Dealing with Ambiguity in Plan Recognition under Time Constraints

Moser S. Fagundes, **Felipe Meneguzzi**, Rafael H. Bordini, Renata Vieira
Pontifical Catholic University of Rio Grande do Sul

Plan Recognition

- Broader Context: Plan, Activity and Intent Recognition
  - Activity Recognition - deals with current (often low-level) actions
  - Plan Recognition - deals with high-level complex goals
  - Intent Recognition - deals with the relation between current plans and the plan library
- In this paper, we talk (mostly) about the latter two areas

Plan Recognition - Terminology

- Observation - input from the environment
- Plan Library (PL) - domain knowledge about the subject being observed, often represented as a directed (possibly cyclic) graph
- Plan Step - one node in the plan library graph
- Plan Hypothesis - a sequence of plan steps consistent with both the Plan Library and the Observations

Motivation for our Work

- Recognition often tied to **doing something** about recognized plans (or plan hypotheses)
  - Assistance (when observed subject is benign)
  - Countermeasures (when observed subject is adversarial)
- Responses usually not instantaneous
  - Observer agent needs to reason about plan hypotheses and time

Background: Symbolic Plan Recognition

- Symbolic Behavior Recognizer (SBR) Avrahami-Zilberbrand and Kaminka
    - Hybrid plan recognition approach
    - Uses a decision tree (FDT) to map observations into plan-steps in the PL
    - Allows quick response for plan-library membership queries
    - Used for anomalous behavior identification

