



Constructing Belief: Using Bayesian Belief Networks to Measure and Manage Uncertainty in Digital Design

Daniel Hambleton^(✉) and Elissa Ross

MESH Consultants Inc., Toronto, ON M5V1E7, Canada
info@meshconsultants.ca

Abstract. A Bayesian Belief Network is a diagrammatic way to reason probabilistically and understand causal inference in complex systems. We propose using Bayesian Belief Networks (BBN) in the early stages of design projects to highlight components with high risk of failure. Identifying these components of high risk can inform how resources should be best used on costly modelling tasks. In addition, high risk components may impose functional modelling requirements, which in turn will inform the design of flexible systems for critical areas. This approach has the potential to significantly reduce risk by focusing and informing modelling efforts, which in turn increases the chance of success of the project and lowers costs for all stakeholders involved.

Using a prototype software application developed to quickly create BBNs and calculate a final probability value of a specific outcome (the “work product”), we test different project scenarios collected through three interviews with industry professionals. In each case, we identify an aspect of the project that changed during the course of the project with far reaching implications. By adjusting the values and structure of these networks we formulate specific functional requirements for digital models and in some cases, the associated construction systems. We find that these requirements would have increased the overall value of their respective projects by directly addressing the areas of strong influence and uncertainty identified in the BBN.

Keywords: Bayesian Networks · Architecture · BIM · Management · Modelling

1 Introduction

The design and construction industry today has access to an unprecedented amount of digital technology. From artificial intelligence based simulation engines to augmented reality helmets for construction workers, the range and depth of technological innovation in digital design is astounding (Umetani and Bickel 2018). The availability of this technology combined with a growing shift in the makeup of a professional design office towards staff with software engineering skills, both formal and self-taught, means that virtually anything can be modelled, documented, published and shared. Yet all design concepts that are intended to be realized as physical objects (buildings, consumer products) must be ushered through a series of stages involving design iteration,

expert consultation, value engineering, quantity estimation, shipping, manufacturing, to name a few. In contrast to the automotive or aerospace industries, the architecture, engineering, and construction (AEC) industry is confounded by the prototypical nature of their work product: rarely are two buildings exactly alike. Without the ability to incrementally improve the production process of any individual building project, stakeholders face many difficult decisions that, especially in the early stages of concept design, can have far reaching and often unexpected implications.

The pressure to reduce the amount of uncertainty during the initial phases of an architecture project is enormous. This challenge is addressed by vast amounts of modelling and documentation work (Building Information Modelling (BIM)) with the goal of creating a complete picture of every aspect of a project. Although daunting in its complexity, BIM has proved itself to be a positive and unifying force in the AEC industry. What remains difficult, however, is how to identify what to model, when to model it, and in how much detail. In the initial stages of design, the need to iterate quickly and test many design alternatives is at odds with the pressure to understand the implications of each of those iterations. In this paper, we propose Bayesian Belief Networks (BBN) as a lightweight but complete framework for directly engaging all project stakeholders during the design process.

A Bayesian Belief Network is a diagrammatic way to reason probabilistically and understand causal connections between different factors, models or stages of a project. Through a series of case studies with building professionals, we find that Bayesian Belief Networks are a useful tool for highlighting critical stages of the design process. They provide a systematic approach to project level optimization without detailed models or documentation. The use of Bayesian Networks to represent the design process may decrease the cost of modelling by focusing on areas of strong influence over the success or failure of the project. By highlighting areas of high risk, Bayesian Networks produce functional requirements on models, which in turn inform the design of flexible systems in critical areas. This has the potential to substantially decrease the cost and risk of a project, which increases the overall value of a building project.

2 Background and Related Work

A Bayesian Network (or Bayesian Belief Network) is a way of diagrammatically representing a set of factors and their influence on one another. The factors are represented by elliptical nodes, and their relationships to one another are indicated by arrows (directed edges). Formally, a Bayesian Network is a directed acyclic graph in which the nodes represent random variables and the directed edges represent conditional dependence. Each node has an associated conditional probability table (CPT), which describes the likelihood of various outcomes resulting from the different influences upon it (the upstream nodes in the Bayesian Network) (Constantinou and Fenton 2018). A strength of Bayesian Networks is that once constructed, the conditional probabilities may be easily updated as new information is added.

