The game of Bridge: a challenge for ILP

S. Legras, C. Rouveirol, V. Ventos

Véronique Ventos
LRI Univ Paris-Saclay
vventos@nukk.ai

Games

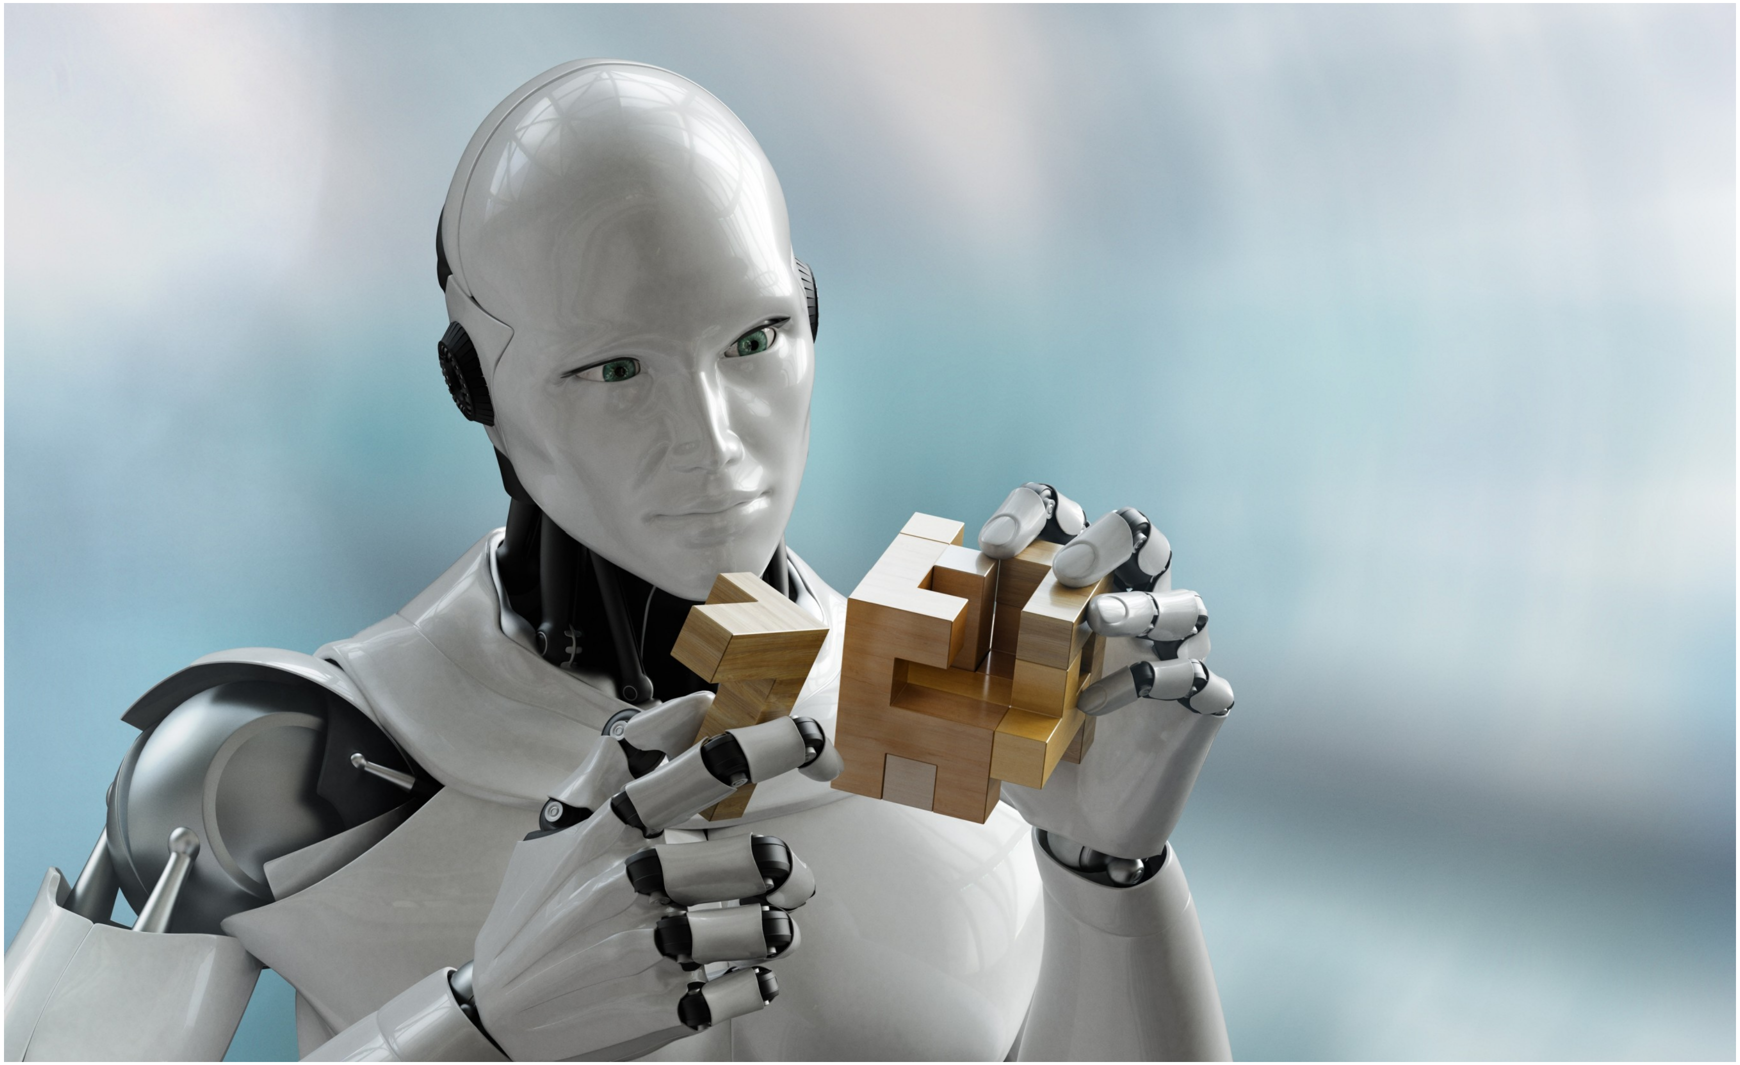

Interest of games for AI

Excellent field of experimentation

Problems are easier to understand and to model than in real life (limited number of simple rules, in-depth human analysis over time, ...)

