> Artificial Intelligence: why should firms care?

Franck.Dernoncourt@gmail.com
http://francky.me

May 30th, 2012
@Swedish Chamber of Commerce in Paris

0) Foreword

Video http://youtu.be/A9pMbRKq8Ro

Why this presentation?

Article on IBM Watson

http://www.theinquirer.net/inquirer/news/2171838/ibm-protects-watson-ai-hungry-firms:

"Eighty or 90 per cent of these requests don't need Watson anyway, technology already exists for what they need."

Table of Contents

Intelligence

The game of chess has always been viewed as an intellectual game par excellence:

"a touchstone of the intellect" (Goethe)

... IBM Deep Blue beat world champion Garry Kasparov in 1997.

Intelligence

No consensus on a formal definition of intelligence. Consensus on:

- Intelligence is the computational part of the ability to achieve goal**s** in the world.
- Varying kinds and degrees of intelligence occur in people, many animals and some machines.

Artificial Intelligence (AI)

3 definitions:

jmc: It is the science and engineering of making intelligent machines, especially intelligent computer programs. It is related to the similar task of using computers to understand human intelligence, but AI does not have to confine itself to methods that are biologically observable.

Artificial Intelligence (AI)

3 definitions:

jmc: It is the science and engineering of making intelligent machines, especially intelligent computer programs. It is related to the similar task of using computers to understand human intelligence, but AI does not have to confine itself to methods that are biologically observable.

Peter Norvig: We think of AI as understanding the world and deciding how to make good decisions. Dealing with uncertainty but still being able to make good decisions is what separates AI from the rest of computer science.

Artificial Intelligence (AI)

3 definitions:

jmc: It is the science and engineering of making intelligent machines, especially intelligent computer programs. It is related to the similar task of using computers to understand human intelligence, but AI does not have to confine itself to methods that are biologically observable.

Peter Norvig: We think of AI as understanding the world and deciding how to make good decisions. Dealing with uncertainty but still being able to make good decisions is what separates AI from the rest of computer science.

comp.ai: AI can mean many things to many people. Much confusion arises because the word 'intelligence' is ill-defined. The phrase is so broad that people have found it useful to divide AI into two classes: strong AI and weak AI.

Artificial Intelligence (AI)

- Weak AI simply states that some "thinking-like" features can be added to computers to make them more useful tools... and this has already started to happen (witness expert systems, drive-by-wire cars and speech recognition software).

- Strong AI makes the bold claim that computers can be made to think on a level (at least) equal to humans and possibly even be conscious of themselves.

AI goal: Achieve strong AI (i.e. singularity).

A dream of robots

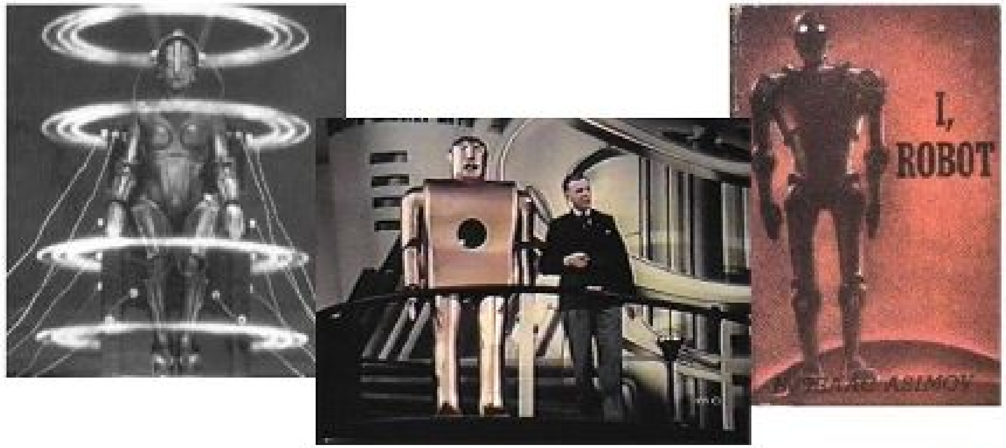

Commercial reality

Research frontier

Artificial Intelligence (AI) – a brief history

After WWII, a number of people independently started to work on intelligent machines.

The English mathematician Alan Turing may have been the first. He gave a lecture on it in 1947. He also may have been the first to decide that AI was best researched by programming computers rather than by building machines.

By the late 1950s, there were many researchers on AI, and most of them were basing their work on programming computers.

Artificial Intelligence (AI)

Artificial Intelligence (AI)

The big picture: induction vs. deduction

Statistical AI, arising from machine learning, tends to be more concerned with "inductive" thought: given a set of patterns, induce the trend.

Classical AI, on the other hand, is more concerned with "deductive" thought: given a set of constraints, deduce a conclusion.

Table of contents

1. AI > definitions
2. AI > induction
3. AI > deduction
4. Conclusion & perspectives

Evolutionary algorithm:

based on Darwin’s natural selection theory.

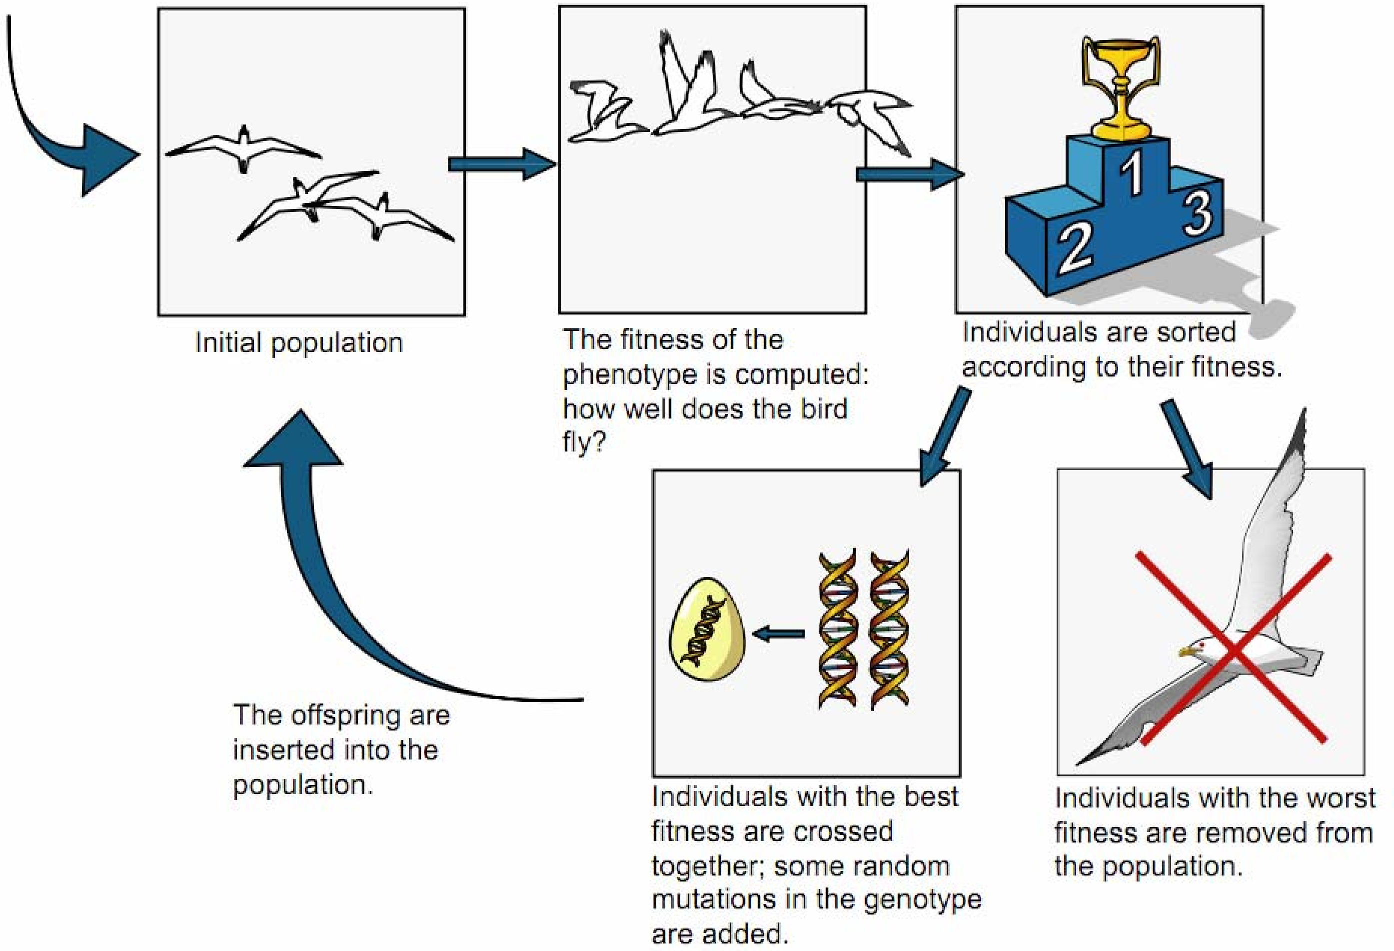

Evolutionary algorithm:

Case study:

SSgA manages institutional portfolios of very high value (from high millions $US to billions $US).

Exploiting evolutionary algorithms to develop stock future valuation models in quantitative asset management.

Members of the population are candidate stock future valuation models.

Evolutionary algorithm:

Case study:

SSgA manages institutional portfolios of very high value (from high millions $US to billions $US).

Exploiting evolutionary algorithms to develop stock future valuation models in quantitative asset management.

Members of the population are candidate stock future valuation models.