For example, consider a system that controls the brightness in a room. Room brightness is dependant on the lighting system and also daylight. However, the lighting system itself *also* depends on daylight. To complete the BBN, we fill out a conditional

probability table (CPT) for each node. Each entry in the table corresponds to a distinct combination of conditional probability scenarios and depends on the number of dependent nodes and outcomes at each node. The value, or probability, associated with each scenario is then used to compute the posterior probability table (the conditional probability of the parameter of interest given the evidence of connected nodes).

2.1 Example: A Simple Bayesian Belief Network

Consider the network below which illustrates the likelihood that a room is bright (illuminated) given two possible causal factors: it is daytime, or the lighting is on (Fig. 1).




Fig. 1. A simple Bayesian Belief Network

Using the connectivity of the resulting BBN, we can work through the CPT for each node. The values for the CPT can be generated by expert knowledge or previously collected data. In this paper, we focus on the knowledge driven approach. Notice that the CPT for “Room is Bright” has four entries, covering all possible states of the connected nodes (Table 1).

With the CPTs complete, we can compute the posterior probability table for the “Room is Bright” node by means of an inference algorithm, specifically the junction tree algorithm. For a full discussion of this algorithm and other inference approaches, such as Markov Chain Monte Carlo (MCMC), (Barber et al. 2004). In this example, note that the probability of success for “Room is Bright” is good but by no means assured. Strategies for improving the performance of the target node can be quickly hypothesized and simulated by updating the values of the CPTs or the connectivity of the network (Table 2).

Judea Pearl developed the term “Bayesian Network”, and established their relevance to probabilistic reasoning throughout the 1980s (see e.g. (Pearl 1985)). Because Bayesian Networks capture the relationships between variables under causal influence, they are a critical tool to facilitate the distinction between correlation and causation. Pearl describes the relevance of Bayesian Networks to the topic of causation and outlines the history of the study in “The Book of Why”, his recent book for a popular audience (Pearl and MacKenzie 2018). Pearl developed the theory of Bayesian

Table 1. Example of conditional probability tables

It is Daytime		Reasoning
0	1	
0.5	0.5	Assume night and day are equally long

	The light is on		Reasoning
It is daytime	P(node=0)	P(node=1)	
0	0.15	0.85	If it is not daytime, light is very likely on
1	0.40	0.60	If it is daytime, it is possible light is on

		Room is bright		Reasoning
It is daytime	The light is on	P(node=0)	P(node=1)	
0	0	1.0	0.0	If it is not day and lighting is off, room is not bright
1	0	0.4	0.6	If it is day, and light is off, it is possible room is bright
0	1	0.1	0.9	If not day, and light is on, room is very likely bright
1	1	0.1	0.9	If both day and light is on, room is very likely bright

Table 2. Example posterior probability table of a single node

P (node = 0)	P (node = 1)	Conclusion
0.2275	0.7725	It is possible for room to be incorrectly lit. Lighting system could be improved

Networks in the study of artificial intelligence, specifically attempting to address the question: how do machines reason? This is a response to the fact that prior to Bayesian Networks, scientists lacked a language for codifying the process through which humans reason about seemingly intuitive questions. For instance, achieving the most comfortable balance between daylight and lighting system is not all obvious and is something that every office worker is quite familiar with.

Pearl outlines a “Ladder of Causation” (Pearl and MacKenzie 2018), which he asserts is the pathway to achieving true artificial intelligence: The first rung is statistical associations (correlations that can be learned from observational data), the second rung is interventions (if-then reasoning), and the third consists of counterfactuals (the “what if” rung). These could also be expressed as “seeing”, “doing” and “imagining”. The goal of Bayesian Networks is to move beyond the first rung into questions of causation and counterfactuals. These distinctions are critical for progressing beyond an analysis of what happened to a prediction of what could be.

The construction of Bayesian Belief Networks may be approached in two ways: either taking a data-driven approach in which the network defines its structure through machine learning on data; or the BBN can be constructed by collecting information

from domain experts (the knowledge-driven approach) (Constantinou and Fenton 2018). It is also possible to use a mix of these two approaches. Once constructed, the key benefit of BBNs is their ability to support complex decision problems under uncertainty. For this reason, BBNs have been utilized in a variety of contexts, from generic project management to highly specialized tasks, for instance the diagnosis of medical disorders (Onisko et al. 1999).