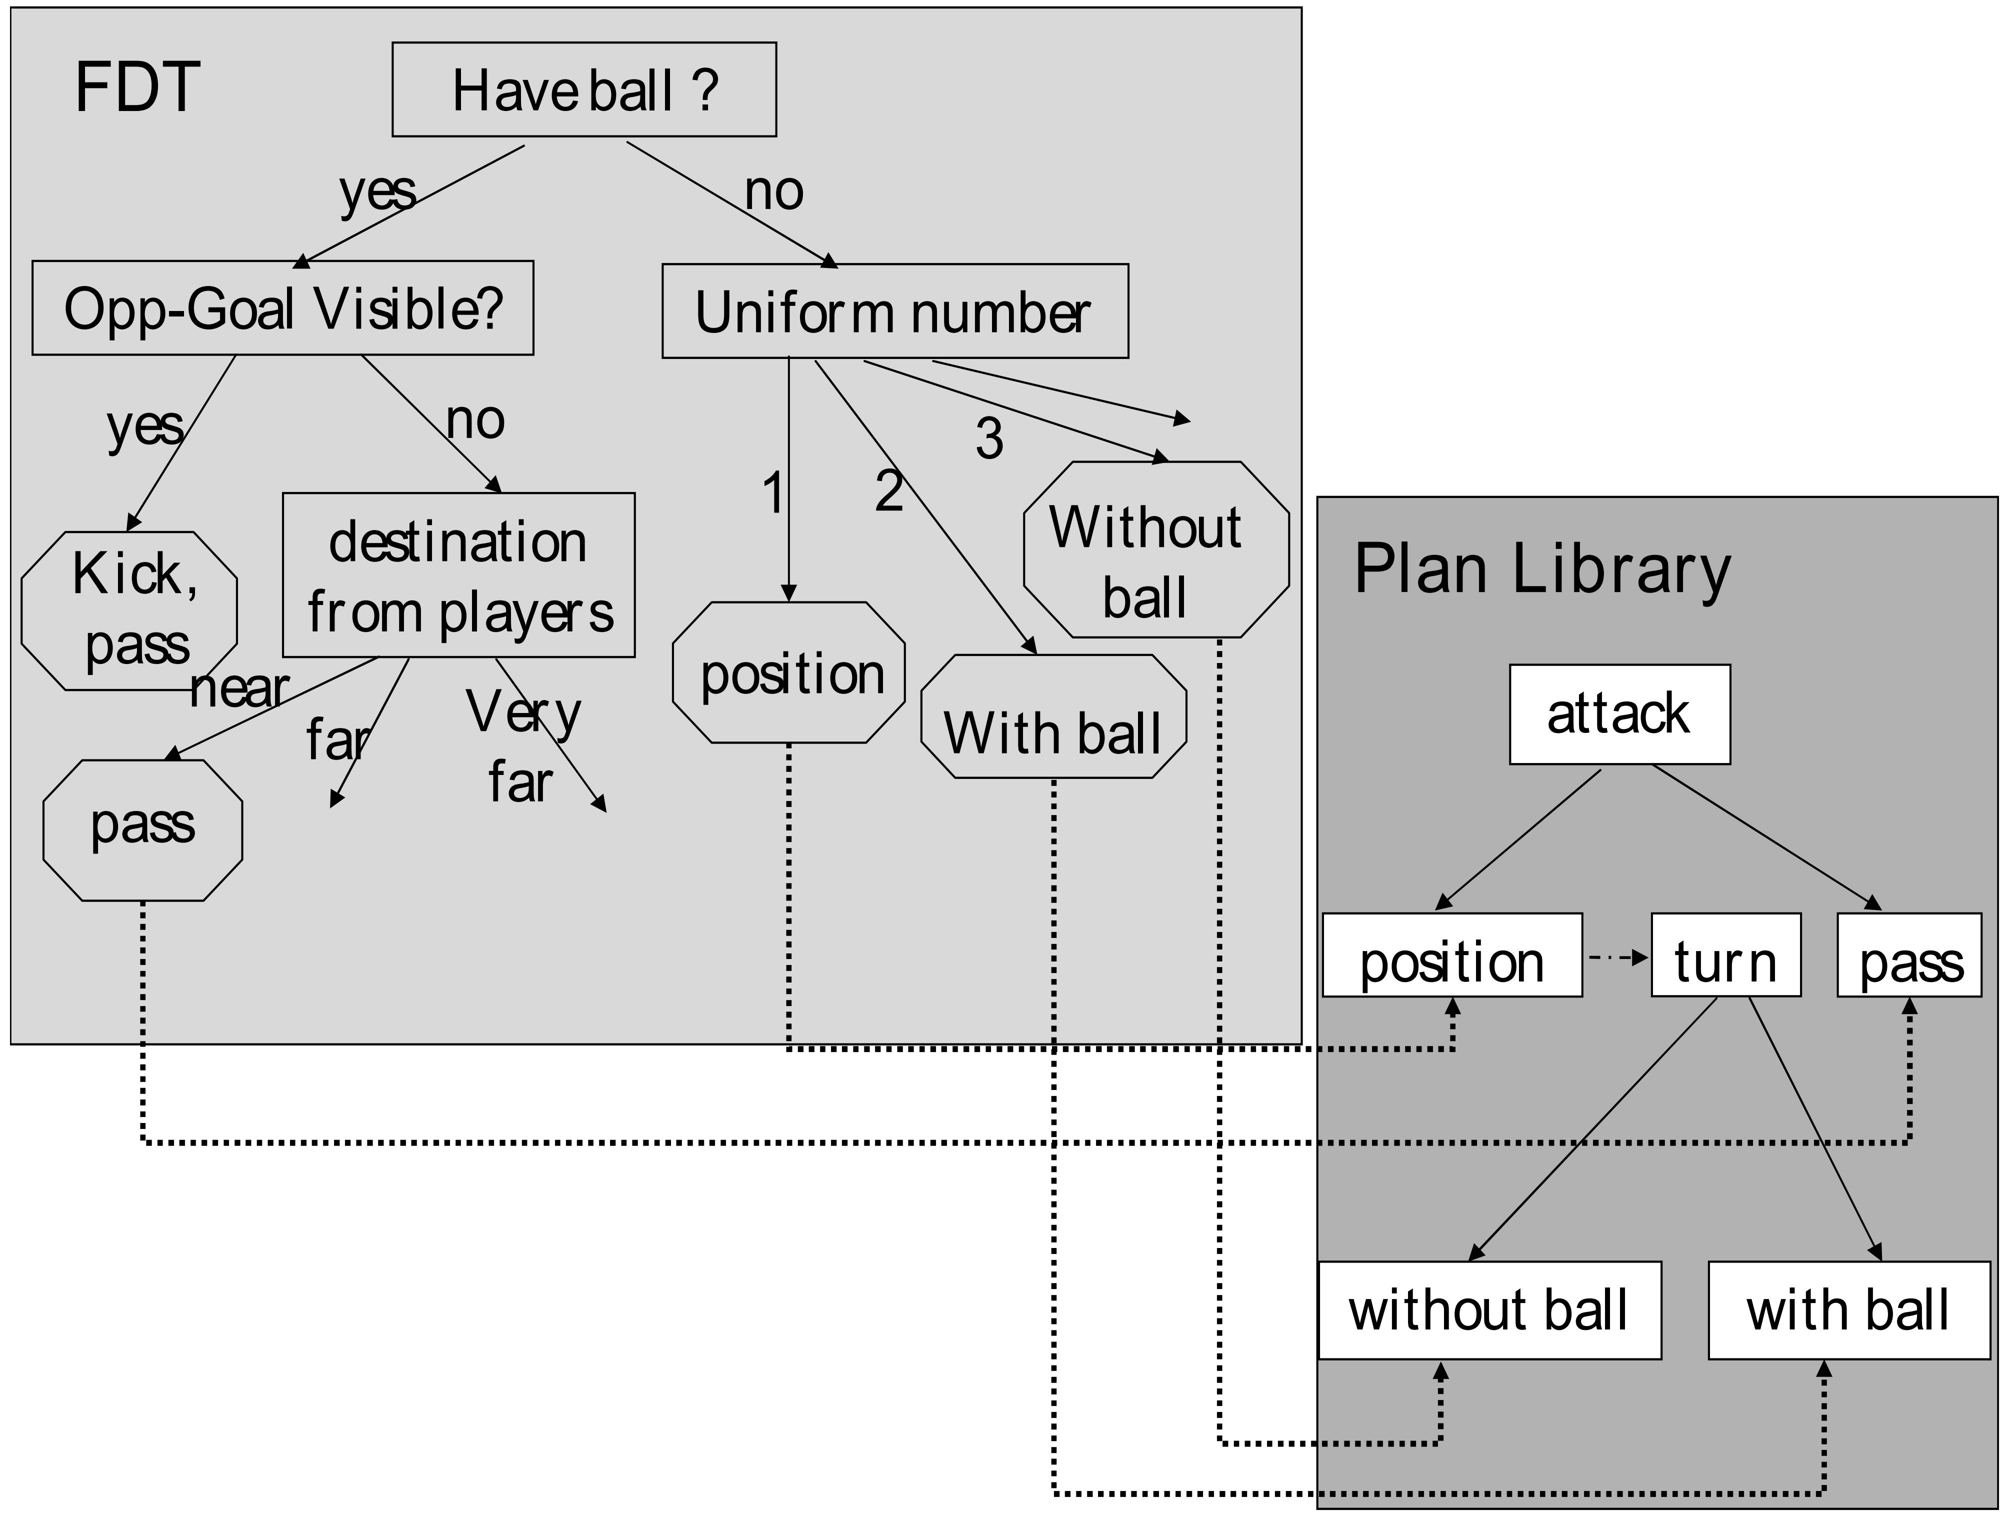

Recognizer Architecture

- We leverage SBR into an overall recognizer architecture, including
  - Actual plan recognition
  - Interaction for disambiguation
  - Response to recognition
  - Estimation of recognition time
  - Assessing plan likelihood