E: analyst variables
P: price momentum variables
Q: quality variables
V: valuation variables

2) AI > induction

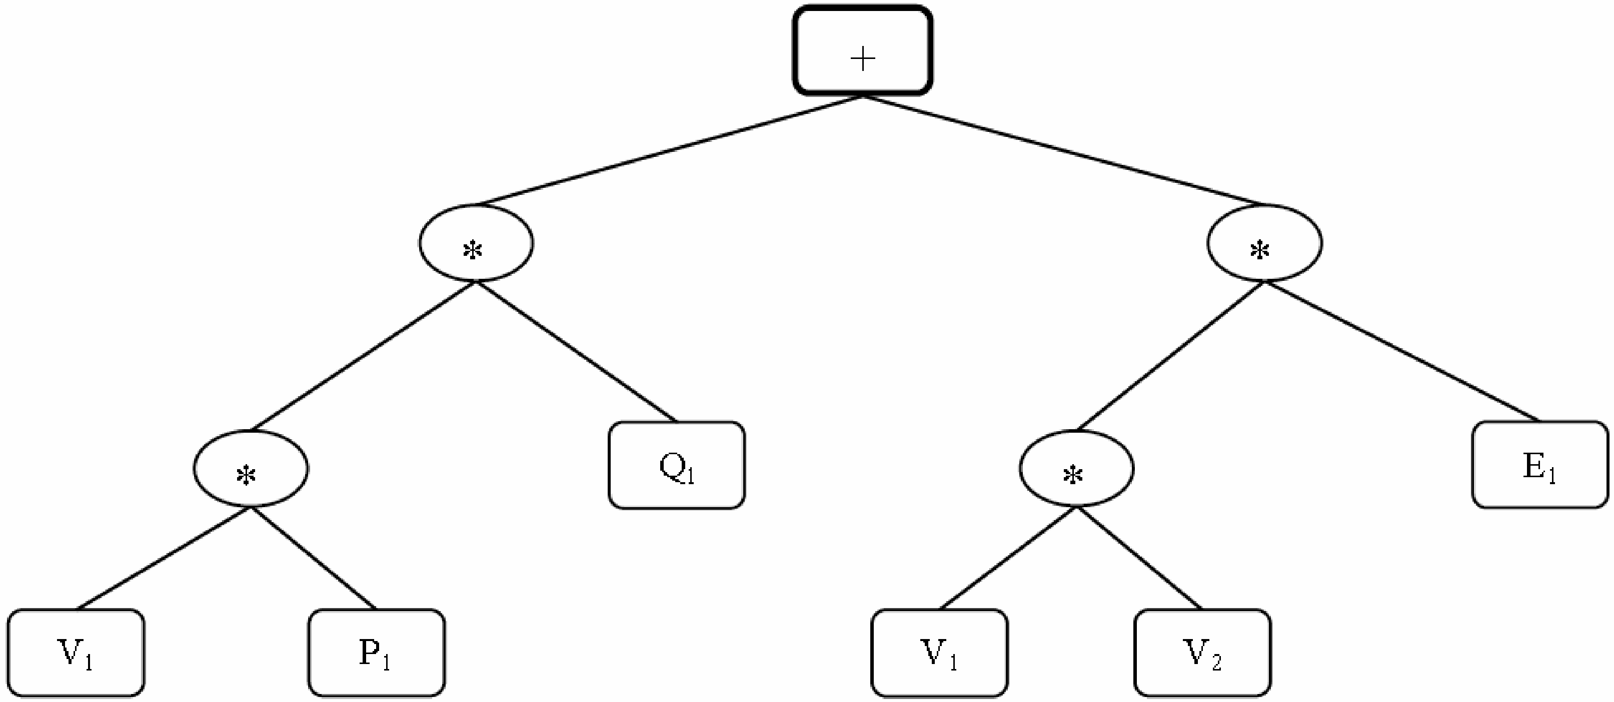

$$Model\,A = V_1 * P_1 * Q_1 + V_1 * V_2 * E_1$$

E: analyst variables
P: price momentum variables
Q: quality variables
V: valuation variables

> Models assess stock future valuation

Evolutionary algorithm:

based on Darwin's natural selection theory.

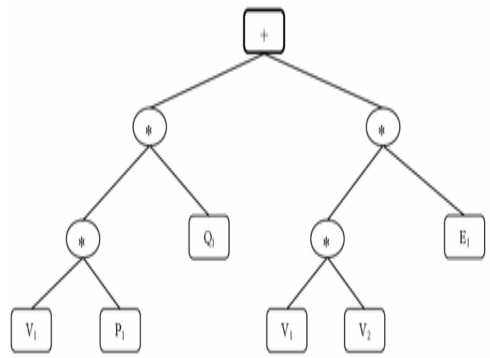

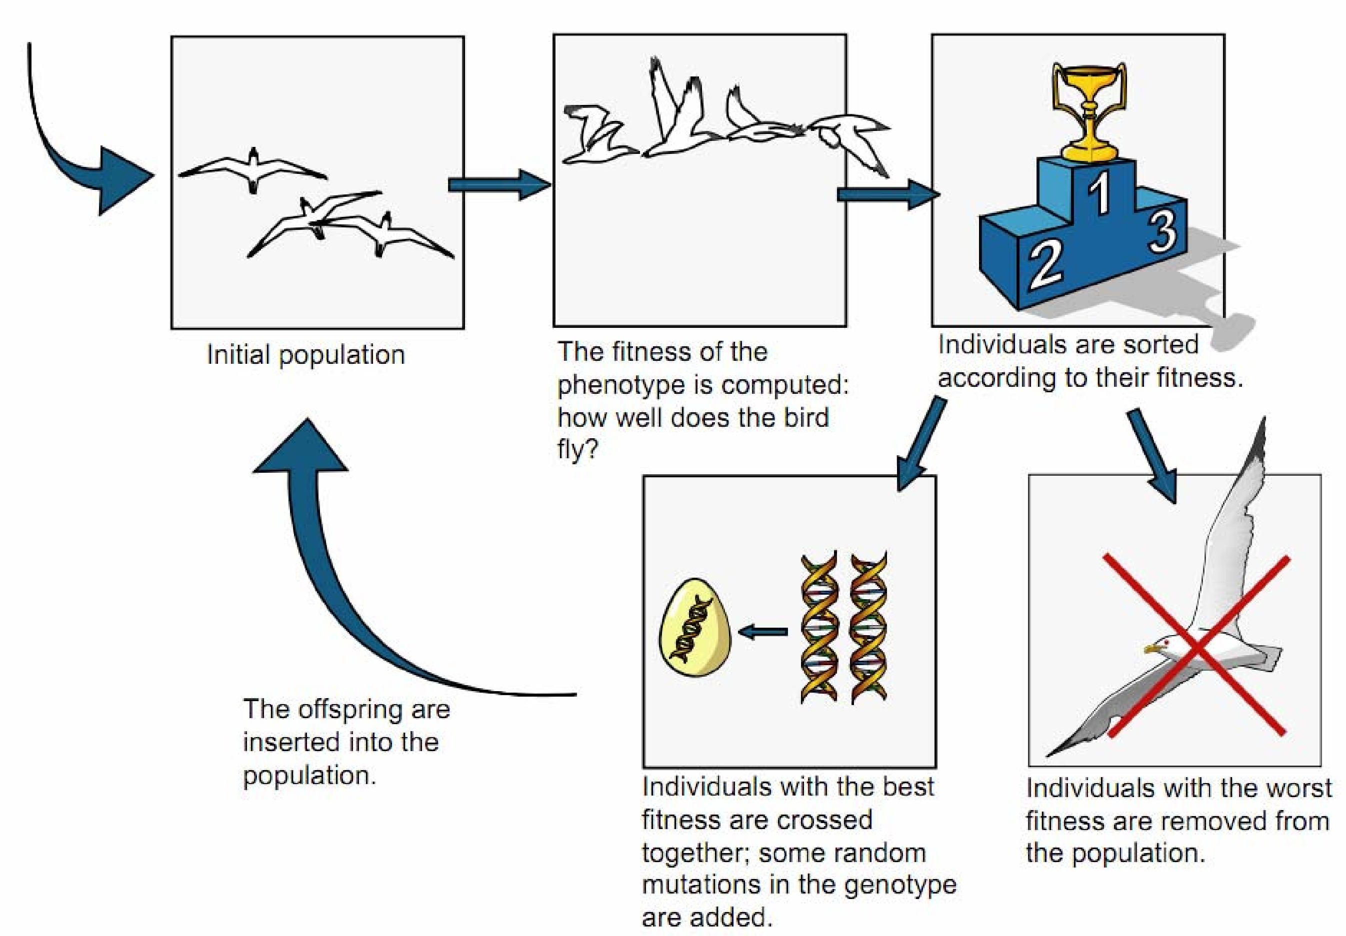

Multi-objective:

- Predict precisely stock future valuation
- don’t take analysts too much into account
- limit tree size (i.e. model complexity)
- …

Multi-objective:

- Predict precisely stock future valuation
- don't take analysts too much into account
- limit tree size (i.e. model complexity)
- …

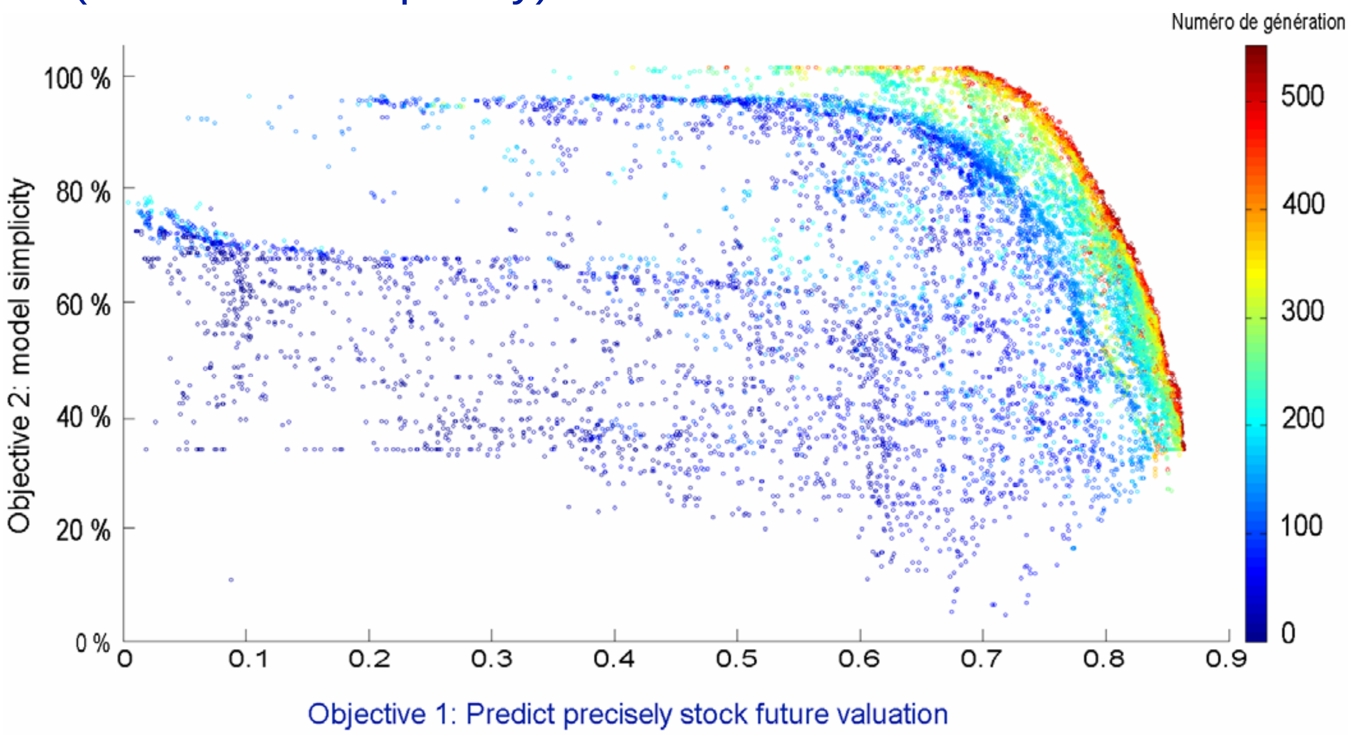

Black box optimization:

