Corpus-Based Incremental Intention Recognition via Bayesian Network Model Construction

Han The Anh (h.anh@fct.unl.pt)
Luís Moniz Pereira (lmp@di.fct.unl.pt)

CENTRIA /UNL

ICAPS "Goal, Activity and Plan Recognition" workshop
Freiburg, Germany, 12 June, 2011

Outline

Why Intention Recognition?

- Intention recognition (IR): inferring intentions of a single agent (individual) or a group of agents (collective) based on their observed actions.
- Why IR is necessary? Acting on environment, agent may have to deal with other agents
  - Ease interactions, improve cooperation and coordination, especially when communication is limited.
  - Defend from potential hostile behaviors, plan in advance to take advantage.

Intention recognition systems: Typical Components

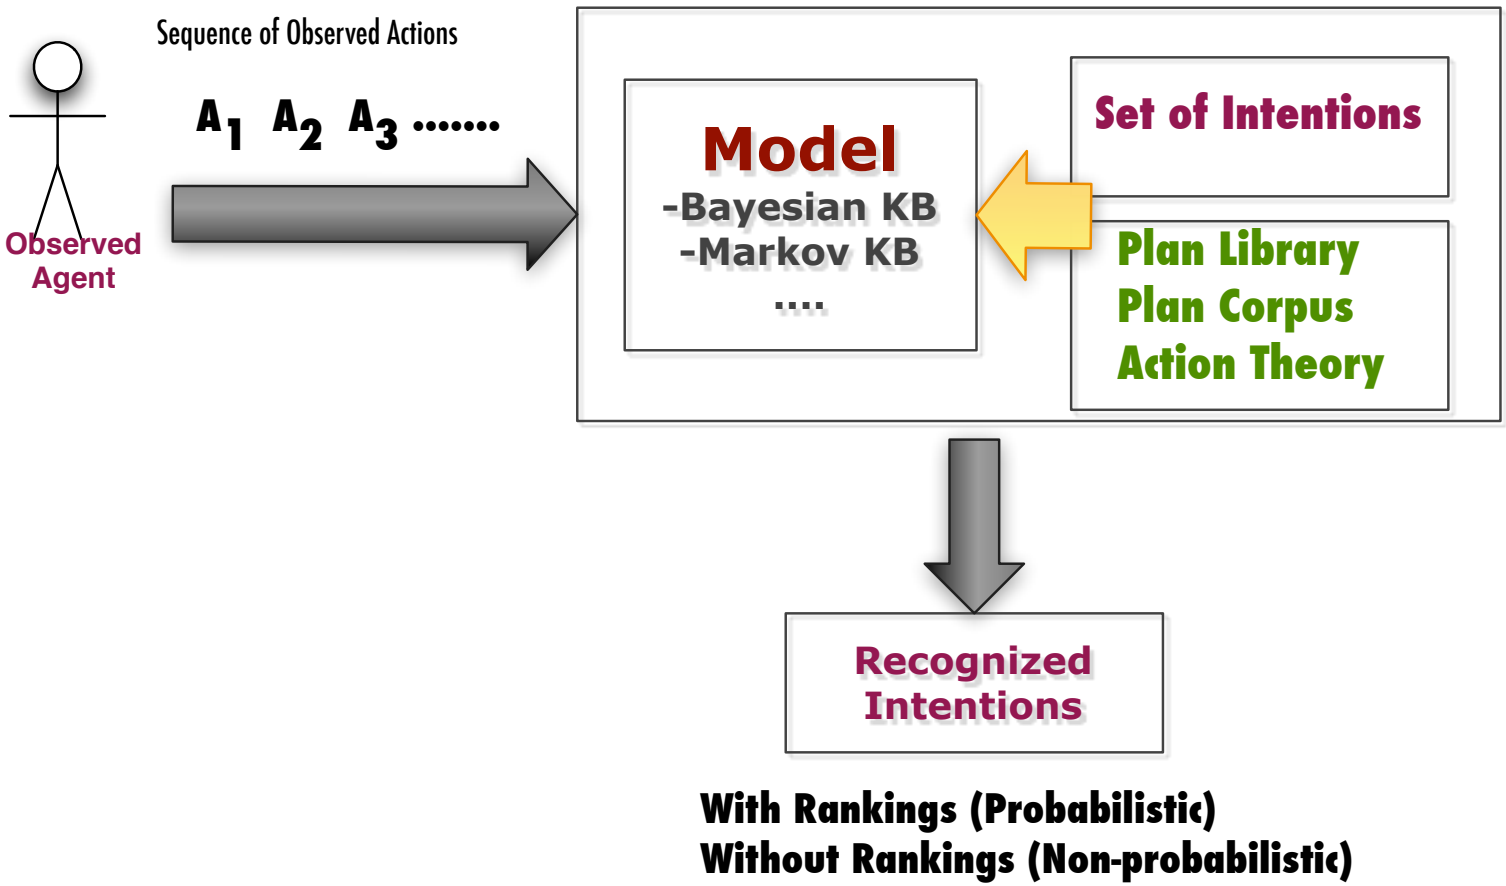

System Description

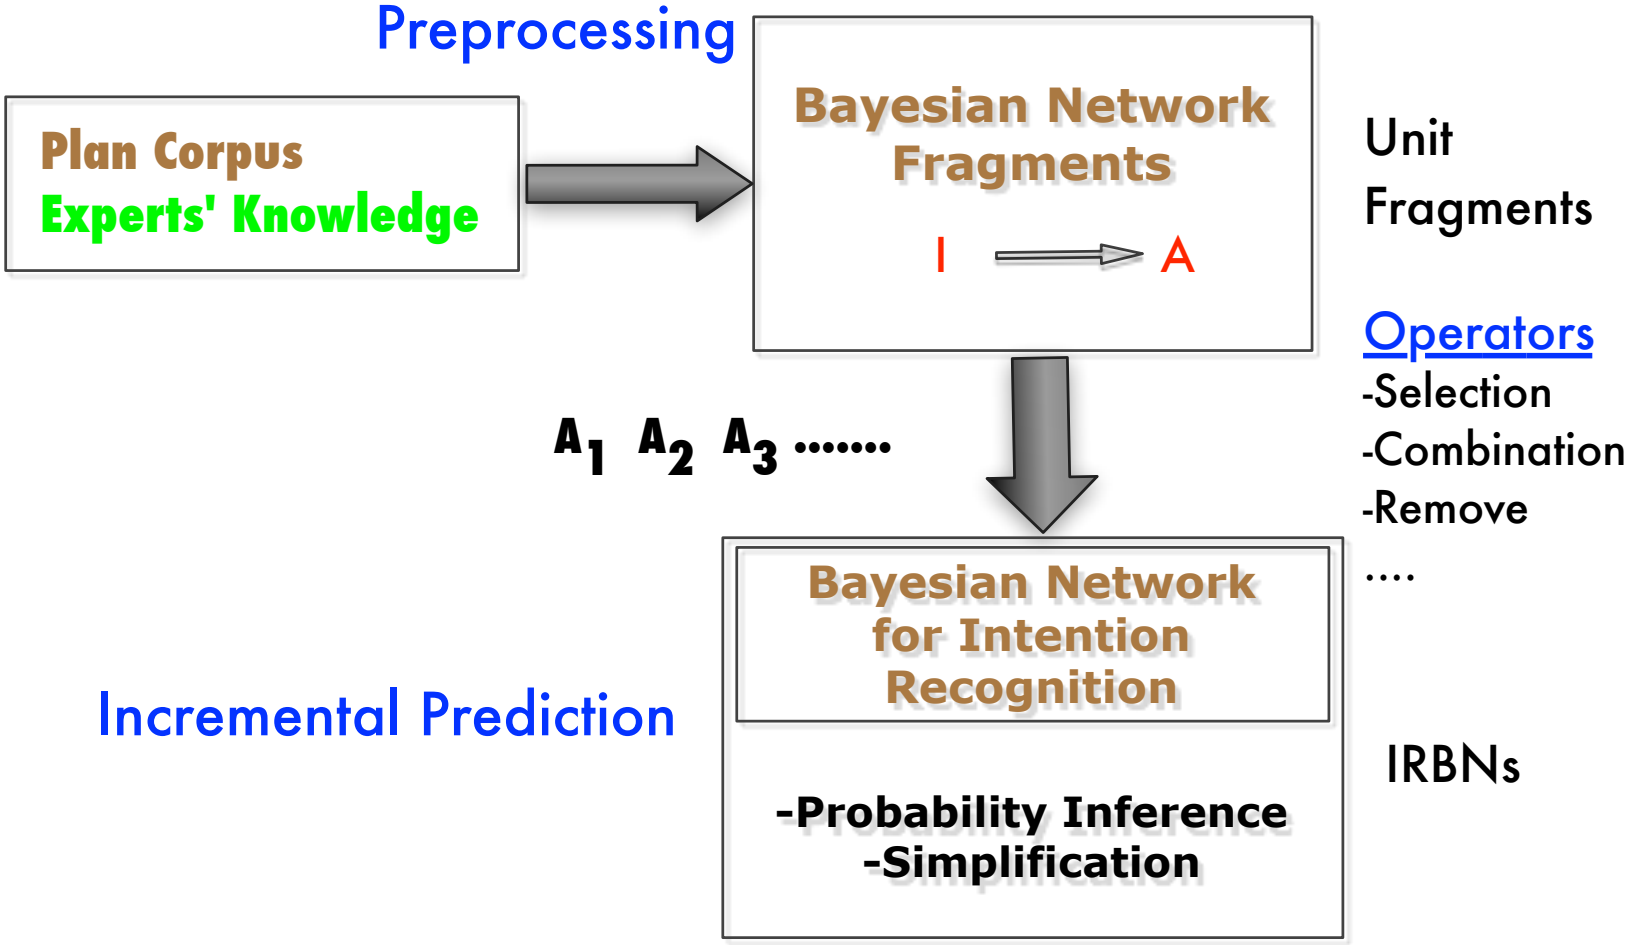

Bayesian Network for Intention Recognition (IRBN)