Bayesian Belief Networks have been applied in several areas related to architecture, engineering and design. A literature review of this application can be found in (Okhoya 2015). The author describes two main application areas for BBNs in architecture: construction management, e.g. logistics, performance, project cost and safety; and building performance, e.g. thermal comfort, operational costs and energy consumption. An overview of artificial intelligence models in construction more broadly, including BBN approaches, appears in (Kulkarni et al. 2017).

(Cardenas et al. 2012) consider Bayesian Belief Networks applied to risk control in construction projects. The authors note that failures occur “in situations in which tacit or explicit information, such as risks and uncertainties, technological knowledge, design assumptions, monitoring records, thresholds and tolerances, was either ignored, improperly used, rejected or not passed on by someone in the project.” (Cardenas et al. 2012, p. 340). This observation suggests that if a BBN is able to codify some of these assumptions, it may act as a unifying communication tool for project stakeholders. The paper sets up a rather detailed BBN, but stops short of a robust analysis of its utility. The authors also note that cost information could be layered on the network to capture broader project information.

Another approach is (Chen and Liu 2015), who consider a Bayesian Belief Network representing elements of construction safety in urban subway engineering. They use a mixed approach of data and expert knowledge to construct the Bayesian Network. One of the key findings of the study is that investing in BIM leads to a high chance of success of a project. In the present paper we hope to draw out this observation and further utilize Bayesian Belief Networks to determine where and when modelling is most crucial to the success of a project.

3 Case Studies

To understand how an BBN might be constructed and used in a production setting, we conducted a series of case studies with industry professionals.

3.1 Methodology

Each case study is taken from a completed or ongoing project that the interviewees had expert knowledge of. Our interview process was structured around a set of questions ranging from estimations on the total cost of modelling for the project to specific questions on fabrication technology. In particular, we were interested in finding aspects of the projects that changed substantially during the design and construction phases and what effect those changes had. Once we had identified such an aspect, we conducted an interactive session with the interviewee with the goal of collaboratively creating a

BBN. To minimize the friction of this process, we used pen and paper to sketch a few iterations of the actions (nodes) and causal relations (links) of the BBN. Once we had a good candidate BBN, we worked through the conditional probability tables (CPTs) with the interviewees, placing more emphasis on coarse estimates and plausible relative values than accurately sourced data. We then entered this data in to our prototype application (see below) and computed the probability of the final “work product”. Where possible, we presented an interpretation of the results to the interviewee and suggested where more focused modelling efforts might have minimized the negative impact of the scenario, or how certain functional requirements, either digital or physical, could have had an impact.

3.2 Qrisko - Quick Risk Operations Application

To facilitate the creation of Bayesian Belief Networks in a design setting, we created a simple prototype application. Written in C++ and using Dlib (King 2009), *Qrisko* is capable of setting up and analyzing BBNs of arbitrary complexity. A minimal interface allows the user to graphically define nodes and links. When a node is selected, the CPT is generated with default values which are then modified by the user. When all the CPTs are complete, all prior probabilities are calculated and presented alongside the CPT.

3.3 Case Study 1: Roof System

This study concerns a multi-layer roof system that is installed on a timber structure that is shaped by a doubly curved design surface. This project is confidential and we have therefore anonymized any identifying information and omitted images. Nonetheless, the structure of this network should be quite familiar as it represents a common approach for creating models that involve different construction systems; that of using the output of one model as input to another. Although a rational approach, and one that captures the dependency relations well, it creates a **causal chain** (variables that have a knock-on effect on one another) (Fig. 2).




Fig. 2. A causal chain of models

Taking the design surface as the starting point, the timber model is created. The timber model is then used as the basis for the roof system models, which in turn creates the facade panel model. We assume that upon receipt of the upstream model, each

intermediate model has a very high chance of success, in this case 0.9. Since the dependencies in this network are quite simple, it is straightforward to compute the final probability of the work product which is 0.7. It is worth emphasizing this result: despite each stage having a high chance of success, the combined chance of success is much lower. We find that this simple fact suggests that this way of organizing such models can be improved.

A possible scenario is when the output of an intermediate stage is found to have failed. For example, the roof system might have a physical limitation that makes its application to some design surfaces impossible. We can simulate this with our BBN by providing evidence of failure to one of the roof system nodes and updating the computation. Not surprisingly, we find that the chance of success of the work product to fall dramatically to 0.18. Not only does this network yield a much lower chance of success for the work product than its intermediate stages, but it is also highly susceptible to failure (Fig. 3).