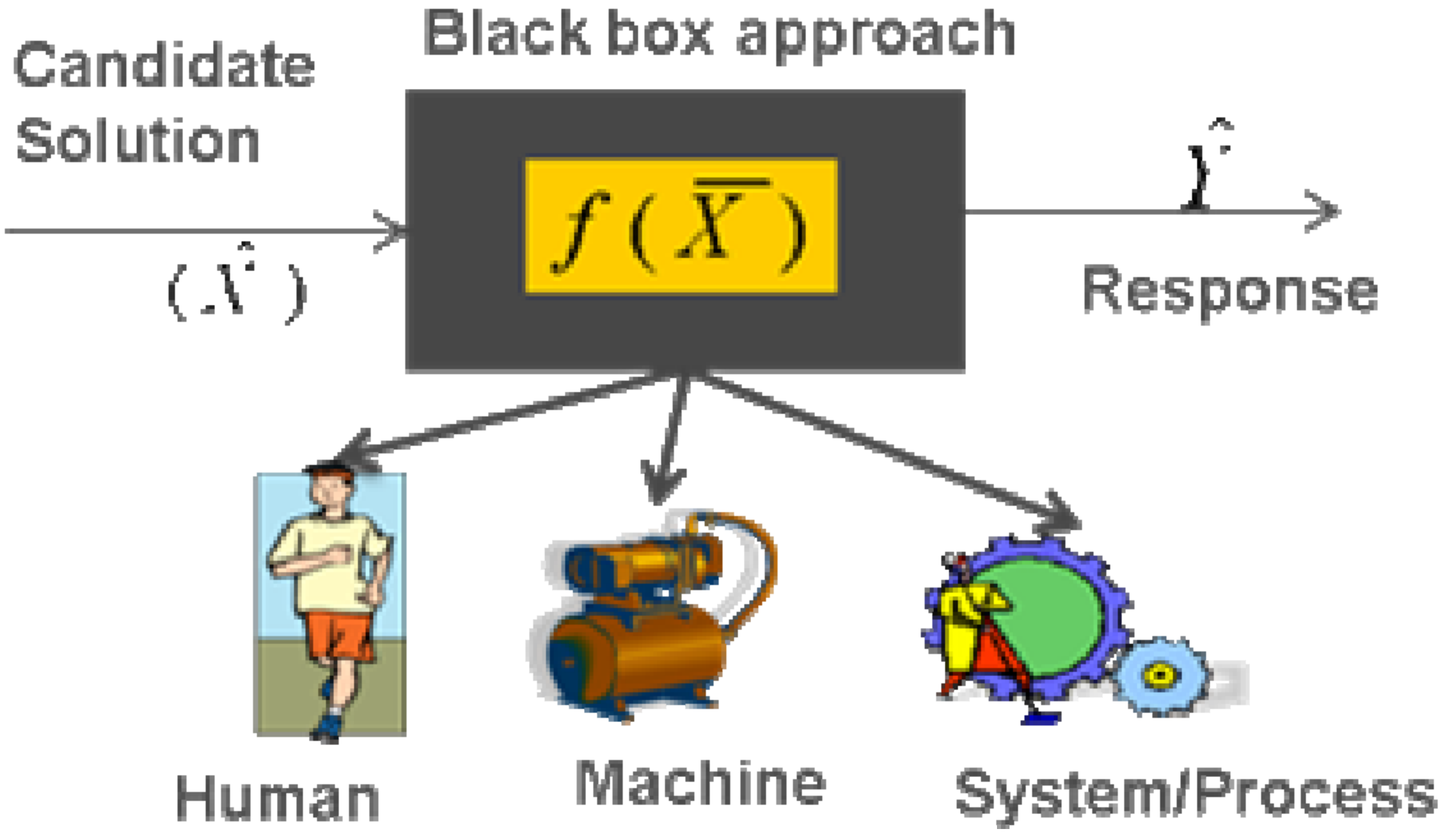

Table of Contents

AI > deduction:

Case study: fuzzy logic

AI > deduction:

Case study: fuzzy logic

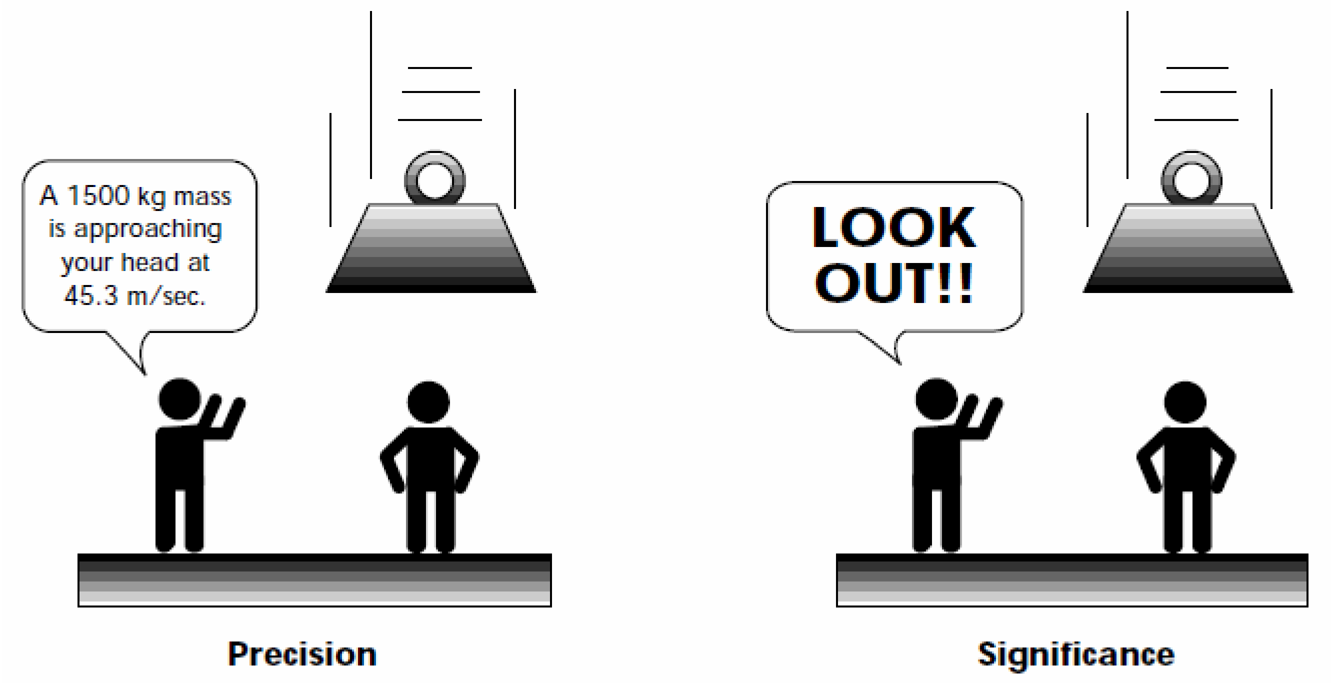

Observation:

Knowledge available to humans are virtually never perfect. These imperfections can be distinguished into two classes:

- **Uncertainties**, which refer to knowledge whose validity is subject to question. For example, if we know someone bumped his head on a ceiling, we can guess that he is likely to be very tall.

- **Inaccuracies**, which refer to knowledge that is not clearly perceived or defined. For example, instead of saying someone is 2 feet and 3 inches, we usually say that person is very tall.

Observation:

Knowledge available to humans are virtually never perfect. These imperfections can be distinguished into two classes:

- Uncertainties, which refer to knowledge whose validity is subject to question. For example, if we know someone bumped his head on a ceiling, we can guess that he is likely to be very tall.

- **Inaccuracies**, which refer to knowledge that is not clearly perceived or defined. For example, instead of saying someone is 2 feet and 3 inches, we usually say that person is very tall.

Observation:

Knowledge available to humans are virtually never perfect. These imperfections can be distinguished into two classes:

Probability!

- Uncertainties, wh[illegible]owledge whose validity is subject to question. For example, if we know someone bumped his head on a ceiling, we can guess that he is likely to be very tall.

- **Inaccuracies**, which refer to knowledge that is not clearly perceived or defined. For example, instead of saying someone is 2 feet and 3 inches, we usually say that person is very tall.

Observation:

Knowledge available to humans are virtually never perfect. These imperfections can be distinguished into two classes:

- Uncertainties, which refer to knowledge whose validity is subject to question. For example, if we know someone bumped his head on a ceiling, we can guess that he is likely to be very tall.

- Inaccuracies, which refer to knowledge that is not clearly perceived or defined. For example, instead of saying someone is 2 feet and 3 inches, we usually say that person is very tall.

Observation:

Knowledge available to humans are virtually never perfect. These imperfections can be distinguished into two classes:

Probability!

- Uncertainties, which refer to knowledge whose validity is subject to question. For example, if we know someone bumped his head on a ceiling, we can guess that he is likely to be very tall.

Fuzzy logic!

- Inaccuracies, which refer to knowledge that is not clearly perceived or defined. For example, instead of saying someone is 2 feet and 3 inches, we usually say that person is very tall.

