Follow-Up Survey (EFUS). The same sets of data were used in the analysis by Petrou et al. (2019). The 2011 EFUS contains sub-hourly measurements of indoor temperature from approximately 800 English homes (Hulme et al., 2013). The monitored homes are a subset of the EHS, a national survey that takes place every two years and consists of household interviews and physical surveys (DLUHC, 2021). It is useful to consider the available data at this stage, because features (in this context) such as the temporal and spatial resolution of the dependent variable can influence the model structure. This helps participants to make decisions about which variables to include in the model and how they are related.

Initially, a single three-hour workshop was planned. However, it was not feasible to finalise a model within the three-hour time frame and a follow-up was arranged.

The initial workshop was composed of three sessions: informing, breakout group DAG co-creation, and whole group consensus checking and evaluation in the context of available data. The second workshop was composed of two sessions, the first focussing on reminding participants of the premise and progress of the project, and the second about whole group consensus.

Participants were composed of department members with expertise suited to input into the problem, along with two facilitators. The intervention of interest was the effect of energy efficiency on summer indoor temperatures in UK residential building stock, hence invitees were composed of eleven academic experts in building physics, indoor overheating and air quality, climatology, and urban health, all based within University College London. Although the majority of attendees were able to join in person, accommodation was also made for those who were available online only. Prior to the workshop, the workshop agenda, and information about the research project was shared with prospective participants, and low risk ethics approval was obtained.

Prior to the second workshop taking place, feedback was


https://doi.org/10.26868/25222708.2023.1588

sought from both experts in causal inference, and in using workshops for expert elicitation. It was suggested that important factors to account for in determining the approach are time constraints, the pre-existence of models, the purpose and scope of the model, and the desire to compare an existing model to one that is developed by stakeholders. The questions developed were thought to be appropriate for the task, if ambitious, for the time allowed.

**Informing**

The initial workshop opened with a short overview of the project, and an introduction to the basic principles and rationale for Causal Inference. Following this was an introduction to DAG theory, such that participants had sufficient knowledge and vocabulary to engage with DAG co-creation. A simple synthetic example of confounding was demonstrated to highlight how DAG structure leads to the causal effect estimates through the addition of controls (based on Figure 1(b) of Cinelli et al. (2022)). The intended message for this exercise was to show that assuming the DAG structure is correct, the issue of confounding can be addressed. However, if the DAG is misspecified, it could also lead to incorrect or biased estimates.

There was also a brief overview of the temperature time series data available for each dwelling, its distributional characteristics, and the time period it covered. However, following the recommendation of Rodrigues et al. (2022), the participants were not provided with detailed information about the available household level data until the "Consensus checks" stage, in order to avoid undue focus on available, rather than causally relevant, variables in the initial model development.

**DAG co-creation**

Following the informing segment of the initial workshop, the participants were divided into two smaller groups for co-creation. The subgroups were determined in advance such that they were composed of participants from a range of disciplines, and with an even distribution of seniority in each group.

Before each co-creation session, participants were provided with a task and some initial content questions, as shown in Table 1. Participants were also prompted to ensure they were clear about the activity. During the first session, participants in each group were able to freely develop their DAGs, with the support of a facilitator who would record variables and relationships on DAGitty, while also noting key discussion points. The facilitator was not there to engage in discussion, other than to support the recording of the DAG. DAGitty (Textor et al., 2016) is a piece of software developed by researchers for the automated analysis of DAGs, facilitating quick and easy testing of whether a set of variables can be used to control for confounding in a causal relationship, and visualisation of the probabilistic relationships between variables.

**Consensus checks**

Following the DAG co-creation section of the first workshop, and during the second workshop, the primary objective was to develop the DAG as a whole group and obtain consensus on its structure and content. There was therefore a group discussion to compare the approach taken by each group, with a focus on drawing out differences in modelling assumptions and the key focus between groups. This discussion was guided using prompts contained in Table 1.

A secondary objective of this step was to determine the DAG variables that data exists for within the EHS-EFUS surveys, other relevant datasets and any potential data gaps. Participants were asked to keep in mind that the model structure is what justifies the causal claims we wish to make, and that the strength of the analysis lies in the validity of the model according to their expertise.