Fig. 3. Causal chain with evidence of failure

Prior to installing the roof system, a laser scan study showed that the as built condition of the timber structure deviated substantially from the digital model. The amount of this deviation varied across the structure but ranged from 100 mm to 500 mm and had the primary effect of increasing the curvature of the roof surfaces. Given that the model of the roof and panel systems were based off of the design surface, this meant that the impact of this change in curvature had to be assessed, quantified and applied to the model. We can capture this scenario in the following BBN (Fig. 4).




Fig. 4. The as built condition diverges from model, leading to a masked dependency.

Once the as-built condition diverges from the design model, the causal chain network is incorrect. The problem lies in the fact that what is being modelled is different from what is being built, creating uncertainty in the model. For instance, in the network above, the success of installing the panels is now dependent on both the model of the panels and the built condition. In the CPT for this network, we maintain the same high probabilities for creating the models as in the first study, however the additional dependency on the built condition dominates the behaviour of the final chance of success. To see this, we simulate the findings of the laser scan study by providing evidence of failure to the “Verify Timber”. Between the additional dependency on the built condition and the fact that there is a lower chance of creating a viable model of the roof system from the information provided by the scan data, the chance of success for the work product falls to 0.4.

The sensitivity of this system to any deviation between the design model and the as built condition suggests two ways to improve this situation. First, we could try to improve the predictions for how the design model will behave over the course of construction. In theory, a complete understanding of this behaviour would allow any problematic scenarios to be identified and dealt with. This would require at best a large increase in modelling efforts, and at worst simply be not possible due to intractable nature of complex physical processes. Secondly, we could make the roof system less sensitive to changes in the timber system. In the context of this project, a roof system and an associated model that was designed to accommodate a range of surface conditions would improve or possibly remove the dependency on the as-built condition and also make the “Verify Timber” action much more likely to succeed.

3.4 Case Study 2: KPSB Public Art

The KPSB Public Art installation was a competition that ran during 2018 in Kelowna, BC, Canada. MAKE Studio won the competition with their “Canopy of Columns” proposal: a tall, curved timber structure. The team is currently in the production phase of the project with a target installation date of October 2019. In addition to the design, The firm is committed to manufacturing all components of the structure and installing them on site. This full-service approach, the ambitious design and the tight budget make this project an excellent study in uncertainty in the business of design (Fig. 5).

As the team progressed through the detail design, they received the results of a geotechnical report conducted at the site. The report determined that the ground condition was poor and required significant changes to the foundation design of the structure. The team had to choose between maintaining their original design intent and adjusting the foundations or maintaining the proposed foundation and adjusting their design. In both cases, the team faced a considerable amount of additional model. We created the following BBN with MAKE Studio co-founder Mani Mani during the interactive session of the interview to capture, in a very general sense, this situation (Fig. 6).

Although only four nodes, this BBN captures the “masked” dependency of the design model on the site condition. Often when creating a parametric model, some aspects of the design are assumed to have a certain state. In this case, instead of creating a model of the site that is able to somehow model changes to soil condition, it is



Fig. 5. Winning entry of the KPSB public art competition




Fig. 6. Design is dependent on estimation and site condition

assumed that the site condition is favourable. This is even more deceptive than the causal chain above since not only does it create a dependency, but it also assigns a biased probability value to the Site Condition node. As more things are omitted from the model, more causal relations are masked, leading to a design model that is a very idealized version of the work product. In our BBN, we can simulate the results of the geotechnical report by finding evidence of failure in the Site Condition (Fig. 7).

The effect is that the chance of success of the final work product is 0.3. Like the causal chain example, this situation is very common. Specifically, all of the professionals that we interviewed state that their current modelling process failed to capture various aspects of the project and that on most if not all of their projects, new information is obtained during the lifetime of the project. The question is, what specific recommendations can we make to increase the chance of a successful work product in this case? Setting aside the role of the “Estimation” node here, we can improve the value of this project by making the design model more robust to changes in the site condition. A concrete way to do this is to obtain a geotechnical report before any



Fig. 7. The results of a geotechnical report should the site condition to be significantly different than expected and had serious implications on foundation design.

significant modelling effort. Unfortunately, the timing of this request is often not feasible, as was the case here, but for larger commercial projects, this is often standard procedure. A more relevant recommendation would be to impose the functional requirement on the design model that would allow it to be adjustable in height. Ranging from a simple scaling to more complex operations involving ratios of beam sizes and variations in design intent, such a feature would allow the design model to accommodate different soil conditions by requiring less or more foundation structure. Indeed, this is the path that was selected by the design team and has allowed them to move past this unexpected constraint.

