

# Towards an Integration of Learning and Reasoning in Agent and Multi-agent Systems

Candidate:

Andrea Rafanelli

*National Ph.D. in Artificial Intelligence  
XXXVII cycle*

Supervisors:

Stefania Costantini  
Andrea Omicini



NATIONAL PH.D. PROGRAM  
IN ARTIFICIAL INTELLIGENCE

**Ph.D. Dissertation, 19<sup>th</sup> February 2025, Pisa**



UNIVERSITÀ  
DEGLI STUDI  
DELL'AQUILA



UNIVERSITÀ DI PISA

# Introduction



## Narrow AI

Single task,  
Single domain,  
High accuracy,  
Speed for certain  
tasks



D'ou venons-nous

## Broad AI

Multi-domain,  
Multi-modal,  
Distributed AI,  
Explainable



Que sommes-nous

## General AI

*"is a **hypothetical**  
stage in which an AI  
can **match** or  
**exceed** the  
cognitive abilities of  
human beings."*



Ou allons-nous

# The path to General AI

- **New disciplines**

- Explainable AI (XAI)
- Trustworthy AI (TAI)

- **New legislations**

- Right for Explanation
- AI Act
- AI in the EU Green Deal

*Researchers are exploring new techniques and pathways to create more reliable and efficient AI systems.*



The New York Times

## *A Wellness Chatbot Is Offline After Its 'Harmful' Focus on Weight Loss*

Stanford University

Algorithms do not treat everyone equally. Black Americans are more likely to be denied loans, jobs, and more when using some AI systems.

News

Opinion

Sport

Culture

Lifestyle



UK UK politics Education Media Society Law Scotland Wales Northern Ireland

Law (Australia)

Fake cases, judges' headaches and new limits: Australian courts grapple with lawyers using AI

CTO Magazine

Trending

About

Contact



AI & MACHINE LEARNING, ETHICS



The Dark Side of AI: Its Growing Environmental Footprint




# Neuro-Symbolic (NeSy)



Learning and Reasoning integration



- MIT-IBM AI Watson Lab<sup>1</sup>
- Neuro Symbolic Learning and Reasoning (ANSR)<sup>2</sup> project.





<sup>1</sup> <https://mitibmwatsonailab.mit.edu/>

<sup>2</sup> <https://www.darpa.mil/program/assured-neuro-symbolic-learning-and-reasoning>



# Scope and Research Flow



# Taxonomy


## Foundation



Which kind of interactions and translation mechanisms can be adopted ?

### *Translation Mechanisms*

Methods used to integrate sub-symbolic and symbolic components.




# Taxonomy

## Foundation



Which kind of interactions and translation mechanisms can be adopted ?

### *Interaction*



#### Consecutive Interaction

Information flows in one direction, with components operating in sequence.

#### Concurrent Interaction

Components operate simultaneously, enhancing parallel processing.

#### Cyclical Interaction

Iterative cycles allow for co-refinement of systems.



# Collaborative Networks

## Applications




## How to integrate heterogeneous agents in multi-agent systems?

- *Leverage learning and reasoning in a MAS.*
- When agents have the **same** knowledge, skills, or way of thinking
  - Groupthink
  - Limited problem-solving
  - Lack of innovation
- The benefits of **diversity**:
  - In psychology, researchers refer to **cognitive diversity**.




# Enhanced Reasoning (Learning for Reasoning)

- Incorporates the output of sub-symbolic systems into *logic programs* or *knowledge graphs*.
- **Translation:** *encoding* sub-symbolic representation into a symbolic one.
- **Interaction:** one-way flow of information, *consecutive interaction*.



# Hybrid MAS for Flooding Events



- **Neural Agent** encodes images to **logic descriptions**.
  - **Logic Image Descriptor (LID)**
- **Perceptor Agent** determines flooding.
- **Weather Agent** provides weather information making use of an API.
- **Alert Agent**, examines all the information and, if needed, provide alert levels.

Neural-logic multi-agent system for flood event detection, A. Rafanelli, S. Costantini, G. De Gasperis.