## Workshop reflections

Attempting to develop a causal model for the effect of energy efficiency on summer indoor temperatures has been instructive. Firstly, and perhaps most obviously, it has had the effect of exposing the complexity of the issue to a level that is beyond what one might expect. Highlighting the multiple pathways and the conditional relationship between variables revealed disagreements about the direction of effects and their relative importance, especially at different spatio-temporal scales of analysis. This has implications for analysing the EHS-EFUS survey data and the relative importance of different controls that have been applied in the past. In addition, this exercise offers a set of learnings applicable to the broader application of Causal Inference methods in Built Environment research. These are detailed in the following sections.

**Timing**

The time required to elicit expert input and develop a DAG will largely depend on the familiarity of the experts with Causal Inference methods, and the complexity of the system being studied. While system complexity will vary, it is expected to remain high across most areas of built environment research. For the case study described in this paper, two workshops of two and three hours long were inadequate to finalise the DAG model. This highlights how difficult it is to conceptualise a model of this nature, not least because experts will have very deep understandings of the problem space, which leads to the addition of a lot of details and the consideration of multiple relationships. Underestimating the time needed to complete the DAG development was also noted as an issue by Rodrigues et al. (2022).

While Koller (2009) noted that the time of experts needed to construct the model can sometimes be considered too valuable, participant feedback revealed a preference for a longer time period to develop the DAG. Thus, a recommendation of at least two four-hour sessions is made, with a gap in-between to allow researchers time to reflect and consolidate information.

**Participants**

As expected, experts in a specific area provided a more nuanced picture of their domain, which validates the ap-


https://doi.org/10.26868/25222708.2023.1588

| Rationale | Prompt |
| --- | --- |
| DAG co-creation | Is everyone clear about the intervention we are interested in?<br>Develop a DAG with key variables and direction of effect.<br>What are the key factors driving the assignment of the intervention?<br>Can we coherently combine our confounds into fewer variables? |
| Consensus checks | How confident are you that $x$ and $y$ are causally/not causally related?<br>Do you think we have missed any important parameters or arrows?<br>What assumptions/knowledge underlie the causal relationships?<br>Does the DAG represent the model structure, and do you have any observations? |

*Table 1: Selected prompts used to elicit expert knowledge and develop consensus during DAG co-creation.*

proach of having experts from several relevant disciplines. It is interesting to note that despite efforts to ensure a similar distribution in expertise between the subgroups, the initial models produced were quite different. This could potentially be due to the starting point chosen by each subgroup or the presence of authoritative voices. It could also be the case that experts drawn from the same discipline and institutional department can have different mental models for the same system.

While the use of subgroups of mixed expertise in this case study was preferred, an alternative would be to create subgroups focused on each domain of expertise. This is expected to result in the development of different DAGs, biased towards particular model framing, that would then be merged during the whole-group discussion.

**Background theory**

It is essential to provide participants with an introduction to the key aims of Causal Inference, and especially the key components and principles of a DAG. Working through an example, such as the one provided in Figure 1, can be especially helpful to illustrate the difference between correlation and causality.

The theory of causal graphical models needs to be communicated carefully to participants in advance of developing the DAG, including key vocabulary and the representation of relationships in the casual graph. In particular, participants need to be aware that arrows denote a causal direction, and that a bidirectional arrow indicates a latent variable (Pearl, 2009).

It is also helpful to communicate the use of the backdoor criterion, such that participants are able to follow the implications of their modelling choices on the sufficient sets implied by the model if using DAGitty. This can have important consequences for the data and analysis that is available once the DAG is developed.

Next, it is important to highlight to participants that it is the absence, not the presence, of an arrow that encodes the model assumptions. Hence, the lack of an arrow is a much stronger statement than the presence of one, since this asserts the conditional independence of the variables (Pearl, 2009). In the absence of a causal model, all variables are assumed to be related (and there would therefore be arrows between all variables).

Lastly, it is important that participants understand that the DAG cannot contain cycles, and to acknowledge that this imposes temporal limitations, as well as having consequences for the structure of variables and their content. Some participants may find this too restrictive, and depending on the level of mathematical expertise it may be helpful to discuss the possibility of state space modelling, or Dynamic Bayesian Networks to overcome the temporal limitations of a static DAG (for example, see Shiguihara et al. (2021)).