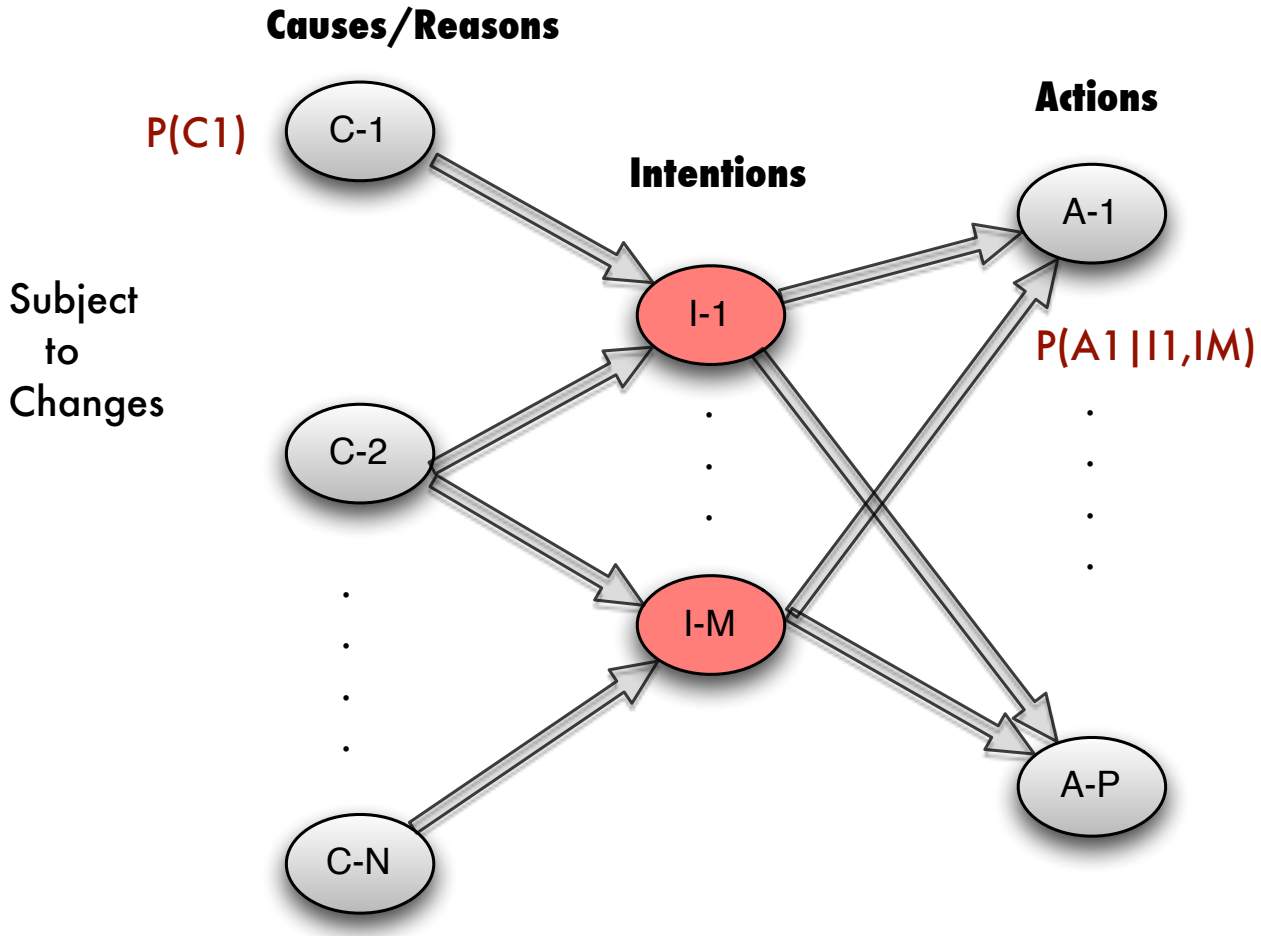

CPD table for each node X
P(X|parents(X))

IR: Compute P(I-i|obs)
i = 1,...,M

Unit Fragments ⟶ Unit IRBN

Figure: Combine context-dependently selected unit fragments of a new action A.

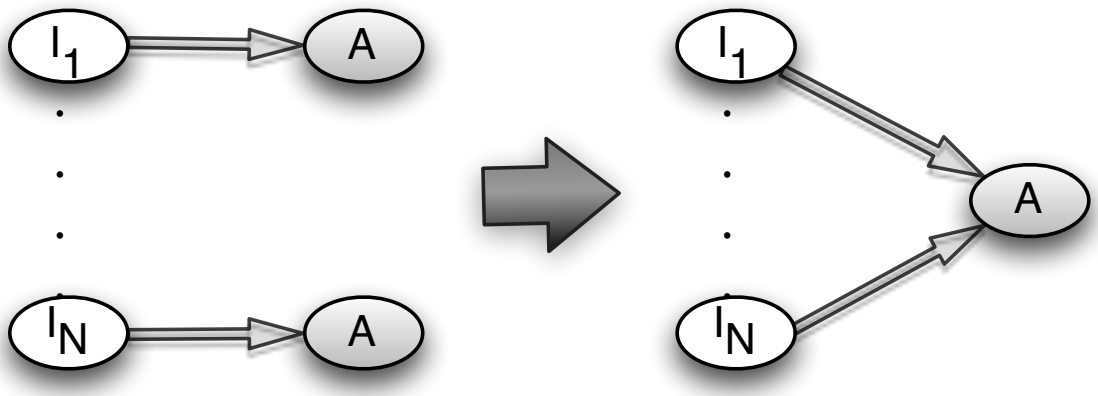

Integrate into the current model ...