3) AI > deduction

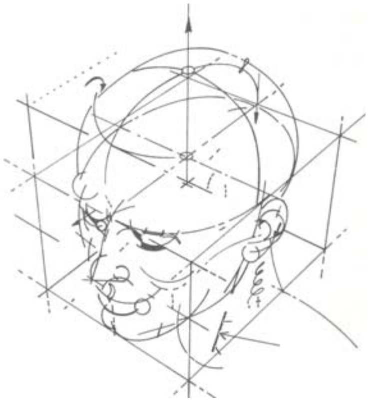

How to express these inaccuracies in logical terms?

In **classical logic**, a proposition is **true** or **false**.
Example: This person is tall. *True* or *false*?
→ Not flexible…

In **multivalued** logic, a proposition may have multiple values.
Example(ternary): This person is tall. *True*, *half true* ou *false*?
→ Slightly more flexible…

In fuzzy logic, a proposition can have as many values as one wants.
Example: This person is tall. This is 30% true.
→ Flexible!

Lofti Zadeh, researcher in systems theory, laid the foundations of fuzzy logic in an article in 1965.

Fuzzy logic is an **extension of the Boolean logic** based on the mathematical theory of fuzzy sets, which is a generalization of the classical set theory.

By introducing the notion of degree in the verification of a condition, thus enabling a condition of being in a state other than true or false, fuzzy logic provides a highly valuable flexibility to reasoning models, making it possible taking into account inaccuracies

3) AI > deduction

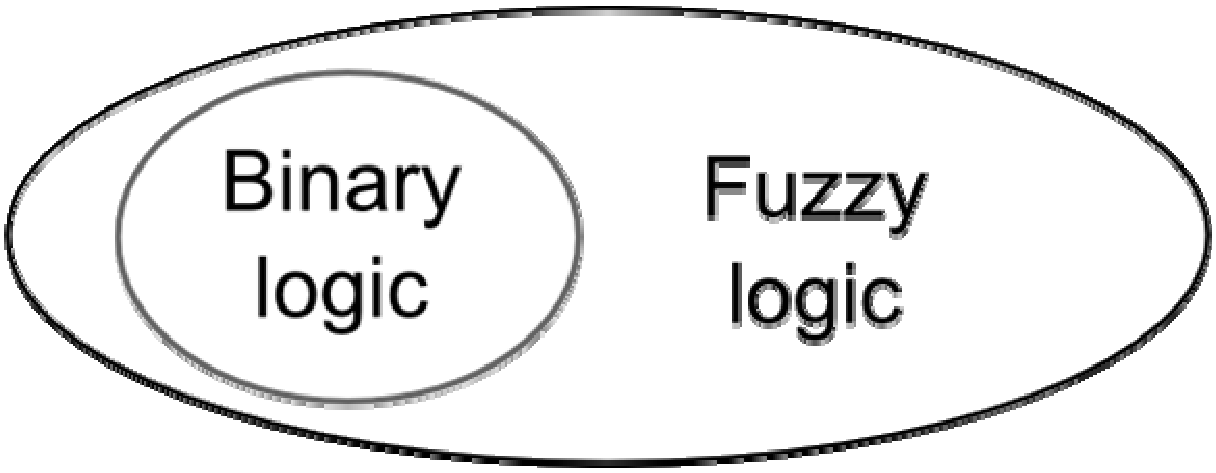

3) AI > deduction

Applications:

- Decision support system (e.g. in healthcare),
- Database (e.g. fuzzy queries),
- Fuzzy commands (e.g. subway line M14 in Paris),
- Data mining (e.g. clustering),
- Pattern recognition,
- …

Let's define the **key concepts of fuzzy logic** through an example of image processing: increasing the contrast of an image

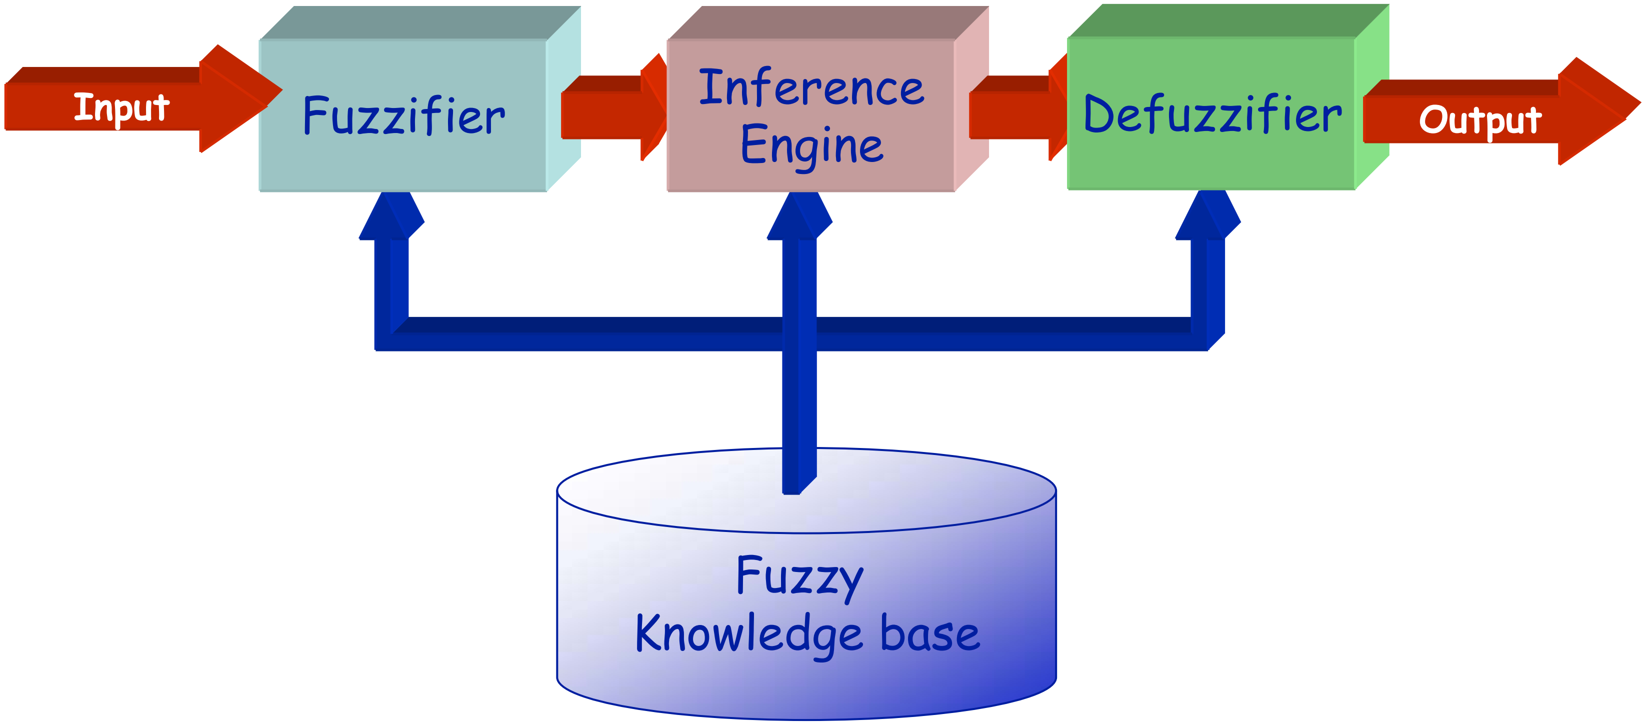
Input
Fuzzifier
Inference Engine
Defuzzifier
Output
Fuzzy Knowledge base

Fuzzifier

1) **Fuzzy set and membership degree**

Classical set	Fuzzy set
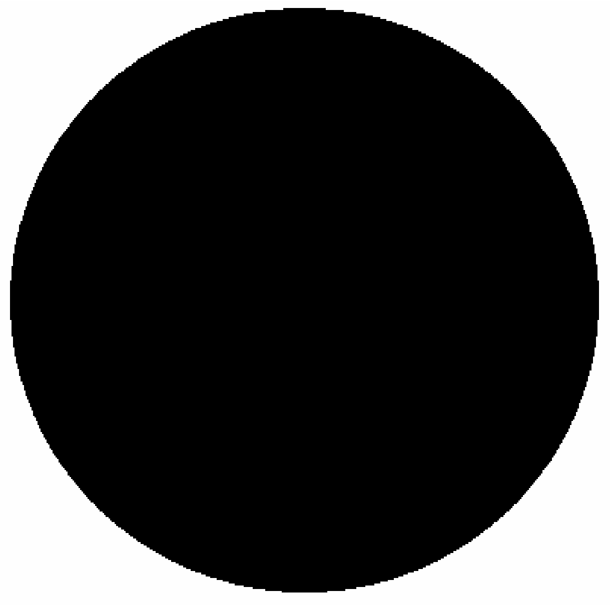	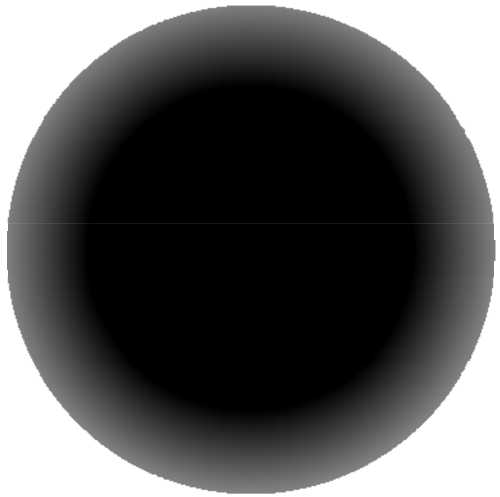
Indicator function: 0 or 1 → **Binary** in classical logic	Membership degree: any real between 0 et 1 →**Membership degree** in fuzzy logic (eg 0.867)