# Experiment Results

Neural Agent



LID



Perceptor Agent

```
1 :- dynamic region/4.
2 region(0, 1, 15.81868131868132, 18.905494505494506).
3 region(1, 4, 125.15170278637771, 19.60061919504644).
4 region(2, 4, 249.5071942446043, 11.720863309352518).
5 region(3, 5, 349.2767857142857, 50.330357142857146).
6 region(4, 6, 12.10117838277123, 113.6176351076798).
7 region(5, 6, 187.05168765367708, 205.03677073243423).
8 region(6, 7, 132.56962025316454, 2.1139240506329116).
9 region(7, 7, 381.81188118811883, 13.386138613861386).
10 region(8, 7, 375.913612565445, 95.30366492146597).
```

```
1 :- dynamic adjacent/2.
2 adjacent(0, 5).
3 adjacent(1, 5).
4 adjacent(1, 6).
5 adjacent(2, 5).
6 adjacent(3, 5).
7 adjacent(3, 7).
8 adjacent(3, 8).
9 adjacent(4, 5).
```

```
water(Region) :- region(Region,3,_,_).
```

```
flooded(Region,Label) :- region(Region,Label,_,_),
    (Label \= 7),
    findall(Adjacent, (adjacent(Region,
        Adjacent),water(Adjacent)),
        AdjWaterRegions, length(AdjWaterRegions,
        N), N >= 2.
```

```
has_flooded_regions :- flooded(_,_).
```

|           | Train | Validation | Test |
|-----------|-------|------------|------|
| Precision | 57 %  | 58 %       | 58 % |
| Recall    | 83 %  | 77 %       | 73 % |



[https://github.com/andrearafanelli/MAS\\_Py\\_FLOOD](https://github.com/andrearafanelli/MAS_Py_FLOOD)




# Knowledge Injection

## Applications

### How to use symbolic knowledge to improve agent training?



- Needs **huge datasets** to generalize well.
- Training is **slow and resource-intensive** (days, weeks, even months).  
*“Training large systems can result in hundreds of metric tons of CO2 emissions”<sup>1</sup>.*



<sup>1</sup> <https://www2.deloitte.com/content/dam/Deloitte/us/Documents/consulting/us-nvidia-gpu-vs-cpu.pdf>

# Guided Learning (Reasoning for Learning)

- Use symbolic representations to guide the sub-symbolic learning process (symbolic knowledge injection).
- **Translation:** *constraining* of symbolic knowledge into sub-symbolic model.
- **Interaction:** symbolic component interacts with the sub-symbolic component in real-time as it operates, i.e. *concurrent interaction*.





# (A) Learning to Search (faster)

- The agent can choose from 3 *discrete action primitives*:

- ❖ **Rotate left** by a fixed angle
- ❖ **Move forward** by a fixed distance.
- ❖ **Rotate right** by a fixed angle.


- The symbolic module overrides the neural controls when needed:
  - **If** Obstacle detection **Then** avoid collision
  - **If** Room Lingering **Then** towards door




# Experiment Results

- The **room change** rate (left) quantifies the agent's exploratory drive.
- The *RL agent lingered in the same room* ( $> 85\%$  of all episodes).
- Non-logic agent scored zero in more than 80% of episodes.
- The hybrid agent scored zero in 60% of episodes .

Room change percentage



Score percentage




<https://github.com/andrearafanelli/Neural-Logic-Reinforcement-Learning>



# (B) Learning to Substitute (efficiently)

## Approach:

- Agent to learn interactively from a few examples.



## Substitution Example:

“What can I use instead of butter, for this cake recipe?”

### ❖ Query:

- Recipe
- Source Ingredient

“You can use oil”

### ❖ Target:

- Substitute Ingredient


“Stop replacing salt with sugar!”: Towards intuitive human-agent teaching. N. Kondylidis, A. Rafanelli, I. Tiddi, A. ten Teije, F. van Harmelen.



# Substitution Query Representation




## Recipe Ingredients:

- Salted Butter
- Almond
- Almond extract
- Powdered sugar
- Vanilla extract
- **Salt** - (To Replace)
- Flour





# Experiment Results



| Learning   Ingr. Representation          | 100 Examples |        |       | 10k Examples |        |       |
|------------------------------------------|--------------|--------|-------|--------------|--------|-------|
|                                          | Hit@1        | Hit@10 | MRR   | Hit@1        | Hit@10 | MRR   |
| Accumulative   1-hot                     | 3.07         | 6.72   | 4.56  | 16.32        | 41.80  | 24.87 |
| Accumulative   1-hot & FoodOn            | 3.82         | 10.49  | 6.25  | 16.90        | 45.52  | 26.43 |
| Balanced   Accumulative   1-hot & FoodOn | 10.99        | 20.82  | 14.59 | 18.97        | 51.19  | 29.67 |

<https://anonymous.4open.science/r/intuitive-human-agent-teaching/README.md>



# Dynamic Agents

## Applications




## How to create agents that adapt dynamically?

As we integrate knowledge, we need to ensure our agents can continuously adapt and refine their understanding.

