BDI Agents in Natural Language Environments

Alexandre Yukio Ichida[1]
Rafael C. Cardoso[2]
Felipe Meneguzzi[1 2]

Pontifical Catholic University of Rio Grande do Sul[1]
University of Aberdeen[2]

Introduction - Contributions

- **Rational Agents to deal with natural language using BDI**
- **Reasoning over natural language information**
  - Natural language beliefs
  - Plan selection through Natural Language Inference
  - Fallback policy

Background - Belief Desire Intention model

BDI - Belief Desire Intention

- Beliefs – Agent's internalised perceptions, or prior knowledge about the world.
- Desire – Objectives to be achieved.
- Intention – Commitments to achieve goals.

Background - Natural Language Inference

Infer logical relation (entailment) between two sentences

- Premise: You are in the kitchen
  - Hypothesis 1: You are not in the bedroom (true - entailment)
  - Hypothesis 2: You are in the living room (false - contradiction)
  - Hypothesis 3: You see a door to the hallway (false - neutral)

Background - Scienceworld

- **Textual Environment: ScienceWorld**[1]
  - Interactive textual environment simulating engines for thermodynamics, electrical circuits, matter and chemistry reactions.
  - States and actions are encoded in natural language texts.

[1]Wang, Ruoyao, Peter Jansen, Marc-Alexandre Côté, and Prithviraj Ammanabrolu. "Scienceworld: Is your agent smarter than a 5th grader?." (2022).

Our Approach

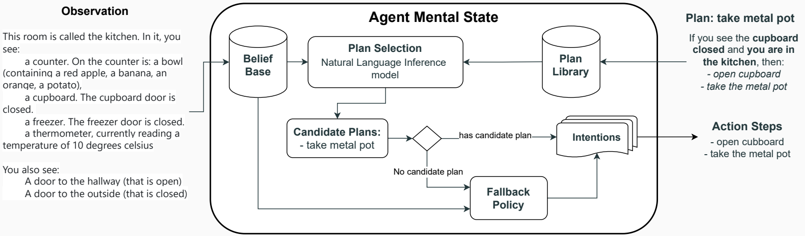

Figure 1: Diagram illustrating the NatBDI architecture to handle and actuate over natural language environments.

Natural Language Beliefs

- **Sentences describing the agent perception about the environment**
  - Objects seen
  - Objects taken (inventory)
  - Location information
- **Beliefs encoded in a human understandable language**
  - Each belief are represented in a descriptive sentence

Natural Language Plans

- **Allow humans to include instructions through plan rules**
- **Plans written in a controlled natural language representation**
- **Plan rules inspired by AgentSpeak language**
  - Triggering event (goal addition)
  - Plan contexts as natural language sentences
  - Sequence of actions or subgoals in the plan body
  - Uses reserved keywords to simplify the language parser

Natural Language Plans

Listing 1: Plans in natural language to pick the metal pot and melt water in ScienceWorld.

```
IF your task is to get the metal pot
CONSIDERING you are in the kitchen
        AND you see the cupboard closed
THEN:
open the cupboard,
take the metal pot

IF your task is to melt water
THEN:
PLAN TO get the metal pot
pick up thermometer
...
```

Natural Language Inference - Inference Steps

1. Create a matrix M to evaluate all belief, context sentence pairs

$$M_{i,j} = \mathcal{C} \times \mathcal{B} = \{(c_i, b_j) \mid c \in \mathcal{C} \wedge b \in \mathcal{B}\} \tag{1}$$

2. Create an entailment matrix E with all NLI result obtained from M

$$E_{i,j} = \{nli(c_i, b_j) \mid (c_i, b_j) \in M_{i,j}\} \tag{2}$$

3. Check whether the plan context are entailed by the belief base ($\mathcal{B} \models \mathcal{C}$) using matrix E as follows:

$$\mathcal{B} \models \mathcal{C} \doteq \bigwedge_{c_i \in \mathcal{C}} c_i \bigvee_{b_j \in \mathcal{B}} b_j \tag{3}$$

Natural Language Inference - Inference Steps - Matrix M

1. Create a matrix M to evaluate all belief, context sentence pairs

$$M_{i,j} = \mathcal{C} \times \mathcal{B} = \{(c_i, b_j) \mid c \in \mathcal{C} \wedge b \in \mathcal{B}\} \tag{4}$$

Plan Context

Belief Base	you see the cupboard closed	you are in the kitchen
This room is called the kitchen		
In it, you see a cupboard. The cupboard is closed.		
You see a door to the hallway (that is open)		

Natural Language Inference - Inference Steps - Matrix E

2. Create an entailment matrix E with all NLI result obtained from M

$$E_{i,j} = \{nli(c_i, b_j) \mid (c_i, b_j) \in M_{i,j}\} \tag{5}$$

Belief Base \ Plan Context	you see the cupboard closed	you are in the kitchen
This room is called the kitchen	False	True
In it, you see a cupboard. The cupboard is closed.	True	False
You see a door to the hallway (that is open)	False	False