Game successes have always been milestones for AI

Go = major challenge

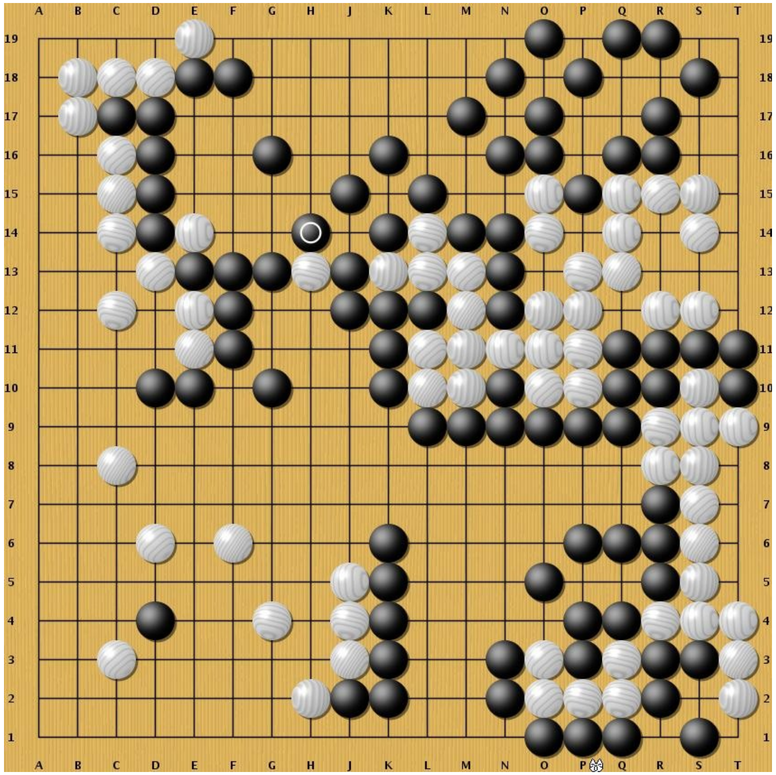

Until 2006 : level of an average amateur player

Crazy Stone, Mogo : Go AI with strategies combining several ML methods

AlphaGo (Deep Mind, google)

March 2016 : alphaGo won 4 to 1 against Lee Sedol

May 2017 : alphaGo Master has defeated Ke Jie, the world's number one Go player

October 2017 :

Zero vs Lee : 100-0

Zero vs Master : 89-11

Next Step ?

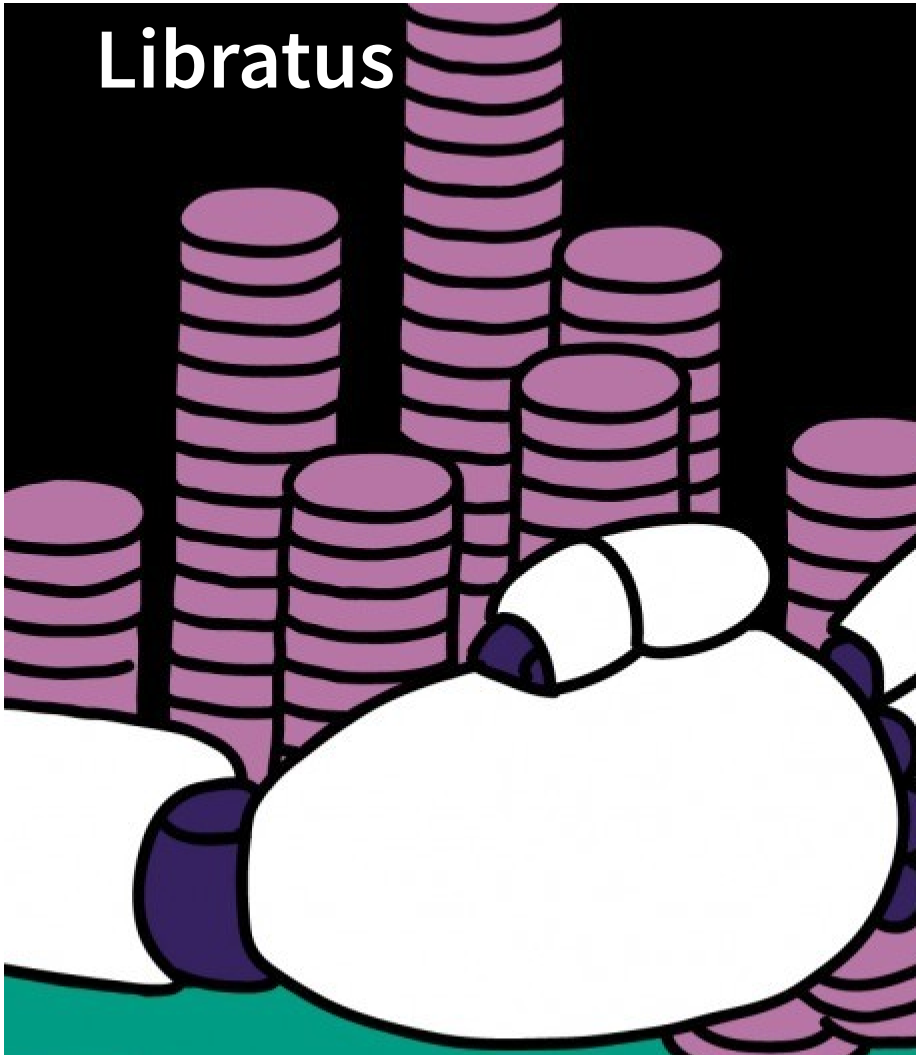

In January 2017, the Poker AI Libratus developed by Carnegie Mellon University won a heads-up no-limit Texas hold'em poker event against four of the best professional players

Poker vs ...

