SESOC CONFERENCE 2021
A New Beginning
CLAUDELANDS, HAMILTON | 5 - 6 JULY 2021

OUR USE OF ENGINEERING MODELS

JOHN HARE

Holmes Group Limited

SUMMARY

Models are an everyday part of engineering life, often in more ways than we recognise.

We take for granted the computer models that we use almost constantly to facilitate design. More critical are the conceptual models that we have in our mind. It is critical that we constantly check the work that we do against these mental models, and don't be misdirected by our computer modelling. If there is disagreement, we should not just assume that the computer is always right. There are many reasons why our models may not be giving us the expected answers and we need to make sure that we understand these before adjusting our mental models.

Common problems are with the assumptions and approximations that we must use in order to make our models manageable. These should be approached with adequate consideration of what level of verification we need to follow up with, to ensure that our assumptions and approximations were reasonable.

This paper explores the different types of model and provides some guidance on their appropriate use. It presents some important principles that we should keep in mind when working with models and some common traps in the modelling processes, with examples of what to look out for and how to avoid them. Finally, it recommends some simple steps to consider in order to ensure successful outcomes from our models.

INTRODUCTION

This paper has been prompted by observations and anecdotes of problems that have repeated themes—typically of too great a level of reliance on computers to 'do the thinking for us' and of increasing complexity in our modelling, frequently working against our design objectives.

Models are an everyday part of engineering life, often in more ways than we recognise. This paper explores the various types of model that we commonly use and provides some guidance on their appropriate use. It presents some important principles that we should keep in mind when working with models and goes through some common traps in the modelling processes, with examples of what to look out for and how to avoid them.

The use of models is both explicitly and implicitly fully integrated into our design office practice. Explicitly, in that we routinely use computer analysis models in the design of new buildings and in the evaluation of existing buildings. In many cases, this is the limit of conscious consideration of models. Implicitly, there are many more types of model used than that. It is the use and integration of both the implicit and explicit models that makes engineering designs complete.

This paper sets out to explore the types of model that we use both formally and informally and to consider the role that they have in determining our design approach. This is followed by some recommendations on how we might apply this thinking to our everyday work practices.

All models are wrong but some are useful. – George E P Box, statistician, 1976

CLASSIFICATION OF MODELS

The term 'model' needs considerable thought. Many engineers restrict their view of what a model is to the image they see on the computer screen in front of them. However there are many more models in use than that, including those that are formed in our own minds.

Elms & Brown (2011) proposed a taxonomy of model types that identified ten main types of model, as in Table 1 below.

Table 1: Taxonomy of Models

	Model Type	Purpose
1	Perceptual	Seeing, understanding and communicating
2	Framing	Fundamental framing
3	Learning	Developing understanding
4	Prediction	Predicting outcomes
5	Comparison	Comparison of Alternatives
6	Conceptual	Developing ideas and understanding
7	Analytic	Checking
8	Blueprint	Specifying what is to be constructed
9	Communication	Communicating between stakeholders
10	Organisational	Ensuring the correct and efficient functioning of an organisation

We do not typically use all of those, and the need to have a name for each of the model types is not a daily occurrence in an engineering office. However, there is a need to understand the characteristics of the model types we most commonly use, what they will do for us and importantly, what they will not do.

The most important model types and uses in our design office practices are as follows:

Table 2: Models in general use in structural engineering practice

Model type	Purpose
Conceptual	This is the picture that flows from our mind's eye to the 'paper napkin' sketch, where our initial ideas of what may be required to solve a perceived problem begin to take form.
Learning	These are the models of that we develop to assist our understanding of the problems we are studying. They may comprise all or part of a system and should allow us to test different parameters. This type of model may include computer models which assist in the design of elements of structures.
Analytic	These are models that we use to validate our proposed solutions to a problem. (Note this is a little different from our general view of 'analysis', which is the term we use for determining design or assessment actions. This in fact falls under Learning models.)
Blueprint	These are the models that we develop in order to represent our solutions so that others can put them into action. Typically, these are the drawings and specifications that we produce.

Most importantly, you don't see the term 'Prediction model' in Table 2. That is a very deliberate omission.