When a new action A_3 arises, integrate it into the current model

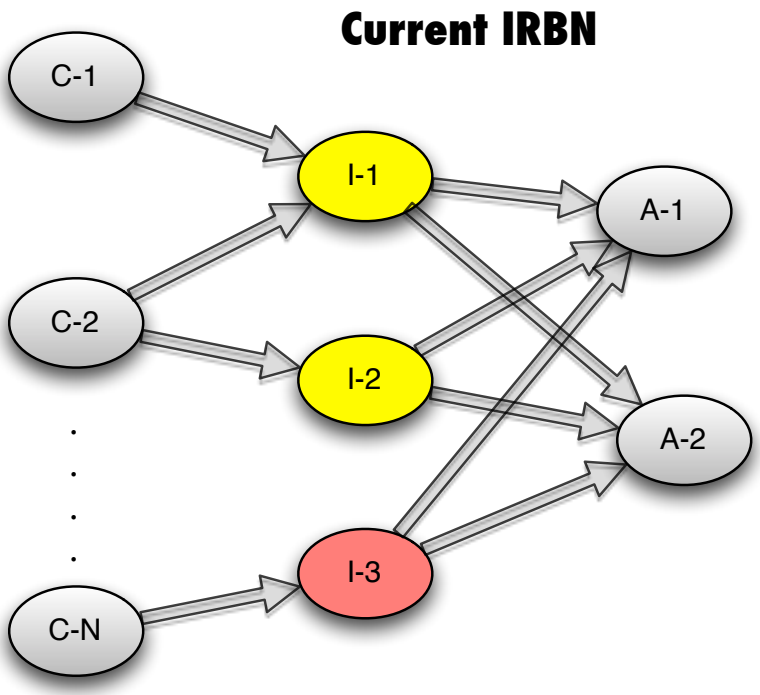

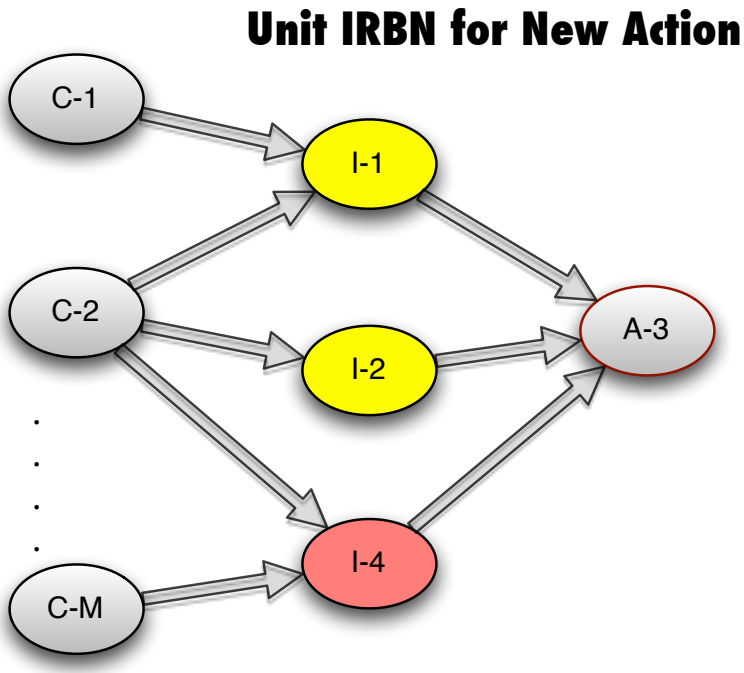

Remove Irrelevant Intentions

If some intentions are found out to be irrelevant, e.g. when their probability is very small, they are removed from the model.