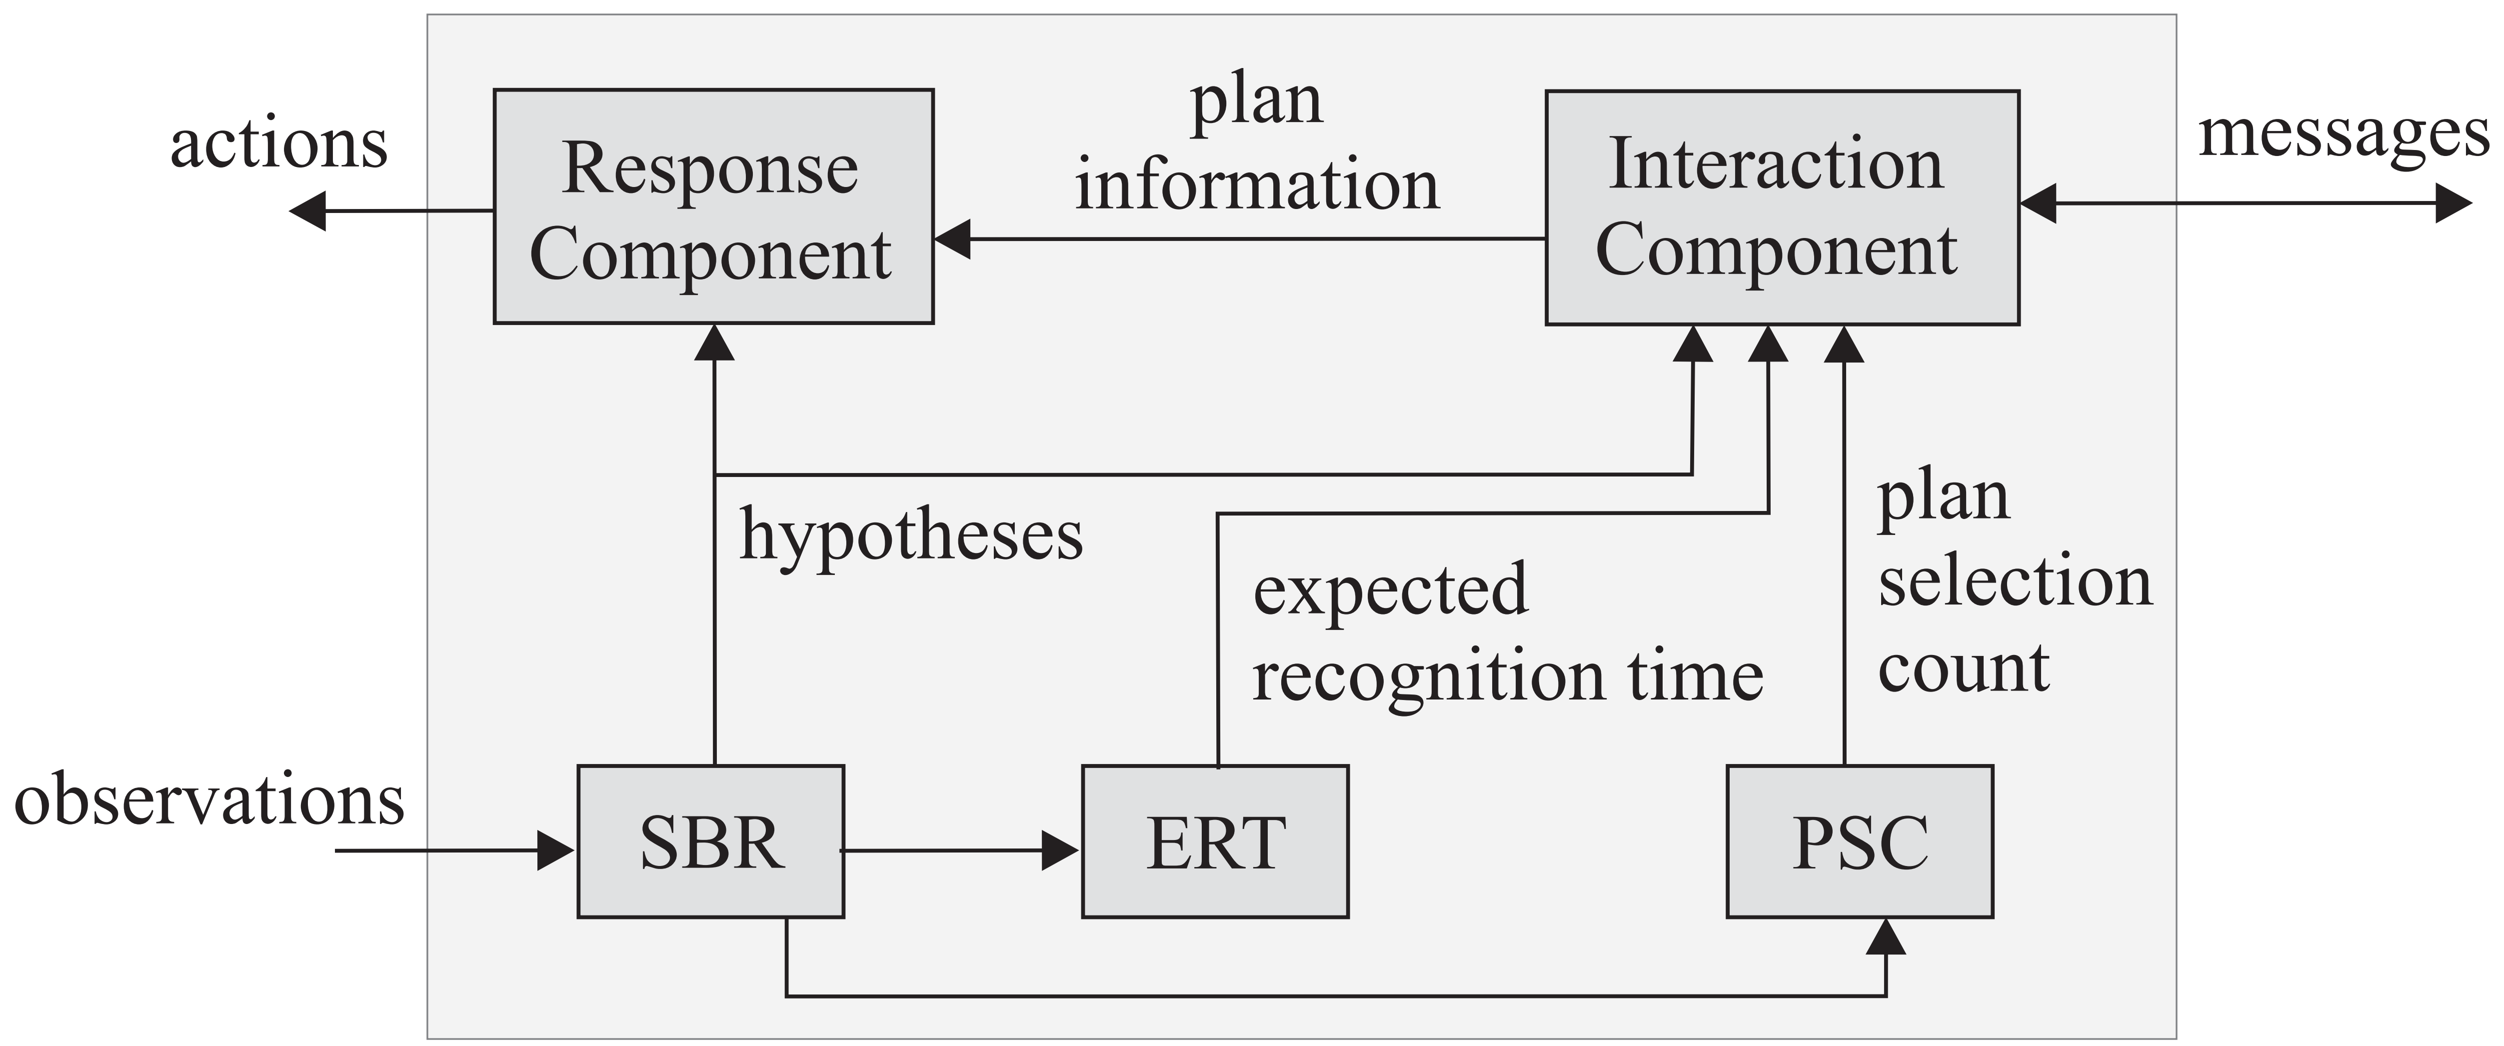

Assessing Time to Recognize

- Assumption: observations are made at regular time intervals
- Basic approach, at every time step:
  - Collect observations and average times (CE Table)
  - Match observations to plan library nodes (via FDT)
  - Tag plan steps with time stamp and actual observation
- When only one hypothesis remains, update ERT Table using a *reinforcement* update $e[\text{``}ert\text{''}] \leftarrow (1 - \alpha(e[\text{``}nupd\text{''}]))e[\text{``}ert\text{''}] + \alpha(e[\text{``}nupd\text{''}])avg$

Assessing Time to Recognize

- ERT Table associates, for each "initial observation", an average recognition time

- Example:

    - In a single episode observations "`location(2,3)`" mapped to "`position`" action in the PL averaged **15 time steps** before recognition

    - Over many episodes, this average resulted in an expected recognition time of **13.15 time steps**

position

CE table
(compact view)

observations	avg
(*location*(2,3), ...)	15.0
(*location*(2,4), ...)	12.5
(*location*(3,4), ...)	10.5

(a)

ERT table

(b)