Fuzzifier

1) **Fuzzy set and membership degree**

Classical logic

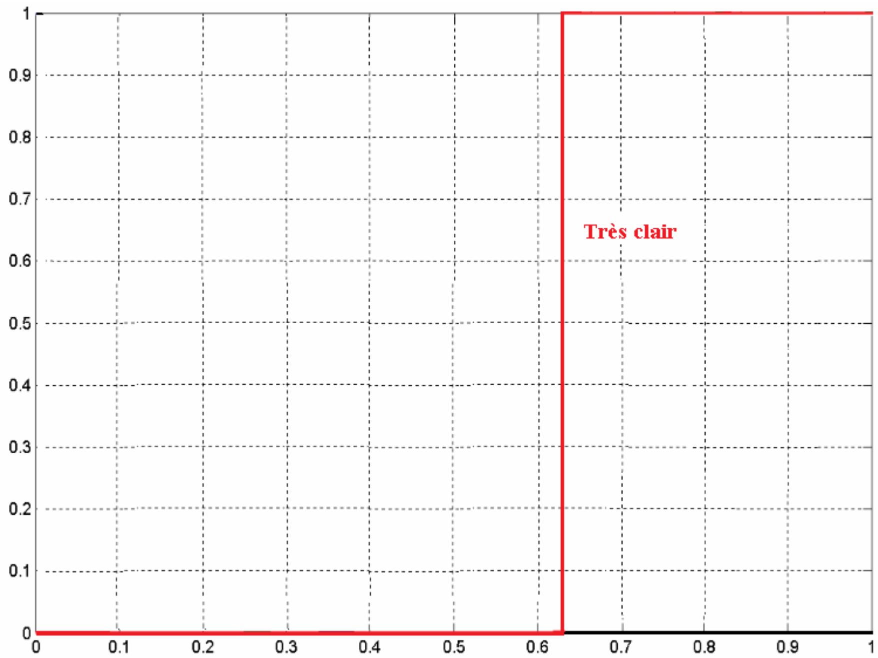

Indicator function: 0 or 1
→ **Binary** in classical logic

Fuzzy logic

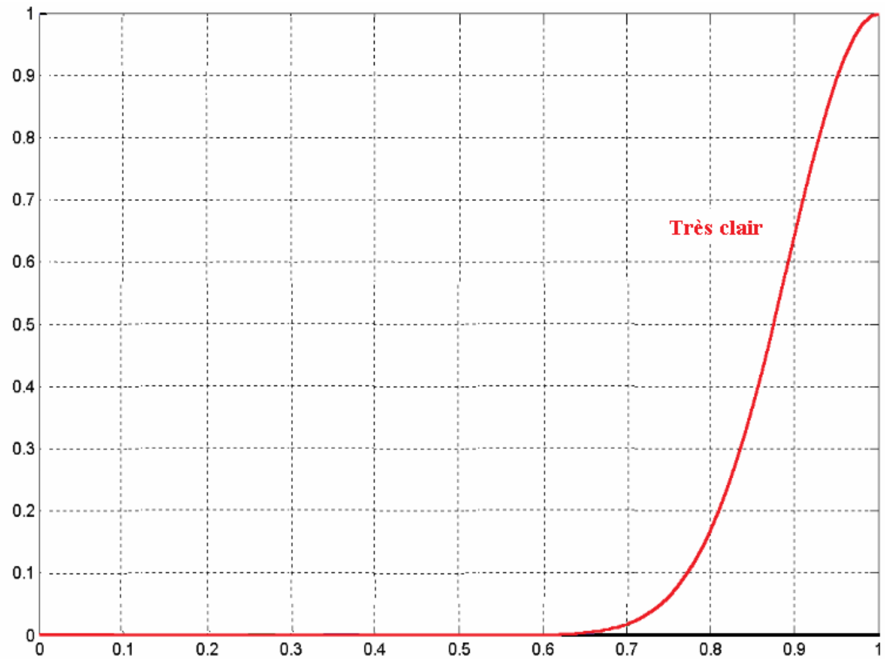

Membership degree: any real between 0 et 1
Membership degree in fuzzy logic
(eg 0.867)

3) AI > deduction

1) **Membership functions (*fuzzification* step) in input**

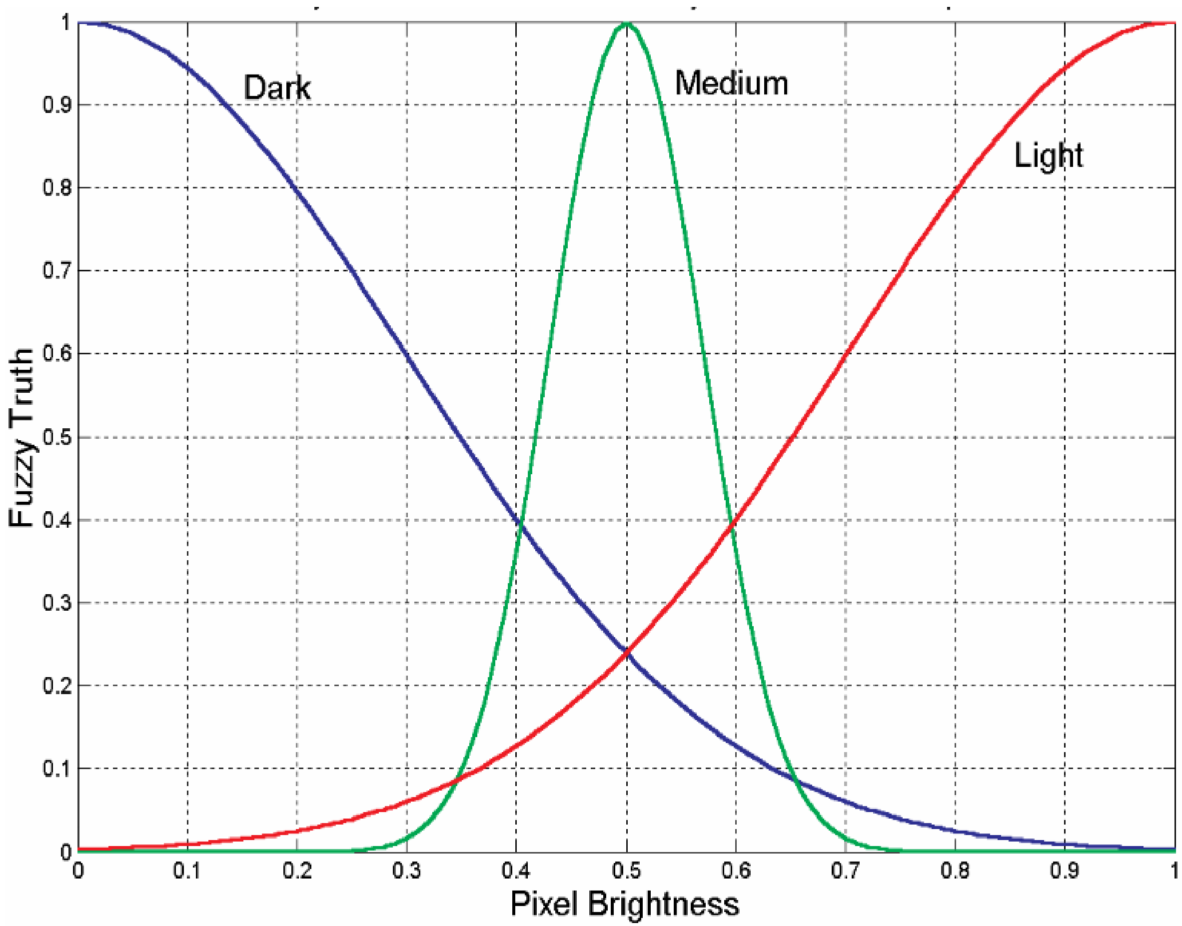

Fuzzifier

1) **Membership functions (*defuzzification* step) in output**

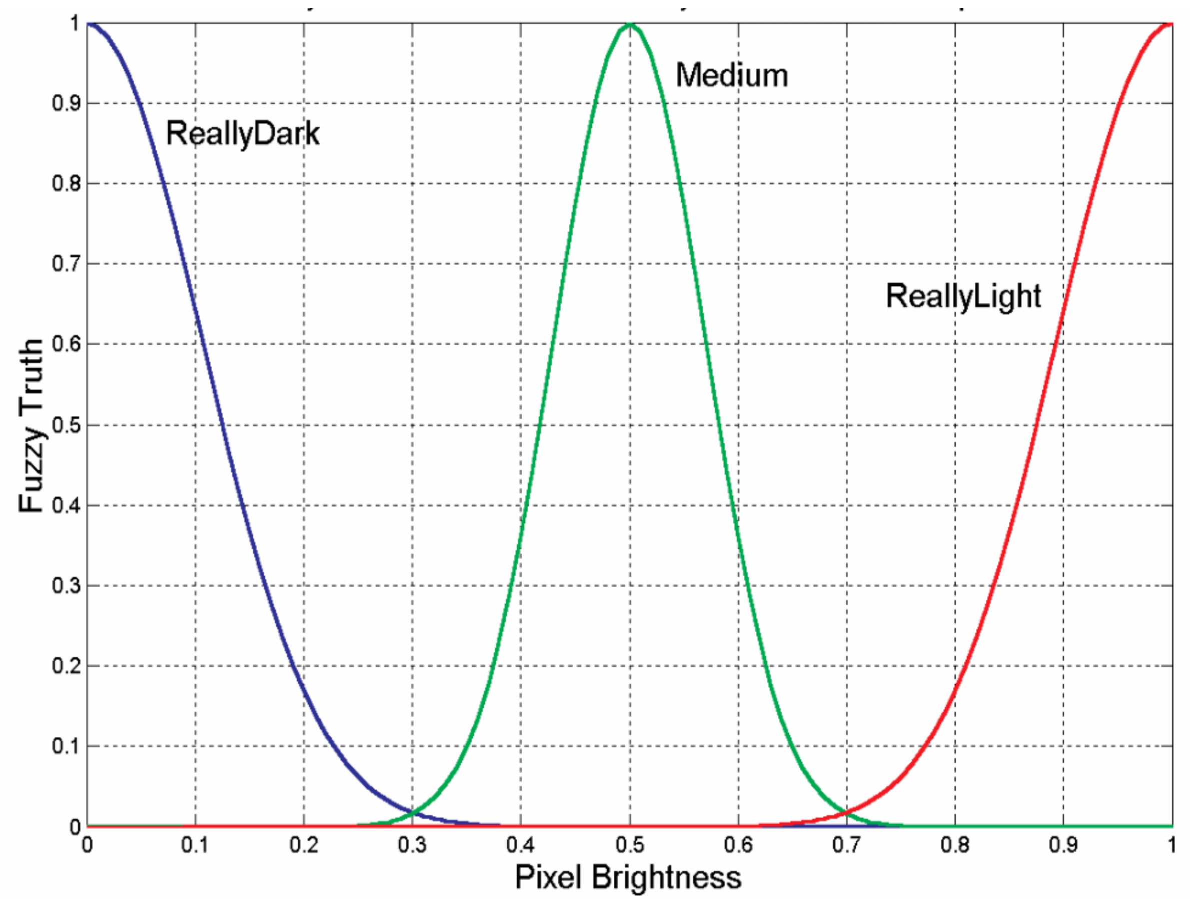

Inference Engine

2) **Redefinitions of basic operations**

Classical logic	Fuzzy logic
• Negation	
• And	
• Or	
• Implication	
• Modus ponens	