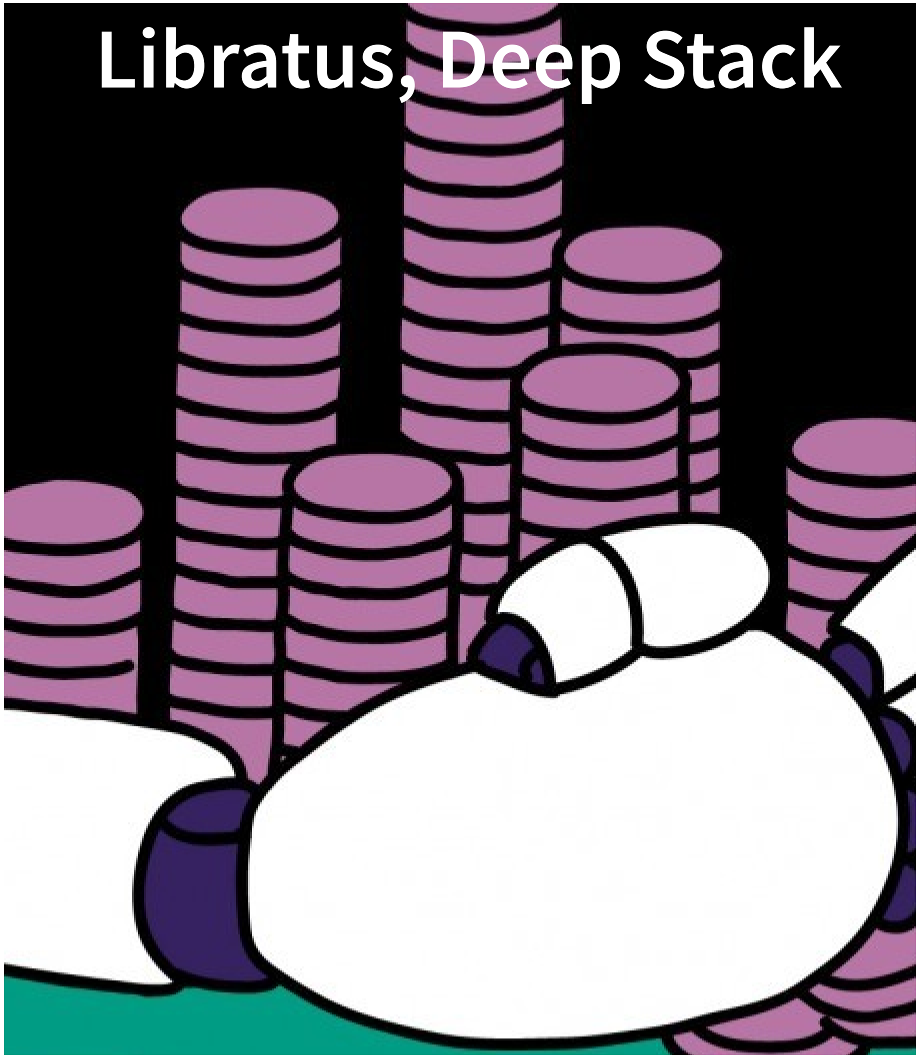

Poker vs bridge

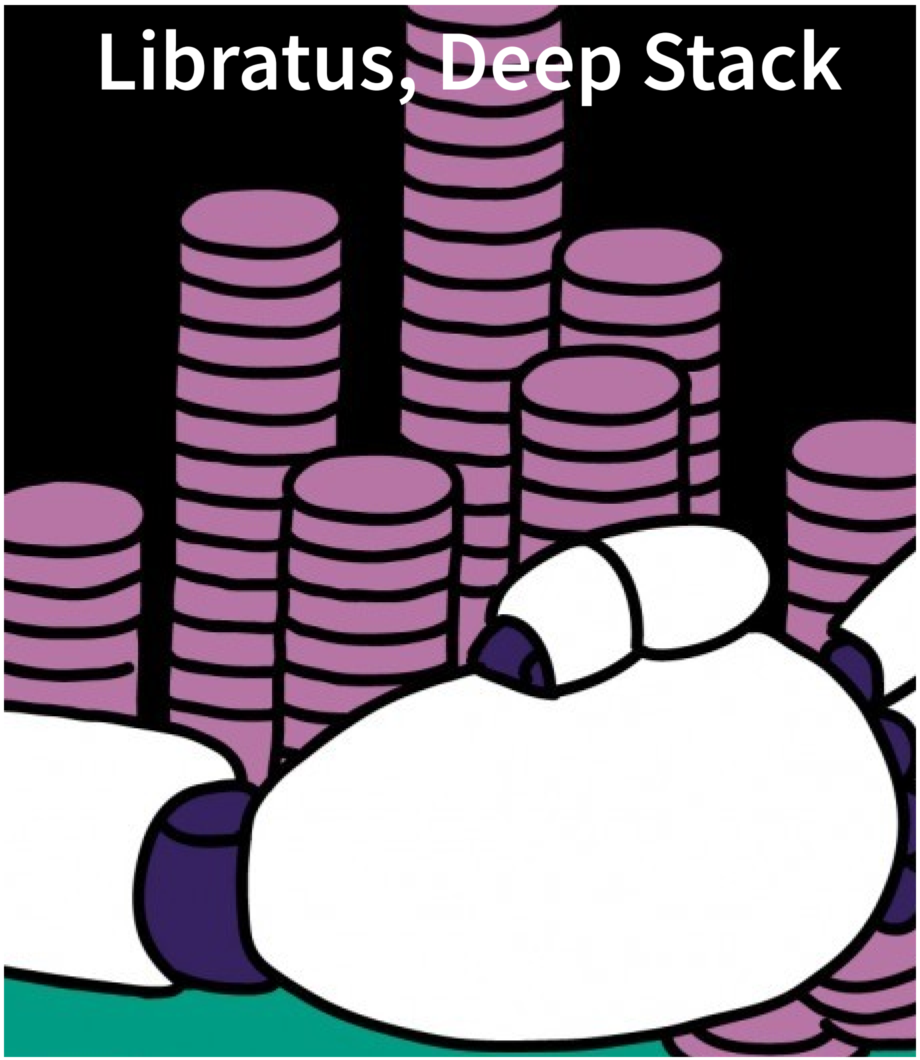

Bridge is the next challenge for AI

Bridge robots : far from best human players (quite similar to go programs before 2006)

Our conviction : « solving » Bridge is a big step between AI such AlphaGo and a General Artificial Intelligence

Bridge needs symbolic approaches

The game of Bridge is an application needing more than black box approaches

Need of explanations: at some point players **must** explain their actions

To "crack" a game, a program needs to play optimally

but …

To "solve" it the program's play must also be explainable in human understandable terms