Adjusts to new inputs, and improves over time.


- Adapts responses: **human centered** [Akata, et al.]




*A research agenda for hybrid intelligence: augmenting human intellect with collaborative, adaptive, responsible, and explainable artificial intelligence., Akata, et al.*

# Iterative Refinement (Learning - Reasoning)

- **Translation:** from symbolic to sub-symbolic , and vice versa.
- **Interaction:** bi-directional, iteratively improving and updating each other, *cyclical interaction*.



# A proposal for human-AI teaming



## Agent

- Sensor
- Sub-Symbolic module for emotion detection
- A behavior tree (BT)

## Knowledge Graph

- User information
- Domain Knowledge

*NEMO - A neural, emotional architecture for human-ai teaming. S. Costantini, P. Dell'Acqua, G. D. Gasperis, F. Gullo, A. Rafanelli.*



# Metrics


## Validation

## How to define metrics for hybrid systems?

Value of Added Knowledge :

Performance with Knowledge - Performance without Knowledge

- Agents in resource-constrained environments
  - Optimize for **minimal energy and computational use.**
- Agents in sensitive domains
  - Prioritize and maximize **robustness.**



# Experiment Results

- Data efficiency gain.
- Energy and latency are almost the same.

| Dataset | Model | Set           | Data efficiency (KB) | Energy (mWh)   | Memory (FLOPs) | Latency (ms) |
|---------|-------|---------------|----------------------|----------------|----------------|--------------|
| BCW     | KBANN | train<br>test | 35.89                | -1.47<br>-0.10 | 3933           | -1.70        |
|         | KILL  | train<br>test | 4.09                 | -0.99<br>0     | 0              | 0.35         |
|         | KINS  | train<br>test | -9.97                | -1.22<br>-0.09 | -559           | -1.41        |
| PSJGS   | KBANN | train<br>test | -4946.81             | -4.67<br>-0.22 | -66944         | -2.56        |
|         | KILL  | train<br>test | 553.89               | -3<br>0        | 0              | 0.04         |
|         | KINS  | train<br>test | -954.80              | -6.53<br>-0.51 | -161779        | -4.75        |
| CI      | KBANN | train<br>test | 1653.79              | -1.41<br>-0.02 | -2468          | -0.43        |
|         | KILL  | train<br>test | 4016.90              | -0.70<br>0     | 4200           | 0            |
|         | KINS  | train<br>test | 4263.50              | -1.41<br>-0.02 | -6220          | -0.44        |



<https://github.com/pikalab-unibo/ski-qos-jaamas-experiments-2022>

*Symbolic knowledge injection meets intelligent agents. A. Agiollo, A. Rafanelli, M. Magnini, G. Ciatto, and A. Omicini*



# Quality Metric


**Robustness** *i.e., how much the predictor is affected by perturbations in the training dataset;*

- Alter the dataset by **adding**, **removing**, or **editing** its entries denote the perturbed dataset .
- Take into consideration also the *magnitude of the perturbation*.

$$\rho_{N,D}(\mathbf{D}) = \frac{1}{n} \sum_{\Delta D \in \mathbf{D}} \|\Delta D\| \cdot \frac{\pi(N_{\Delta D}, D \circ \Delta D)}{\pi(N, D)}$$

$$R_{N,D}(\mathcal{I}) = \frac{\rho_{\hat{N},D}(\mathbf{D})}{\rho_{N,D}(\mathbf{D})}$$

$R_{N,D}(\mathcal{I}) > 1$  indicates the symbolic knowledge create a more robust predictor.



*An empirical study on the robustness of knowledge injection techniques against data degradation. A. Rafanelli, M. Magnini, A. Agiollo, G. Ciatto, and A. Omicini.*



# Experiment Results



## Data drop

agent lacking training data.



## Noise addition

data acquisition process affected by error.