3.5 Case Study 3: Toronto Centre for the Arts

As part of a larger renovation of the Toronto Centre for the Arts, Eventscape Inc. was contracted to design, detail, and build a multifaceted wall treatment of the Lyric theatre with the goal of creating a more intimate and dramatic space. Although integral to the design from the start, the architect felt that significant input from Eventscape was required in order to push the development of the wall system forward. To this end, Eventscape proposed a framing system that would follow a conceptual design surface and serve as the structure for the panelization. The panels themselves house lighting and perform acoustically by being reflective in some parts of the theatre and absorptive in other parts. Leading this effort was Eric Bury and Jonathan McGregor, who were the participants of this interview and case study (Fig. 8).

At a design review with their client, an unlikely actor emerged that had far reaching implications: the vastly complex pulley mechanism found above the stage. The function of the pulley system is to create different environments on stage and is therefore fundamental to the theatre. The difficulty lay in understanding the impact that the full range of set designs might have on the design surface for the panel system. As the conceptual design surface was updated to accommodate a new configuration of the pulley system, so too was the model of the framing and panels. At this stage in the project, propagating these changes through the different models being maintained by Bury and McGregor was no small task. With this situation in mind, we created the following BBN during the interactive part of the interview (Fig. 9).

Here we see that the Panels node is strongly influenced by two factors: Conceptual Mass and Installation. The Conceptual Mass is further influenced by a number of




Fig. 8. The interior panel system in the lyric theatre, Toronto centre for the arts




Fig. 9. The BBN of the panel design and installation

factors, but of note is the Client Requirements node, which represents the different set designs that the pulley system must accommodate. The dependency of a successful Work Product on Installation is clear: if the panels cannot be installed, the project cannot be completed. Finding evidence of failure to install in this BBN yields a final probability of 0.001, which is no surprise. Of note is that finding evidence of successful installation yields a final probability of 0.7 (Fig. 10).

This suggests that while installation is required, it is not sufficient for further increasing of the value of the project. To do this, we must look at the Conceptual Mass node. Firstly, we note the high degree of dependency of this node. Site Condition, Client Requirements, Estimation, and Installation are all factors in creating a viable conceptual mass. Specifically, the CPT for the Conceptual Mass node contains sixteen elements, each of the representing a different scenario and having distinct implications on the conceptual mass. Since the values of the CPT depend on how well these relationships are understood, a clear way to increase the overall value of the project would be to spend more time modelling and refining the behaviour of the conceptual mass. Secondly, we note that finding evidence of a viable conceptual mass results in a target probability of 0.98, whereas finding failure results in a target probability of 0.5. This reinforces the well-known design principle that understanding the needs of the client is fundamental to a successful project. It also suggests that this discussion has




Fig. 10. With successful conceptual mass and installation is the work product probability very high.

hard boundaries: not every configuration or requirement is viable and too many constraints at that programmatic stage has real impact on the final work product. It is simply not always possible to “make it so”.

4 Discussion

A theme of the case studies is that current practices often mask dependencies, hiding modes of failure (e.g. lack of understanding of the site condition in Case Study 2, or the divergence of the as-built condition from the design model in Case Study 1). Even simple BBNs can highlight important dependencies and strong actors. In this way, BBNs may decrease the cost of modelling by focusing on areas of strong influence. Moreover, once a BBN is created and the areas of strong influence identified, it becomes much easier to propose specific functional requirements on the modelling approach and even the modelling technology itself. An example of this is the role of laser scanning as-built conditions. Combined with robust procedural modelling capabilities, laser scanning could provide a way to break the causal chain model illustrated in Case Study 1. Instead of relying on the idealized output of another stakeholder, it would be possible to continuously refine each distinct model based on the physical measurements obtained by laser scanning. In this way, these functional requirements revealed by the BBN would have significant and measurable effects on model robustness and efficiency, further increasing overall project value. These requirements might also help to guide firms looking to develop new, in-house capabilities. Rather than another productivity tool or generic design option generator, firms could allocate resources on targeted features whose value has been established through examination of the BBN.