A prediction model would require us to consider much more deeply the confidence limits of the data that we use, the completeness of the model and the ways that we may manage uncertainty. A prediction model, by definition, sets out to forecast what will happen, when it will happen and the magnitude of any effects. Engineers may sometimes represent the output of learning models as predictions. However, a learning model cannot provide predictions, as all of the required inputs are never known in the design stages of a project. In particular, designers seldom have full knowledge of the variance in the inputs to their models.

Our learning models frequently take the form of scenarios—they help us form a picture of what might or might not happen, within the limits of the accuracy of the model and data. **Learning models inform our judgement but they should not direct it**. That is, we use them to help build a picture of the problems that we are evaluating but we must not accept and use their output with blind faith.

It is sometimes overlooked but it is critical to start with a well-formed conceptual model. Without a sound conceptual model, designers have no basis from which to assess whether the outputs from their learning and analytical models are accurate or not.

It is critical that we constantly check the work that we do against these mental (conceptual) models, and don't be misdirected by our computer modelling. If there is disagreement, we should not just assume that the computer is always right. There are many reasons why our models may not be giving us the expected answers and we need to make sure that we understand these before adjusting our mental models.

IMPORTANT PRINCIPLES

There are three generally useful principles (Brown & Elms 2015), (Elms 1985) that engineers should be mindful of when considering models:

1. The *Principle of Requisite Detail*, which states that there is a minimum level of detail necessary in a (system) model for adequately emulating the reality which is intended to be modelled. In other words, it is important that assessors do not over-simplify the assessment to the extent that what is being modelled is not captured.

2. The *Principle of Decision Invariance*, which states that the system should be sufficiently detailed that the addition of further refinement will not affect the decision. There is no value in making models ever more complicated or comprehensive in the name of precision, if the additional detail makes no difference to the outcome; in fact, it may serve to obscure the outcome and simply add time and cost to the process.

3. The *Principle of Consistent Crudeness*, which states that the choice of the level of detail of the parts of an engineering system must, to some extent, be governed by the crudest part of the system.

   This leads us to two further points—the greatest effort should be put into the variable that has the greatest impact on the outcome. Equally, there is little use in refining our knowledge of a single variable in a problem if there are other variables with equal influence that we have little knowledge of.

THE ROLES OF APPROXIMATION AND ASSUMPTION

It is rarely possible to develop a complete model. As a result, we need to use approximation and assumption to make the modelling process timely and efficient.

Underlying this, we should remember that most structural design is to some degree empirical. Consider for example shear design of structural elements. We routinely approximate this to a uniform stress over the depth of a section, for simplicity. This works when the sections we use and the actions we design for fall within acceptable limits. However we need to make sure these limits are satisfied or our results may prove to be flawed. The point is that most design is a means to an end and implicitly includes assumptions and approximations even before we add our own to the process.

Problems arise in modelling when the assumptions and approximations that we use prove to be inappropriate. It is critical that assumptions and approximations are appropriately validated.

Addressing each in more detail:

Assumption

Assumptions are nearly always required in the modelling process. We seldom have all the detailed knowledge that we need to complete a full model and so we need to make assumptions to progress the design.

We can manage the risk of making bad assumptions by being appropriately conservative. An example of this is our use of lower characteristic steel properties in the design of reinforced concrete elements. By adopting a value that will be exceeded most of the time (95%), we avoid the need to go back and validate the design based on the measured properties of the steel that is actually supplied to site. It is implicit that even if the steel has a lower actual strength, the overall impact on the design will not be significant.

In other cases, we need to make assumptions that will (or should) be verified later—for example we may assume geotechnical properties based on the geotechnical engineer's experience and opinion, prior to a full investigation being completed. We can then review the model later when we have the completed geotechnical report and recommendations and decide whether we need to update it.

It is important for designers to consciously identify which assumptions require later validation. Where there is significant uncertainty, we can bound the possible range of the input values and perform sensitivity analysis to better understand the implications of variation. This may inform designers as to whether it will be critical to check the assumption or whether the design is in fact insensitive to variation of that input. This is the principle of consistent crudeness at work.

Approximation

Similarly, we need to make approximations in order to limit the level of detail of our model. This is in order that processing time and effort are kept proportional to the value of the knowledge gained and both the input and output data are kept to manageable volumes.

We need to use judgement in deciding what approximations are appropriate for a given situation. Sometimes this will be relatively easy, based on convention—for example, 'LOD 300' defines a standard level of detail that we work to in BIM modelling, that is universally understood. But more often, we have to use more judgement in assessing this.