Inference Engine

3) **Decision matrix**

Input	Output
Light	Very light
Average	Average
Dark	Very dark

Power of fuzzy logic!
A fuzzy system is expressed directly in natural language

4) **Defuzzification !**

3) AI > deduction

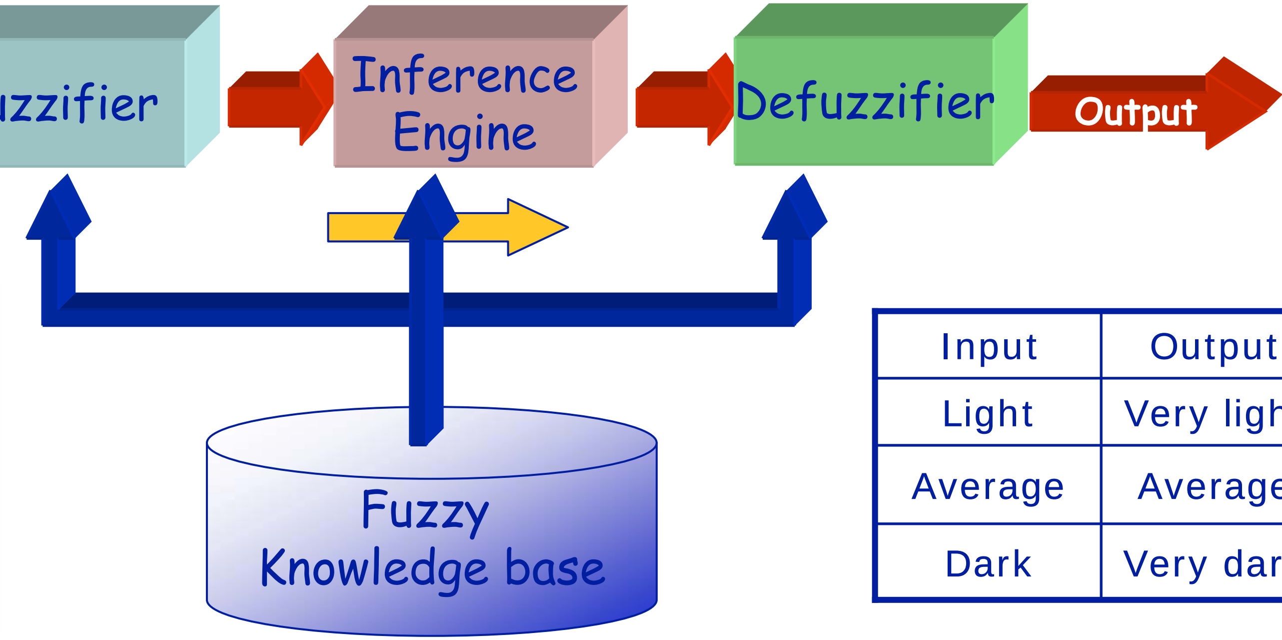

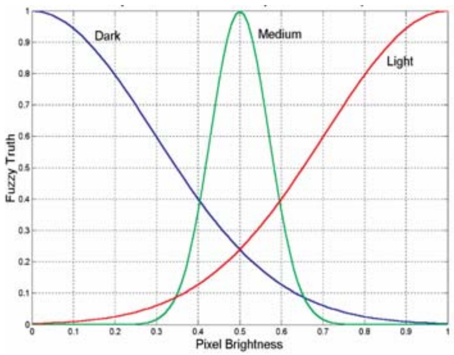

Input	Output
Light	Very light
Average	Average
Dark	Very dark

Exemple 2 : Décision du montant du pourboire à l'issue d'un repas au restaurant, en fonction de la qualité du service ressentie ainsi que de la qualité de la nourriture.

3 variables :

- Input 1 : qualité du service. Sous-ensembles : mauvais, bon et excellent.
- Input 2 : qualité de la nourriture. Sous-ensembles : exécrable et délicieux.
- Output : montant du pourboire. Sous-ensembles : faible, moyen et élevé.

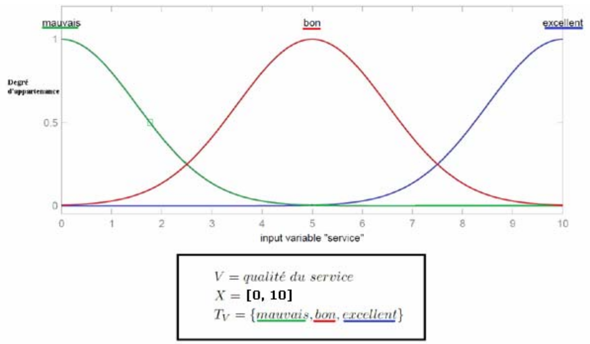

Input 1

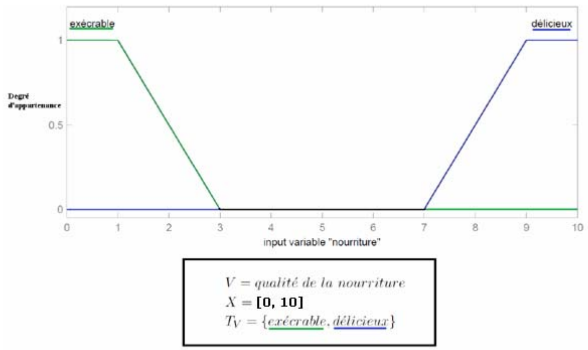

Input 2

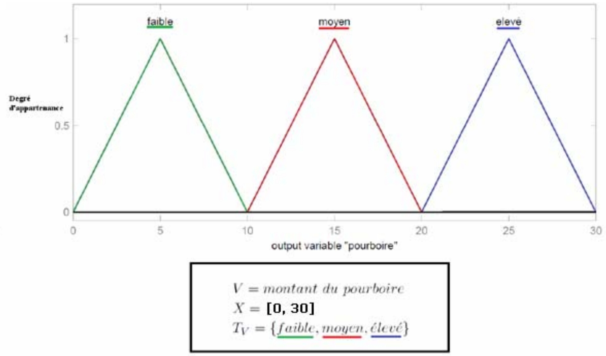

Output 1

3) AI > deduction

3 variables :

- **<u>Input 1 : qualité du service. Sous-ensembles : mauvais, bon et excellent</u>.**
- Input 2 : qualité de la nourriture. Sous-ensembles : exécrable et délicieux.
- Output : montant du pourboire. Sous-ensembles : faible, moyen et élevé.

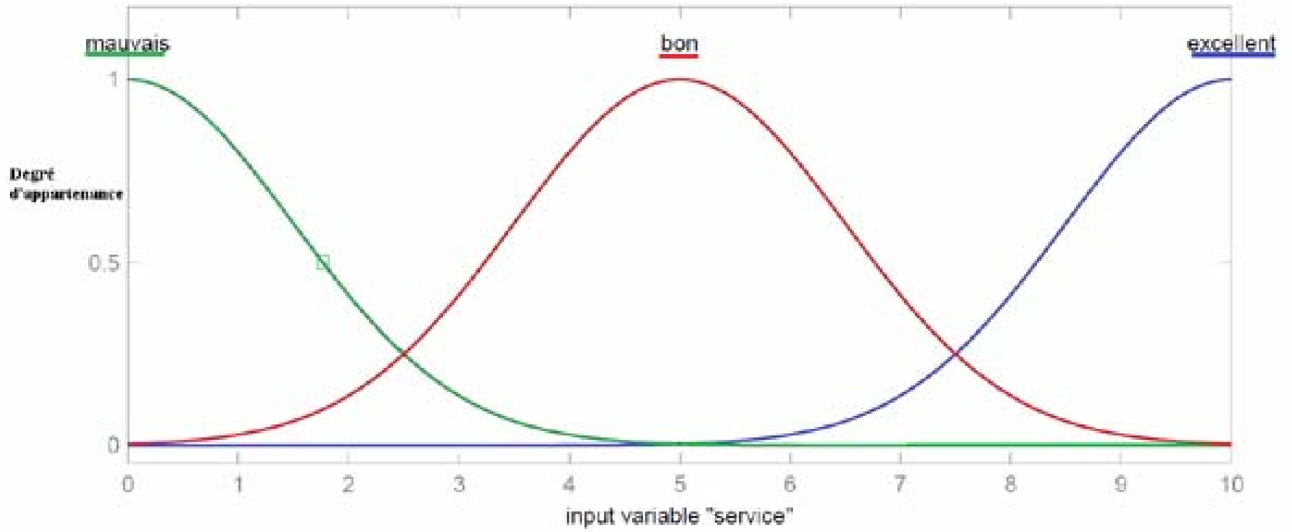

$$V = \textit{qualité du service}$$
$$X = [0, 10]$$
$$T_V = \{\underline{mauvais}, \underline{bon}, \underline{excellent}\}$$

3) AI > deduction

3 variables :

• Input 1 : qualité du service. Sous-ensembles : mauvais, bon et excellent.
• **Input 2 : qualité de la nourriture. Sous-ensembles : exécrable et délicieux**.
• Output : montant du pourboire. Sous-ensembles : faible, moyen et élevé.