## Label flipping

data labelling process affected by error.


| Dataset | $R_{N,D}(\mathcal{I})$ drop |        |        | $R_{N,D}(\mathcal{I})$ noise |        |        | $R_{N,D}(\mathcal{I})$ flip |        |        |
|---------|-----------------------------|--------|--------|------------------------------|--------|--------|-----------------------------|--------|--------|
|         | KINS                        | KILL   | KBANN  | KINS                         | KILL   | KBANN  | KINS                        | KILL   | KBANN  |
| BCW     | 1.0493                      | 1.0318 | 1.0382 | 0.9960                       | 0.9985 | 1.0109 | 0.9994                      | 1.0184 | 0.9520 |
| PSJGS   | 1.0045                      | 0.9968 | 0.8425 | 0.9950                       | 0.9984 | 1.0145 | 0.9962                      | 1.0026 | 1.6749 |
| CI      | 0.9998                      | 1.0039 | 1.0043 | 0.9992                       | 1.0012 | 0.9965 | 0.9897                      | 1.1703 | 0.9815 |

<https://github.com/pikalab-unibo/ski-data-robustness-experiments>



# Conclusion

- Learning and reasoning integration help training with
  - less data
  - less time / energy to spend
- Importance of considering efficiency
  - symbolic knowledge injection brings data efficiency benefits reducing data for train
    - robustness
    - and data efficiency metric



## ▶ Learning Efficiency

The ability to learn from minimal data inputs

## ▶ Energy Efficiency

Reducing energy consumption in AI processes

## ▶ Human Adaptability


AI systems adjusting to human needs

## ▶ Interpretability

Making AI decisions understandable



# Thank You



# List of Published Papers (I)



A. Rafanelli, S. Costantini, and G.D. Gasperis. A multi-agent-system framework for flooding events. In A. Ferrando and V. Mascardi, editors, Proceedings of the 23rd Workshop "From Objects to Agents", Genova, Italy, September 1-3, 2022, volume 3261 of CEUR Workshop Proceedings, pages 142–151. CEUR-WS.org, 2022.

A. Rafanelli, S. Costantini, and G. D. Gasperis. Neural-logic multi-agent system for flood event detection. *Intelligenza Artificiale*, 17(1):19–35, 2023. doi: 10.3233/IA-230004.

A. Rafanelli, S. Costantini, and G. D. Gasperis. Experimenting an approach to neuro-symbolic RL. In R.D. Benedictis, M. Castiglioni, D. Ferraioli, V. Malvone, M. Maratea, E. Scala, L. Serafini, I. Serina, E. Tosello, A. Umbrico, and M. Vallati, editors, Proceedings of the the Italian Workshop on Planning and Scheduling, RCRA Workshop on Experimental evaluation of algorithms for solving problems with combinatorial explosion, and SPIRIT Workshop on Strategies, Prediction, Interaction, and Reasoning in Italy (IPS-RCRA-SPIRIT 2023) co-located with 22nd International Conference of the Italian Association for Artificial Intelligence AIxIA 2023, November 7-9th, 2023, Rome, Italy, volume 3585 of CEUR Workshop Proceedings. CEUR-WS.org, 2023.

G. De Gasperis, S. Costantini, A. Rafanelli, P. Migliarini, I. Letteri, and A. Dyoub. Extension of constraint-procedural logic-generated environments for deep Q-learning agent training and benchmarking. *Journal of Logic and Computation*, exad032, 06 2023. ISSN 0955-792X. doi: 10.1093/logcom/exad032. URL <https://doi.org/10.1093/logcom/exad032>.

S. Costantini, P. Dell'Acqua, G. D. Gasperis, F. Gullo, and A. Rafanelli. NEMO - A neural, emotional architecture for human-ai teaming. In E. D. Angelis and M. Proietti, editors, Proceedings of the 39th Italian Conference on Computational Logic, Rome, Italy, June 26-28, 2024, volume 3733 of CEUR Workshop Proceedings. CEUR-WS.org, 2024.

A. Rafanelli, M. Magnini, A. Agiollo, G. Ciatto, and A. Omicini. An empirical study on the robustness of knowledge injection techniques against data degradation. In M. Alderighi, M. Baldoni, C. Baroglio, R. Micalizio, and S. Tedeschi, editors, Proceedings of the 25th Workshop "From Objects to Agents", Bard (Aosta), Italy, July 8-10, 2024, volume 3735 of CEUR Workshop Proceedings, pages 20–32. CEUR-WS.org, 2024. URL [https://ceur-ws.org/Vol-3735/paper\\_02.pdf](https://ceur-ws.org/Vol-3735/paper_02.pdf).