Part 1: Bridge

Part 2: Opening bid problem

Part 3: ML settings and experiments

Part 4: Brief conclusion

Part 1: Bridge

Usual vision of bridge

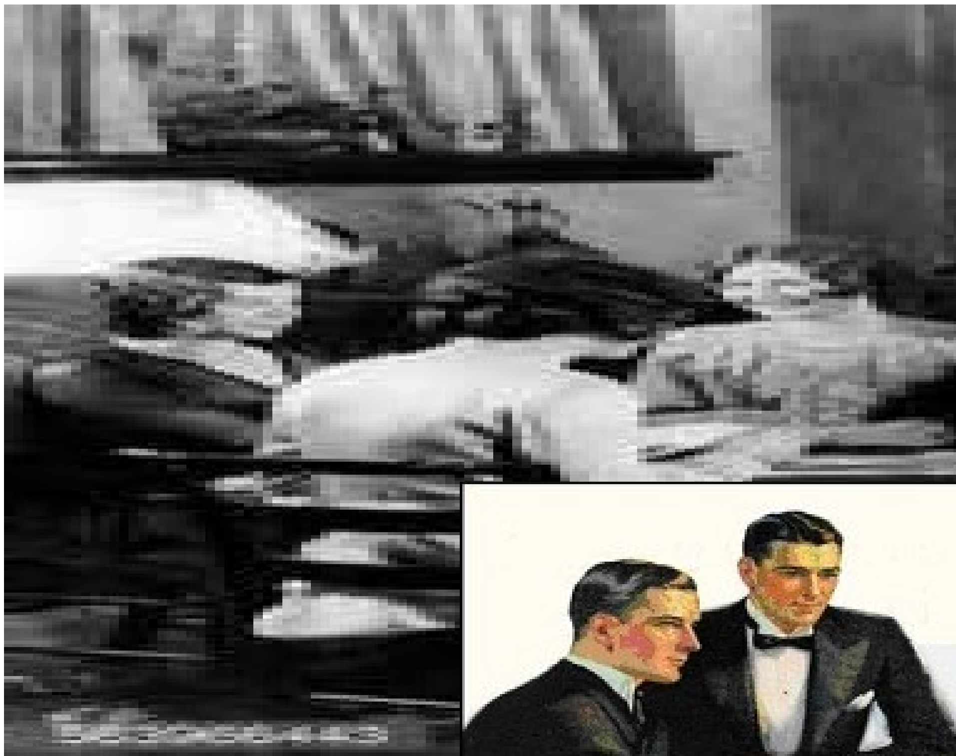

Bridge in 2018

World championships

Wroclaw 2016

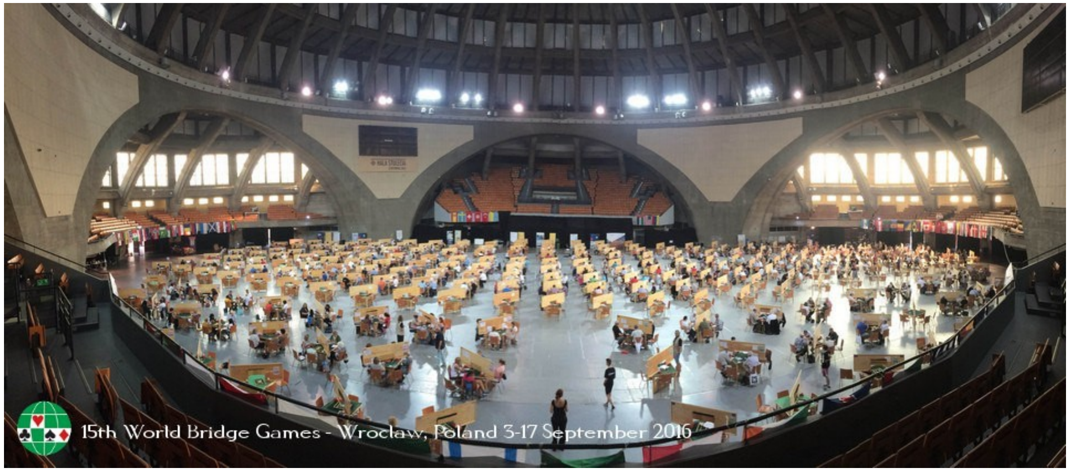

Lyon 2017

Bridge is tough but …

Bridge in short

Trick-taking game, played with 52 standard cards opposing two pairs of players

Cards are dealt randomly to the four players

Each of them only sees his hand (13 cards)

Incomplete information game : players do not have common knowledge of the game being played

Two steps: the bidding phase then the card play

Bidding phase

Coded language used by players to pass information to their partner about their hand

Goal : reach an optimal contract. The contract specifies the minimum number of tricks among the thirteen to be won in the second phase

Card play

Goal : to fulfill (or to defeat for the opposite side) the contract reached during the bidding phase

Part 2: Opening bid problem

Set of bidding cards

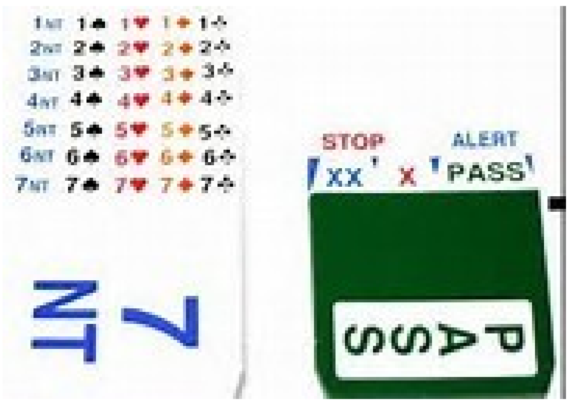

35 symbols of bid : from 1 to 7NT

Cards for other calls :

Pass, X, XX

Stop, Alert

There exist many bidding systems assigning meanings to bids : e.g. Acol , Standard American, Precision Club , Polish Club

Standard American Yellow Card

SAYC (Standard American Yellow Card) is a bidding system which is prevalent in online bridge games

My hand : ♠AK83 ♥QJ2 ♦1076 ♣AJ8

Pass ? 2NT ? 1NT ?

My bid :

1. Counting the high card points (HCP) of my hand

with Ace : 4, King : 3, Queen : 2, Jack : 1

♠AK83 ♥QJ 2 ♦1076 ♣AJ8

1. Counting the high card points (HCP) of my hand

with Ace : 4, King : 3, Queen : 2, Jack : 1

♠AK83 ♥QJ 2 ♦1076 ♣AJ8 ➡ 15 HCP

1. Counting the high card points (HCP) of my hand

with Ace : 4, King : 3, Queen : 2, Jack : 1

♠AK83 ♥QJ 2 ♦1076 ♣AJ8 → 15 HCP

2. Determining the hand pattern: distribution of the thirteen cards in a hand over the four suits

♠AK83 ♥QJ 2 ♦1076 ♣AJ8

1. Counting the high card points (HCP) of my hand

with Ace : 4, King : 3, Queen : 2, Jack : 1

♠AK83 ♥QJ 2 ♦1076 ♣AJ8 → 15 HCP

2. Determining the hand pattern: distribution of the thirteen cards in a hand over the four suits

♠AK83 ♥QJ 2 ♦1076 ♣AJ8 → 4-3-3-3

1. Counting the high card points (HCP) of my hand

with Ace : 4, King : 3, Queen : 2, Jack : 1

♠AK83 ♥QJ 2 ♦1076 ♣AJ8 ➡ 15 HCP

2. Determining the hand pattern: distribution of the thirteen cards in a hand over the four suits

♠AK83 ♥QJ 2 ♦1076 ♣AJ8 ➡ 4-3-3-3

3. Classifying my hand : balanced (no short suit) or unbalanced ?

1. Counting the high card points (HCP) of my hand

with Ace : 4, King : 3, Queen : 2, Jack : 1

♠AK83 ♥QJ 2 ♦1076 ♣AJ8 → 15 HCP

2. Determining the hand pattern: distribution of the thirteen cards in a hand over the four suits

♠AK83 ♥QJ 2 ♦1076 ♣AJ8 → 4-3-3-3

3. Classifying my hand : balanced (no short suit) or unbalanced ? → balanced

Using SAYC opening rules

Finally :

Choosing a rule

→ Bid 1NT with 15-17 HCP, balanced

1NT :)

♠AK83 ♥QJ2 ♦1076 ♣AJ8

Opening problem in Bridge

'Should I bid or pass with a limit hand ?'

The first bid is called the opening

In SAYC, 1-of-a-suit opening requires at least 12 HCP

but …

Opening problem in Bridge

'Should I bid or pass with a limit hand ?'

The first bid is called the opening

In SAYC, 1-of-a-suit opening requires at least 12 HCP

but … experts allow themselves to deviate slightly from the rule by opening some 11 HCP hands

This decision is very important (big impact on the final scoring)

Part 3: ML settings and experiments

Machine Learning setting

The opening bid problem is a binary classification problem where Task T consists in predicting if a given expert opens or passes with a 'limit' hand according to a bridge situation.