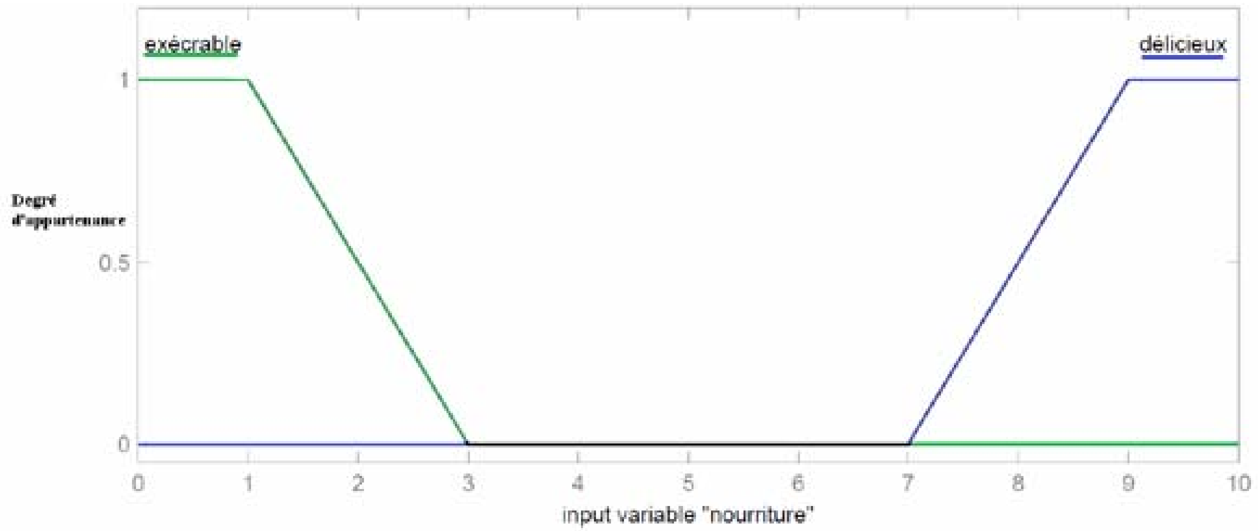

$V = qualit\acute{e}\ de\ la\ nourriture$
$X = [0, 10]$
$T_V = \{ex\acute{e}crable, d\acute{e}licieux\}$

3) AI > deduction

3 variables :

- Input 1 : qualité du service. Sous-ensembles : mauvais, bon et excellent**.**
- Input 2 : qualité de la nourriture. Sous-ensembles : exécrable et délicieux.
- **Output : montant du pourboire. Sous-ensembles : faible, moyen et élevé**.

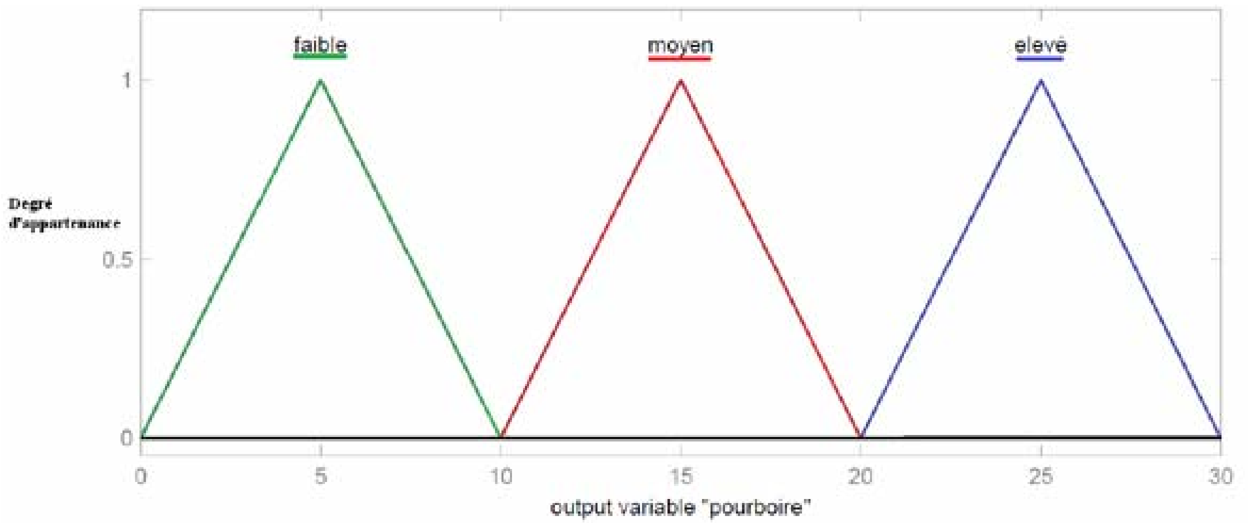

V = *montant du pourboire*
X = [0, 30]
T_V = {*faible*, *moyen*, *élevé*}

3) AI > deduction

Règle 1	Si le service est mauvais ou la nourriture est exécrable	alors le pourboire est faible.
Règle 2	Si le service est bon	alors le pourboire est moyen
Règle 3	Si le service est excellent ou la nourriture est délicieuse	alors le pourboire est élevé.

Service = 7.83
Nourriture = 7.32

Pourboire ??

Règle 1	Si le service est mauvais ou la nourriture est exécrable	alors le pourboire est faible.
Règle 2	Si le service est bon	alors le pourboire est moyen
Règle 3	Si le service est excellent ou la nourriture est délicieuse	alors le pourboire est élevé.

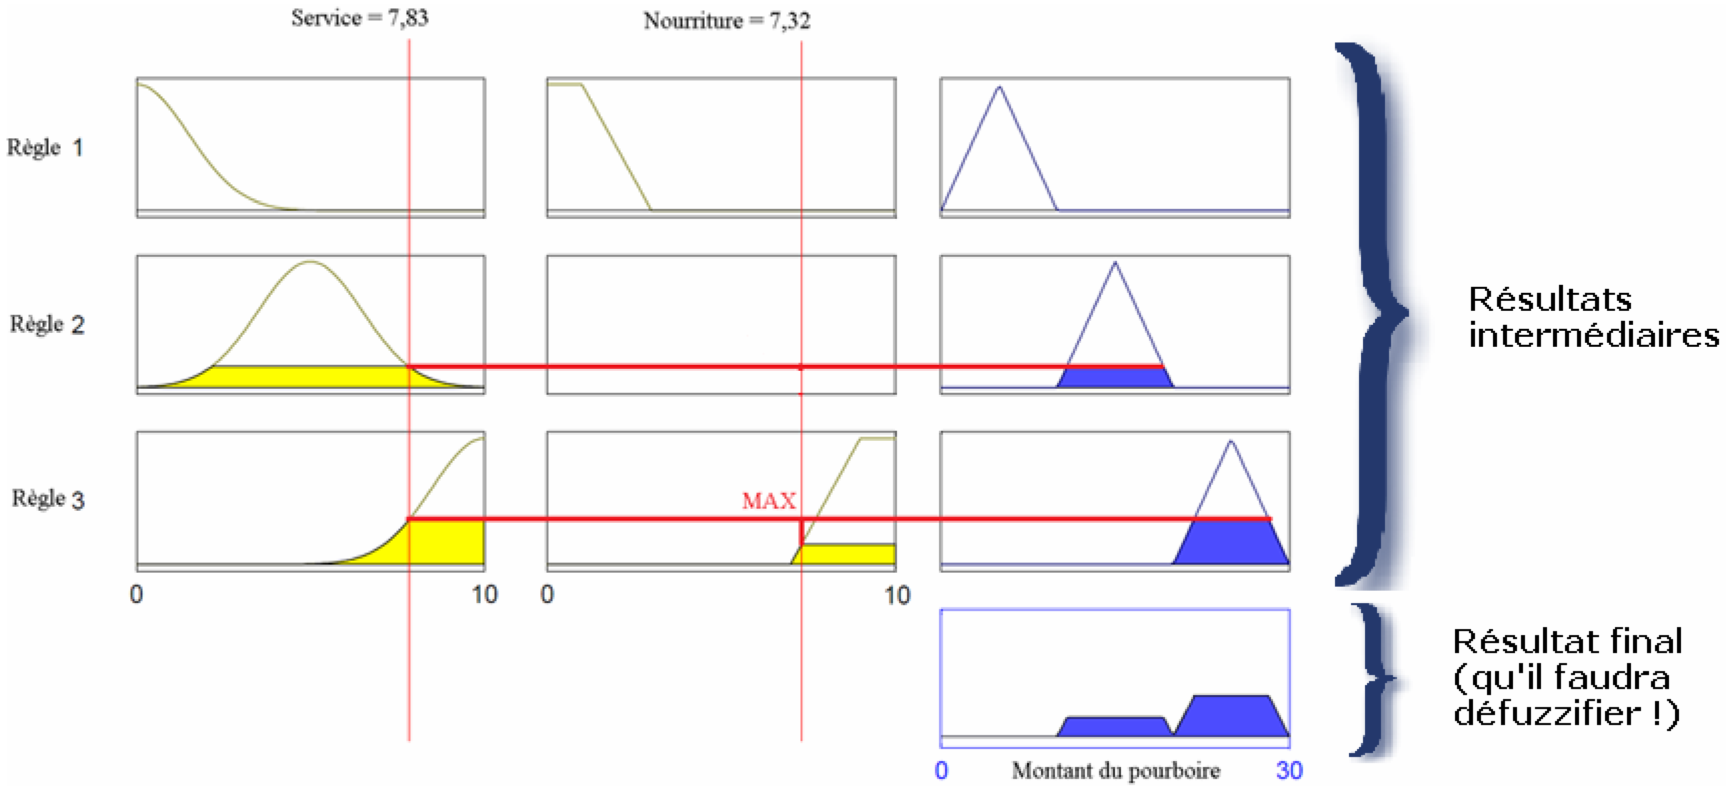

3) AI > deduction

Règle 1	Si le service est mauvais ou la nourriture est exécrable	alors le pourboire est faible.
Règle 2	Si le service est bon	alors le pourboire est moyen
Règle 3	Si le service est excellent ou la nourriture est délicieuse	alors le pourboire est élevé.

