The NEMO co-pilot

Stefania Costantini [a] Pierangelo Dell'Acqua[b] Giovanni De Gasperis[a]
Francesco Gullo[a] Andrea Rafanelli[a]

[a]University of Aquila, Italy
[b]University of Linkoping, Sweden

CILC 2024, 39th Italian Conference on Computational Logic, 26-28 Giugno 2024, CNR - Roma

Introduction

A recent focus in Artificial Intelligence (AI) is the development of intelligent systems in which humans and AI Systems work together as teams (human-AI teaming, HAI).

When working together, humans and AI can produce results that exceed what they can achieve alone, whereas they can control and improve each other. A particularly important application of human-AI teaming pertains critical tasks which are burdensome and error-prone for humans.

AI Agents can thus be endowed with emotion recognition and be capable of empathy and modeling aspects of the Theory of Mind (ToM), in the sense of being able to reconstruct what the human is thinking or feeling.

Introduction (cont'd)

Affective computing refers to the domain of computing focused on leveraging emotions, feelings, moods, and other aspects of human psychology to accomplish specific tasks

Empathy brings benefits to social interactions, fostering better relationships. In fact, endowing artificial agents with empathetic capabilities enhances human-agent interactions. So, computational modelling of empathy has become an active research area.

This work focuses on modelling empathetic behaviour in virtual agents to be employed for user support in critical fields.

An Envisaged Application

In The COST Action DigForASP, we explored methods for aiding investigators by using AI, particularly computational logic and computational intelligence applied to digital forensics.

Digital forensics experts still face high pressure in high-stakes cases involving serious crimes. Long hours poring through volumes of data on multiple devices, though with the help of AI tools, can take a physical and mental toll.

The same can happen to medical doctors, patients and caregivers, who face long hours with difficult and tiresome tasks and critical situations. (PRIN Projects TRUSTPACTX and DECODER)

Envisaged Application to HAI

The framework proposed in this work integrates emotional intelligence and empathy in a rigorous formal setting, aiming to provide adequate support to human users in high-pressure, high-stakes tasks.

By monitoring and analysing a human's affective and physiological states, intelligent personal agents can adapt their behaviour to provide empathetic support. Thus, they can help humans maintain their adequacy and accuracy in the task/situation at hand and preserve their physical and mental well-being.

A Formal Tool: Behaviour Trees (BTs)

A behaviour tree is a mathematical model of plan execution, widely adopted, for example, in videogames to model the behaviour of non-player characters (NPCs)

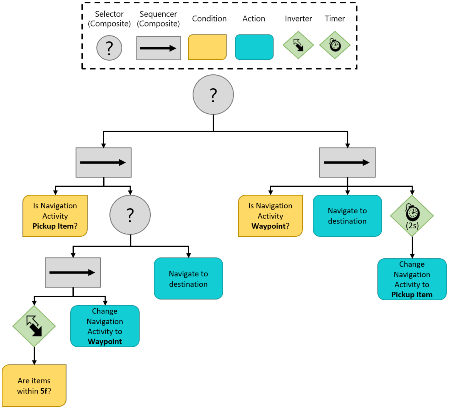

Our Proposal for HAI Agents: Affective States in BTs

Implemented in Prolog

We first enhanced BTs with an "affective state", composed of a set of variables representing the user's features and the agent's "emotions", i.e., how the agent "feels" in a human-relatable sense.

The agent elaborates the affective state during repeated interactions with the user and can thus tune its reaction accordingly.

Our Proposal for HAI Agents: Neural Empathy-Aware Behavior Trees (NEABT)

We enhance the basic BT definition with various kinds of nodes:

Neural Node
A neural node draws conclusions about a user's emotional state through a deep learning model. It takes sensor input from the environment (e.g., wearable devices, videocameras, etc.) catching sensory data like facial expressions, vocal tones, and body language, and the agent's previous emotional state. The node's outcome contributes to updating the agent's affective state variables.

By continually updating the agent's internal emotional state, the neural node allows dynamic adaptation of the NEABT to the emotional context.

Our Proposal for HAI Agents: Neural Empathy-Aware Behavior Trees (NEABT)

We enhance the basic BT definition with various kinds of nodes:

Emotional Selector

The emotional selector is a node that orders its child sub-nodes based on a set of relevant decision factors and the agent's current "affective state".

Once the ordering has been established, the emotional selector behaves as a priority selector.

Our Proposal for HAI Agents: Neural Empathy-Aware Behavior Trees (NEABT)

We enhance the basic BT definition with various kinds of nodes:

Empathy Node An empathy node provides an emotional characterization of its single child node. An empathy node can only be a child of an emotional selector.

Its child can be a leaf or an inner node; where empathy nodes cannot be nested. An empathy node is represented by a dashed line circle with the name of the empathy emotion in it.