Input : set of *n* labeled examples (x_i,$class_i$)

Output : f(x) assigning each example x to its class + (open) or - (pass)

DataSets

The goal is to learn rules linked to experts' decisions

Random generation of 6 sets of unlabeled examples

Labeling by 4 Bridge experts (among the best 100 players of their country) using a system requiring 12 HCP for opening

Important remarks

Experts have the same level but different styles

Decisions vary a lot from an expert to another

Learning of 'personal rules', different learning tasks

Consistency : the same expert can make different decisions facing the exact same situation

Tagging Interface

To bid or not to bid ?

♠ A 10 6 4
♥ Q 3 2
♦ K Q 10 2
♣ 9 6

North	East	South	West
PASS	PASS	PASS	?

Yes No Oops

Summary and statistics

6 samples sets, 4 experts, aggressiveness

Labeled set	S_1	S_2	S_3	S_4	$S_{5,1}$	$S_{5,2}$	S_6
Unlabeled set	R_1	R_2	R_3	R_4	R_5	R_5	R_5
Expert	E_1	E_2	E_2	E_3	E_4	E_4	E_4
Size	1000	1000	970	790	1222	1222	1079
Pos./Neg.	768/232	681/319	540/430	603/187	681/541	582/640	560/519
Pos rate (%)	76.80	68.10	55.67	76.33	55.72	47.63	51.90

Table 1. Samples sets

Experts' consistency

Labeled set	S_1	S_2	S_3	S_4	$S_{5,1}$	$S_{5,2}$	S_6
Unlabeled set	R_1	R_2	R_3	R_4	R_5	R_5	R_5
Expert	E_1	E_2	E_2	E_3	E_4	E_4	E_4
Size	1000	1000	970	790	1222	1222	1079
Pos./Neg.	768/232	681/319	540/430	603/187	681/541	582/640	560/519
Pos rate (%)	76.80	68.10	55.67	76.33	55.72	47.63	51.90

Table 1. Samples sets

3 ML systems

The Support Vector Machine (SVM) learner and the ILP systems (Aleph and Tilde) used in the experiments are both state of the art ML systems

Aleph : learning from entailment (set of prolog rules)

Tilde : learning from interpretations (relational decision tree)

Background knowledge : set of definite clauses

Expected ILP added value

Flexibility : allows experimenting with various abstractions of examples description through the use of background knowledge

Explainability : learned models are readable by experts who can then help us update current BK

Designing BK

Designing the BK stems from a joint work between experts and us in order to achieve both an acceptable bridge-wise representation and an acceptable learning performance

First representation (propositional)

BK_0: The following predicates form the common background for all BKs (E represents the example primary key).

- $open(E, Class)$ with $Class = pos$ or neg
- $has\_hand(E, H)$ where H is a list of card constants representing the hand's example
- $has\_card(E, C)$ is $true$ if C is a card occurring in E's hand
- $position(E, P)$ with P a digit between 1 and 4 representing the position of the expert
- $vuln(E, V1, V2)$ with $V1$ and $V2 = g$ (not vulnerable) or r (vulnerable) representing the vulnerability of the player and of its opponents.

Example 1 using BK0

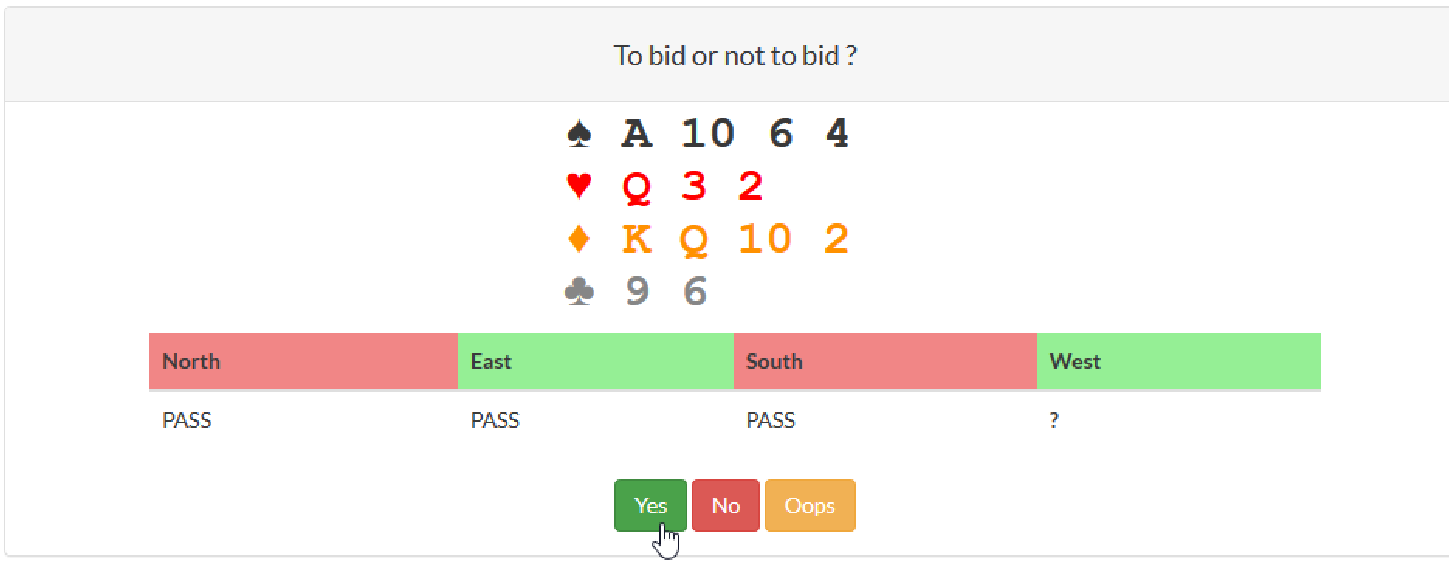

$open(e1, pos).$ $has\_hand(e1, [c6, c9, d2, d10, dq, dk, h2, h3, hq, s4, s6, s10, sa]).$

$position(e1, 4).$ $vuln(e1, g, r).$

$has\_card(e1, c6).$ $has\_card(e1, c9).$ ... $has\_card(e1, sa).$

King of heart description

has-card(h1, hk)

card(hk)

has_suit(hk,heart)

has_rank(hk,k)

card (X) $\wedge$ has_suit(X,heart) $\rightarrow$ major(X)

card(X) $\wedge$ has_rank(X,k) $\rightarrow$ honor(X)

Saturation : major(hk), honor(hk)

Relational representation

BK1 extract (card is structured and abstracted)

has_suit(Card,Suit), has_rank(Card,Rank)

honor(Card) / small card(Card)

minor(Card) / major(Card)

nb(E,Suit,Num)

lteq(Num, Num), gteq(Num, Num)