In less clearly defined situations, the approximations we adopt may have a significant impact on our design or analysis output. Hence, the judgement required is more significant and we may need to test this by iterating the design and by modelling a range of inputs or boundary conditions.

An example of this is the way that we model a building response to seismic load with differing foundation fixity. An explicit model might be too complex and may require us to have vastly more information than we can practically obtain. In theory this may need a full non-linear model, but the time and effort required for this may be too great. So, we may choose to envelope the building actions using a fixed base at the one extreme and a fully pinned base at the other, both linear models. We may then envelope the design with confidence, knowing that the design actions of a more complete data set would fall between those extremes. However, we should consider also whether the range of the extremes was so high as to no longer be useful.

In considering this, remember Occam's razor (refer Figure 1 below).

Occam's razor (sometime also known as the law of parsimony): when presented with competing hypothetical answers to a problem, one should select the answer that makes the fewest assumptions. Otherwise sometimes re-stated more simply—the simplest explanation is usually the best one!

William of Ockham, a Franciscan friar who studied logic in the 14th century, first made this principle well known.

Figure 1: Occam's razor

WHEN GOOD MODELS GO BAD

When we have problems with our work, it often arises from a misalignment of the models that we use. Sometimes problems arise from the necessities of approximation and assumption in the modelling process. Sometimes it is random (stochastic) error. More often, it is in the process that we follow.

Too Complex?

Experience shows us that in many cases, the source of error is a failure at the outset to develop a full enough understanding of what it is that is being modelled. That is, we have not formed a clear concept of what we need to learn from our model and therefore what we are modelling and how we model it. Our default setting for this is often to opt for the most complex form of model that we can envisage for a given problem in the belief that what we need to know will emerge from the mist.

At risk of generalising, one of the most common reasons for adopting a process that is too complex is a failure to have first identified the key questions that the modelling process is intended to answer. This points to a need to clearly identify the success factors for the 'project', both at a high level (overall) and at a detailed level with respect to individual components or joints.

Too Precise?

Confusing precision with accuracy is another trap that we should avoid. Often we assume that the most precise (and generally complex) form of model will give the 'best' answer. This is not a given and in fact, complex models can often obscure the underlying issues that we are trying to resolve. Conversely, a simple model may often be the quickest way to get to an answer that is accurate enough for the ensuing design process.

Part of the reason that we tend to gravitate towards complex models may be a confusion of precision and accuracy. As **Error! Reference source not found.** illustrates, it is possible to be precise but inaccurate.