NEABT of the personal assistant agent for intellectual tasks

The objective is to monitor human's affective states and modulate the agent's behaviour accordingly. We adopt a model characterising the human affect system across six primary emotive states– *anger*, *fear*, *surprise*, *sadness*, *disgust*, and *joy*.

By formalising the agent's behaviour via a NEABT guided by affective factors, the agent can match its collaboration to the human's needs. For human experts, this empathetic support keeps their work productive without compromising human well-being.

NEABT of the personal assistant agent for intellectual tasks

The empathy state variables are:
$\widehat{E} = \{\widehat{\text{annoyance}}, \widehat{\text{interest}}, \widehat{\text{distraction}}, \widehat{\text{apprehension}}\}$

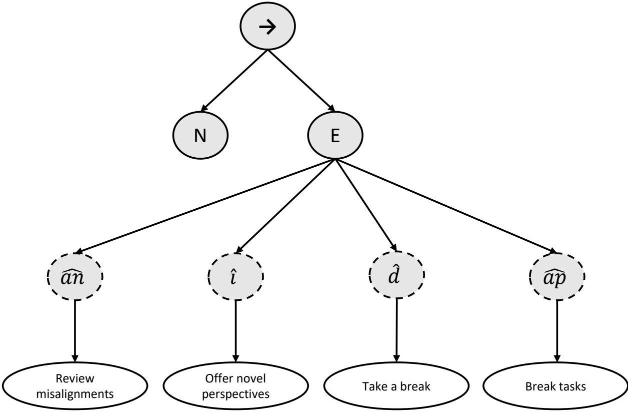

Future Directions

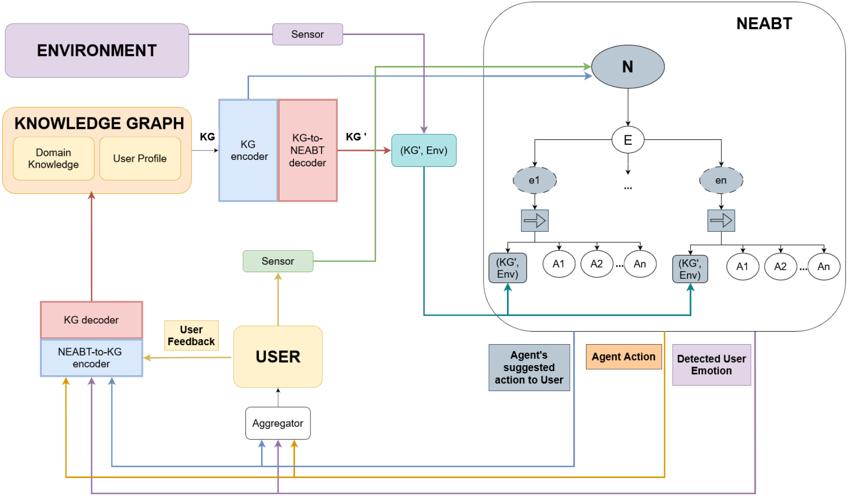

Prolog code for NEABT

```
neabt_structure(E, Children) :-
    % output: Children
    append(Action_nodes, Emphatic_nodes, All_nodes),
    findall(Child, child(E, All_Nodes, Child), Children),
    member(C, Children), C \= E.

neural_node(UserSensor, KGEmb, Prob, EmoState) :-
    % output: EmoState
    nn(UserSensor, KGEmb, Prob, EmoState).

emo_selector(E, EmoState, Prob, KGd, Env, SuccessNodes) :-
    % output: SelectedNodes
    neabt_structure(E, Children),
    select_relevant_nodes(EmoState, Prob, Children, SelectedNodes),
    % select_relevant_nodes is an external algorithm
    findall(Node, (member(Node, SelectedNodes),
                   empathy_node_success(Node, kg_condition(KGd, Env))),
            SuccessNodes).
    % kg_condition(KBd, Env) are facts from the KG

empathy_node_success(N, Condition) :-
    % BT implementation, extract the first that succeed
    neabt_structure(N, Children),
    (
        call(Condition) ->
        foreach(A, Children, action(A),
                execute_actions(A)), % external plugin
        foreach(SubN, Children, node(SubN),
                empathy_node_success(SubN, Condition)),
    ;
    true
    ).

neabt(UserSensor, KG, Env, E, SuccessNodes) :-
    neural_codec(KG, KGd, KGEmb),
    neural_node(UserSensor, KGEmb, Prob, EmoState),
    emo_selector(E, EmoState, Prob, KGd, Env, SuccessNodes).
```

Conclusion

The research that we just illustrated is a follow-up of DigForASP, and falls within the activities of the Italian PRIN Projects:

TrustPACTX - Design of the Hybrid Society Humans-Autonomous Systems: Architecture, Trustworthiness, Trust, EthiCs

ADVISOR - ADaptiVe legIble robotS for trustwORthy health coaching

Participants are the Universities of L'Aquila, Messina, Naples "Federico II", Palermo, and the National Research Council.

Thank you for your Attention!