Relational representation

BK1 extract (abstraction of Hand description)

distribution(E, [Num,Num,Num,Num])

balanced(E) / semi_balanced(E) / unbalanced(E)

plusvalue(E)/moinsvalue(E) (e.g. at least two honors in a suit with at-least 5 cards)

BK2: all BK1 predicates + list_honor(E, Suit, ListH)

Partial relational description of example 1

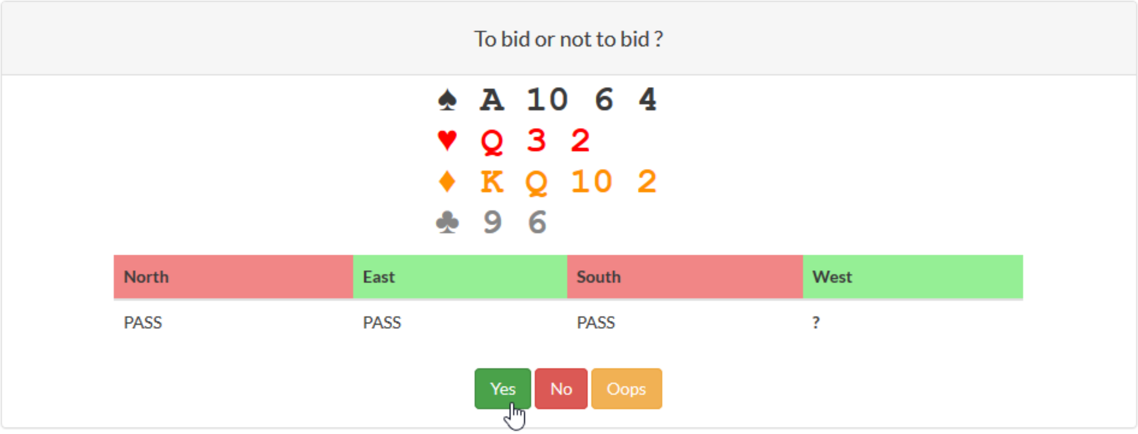

nb(e1,Spade,4)

nb(e1,Heart,3)

distribution(e1,[4,4,3,2])

balanced(e1)

plusvalue(e1)

Experiments

We have made experiments on labeled sets with several BK of increasing expressivity using SVM, Aleph and Tilde

Accuracy comparaison of SVM, Aleph and Tilde

For ILP systems :

Complexity of the learned models

Relevance according to experts' feedback

Accuracy of learned models

10-fold cross validation

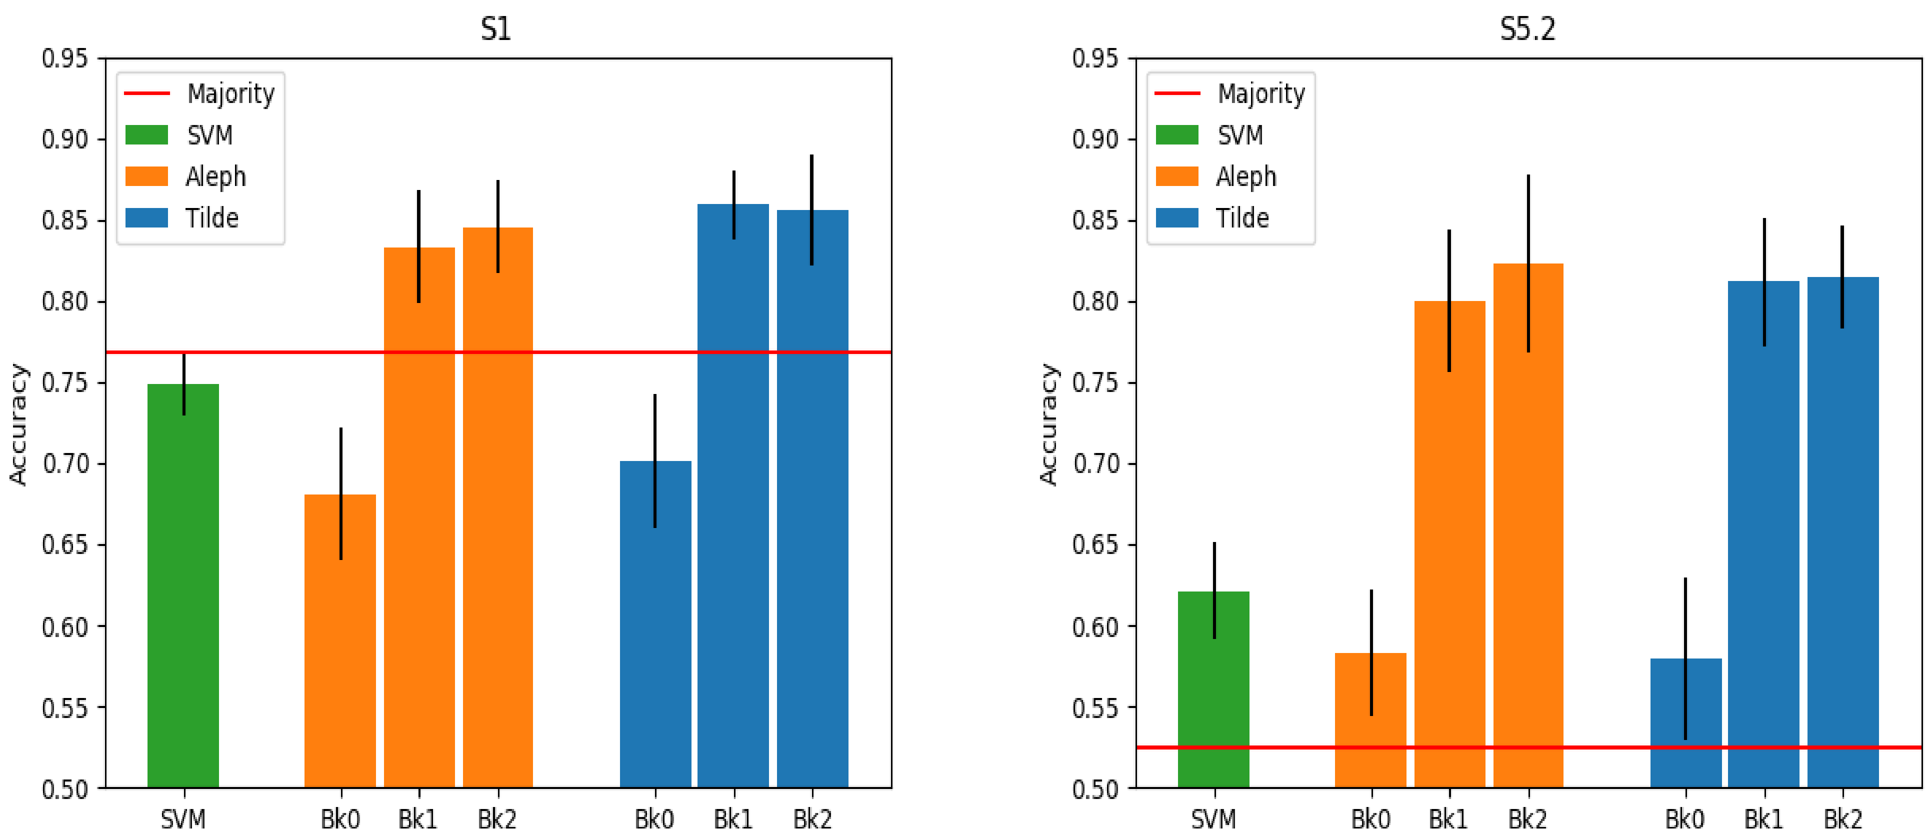

Accuracy of learned models

The performance with propositional BK (BK0) is low as expected

Models learned with BK1 and BK2 have significant better results

No significant difference between BK1 and BK2

Performance of Aleph and Tilde are close

Similar conclusions on other datasets (results available on our website)