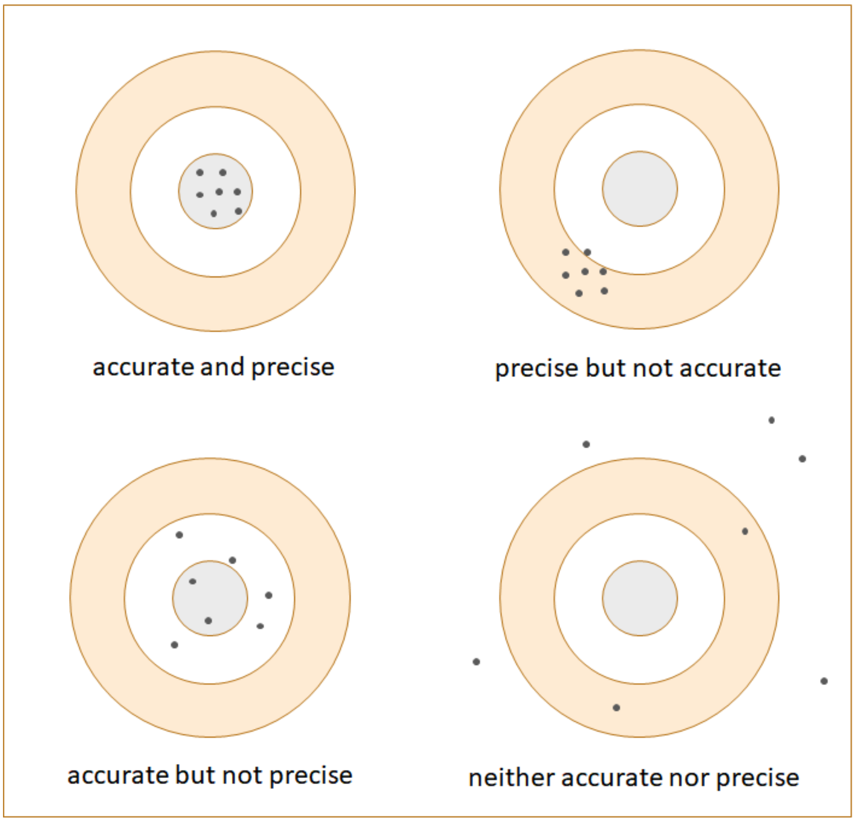

Figure 2: Precision vs. accuracy

Sometimes we fall into a trap of thinking that a qualitative or approximate model of a given problem may be inaccurate, whereas it may be simply imprecise. But provided we understand how imprecise, we can manage that in the way that we use the outputs. The principle of consistent crudeness is important here—it helps us decide how precise we need to be. There is no value in a precise modelling process to be followed by an approximate empirical design process.

The corollary is that we can convince ourselves that a precise but inaccurate set of outcomes is more ‘correct’ than an imprecise but accurate model of the same thing. We are wired to trust the more complex and precise models over the simple models but this is often wrong.

An common example of this in our industry is in the confusion that often arises between Initial Seismic Assessment (ISA), an empirical method with judgement applied; and a Detailed Seismic Assessment (DSA), which generally uses more explicit computer modelling and specific calculations of element and assembly capacity. The common assumption is that the DSA will always provide a more accurate assessment of seismic capacity, but this may confuse accuracy and precision. If the DSA assessor has spent more time to validate assumptions, consider potential mechanisms and understand the as-built and maintained condition of the building, it will probably be more accurate, but if they have missed fundamental vulnerabilities

that the ISA assessor took into account in applying judgement to their empirical assessment, possibly not.
Checking back on our Conceptual Models

Regardless of the implicit precision of the modelling output, we must have an expectation of what it is that we expect the outcome to have been (ie our Conceptual model). This will generally lead us to one of two conclusions. Either:

1. The Learning or Analytical model confirms our Conceptual model, in which case we have been confirmed in our thinking (but should not be blind to the possibility of our own confirmation bias, that is, we accepted the first thing that told us we were right, rather than give it deeper objective investigation); or

2. The Learning or Analytical model gives a very different outcome, from which we can conclude that either our concept was wrong, or there are errors in our model. Further work is required to determine which, but it is not a given that the imprecise concept was wrong.

A failure to check back on our conceptual models will often result in a divergence from what it was that we set out to achieve from the outset and the greater the delay in checking back, the greater the divergence.

MODELING PROCESS

Building a Picture

When approaching a design or analysis project, it is important that we start by planning the process. This in turn requires that we develop a view of what models may be required (beyond the Conceptual model that is already forming in our heads as we consider the problem).

The true magic of the engineer is in the answering of the question that follows—what is it that we need to learn from our more detailed models in order to complete the design or assessment? It is not in the development of the detailed models, which we should not attempt without considering that question.

Know what success looks like

It seems trivial, but the lack of this forethought often leads to divergent outcomes. Unfortunately there are many past examples where designers have developed a complex and comprehensive model and subsequent elaborate design, for significant time and expense, only to find that a simple, repeatable and quick solution was all that was required.

Remember that our design is one step of a much larger process of creating a structure or process, often to be executed by unskilled labourers or in difficult circumstances. So with reference to the Principle of Consistent Crudeness, designers should not exclude the implementation of your design from their considerations. Put simply, if you cannot see how it will be built, finessing the design is the least of your issues!

Choice of Model

The choice of what forms of model to use is entirely subjective. However, the general approach to learning models should be the same. Once we have determined what it is that we need to find out in order to complete our task, the forms of model should be selected that will efficiently and effectively get us to the outcome.

We should start by considering the data set that we have, then the nature of the problem—for example: is it linear or non-linear in behaviour? If non-linear, can an approximation provide us with sufficient understanding? (For example, the basis of ductility and capacity design in our Building Code is that we can use a linear analysis approach and allow for non-linearity through appropriate detailing and using some approximations to check outcomes such as drift).