**Elicitation process**

Following the advice of Rodrigues et al. (2022), DAGitty was used to record participant inputs to the DAG, and notes were taken to record key discussion points. DAGitty is an appropriate choice since it is a straightforward tool, readily understood by participants. It is also helpful because after the workshop, the model can be shared in text format among interested parties, who can then load it themselves and reflect on the model structure and its implications.

Despite the ease of use of DAGitty, it would have been helpful to have a note-taker, as well as someone who is able to facilitate the DAG development, because it is challenging in a free-flowing discussion context to record participant comments as well as ensure that the DAG is accurately specified.

During the second workshop, participants made some useful observations about the communication of the model. Due to its overall complexity, it may be helpful when reminding participants of its contents to break it down into smaller components, which can then be discussed individually and built upon. This would certainly help with clarity of communication and allow participants to incrementally digest the model, rather than all at once.

**DAG Amalgamation**

Key to the process of developing a causal graphical model is knowledge of the variables in the system of interest. This is why the use of experts in the development phase is appropriate, and can lead to confidence in model outputs and consolidation of variable interpretation. It also leads to robust discussion of the effect of interest and exposes the different points of view held between different disciplines.

The iterative process of DAG co-development and consensus checking (highlighted in Figure 2), is crucial to capturing participant knowledge. To facilitate this process and guide participants in responding to the prompts


https://doi.org/10.26868/25222708.2023.1588

in Table 1, a dedicated worksheet was developed. The worksheet contained questions designed to assist participants in assessing their causal assumptions, and the most and least significant variables in generating the data features. It also asked participants whether there were any variables that could be combined into a single indicator, whether they had any observations to add, and whether in their view the DAG accurately represented the model. It was suggested that it may be helpful in a future workshop to get different groups to look at different levels of the problem. This might have the benefit that in ensuring the levels of the model are coherent, high level group structure may emerge, which may be exploitable. In particular, one approach in this context could be to look at building variables, occupant characteristics, and variables that are external to the building.

## Conclusion

Causal Inference techniques, including the use of Directed Acyclic Graphs (DAGs), have the potential to improve current research practices, both in data-driven and building simulation applications. This paper describes the process of developing a DAG using expert elicitation, and reflects on its application in a case study on the causal effect of home energy efficiency on summer indoor temperature. The simplicity of the concept of DAGs means that it is possible to consider structural model effects without the need to have specialised knowledge of statistical methods. This enables researchers who may have highly valuable expertise in one domain to consider the interaction of the system's various components, and co-develop a holistic graphical model, without the need for a background in statistics or machine learning. The process of developing a DAG for this system has revealed several points of interest to others. These are that the time needed to develop these models may be large, yet researchers felt it was a valuable exercise and a good use of their time. Next, that it is necessary to deliver DAG theory to researchers with care, such that they are able to appreciate how the DAG structure leads to the addition of controlling variables and the ability to assert causal, rather than associational, effects. For a situation as complex as an occupied building, the DAG quickly becomes highly complicated. In the context of a workshop, communication of the model should be done gradually such that participants have the time to absorb the graph properties. This is useful to know and can guide further research, as well as supporting effective communication in the modelling community and more widely, building science researchers.

## Acknowledgment

This work is an output of the project titled Pilot study on the application of causal inference methods in complex built environment research (CIMBER), supported financially by the CBES Platform Grant [EPSRC grant number is EP/P022405/1]. This research was also funded in part, by the Wellcome Trust [Grant number 209387/Z/17/Z].

## References

Berliner, L. M. (2003). Physical-statistical modeling in geophysics. *Journal of Geophysical Research - Atmospheres 108*, 8776–n/a.

Chen, X., J. Abualdenien, M. M. Singh, A. Borrmann, and P. Geyer (2022). Introducing causal inference in the energy-efficient building design process. *Energy and Buildings 277*, 112583.