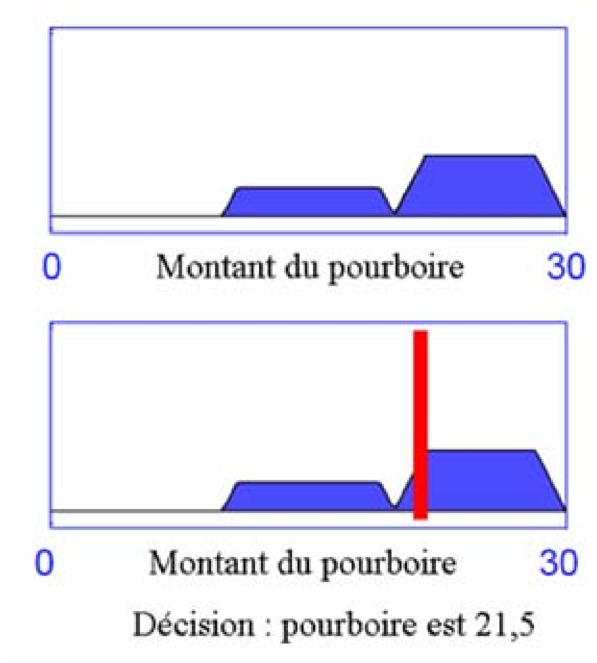

Défuzzification avec la méthode centre de gravité (COG)

$$Decision = \frac{\int_Y y \cdot \mu(y) \cdot dy}{\int_Y \mu(y) \cdot dy}$$

3) AI > deduction

Règle 1	Si le service est mauvais ou la nourriture est exécrable	alors le pourboire est faible.
Règle 2	Si le service est bon	alors le pourboire est moyen
Règle 3	Si le service est excellent ou la nourriture est délicieuse	alors le pourboire est élevé.

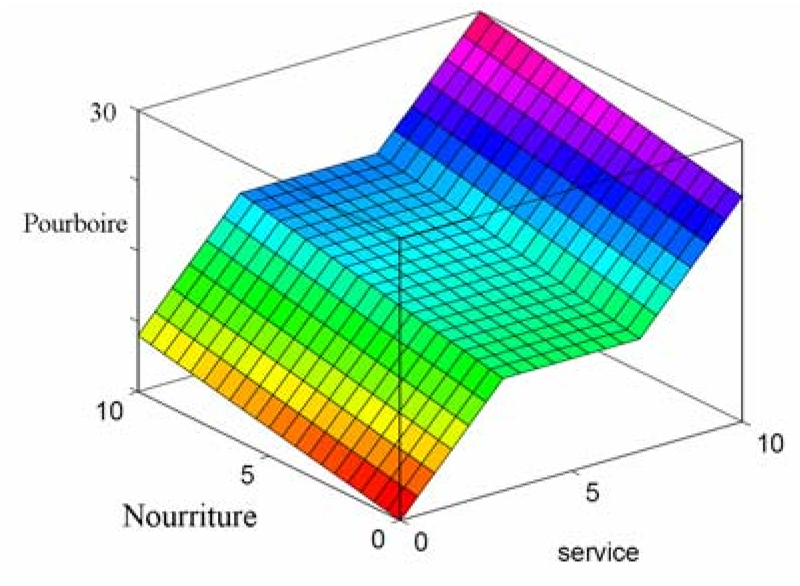

Ensemble des décisions d'un système se basant sur la logique classique

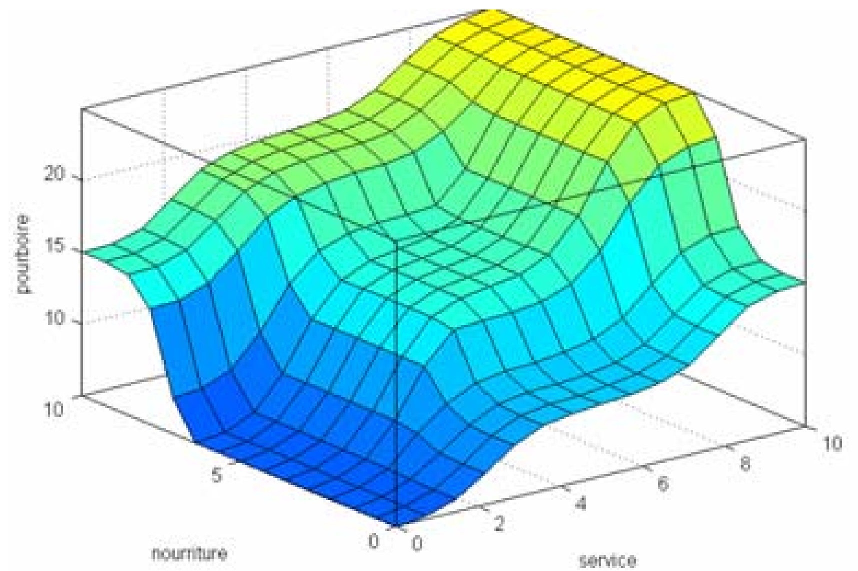

Ensemble des décisions d'un système flou

3) AI > deduction

Exemple 3 : pilote automatique

Règle 1	Si le feu est rouge...	si ma vitesse est élevée...	et si le feu est proche...	alors je freine fort.
Règle 2	Si le feu est rouge...	si ma vitesse est faible...	et si le feu est loin...	alors je maintiens ma vitesse.
Règle 3	Si le feu est orange...	si ma vitesse est moyenne...	et si le feu est loin...	alors je freine doucement.
Règle 4	Si le feu est vert...	si ma vitesse est faible...	et si le feu est proche...	alors j'accélère.

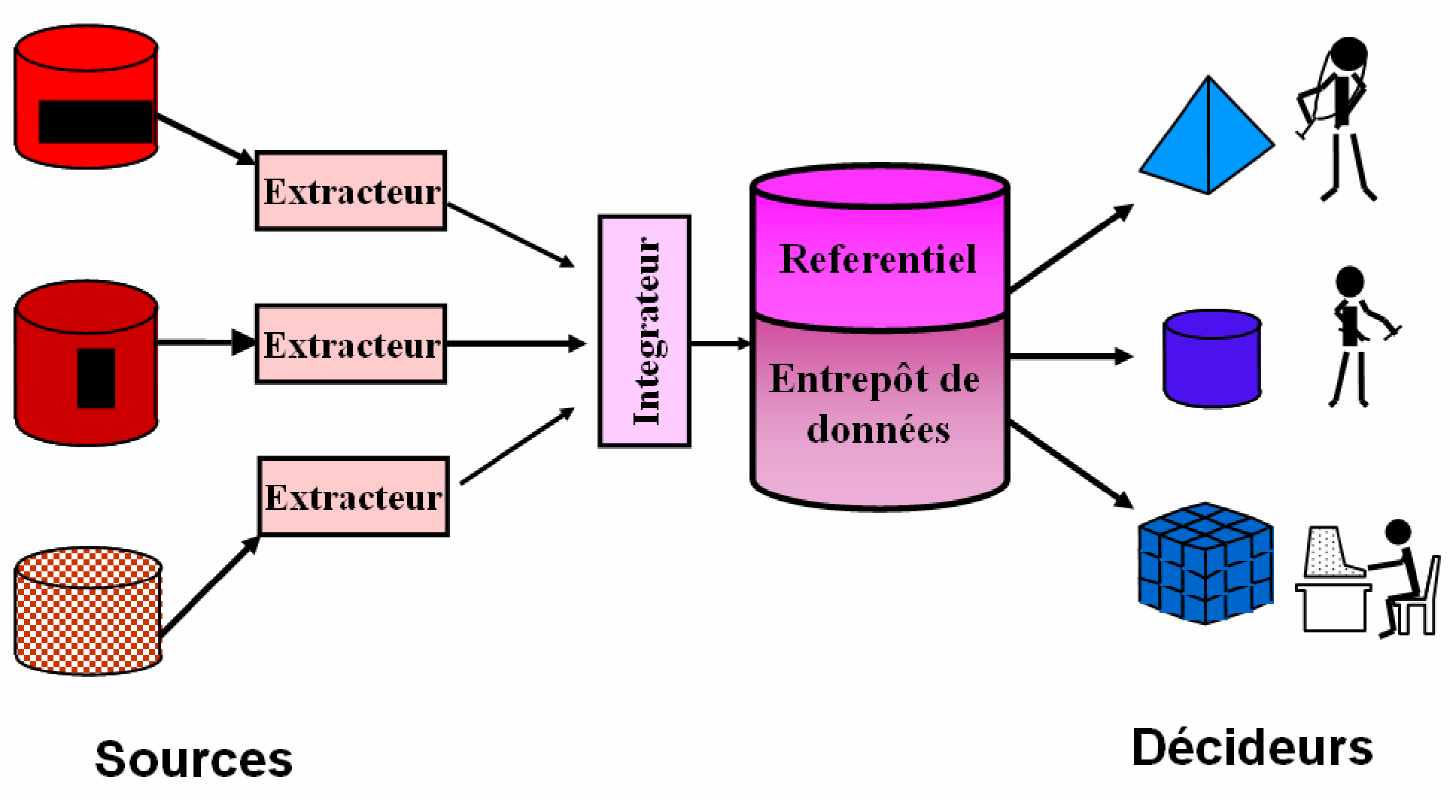

The decision-making chain

Table of Contents

Conclusion:

AI's aim: Integrating the approaches.

4) Conclusion & perspectives

Conclusion:

Regard AI as a toolset for your business needs.

Do computer programs have IQs?

jmc: No. IQ correlates well with various measures of success or failure in life, but making computers that can score high on IQ tests would be weakly correlated with their usefulness.

However, some of the problems on IQ tests are useful challenges for AI.

How far is AI from reaching human-level intelligence?
When will it happen?

jmc: Most AI researchers believe that new fundamental ideas are required, and therefore it cannot be predicted when human-level intelligence will be achieved.

Conclusion:

Are current computers the right kind of machine to be made intelligent?

Yes and no.

Yes: All computations that have been observed so far in human brains can be reproduced on computers.

No: Yet some AI researchers consider the Von Neumann architecture as a computational and architectural bottleneck for AI. For example, IBM's SyNAPSE project aims to design computing chips that map closer human brain functions.

Conclusion:

Are computers fast enough to be intelligent?

We don't know for we haven't yet figured out the right algorithms to compute intelligence. ☹

4) Conclusion & perspectives

10 important differences between brains and computers:

1. Brains are analogue; computers are digital
2. The brain uses content-addressable memory
3. The brain is a massively parallel machine; computers are modular and serial
4. Processing speed is not fixed in the brain; there is no system clock
5. Short-term memory is not like RAM
6. No hardware/software distinction can be made with respect to the brain or mind
7. Synapses are far more complex than electrical logic gates
8. Unlike computers, processing and memory are performed by the same components in the brain
9. The brain is a self-organizing system
10. Brains have bodies

Conclusion:

Can machines think?

Edsger Dijkstra: Whether computers can think is like the question of whether submarines can swim. In English, we say submarines don’t swim, but we say aeroplanes do fly. In Russian, they say submarines do swim.


franck.dernoncourt@gmail.com