Once we have established these background matters, we can select a method. Considering the principles outlined above, the method used should be capable of accurately determining what we need to know (requisite detail), without undertaking too much unnecessary work (design invariance). As a rule, **our modelling processes should generally gravitate to the least complex model that will accurately provide the information that we are looking for, to a level of precision that is appropriate for the nature of the task at hand.** "Because we can" is never a justification for using an unnecessarily complex approach!

Occasionally, a designer's initial modelling may reveal previously hidden aspects of the system and highlight the need to use a more detailed approach. Note that the significance of such hidden aspects must first be confirmed in your understanding and hence your revised mental model. Without taking the time to 'recalibrate' your mental conceptual model, you will not be able to assess the output from your learning model.

Alternatively, the learnings from the modelling outcomes may show that our assumptions were invalid. In these cases, it may be necessary to review and change our approach, still taking that essential step of recalibrating your mental conceptual model.

Returning to the question of accuracy, it is important to keep assumption and approximation in mind, as well as the principles articulated above.

Consider the design of a simple steel portal frame structure with concrete panels. There are several options that we might consider (in order of increasing complexity):

1. We could use a standard solution from a guidance document. This would allow us to quickly pick a uniform section for the portal frame. We would need to calculate roof bracing actions by hand for a loading envelope.

2. We could develop a simple 2D linear model. We would make an assumption about the loading by adding a margin for eccentricity (if necessary). That would allow us to envelope the displacements, based on a worst case. This would allow us to more readily adopt tapered or otherwise variable sections of members within the portal frame. We would need to calculate roof bracing actions independently.

3. We could develop a simple 3D linear model and complete a 3D static analysis. We could model the roof diaphragm explicitly in order to ascertain loadings in the diaphragm and would have to calculate loads at every frame. Again, we could design all elements directly to the actions derived from the output.

4. We could develop a full 3D model that could be used to complete a full 3D dynamic analysis, refining the seismic loading to a degree beyond what we are allowed with conventional static analysis. It would require us to model the roof bracing explicitly as a rigid diaphragm assumption would be inappropriate. We would have to go to some lengths to model mass distribution appropriately. We could design all elements directly to the actions derived from the output.

Each of these has varying levels of precision, but all may be accurate. In fact the first is arguably the most accurate, as it has been through significant validation in order to be included in published guidance. But in either case, the difference in steel weight may be only a few kg/m on the final sections selected. Furthermore, the more elaborate fabrication methods that may be supported by the more complex modelling may in fact cost more than the saving generated in the weight of the steel. Obviously the last is included for completeness – it would not be something that we would normally contemplate.

Importantly, verification of each requires different effort. It is conceivable that the first may in fact be something that you would do to verify the later approaches. But if the first works well enough, why do all the extra work to complete the more complex models? This is the principle of decision invariance at work—the additional effort required to perform complex analysis should only be undertaken if the more complex analysis affords a different and markedly better outcome than can be generated by simple analysis.

However, a more elaborate model may be warranted if it is to be re-used many times or if frequent changes or iterations will be needed.

Completeness

A model is complete when it provides the insights and outputs we need. That does not require everything to be modelled. We may find out enough about the performance of a whole structure by modelling sub-assemblages of key components, provided that the boundary assumptions that we adopt are truly reflective of the rest of the structure. This last point is very important.

Most models have to include boundary assumptions but often these can be trivial. However, some are significant, for example soil-structure interaction at foundation level of a building. When doing an ETABS analysis, it is generally simplest to model a fixed base. For structures with rigid foundations, this may be a valid approximation and we can check the potential influence of this in the subsequent foundation design process. However, if we have, for example, a raft foundation, the flexibility of the foundation may have to be modelled explicitly or we may need to artificially soften structural elements in order to approximate the raft behaviour.

> *A theory has only the alternative of being right or wrong. A model has a third possibility: it may be right, but irrelevant. — Manfred Eigen, Nobel Prize winner, 1973*

The Influence of Models on Design

It is also important to consider the influence of our modelling decisions on the design. If we are not careful, a poor modelling decision may lead to a poor design decision. This can happen when the design is arbitrarily fitted to the assumption, rather than the reverse.

The two examples below illustrate this (both situations that have been observed, more than once), with alternative approaches also offered. An important point about both is that a complex model could specifically address the issue being highlighted, but each has a simple approach that will yield a positive outcome for relatively little effort.