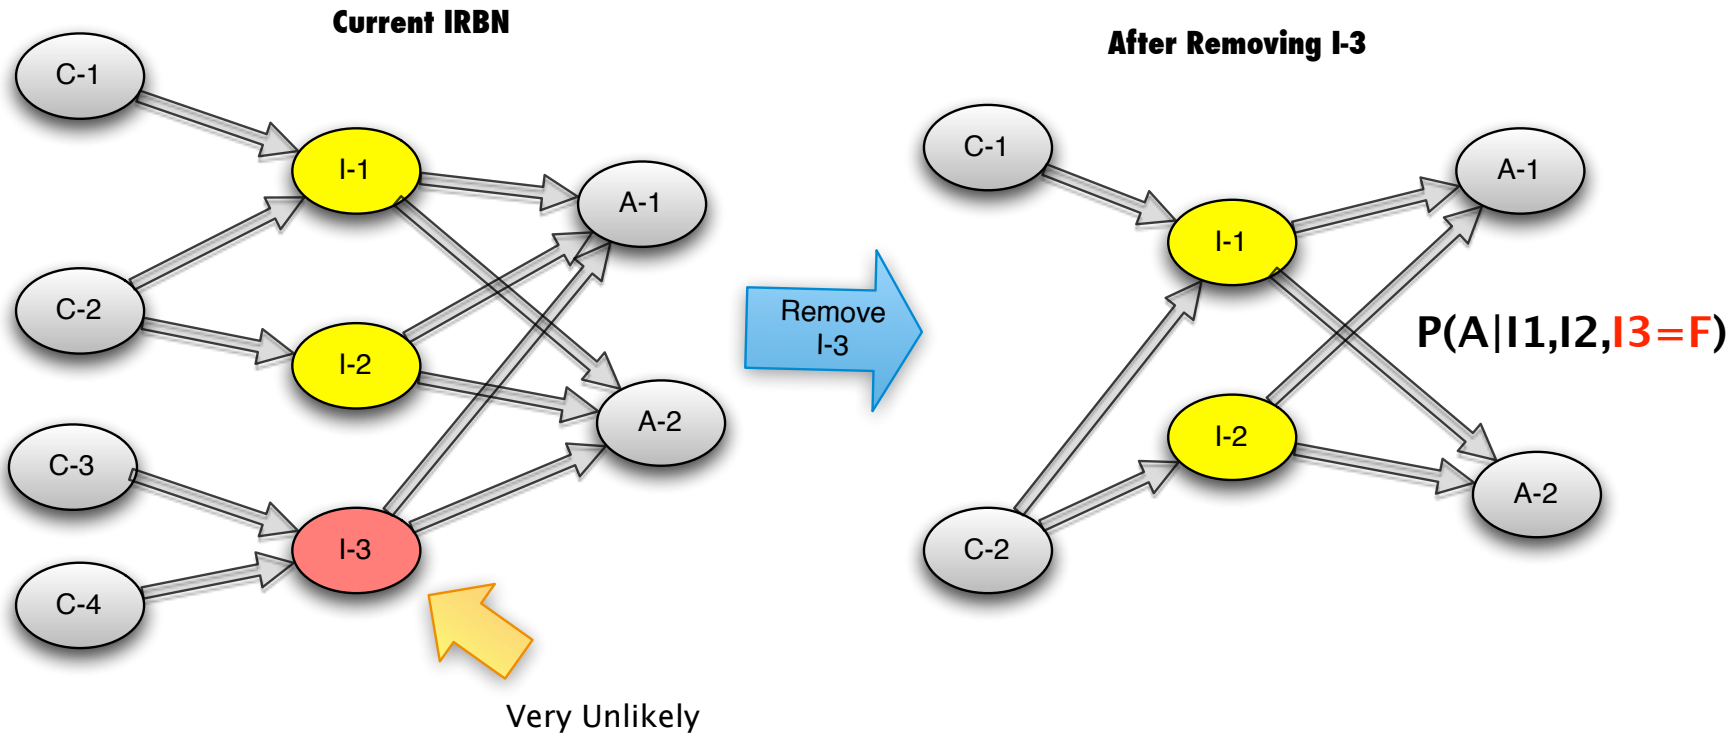

Incremental Intention Recognition Algorithm

Repeat until one intention remains or time limit is reached.

1. If new actions are observed: combine unit IRBNs for them with the current IRBN.
2. Compute conditional probability of each intention given current observed actions. **Remove** irrelevant intentions.

Relations Amongst Intentions: multiple intentions

- In case agent pursues multiple intentions simultaneously, intentions may support or exclude each other.
- Two mutually exclusive intentions cannot be parents of an action: $P(I_1 = T, I_2 = T) = 0 \rightarrow P(A|I_1, I_2, ...)$ is undefined.
- Need to combine mutually exclusive intentions in a single node.

Single Intention Case

- In case agent pursues a single intention (Linux Plan corpus), all intentions are mutually exclusive.
- They are combined in a single node.
- $A_1, ..., A_m$: current observed actions. Then,

$$P(I = I_j | A_1, ..., A_m) = \frac{P(I_j) \prod_{i=1}^{m} P(A_i | I_j)}{\sum_{j=1}^{n} P(I_j) \prod_{i=1}^{m} P(A_i | I_j)}$$

- The recognizer has linear complexity on the number of intentions being modeled.

Evaluation Metrics

- $Seq = a_1, ..., a_n$: sequence of actions achieving intention I.
- N-best prediction: $correct(A) = 1$ if I is one of N most likely intentions, and 0 otherwise.

$$precision(Seq) = (\sum_{i=1}^{n} correct(a_i))/z$$

$$recall(Seq) = (\sum_{i=1}^{n} correct(a_i))/Z$$

$$convergence(Seq) = (z - t + 1)/z$$

z: #predictions; Z: #prediction opportunities; t: smallest number from that on correctly predicts.

Linux Plan Corpus

- One of rare, often used, plan corpora available for evaluating intention/plan recognizers.
- goals = tasks in Linux (e.g., find a file, copy some files); actions = Linux commands (e.g., find, cp, cd, ls).
- 19 goals; 43 actions (commands); 547 sessions. Collected from Linux users.