Cinelli, C., A. Forney, and J. Pearl (2022). A crash course in good and bad controls. *Sociological methods research*, 4912412210995–.

de Wilde, P. (2014, May). The gap between predicted and measured energy performance of buildings: A framework for investigation. *Automation in Construction 41*, 40–49.

DLUHC (2021). English Housing Survey. https://www.gov.uk/government/collections/english-housing-survey.

Fosas, D., D. A. Coley, S. Natarajan, M. Herrera, M. Fosas de Pando, and A. P. Ramallo-González (2018). Mitigation versus adaptation: Does insulating dwellings increase overheating risk? - ScienceDirect. *Building and Environment 143*, 740–759.

Goldberger, A. S. (1972). Structural equation methods in the social sciences. *Econometrica 40*, 979–1001.

Granger, C. W. J. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica 37*, 424–438.

Greenland, S., J. Pearl, and J. M. Robins (1999). Causal diagrams for epidemiologic research. *Epidemiology 10*.

HM Government (2021, October). *Heat and Buildings Strategy*.

Building Research Establishment on behalf of the Department of Energy and Climate Change (2013). *Energy Follow-Up Survey 2011, Report 11: Methodology*.

International Energy Agency (2021, October). *Net Zero by 2050 - A Roadmap for the Global Energy Sector*.

Koller, D. (2009). *Probabilistic graphical models : principles and techniques*. The MIT Press.

Kuhnert, P. M., T. G. Martin, and S. P. Griffiths (2010, 7). A guide to eliciting and using expert knowledge in bayesian ecological models. *Ecology Letters 13*, 900–914. https://doi.org/10.1111/j.1461-0248.2010.01477.x.

Lomas, K. J., S. Watson, D. Allinson, A. Fateh, A. Beaumont, J. Allen, H. Foster, and H. Garrett (2021, August). Dwelling and household characteristics' influence on reported and measured summertime overheating: A glimpse of a mild climate in the 2050's. *Building and Environment 201*, 107986.

Mavrogianni, A., P. Wilkinson, M. Davies, P. Biddulph, and E. Oikonomou (2012, September). Building characteristics as determinants of propensity to high indoor summer temperatures in London dwellings. *Building and Environment 55*, 117–130.

Pearl, J. (2009). *Causality* (2 ed.). Cambridge University Press.

Pearl, J. (2022, 1). Causation and decision: On dawid's "decision theoretic foundation of statistical causality". *Journal of Causal Inference 10*, 221–226.

Petrou, G., E. Hutchinson, A. Mavrogianni, J. Milner, H. Macintyre, R. Phalkey, S.-C. Hsu, P. Symonds, M. Davies, and P. Wilkinson (2022, May). Home energy efficiency under net zero: Time to monitor UK indoor air. *BMJ*, e069435.

Petrou, G., P. Symonds, A. Mavrogianni, A. Mylona, and M. Davies (2019, July). The summer indoor temperatures of the English housing stock: Exploring the influence of dwelling and household characteristics. *Building Services Engineering Research and Technology 40*(4), 492–511.

Rodrigues, D., N. Kreif, A. Lawrence-Jones, M. Barahona, and E. Mayer (2022, August). Reflection on modern methods: Constructing directed acyclic graphs (DAGs) with domain experts for health services research. *International Journal of Epidemiology 51*(4), 1339–1348.

Shiguihara, P., A. D. A. Lopes, and D. Mauricio (2021). Dynamic bayesian network modeling, learning, and inference: A survey. *IEEE Access 9*, 117639–117648.

Textor, J., B. van der Zander, M. S. Gilthorpe, M. Liśkiewicz, and G. T. H. Ellison (2016, 12). Robust causal inference using directed acyclic graphs: the r package 'dagitty'. *International Journal of Epidemiology 45*, 1887–1894.

Tink, V., S. Porritt, D. Allinson, and D. Loveday (2018, August). Measuring and mitigating overheating risk in solid wall dwellings retrofitted with internal wall insulation. *Building and Environment 141*, 247–261.

Zhuravchak, R., N. Nord, and H. Brattebø (2022). The effect of building attributes on the energy performance at a scale: an inferential analysis. *Building Research and Information 50*, 662–680.


https://doi.org/10.26868/25222708.2023.1588