Example 1: Consider the design of a ground floor slab over a basement, with two or more towers on it. Such a slab will tend to act as a transfer diaphragm and the common approach is to assume a rigid diaphragm for efficient modelling. This results in large transfer forces, often requiring large collector elements, thickened slabs and a lot of reinforcement. It may also be accompanied by massive return shears in the primary structure below the diaphragm—the peg in the hole effect. This gets potentially worse in the region between the towers and is made more complex in that the worst actions arise when the towers are out of phase – difficult to model and get good useful output from.

But the consequence of ‘failure’ of the diaphragm may only be a very small movement before alternative load paths develop. Inserting a full seismic joint is expensive and may result in other issues—for example weather-proofing. But what if we simply allow it to fail – either by specifically modelling diaphragm flexibility (which may introduce significant complexity), or even by controlling where the weak point is and ensuring the failure will not result in a more critical issue? The latter can be done by modelling the main structure without the diaphragm and checking if the primary structure can deal with the actions (which conceivably be less than those resulting from return shears); and then checking displacements. The differential displacement may turn out to be little more than a large crack; and may not amount to anything significant at SLS. That is, we could adopt the design solution (with communication to the client) of treating it as monolithic concrete with a form of capacity design to let it fail in a way that will be safe and repairable, if it ever happens.

Example 2: Even more common – we model a structure with a fixed base, generating significant wall/column base actions. We then use the overstrength actions from these elements to determine overall foundation design actions, which show we need a very large pad, raft or even piles in order to resist these overturning actions. So the model drives a design decision.

But if we envelope the design actions on the superstructure by running both a fixed base and a pinned base option, we can then design a flexible foundation which we can ‘tune’ to rock at the level of demand we accept. In this way, we can limit the foundation to a reasonable size and capacity, while ensuring the primary structure is protected and that the increased flexibility of the structure is addressed (by considering drift and its impact on other elements). This will typically be more materially efficient and easier to build.

The downside of this approach may be that this should be treated as an alternative solution, with ensuing difficulties in building consent process, but the technical solution is .

The important step in each case is to step back from the modelling process and ask the question – does this make sense? Without that step, it would be possible to carry on and design something that is precise and correct (for the assumptions adopted in the model) and yet not the best outcome for the project.

Verification

All models require some level of verification, depending on how they were developed and what modelling process was followed.

Verification should always commence with a comparison of the outcomes of the modelling to the conceptual model, the picture that we had in our mind’s eye as we set out on the process. In structural analysis, this may be as simple as looking at the deformed shapes and a quick check of member actions to see that they match our general expectation. For more complex models such as a full ETABS analysis of a building, we should complete a basic check of key outputs, such as:

- Building period
- Base shear
- Displacement (absolute and interstorey)

If the model is viable (that is, it returns results that indicate a structure that is capable of being built and will ‘work’), but is not doing what we expected, we should not proceed further until we know why.

CONCLUSIONS

Developing models is a key part of what we do. However, we need to be careful to understand the intent of our models and to consider carefully what the questions are that we use them to answer. A viable and well-constructed but inappropriate model may cloud us to what is most important.

We do not need to explicitly define the types of model that we use, but we must understand their purpose. We should frequently remind ourselves that models are generally developed to inform our judgement—regardless of what our models tell us, we need to understand why or continue investigating until we do. Without that understanding, we cannot form a judgement as to the appropriateness of the outcomes.

THE KEY STEPS TO SUCCESS

Here are some simple points to consider when embarking on your modelling journey

1. Understand what success looks like – for the client, for you.
2. Have a clear concept in mind from the outset—a mental model that you can use as a benchmark.
3. Understand the problem(s) that you are approaching clearly so that you can define what output you require.
4. Be clear what the objective of any individual model is and how it will inform your judgement.
5. Always use the simplest model that can accurately capture the problem.
6. Don't lose sight of the need for someone to implement the solutions that we develop; and use this to inform the modelling and design.

REFERENCES

Brown CB & Elms DG (2015), Engineering decisions: Information, knowledge and understanding, Structural Safety 52 (2015) 66-77, Elsevier.

Elms, D., & Brown, C. (2011). Professional Decisions: the Central Role of Models. Civil Engineering and Environmental Systems, 2012, 1-11.

Elms DG (1985), The principle of consistent crudeness, Purdue University, IN, 1985 Proc Workshop on Civil Engineering Application of Fuzzy Sets, Purdue University, IN, 1985.