Intention Recognition Results on Linux Plan Corpus

τ – confidence level: make prediction only when probability $\geq \tau$.

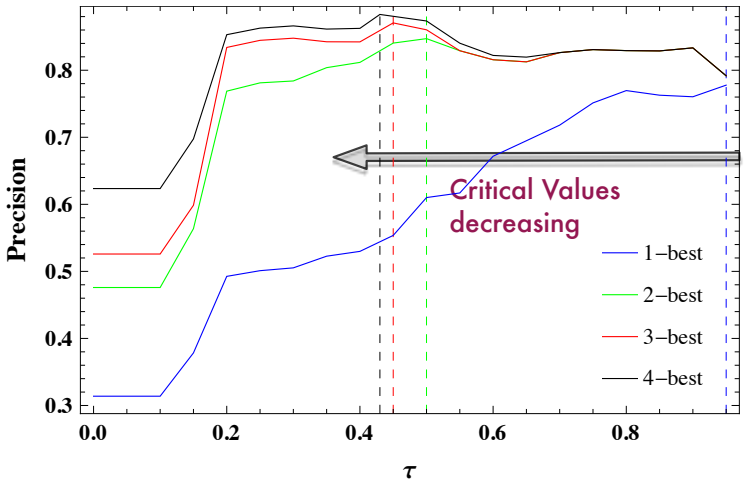

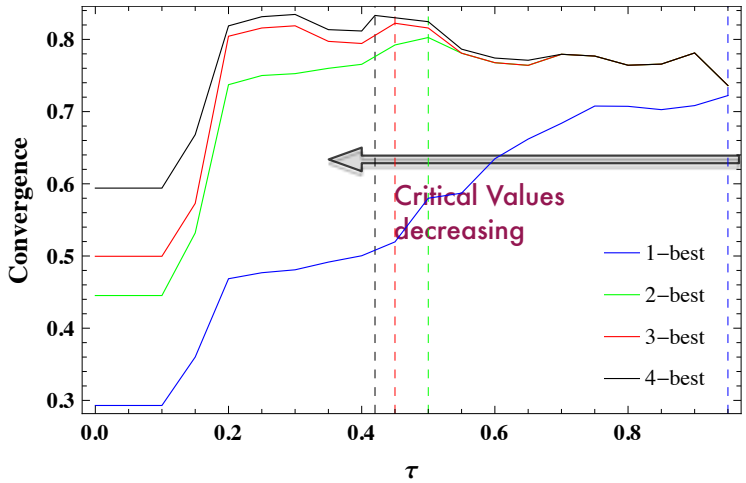

Figure: Precision and Convergence for $\tau \in [0, 1]$; N = 1, 2, 3, 4.

We obtained better results than existent corpus-based intention recognizers.

N-best	1-best	2-best	3-best	4-best
τ	0.95	0.5	0.45	0.42
Precision	0.786	0.847	0.870	0.883
Recall	0.308	0.469	0.518	0.612
Converg.	0.722	0.799	0.822	0.824

Table: Intention Recognition Results on the Linux Plan Corpus

Iterated Prisoner's Dilemma Corpus (IPD)

Prisoner's Dilemma is a symmetric two-player non-zero game defined by the payoff matrix

$$\begin{array}{c} \\ C \\ D \end{array} \begin{array}{c} \begin{array}{cc} C & D \end{array} \\ \begin{pmatrix} R,R & S,T \\ T,S & P,P \end{pmatrix} \end{array}$$

Intentions to be recognized are strategies in IPD: TFT, WSLS, GTFT, GRIM, FBF, etc.