We find the implications to design very compelling. The functional requirements that emerge from a Bayesian Belief Network can extend to physical construction systems and not just digital tools. Most construction systems have some degree of tolerance. For instance the holes for screws are a little bit larger than the screw diameter, or perhaps there are multiple possible assembly configurations of a facade system. This generic approach can be insufficient for more complex designs, in

particular designs involving any amount of curvature. In those cases, there is a strong coupling between the limitations of the hardware components and the degree of curvature that can be accommodated. This suggests that the flexibility of a system (the ability of a system to adapt to a range of conditions) might be more relevant than modularity. Bayesian Belief Networks can establish where such dependencies lie and how important might be to the success of the project. Engaging the experts with relevant knowledge at key points in the lifetime of a project is critical for success and BBNs can create a framework for achieving this success.

In all of the studies we conducted, we were able to quickly create a coarse but more or less complete pictures of the project not by sophisticated modelling machinery, but by using the framework of Bayesian Belief Networks for industry professionals to more directly apply their expertise. We believe this approach can substantially increase the overall value of the project. This value would be distributed to all stakeholders, but the owner of the project would likely benefit most of all.

5 Further Work

A key step in future work is to apply this methodology to a new project. This in turn would be facilitated by the extension of our prototype software to a web application for the development of Bayesian Networks. This could be targeted to the design and architecture industries by the inclusion of common templates (such as those seen in the case studies above). Because we view BBNs are a useful communication tool for different stakeholders in a project, this application would allow multiple users to interact with the network.

Discussion of Bayesian Networks often focuses on the data-driven approach to their construction. In our case studies, we rely exclusively on the expert knowledge method of construction. It is possible, however, that combining expert knowledge with a machine learning approach on certain areas could be beneficial. This would likely involve the creation of a larger and more detailed BBN, which could include collected data about some of the less prototypical aspects of the project pipeline, for instance data about average material wastage by material type.

A further extension of this work would be to explicitly codify the *decisions* involved in the scenarios captured by the BBN. A **Bayesian Decision Network** (BDN) is an extension of a Bayesian Network which includes additional nodes called **decision nodes** (Constantinou and Fenton 2018). These nodes do not have conditional probabilities associated with them. BDNs have the potential to further support complex decision making under uncertainty. While BDNs may further clarify pathways of causality, they are also more time-consuming to create, and would depend heavily on the input of multiple stakeholders in a project.

References

- Barber, D.: Ihler, Fisher, Willsky, 2005. Probabilistic Modelling and Reasoning: The Junction Tree Algorithm (2004)
- Cardenas, C., Al-jibouri, S., Halman, J.: A Bayesian belief networks approach to risk control in construction projects. In: 14th International Conference on Computing in Civil and Building Engineering, Moscow, Russia, pp. 340–351 (2012)
- Chen, F., Liu, Y.: Innovation performance study on the construction safety of urban subway engineering based on bayesian network: a case study of BIM innovation project. *J. Appl. Sci. Eng.* **18**(3), 233–244 (2015)
- Constantinou, A.C., Fenton, N.: Things to know about Bayesian networks: decisions under uncertainty, part 2. *Significance* **15**, 19–23 (2018)
- Fenton, N.E., Neil, M., Constantinou, A.C.: Review of: Judea Pearl and Dana Mackenzie: “The Book of Why: The New Science of Cause and Effect”. Preprint (2018)
- King, D.E.: Dlib-ml: a machine learning toolkit. *J. Mach. Learn. Res.* **10**, 1755–1758 (2009)
- Kulkarni, P.S., Londhe, S.N., Deo, M.C.: Artificial neural networks for construction management: a review. *J. Soft Comput. Civil Eng.* **1**(2), 70–88 (2017)
- Onisko, A.; Druzzdel, M.; Wasyluk, H.: A Bayesian network model for diagnosis of liver disorders. In: Proceedings of the Eleventh Conference on Biocybernetics and Biomedical Engineering, Warsaw, Poland (1999)
- Pearl, J.: Bayesian networks: a model of self-activated memory for evidential reasoning (UCLA Technical Report CSD-850017). In: Proceedings of the 7th Conference of the Cognitive Science Society, University of California, Irvine, CA, pp. 329–334 (1985)
- Pearl, J., MacKenzie, D.: *The Book of Why: The New Science of Cause and Effect*. Basic Books, New York (2018)
- Okhoya, V.: *Bayesian Networks as an Architectural Decision Support Tool* (2015). Unpublished
- Umetani, N., Bickel, B.: Learning three-dimensional flow for interactive aerodynamic design. *ACM Trans. Graph. (TOG)* **37**(4), 89 (2018)