Complexity of learned models

Nb of rules in terms of the size of the training set

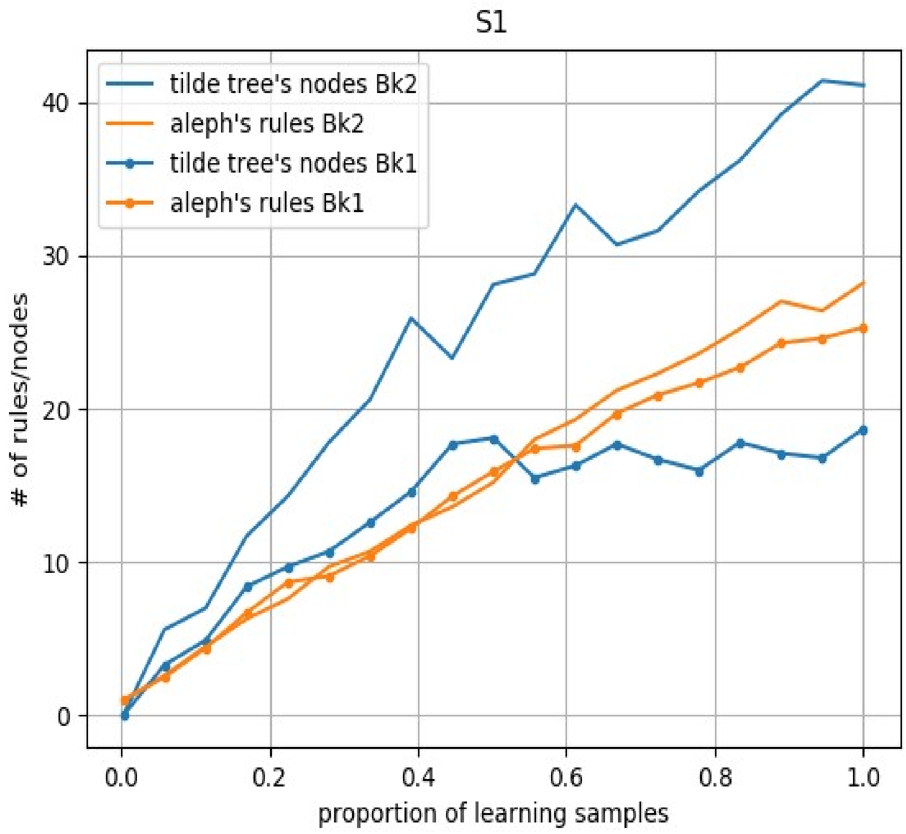

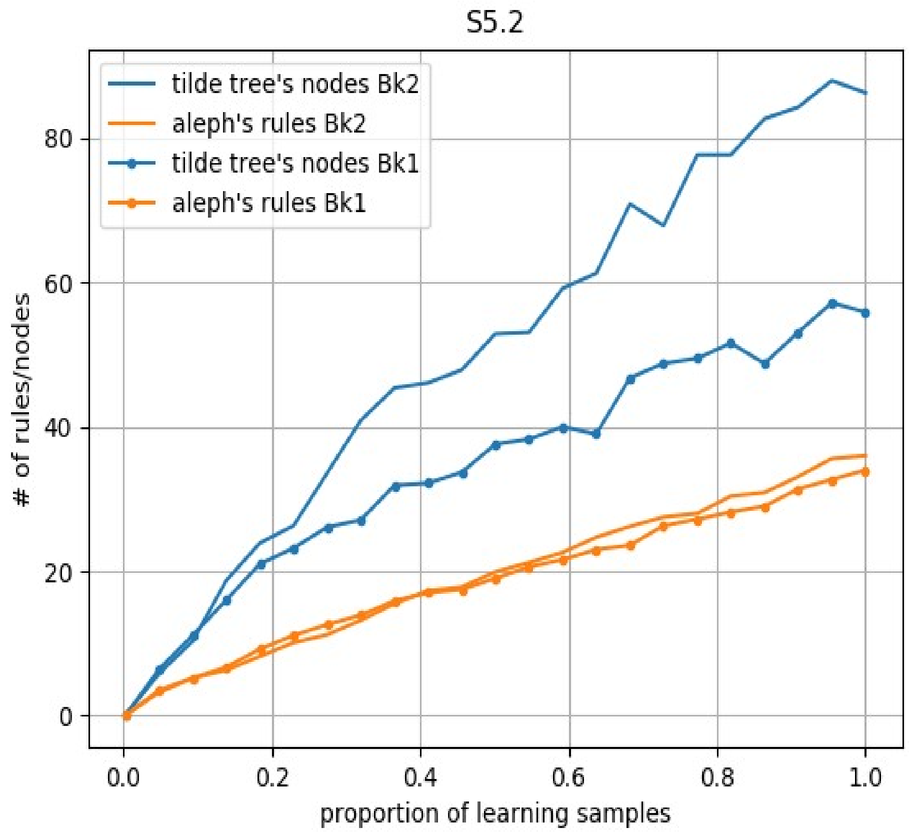

Complexity of learned models

The number of rules regulary increases for Aleph whereas its performance is stable (overfitting?)

The size of Tilde's models stabilizes for BK1 when it nearly reaches its best performance

BK2 seems less adapted for Tilde (bigger complexity with similar performance)

Both ILP systems reach a good performance while seing few examples and with small models

Relevance: Expert feedback

Some of the rules produced are of the 'common bridge knowledge' type whereas the others are more subjective and personal

R1 : open(A) :- plusvalue(A), position(A,3)

R2 : open(A) :- nb(A,spade,B), gteq(B,4), position(A,4)

Famous bridge rule known as ‘the rule of 15’

Intuitive vs analytical mind

Tilde : the complexity of the model learned is significantly different from an expert to another

Relationship between this complexity and the expert’s way of thinking

(e.g. E1 has an analytical mind, his DT is very concise, E4 is more intuitive, he is a slow player, his DT is two times larger and generated rules are too specific)

What's in an expert's Mind ?