Intention Recognition Results on IPD Plan Corpus

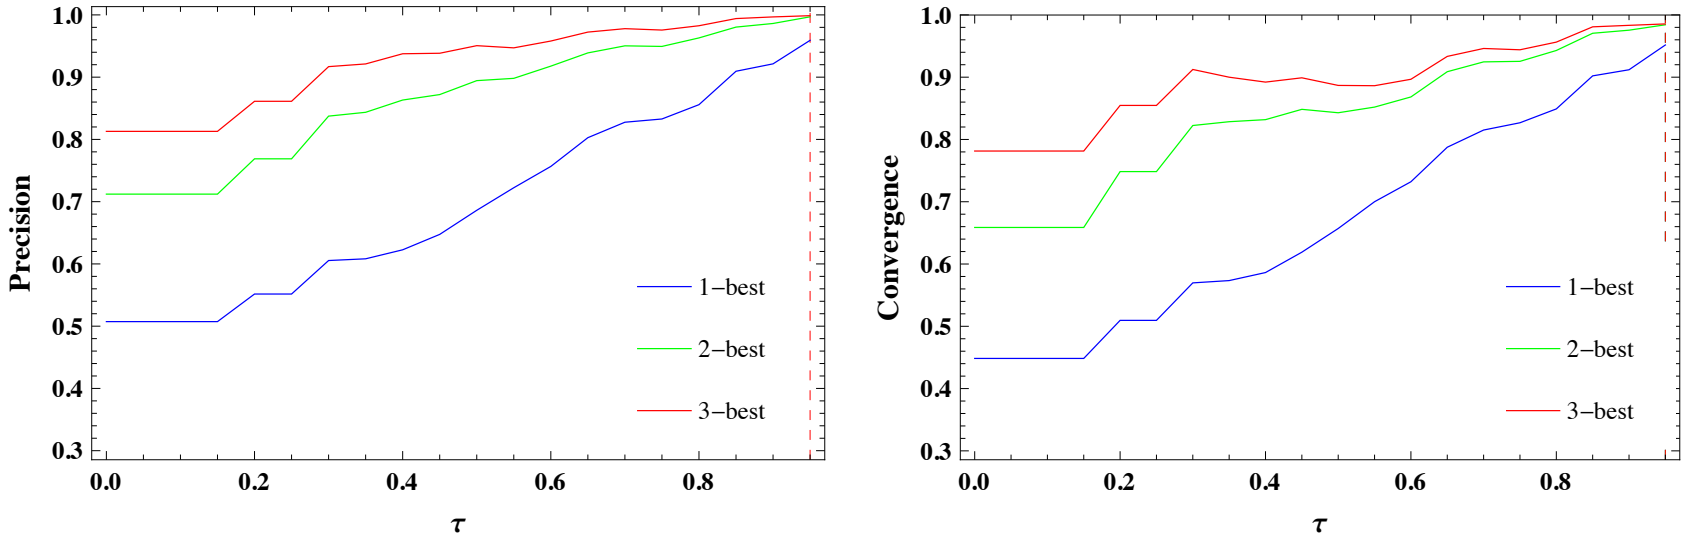

Figure: Precision and Convergence for $\tau \in [0, 1]$; N = 1, 2, 3

Summary

- Incremental IR model via incrementally constructing a Bayesian network model.
- Evaluated on Linux Plan corpus, obtaining better results than existent systems.
- Present a new benchmark for intention/plan recognition based on social dilemmas (Prisoner's Dilemma, Stag hunt,...).
- Able to handle multiple intention case by representing relations amongst intention nodes.

Future Work

- Create plan corpora to evaluate multiple-intention recognition (suggestions of Planners for this purpose are greatly appreciated)
- Apply to explain the role of intention recognition for the evolution of cooperation (see our papers in refs).
- We have provided in another IPD benchmark corpus for testing context-dependent aspect (BMAW@UAI paper).

Thank you!

QUESTIONS?

IPD Plan Corpus

- Intentions are famous strategies in IPD: TFT, WSLS, GTFT, GRIM, FBF. There are infinite number of strategies.
- Corpus actions: $s_1 ... s_M \xi$, where $s_i \in \{E, R, T, S, P\}$ – states of the M last interactions; $\xi \in \{C, D\}$ – current move. $\Sigma_1 = \{EC, RC, TC, SC, PC, ED, RD, TD, SD, PD\}$
- A plan session of TFT: $[EC, RC, SD, PD, TD]$

round :	0	1	2	3	4	5
TFT :	–	C	C	D	D	D
X :	–	C	D	D	C	D
TFT-states :	E	R	S	P	T	P

Generating Training and Testing Datasets

- 7 strategies: *AllC*, *AllD*, *TFT*, *GTFT*, *WSLS*, GRIM and FBF.
- Training corpus is generated by playing with each strategy all the possible combinations 10 times ($r = 5, \ldots, 10$).
- The testing dataset is generated by playing a random choice with each strategy in each round ($r = 5, \ldots, 10$).