observations	ert	nupd
(*location*(1,3), ...)	21.04	20
(*location*(2,3), ...)	12.92	8
(*location*(3,2), ...)	14.65	11
(*location*(3,3), ...)	10.77	7
(*location*(3,4), ...)	7.62	13

ERT-UPDATE

(c)

observations	ert	nupd
(*location*(1,3), ...)	21.04	20
(*location*(2,3), ...)	**13.15**	**9**
(*location*(2,4), ...)	**12.50**	**1**
(*location*(3,2), ...)	14.65	11
(*location*(3,3), ...)	10.77	7
(*location*(3,4), ...)	**7.82**	**14**

Assessing Probability of Plan Selection

- In each recognition episode we keep track of:
  - the number of times a node in the plan library was updated with ERT; and, from this count
  - the number of times a node in the plan library was **actually part** of a successfully recognized plan
- This allows us to estimate how likely a hypothesis leads to a successful recognition using

$$maxChance(t) = \max_{e \leftarrow CE(t,l)} \frac{e[nps]}{\sum_{e_i \in CE} e_i[nps]}$$

Interaction Component

- The Interaction Component uses the probability and the estimated recognition time to:
  - compute the “value” of current plan recognition hypotheses
  - decide whether to disturb the observed subject or not;
- Decision uses a combination of parameters and estimations made by our algorithm

Bringing it all Together

- Given the expected recognition time at a step *ert(t)*,
  a recognition deadline *ρ(t)*,
  a maximum chance for a successful hypothesis *maxChance(t)*
  and a decision threshold *ϕ*,

  - The observer agent can decide whether to interrupt the user based on two criteria:

    - $ert(t) \leq \rho(t)$ - whether the expected time is lower than the deadline; and

    - $maxChance(t) \geq \phi$ - whether the maximum chance is greater than a threshold

Conclusions

- Our main contributions are:
  - A plan recognition algorithm and surrounding architecture that
    - Estimates time until a plan can be recognized in various contexts
    - Provides a probability estimation for plan recognition
  - Providing decision criteria on whether to interrupt a user to disambiguate multiple plan hypotheses

Future Work

- Take into account interleaved plan execution and lossy observations
- Evaluate the architecture with human-generated data

Questions?