# List of Published Papers (II)



A. Agiollo, A. Rafanelli, and A. Omicini. Towards quality-of-service metrics for symbolic knowledge injection. In A. Ferrando and V. Mascardi, editors, WOA 2022 – 23rd Workshop “From Objects to Agents”, volume 3261 of CEUR Workshop Proceedings, pages 30–47, Nov. 2022. URL <http://ceur-ws.org/Vol-3261/paper3.pdf>.

A. Rafanelli, S. Costantini, and A. Omicini. Position paper: On the role of abductive reasoning in semantic image segmentation. In A. Dovier, A. Montanari, and A. Orlandini, editors, Proceedings of the Discussion Papers - 22nd International Conference of the Italian Association for Artificial Intelligence (AIxIA 2022 DP), Udine, Italy, November 28 - December 2, 2022, volume 3419 of CEUR Workshop Proceedings, pages 75–84. CEUR-WS.org, 2022.

A. Agiollo, A. Rafanelli, M. Magnini, G. Ciatto, and A. Omicini. Symbolic knowledge injection meets intelligent agents: QoS metrics and experiments. *Autonomous Agents and Multi-Agent Systems*, 37(2):27:1–27:30, June 2023. ISSN 1573-7454. doi: 10.1007/s10458-023-09609-6.

A. Rafanelli. Beyond traditional neural networks: Toward adding reasoning and learning capabilities through computational logic techniques. In E. Pontelli, S. Costantini, C. Dodaro, S. A. Gaggl, R. Calegari, A. S. d’Avila Garcez, F. Fabiano, A. Mileo, A. Russo, and F. Toni, editors, Proceedings 39th International Conference on Logic Programming, ICLP 2023, Imperial College London, UK, 9th July 2023 - 15th July 2023, volume 385 of EPTCS, pages 416–422, 2023. doi: 10.4204/EPTCS.385.51.

P. Dell’Acqua, S. Costantini, A. Dyoub, G. D. Gasperis, A. Monaldini, and A. Rafanelli. Empathy-aware behavior trees for social care decision systems. In M. M. Cruz-Cunha, D. Domingos, E. Peres, and R. Rijo, editors, CENTERIS 2023 - International Conference on ENTERprise Information Systems / ProjMAN - International Conference on Project MANagement / HCist - International Conference on Health and Social Care Information Systems and Technologies 2023, Porto, Portugal, November 8-10, 2023, volume 239 of *Procedia Computer Science*, pages 1727–1735. Elsevier, 2023. doi: 10.1016/J.PROCS.2024.06.351.

F. Persia, M. Ge, G. Pilato, D. D’Auria, A. Rafanelli, S. Costantini, and G. D. Gasperis. Leveraging DALI to refine route planning by dynamically avoiding risky pois. In 18th IEEE International Conference on Semantic Computing, ICSC 2024, Laguna Hills, CA, USA, February 5-7, 2024, pages 351–354. IEEE, 2024. doi: 10.1109/ICSC59802. 2024.00061.



# List of Published Papers (III)



S. Costantini, P. Dell'Acqua, G. D. Gasperis, F. Gullo, and A. Rafanelli. The NEMO co-pilot. In S. D. Martino, C. Sansone, E. Masciari, S. Rossi, and M. Gravina, editors, Proceedings of the Ital-IA Intelligenza Artificiale - Thematic Workshops co-located with the 4th CINI National Lab AIIS Conference on Artificial Intelligence (Ital-IA 2024), Naples, Italy, May 29-30, 2024, volume 3762 of CEUR Workshop Proceedings, pages 124–128. CEUR-WS.org, 2024. URL

S. Costantini, P. Dell'Acqua, G. De Gasperis, and A. Rafanelli. Empowering emotional behavior trees with neural computation for digital forensic. 15th European Symposium on Computational Intelligence and Mathematics (ESCIM 2024).

G. Amaral, S. Costantini, G. D. Gasperis, P. Dell'Acqua, G. Guizzardi, F. Gullo, and A. Rafanelli. Neural empathy-aware behavior trees meets knowledge graphs for affective human-ai teaming. Under revision in the IEEE Transactions on Affective Computing Journal.

N. Kondylidis, A. Rafanelli, I. Tiddi, and A. ten Teije., and F. van Harmelen, “Stop replacing salt with sugar!”: Towards intuitive human-agent teaching. Under double-blind submission.