Natural Language Inference

3. Check whether the plan context are entailed by the belief base ($\mathcal{B} \models \mathcal{C}$) using matrix E as follows:

$$\mathcal{B} \models \mathcal{C} \doteq \bigwedge_{c_i \in \mathcal{C}} c_i \bigvee_{b_j \in \mathcal{B}} b_j \tag{6}$$

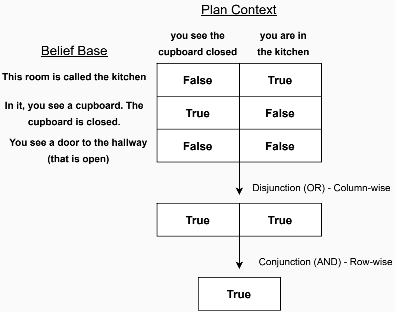

Fallback Policy

- **If there is no plan candidate, the agent triggers the fallback policy to select an action.**
- **Fallback Policy:**
  - The fallback policy π is a function $f : \mathbb{R}^{|\mathcal{B}|} \rightarrow \mathbb{R}^{a}$ that maps the belief base and results into a plan.
  - Learnable (Reinforcement Learning)
- **Reinforcement Learning**:
  - Plan selection as POMDP
  - Off-the-shelf RL approaches

Results

Table 1: Comparison of our natural language BDI agent in two tasks in the ScienceWorld environment with different plan library sizes. We show the average scores obtained and the average number of actions performed in each phase out of all task variations. The bold font identifies which approach (BDI or DRRN) contributed more to the total score.

Task	Variations	Episodes	Number plan-rules	Score (Total)	Score (BDI)	Score (DRRN)	Number BDI actions	Number RL actions
find-non-living-thing	75	242	0	0.66	0.00	**0.66**	0.00	50.00
			8	0.75	0.30	**0.45**	3.33	38.00
			15	0.84	**0.58**	0.26	6.25	24.00
			23	0.91	**0.79**	0.12	7.64	13.33
			30	0.98	**0.98**	0.00	9.19	4.00
melt	9	457	0	0.03	0.00	**0.03**	0.00	50.00
			4	0.14	**0.11**	0.03	5.11	44.44
			7	0.36	**0.34**	0.02	10.89	33.33
			10	0.57	**0.56**	0.01	17.11	22.22
			13	0.67	**0.67**	0.00	20.89	16.67

Results - Plan Library Size

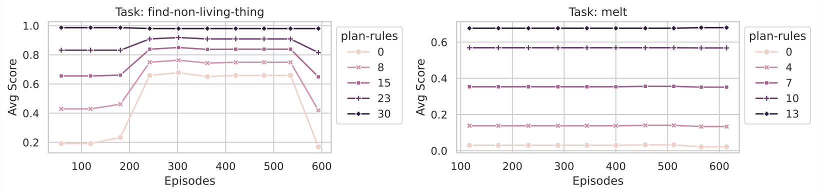

Figure 2: Scores per episodes when scaling the number of plan-rules.

Results - Natural Language Inference

Table 2: Results of using different LLMs for NLI. The following columns describe them: model size (*Params*); accuracy on MultiNLI matched test set (*MNLI*-m); score obtained using NatBDI; average number of *actions* performed, *errors* raised and plan-rules (*Plans*) used; lexical overlap computed on entailment pairs (*LO(E)*); average word number in belief (—*B*—) and context (—*C*—) sentences; and total sentence pairs processed. We highlight the best scores in bold font.

Model	Params	MNLI-m	Task	Score	Actions	Errors	Plans	LO(E)	—B—	—C—	Pairs
MiniLM (L6)	22M	82.2	find-non-living-thing	0.69	7.65	0.37	2.57	0.64	7.84	4.20	1076
			melt	0.23	8.78	1.00	3.59	1.34	10.96	4.83	691
Bert (base)	110M	84.6	find-non-living-thing	0.84	9.61	0.20	2.72	0.60	8.40	4.16	1075
			melt	0.33	11.44	0.67	3.56	1.16	11.12	4.80	690
Roberta (large)	355M	90.8	find-non-living-thing	**0.98**	9.19	0.08	2.84	0.40	7.30	4.19	1076
			melt	**0.67**	20.89	0.33	5.67	1.21	11.06	5.25	790

Conclusion and Future Work

- **We develop an BDI-based agent architecture to reason over natural language environment**
  - We leverage novel NLI language models to infer entailment over natural language sentences
  - We include a plan library consisting in natural language plan-rules
  - We extend the mechanism for plan selection including a fallback policy in our architecture
- **Future Work**
  - Cover more ScienceWorld task
  - Improve the fallback policy component
  - Generating plan-rules from data through learning techniques