CS 4390/CS 5390 Uncertainty in AI: Fall 2025 Class, TR 12-1:20 pm

Vladik Kreinovich

Department of Computer Science, University of Texas at El Paso
El Paso, Texas 79968, USA, vladik@utep.edu
https://www.cs.utep.edu/vladik/cs4390.25/syllabus.html

1. Uncertainty in AI: Main Objective of the Class

- To teach methods of dealing with uncertainty in AI.

2. What Is AI

- In order to talk about uncertainty in AI, we first need to discuss what is AI and what is not AI.
- From the computational viewpoint, an ideal computational problem is when:
  - we have a well-defined problem, and
  - to solve this problem, we apply a well-justified algorithm to objective (measurement-based) data.
- In this form, solving this problem is routine, no intelligence is needed.

3. What Is AI (cont-d)

- In many real-life situations, however:
  - the problem is not well-defined (e.g., how to formalize "comfortable"),
  - some data comes not from measurements but from expert estimates, and
  - to process this data, we use expert-proposed heuristic (= not justified) methods.
- In such non-routine cases, we need to use intelligence.
- When computers are used to solve such intelligence-related problems, this becomes AI.
- So, in a nutshell, AI is using heuristics, using expert knowledge.

4. Why uncertainty?

- Uncertainty is everywhere in AI:
  - heuristics are not precise,
  - expert knowledge is not precise,
  - measurements are not precise.

  So, we need:
  - to *quantify* this *uncertainty*,
  - i.e., describe it in computer-understandable numerical terms.
- We also need:
  - to *propagate* this *uncertainty* through the algorithms,
  - i.e., describe, based on the uncertainty with which we know inputs, what will be the uncertainty of the algorithm's output.

5. Why uncertainty (cont-d)

- In uncertainty propagation, in addition to the general topic, there are also two special topics.
- The first is taking into account additional uncertainty caused by the computational device:
  - from *round-off errors* in the traditional computers
  - to the probabilistic character of *quantum computers*.
- The second is taking into account uncertainty that we introduce ourselves on purpose, e.g., to preserve *privacy*.

6. Why not just use machine learning (ML)

- Current ML tools are so good that it seems like they can compute everything.
- Unfortunately, uncertainty is an exception.
- Current ML tools:
  - do perform most computations very well, but
  - they are not very good in estimating uncertainty,
  - in particular, uncertainty of their own results.
- For example, when they mistake a dog for a cat, they claim it with 99.9% confidence.

7. Why not just use machine learning (cont-d)

- There is a reason for this.
- When we train a system to predict the result, we can (and do) use thousands of examples.
- However:
  - when we train it to get the probability of each result,
  - we need about 100 examples of computations to get one probability estimates.
- So the number of inputs for such training is 100 times smaller, not enough to get good results.

8. Types of uncertainty

- In the ideal case:
  - we know all possible values of the corresponding quantity and
  - we know the frequency with which each value occurs,
  - i.e., we have a *probabilistic* uncertainty.
- In other cases, we have no information about the frequencies.
- This corresponds to *interval* or, more generally, *set-valued* uncertainty.
- These are the two extreme cases.
- In general, we may have some *partial information* about the probabilities.
- This is called *imprecise probabilities.*

9. Types of uncertainty (cont-d)

- Important particular cases are:
  - cases when we know intervals of possible values of probability;
  - this is known as *probability boxes* or, for short, *p-boxes*, and
  - cases when we know the frequency of different possible probability distributions;
  - this is known as the *Bayesian approach*.
- We can also have uncertain expert statements.
- There are two ways to describe their uncertainty.
- We can try to elicit *subjective probability* (also known as *prior distribution*) from the experts.
- We can also simply ask the experts to gauge their own degree of certainty, e.g., on a scale from 0 to 1.
- This is known as the *fuzzy* approach.

10. Types of uncertainty (cont-d)

- Instead of numbers, it often makes sense to elicit intervals – or even fuzzy statements,
- This modification is known as *type-2* fuzzy techniques.

11. Measures of uncertainty

- There are two basic ways to measure uncertainty.
- We can view it from the viewpoint of amount of knowledge.
- I.e., we can describe the amount of effort needed to complete our uncertain knowledge.
- This approach leads to *entropy*.
- We can also explicitly take into account the effect of uncertainty on our decisions.
- I.e., we can gauge *losses* caused by uncertainty.

12. Measures of uncertainty (cont-d)

- For example, if a company does not want to reveal its current sales amount to its competitors:
  - disclosing the least significant binary digit (bit) of this amount and
  - disclosing the first bit both provide the competitors with exactly 1 bit of information.
- However:
  - the first disclosure is harmless, while
  - the second can lead to serious losses.

13. How to come up with algorithms

- Here, we will talk about standard data processing techniques, with an emphasis on the probability-based ones:
  - *Maximum Likelihood*,
  - its most used particular case – *Least Squares*, and
  - *Markov Chain* techniques.
- We will also talk about *simulation*, *interpolation* and *extrapolation*, *robustness*, and *regularization*.

14. How to propagate uncertainty, in particular through Large Language Models (LLMs) and simulation algorithms

- This will be one of the main topics of this course.
- We will learn uncertainty propagation algorithms covering all types of uncertainty: probabilistic, interval, and fuzzy.
- Most of these algorithms are designed for propagating uncertainty through open-source algorithms.
- We will also cover techniques for propagating uncertainty through proprietary *black-box algorithms*.

15. How to combine uncertainty

- Often, different inputs come with uncertainty of different type.
- For example:
  - we know the probability distribution of one of the quantities, and
  - we only have expert (fuzzy) information about another quantity.
- For such cases, we will study two types of methods.
- When most of the information is of the same type, we can simply transform the remaining different-type information to this type.
- We can use *Maximum Entropy* approach (MaxEnt) to transform intervals into probabilities.
- We can use *defuzzification* to transform fuzzy data into numbers.
- We can use midpoint (or *Hurwicz approach*) to transform an interval into a number, etc.

16. How to combine uncertainty (cont-d)

- In some situations, there is no dominant type of uncertainty.
- In this case, we need to use more sophisticated methods for actually combining different types of uncertainty.

17. Reasoning under uncertainty, causality, etc.

- In this section, we will be studying:
  - basic fuzzy techniques,
  - basic statistical testing, and
  - basic ideas of non-monotonic logics.

18. Decision making under uncertainty

- We start with decision making:
  - in situations without uncertainty,
  - i.e., in mathematical terms, with *optimization*.
- We will study *gradient descent*.
- It is the simplest numerical optimization technique that works so well in neural networks.
- Then we will study decision making under the most detailed type of uncertainty – probabilistic uncertainty.
- We will study the *utility approach* – the standard approach of *decision theory*.
- Then, we will deal with the case of interval uncertainty.
- In this case, decision theory leads to the *Hurwicz criterion*.
- After that, we will study *fuzzy decision making*.

19. Decision making under uncertainty (cont-d)

- Finally, we will study *symmetry approach*.
- It is a physics-motivated approach which is useful when we have very little information.
- We will show how this approach explains many successful heuristics in fuzzy and neural networks.

20. Other topics

- At the end of the class, we will briefly mentioned other important uncertainty-related topics, such as:
  - *uncertainty visualization* and
  - dealing with *conflict situations under uncertainty*.

21. Sources

- On this topic, there is no up-to-date textbook yet.
- We will use several papers.
- For some topics, there are no easy-to-read papers.
- We will then try to post easier-to-summaries of the not-so-easy-to-read papers on the class website.