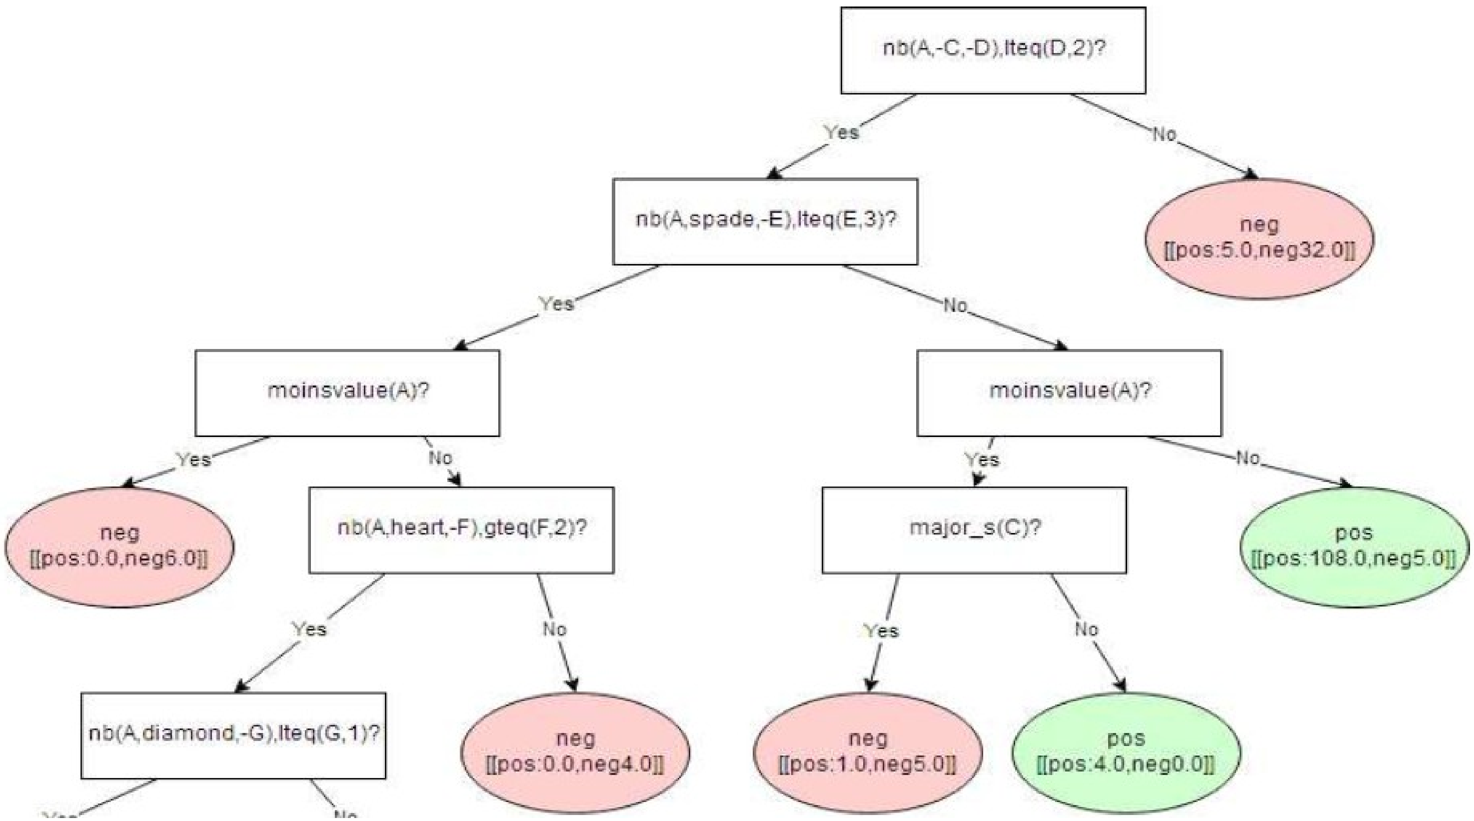

E1 First order logical decision tree

E1 feedback

The first node has been validated by E1 as the first criteria of his decision

Several rules have been described as ‘excellent’

The global vision of the DT appeared to him congruent with his approach to the problem

Before the experiments E1 was not able to explain clearly his decision-making process

Bridge experts have black-box approach :)

Part 4: Brief conclusion

Different skills

Being a good bridge player requires :

- **depth of analysis**
- **reasoning with incomplete information**
- **ability to establish a diagnosis based on different sources**
- **evaluation of opponent's level and psychology**
- **communication with partner etc**

νBridge Project

2015-2017: AlphaBridge academic Project Univ Paris Saclay

(http://vvopenai.monsite-orange.fr/)

2018-… : νBridge project designed by NukkAI to solve the game of bridge by defining a hybrid architecture including recent numeric and symbolic Machine Learning modules

NukkAI : a private AI Lab

Building next generation AI

Cofounded with

JB Fantun in may 2018

Web site :

www.nukk.ai

νBridge architecture

Hybrid architecture combining different AI paradigms: Symbolic Reinforcement Learning, Description Logics, Planning in MDP, POMDP, Deep Learning , (Probabilistic) Inductive Logic Programming

Symbolic modules

Main goal : use formalisms understandable for humans

Bridge Background Knowledge (BK)

Decision making rules

Adaptation, automatic update of set of rules

Transfer Learning

Approaching the real situation

Throughout the game, the hidden information is reduced

The main goal of each player consists in 'rebuilding' the hidden hands in order to make decisions

Bridge is probabilistic

Rebuilding is based on probabilistic reasoning

A= ‘Opponent holds king of club’

B= ‘My partner holds king of club’

C=’Opponent holds 3 cards in club and my partner holds 2 cards in club’

p(A)= p(B)=1/2   P(A/C)=3/5

Each new information modifies the probability of the distribution of the hidden cards and influences the player’s strategy

It was difficult at first to convince people that Bridge was more than juste a game

It is still difficult to convince people that hybrid approach is welcome

But ...

It was difficult at first to convince people that bridge was more than juste a game

It is still difficult to convince people that hybrid approach is welcome

But … Bridge is a killer application for that

NukkAI collaborations

Bridge is a great challenge for AI and much work related to the definition of a Bridge AI remains to be done

Collaborations are welcome


Thanks
For
Listening
Any
Questions
KeepCalmAndPosters.com

http://vvopenai.monsite-orange.fr/

Home | Research Areas | AlphaBridge motivations | AlphaBridge project ▼

Véronique Ventos

PhD in Artificial Intelligence (1997)

Brief bio:

I have been an associate professor at Paris Saclay University, France since 1998.

1998-2015: member of LaHDAK (Large-scale Heterogeneous DAta and Knowledge).

Since 2015, I carry my research inside the team A&O (Machine Learning and Optimization), at Laboratory of Computer Science (LRI).

AlphaBridge Project

I started playing bridge in 2004 and am now 59th French woman player out of 48644 players.

In 2015, I set up the **AlphaBridge** project combining my two passions. AlphaBridge is dedicated to solve the game of bridge by defining a hybrid architecture including recent numeric and symbolic Machine Learning modules.

AI winter is not coming (back) :)