A Modular REST-Based Framework for Human-in-the-Loop Robot-Assisted personalized Rehabilitation in Neurodevelopmental Disorders

Marco De Luca[1,*], Davide De Tommaso[2], Davide Ghiglino[2], Anna Rita Fasolino[1], Porfirio Tramontana[1] and Agnieszka Wykowska[2]

[1]University of Naples "Federico II", Napoli, Italy
[2]Istituto Italiano di Tecnologia, Genova, Italy

Abstract

Socially Assistive Robots (SARs) are increasingly used in therapeutic settings to support cognitive, social, and motor development through personalized, interactive engagement. However, deploying SARs in real-world clinical environments requires software frameworks that are both adaptable to individual patient needs and accessible to non-technical users. This work presents a modular full-stack REST framework designed to operationalize robot-assisted therapy with the iCub humanoid robot. The system supports the rapid deployment of personalized therapeutic workflows—referred to as *iCub Applications*—by composing modular robot actions into reusable templates. These applications are executed through a Wizard-of-Oz (WoZ) paradigm, enabling a human operator to teleoperate the robot online and maintain adaptive, patient-centered control. An intuitive, web-based frontend allows clinicians and researchers to configure and supervise sessions without requiring programming expertise. The interface dynamically generates its content based on the backend configuration exposed via a REST API, allowing multiple applications to be supported without modifying the frontend code, since all application logic and structure are defined on the backend. By integrating flexibility, modularity, and human-in-the-loop interaction, the framework bridges the gap between low-level robotic control and clinical usability, contributing to the advancement of socially-aware, human-centered AI in assistive and rehabilitation contexts.

Keywords

human–robot interaction, software architecture, socially assistive robots, neurodevelopmental disorders

1. Introduction

Integrating robotics into therapeutic care has opened new opportunities to deliver engaging, adaptive, and personalized support for individuals with developmental, cognitive, and motor challenges. In particular, Socially Assistive Robots (SARs) have shown promise in domains such as autism therapy, motor rehabilitation, and cognitive training [1, 2]. Through their physical embodiment and ability to communicate using multiple channels—speech, gesture, facial expressions—SARs can foster social interaction, increase motivation, and enhance patient involvement.

However, the effectiveness of SAR-based interventions depends heavily on the underlying software architecture. Simply deploying robot hardware is not sufficient; what is needed is a system that enables real-time control and adaptation of robot behavior according to individual user responses. This highlights the central role of Human–Robot Interaction (HRI) design, particularly in clinical settings that require personalization, transparency, and therapist supervision [3].

In this context, the iCub humanoid robot[4] stands out as a valuable platform. Its anthropomorphic form, child-like dimensions, and comprehensive sensorimotor capabilities make it particularly well-suited for working with children and vulnerable populations. The software infrastructure behind iCub is built on YARP (Yet Another Robot Platform)[5], a modular middleware designed to enable distributed, real-time communication among robot components. This architecture supports integration of features such as speech synthesis, gaze control, and gesture coordination.

Workshop on Social Robotics for Human-Centered Assistive and Rehabilitation AI (a Fit4MedRob event) - ICSR 2025
[*]Corresponding author.
marco.deluca2@unina.it (M. D. Luca); davide.detommaso@iit.it (D. D. Tommaso); davide.ghiglino@iit.it (D. Ghiglino); fasolino@unina.it (A. R. Fasolino); ptramont@unina.it (P. Tramontana); agnieszka.wykowska@iit.it (A. Wykowska)

Despite this potential, developing therapeutic sessions with iCub is a challenge for non-technical users. Therapists and clinical staff typically require programming skills to define and execute interaction scenarios, which limits the scalability and day-to-day usability of the platform in healthcare environments. This creates a gap between what the robot is capable of and what clinicians are able to deploy independently.

To address this limitation, we propose a new software framework that allows the configuration and execution of structured *iCub Applications*, customizable therapeutic workflows built from modular robot actions. These applications are modeled using Finite State Machines (FSMs) and follow a Wizard-of-Oz (WoZ) execution model [6, 7], where a human operator guides the robot's behavior during the session. This hybrid approach maintains session structure while enabling online adaptation to patient behavior.

The system is built on a full-stack REST architecture that cleanly separates the backend—responsible for protocol logic, application definition, and robot coordination—from the frontend, which provides an intuitive web-based interface for therapists and researchers. Notably, the frontend dynamically generates its content based on definitions received via the REST API, allowing new applications to be deployed or reconfigured entirely from the backend without requiring changes to the frontend code. This design supports accessibility, scalability, and adaptability, making the platform well-suited for both clinical and research settings.

2. Adaptive Therapeutic Sessions with the iCub Robot Platform

The iCub humanoid robot is a flexible research platform for Human–Robot Interaction (HRI), cognitive robotics, and therapeutic settings [4]. Its child-sized body, 53 degrees of freedom, and rich sensor suite—including stereo vision, tactile skin, microphones, and force sensors—make it especially suitable for engaging with children and individuals with neurodevelopmental disorders. These capabilities enable natural, multimodal interactions through gaze, speech, gestures, and touch.

iCub's software is built on the modular YARP middleware [5], which separates low-level hardware control from high-level behaviors. This design exposes core functionalities—like speech output, gaze shifts, or motor control—as independent services that can be flexibly combined. Therapists can therefore activate only the robot components needed for each session, making it easier to adapt interactions to specific therapeutic goals.

Within this framework, we define an *iCub Application* as a structured sequence of robot actions designed to support a specific therapeutic activity. Each application orchestrates the robot's multimodal capabilities into a coherent interaction flow that aligns with clinical objectives, such as promoting joint attention, encouraging verbal interaction, or guiding motor tasks.

Applications are executed using a Wizard-of-Oz (WoZ) approach [8], where the robot appears autonomous to the patient but is in fact remotely controlled by a human operator. This model offers both flexibility and safety—particularly important in pediatric or cognitively sensitive contexts—while maintaining the illusion of autonomy that is often crucial for engagement. In a typical session setup (Figure 1), the therapist and patient are seated across from iCub, with the therapist operating the session from a separate interface, out of the patient's view. The therapist selects an *iCub Application* and configures its session parameters according to the patient's profile. Throughout the session, the therapist assesses the patient's responses and selects the appropriate robot actions, ensuring that the experience is personalized, engaging, and therapeutically relevant.

This framework supports a human-in-the-loop interaction model in which robot behavior is not pre-scripted but flexibly directed by the therapist's clinical judgment. It offers a balance between repeatability—important for structured therapy—and real-time adaptability—essential for personalization. As such, iCub serves not just as a programmable robot, but as a responsive therapeutic partner capable of supporting individualized, interactive interventions.

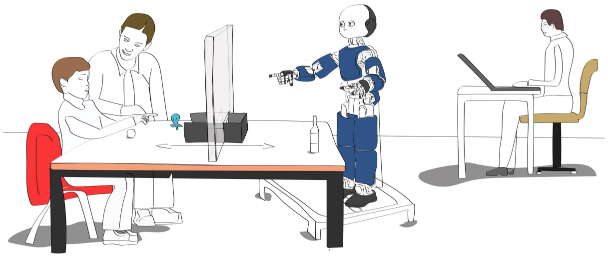

Figure 1: Experimental Session Setup in a Wizard of Oz setting.

3. Proposed Solution

To support the execution of structured yet flexible robot-assisted therapy sessions using the iCub humanoid platform, we introduce a modular, full-stack REST software framework. The proposed framework is designed to abstract the technical complexity of robotic control and provide an accessible interface for clinicians and technical operators to define, configure, and supervise therapeutic protocols without requiring robotics or programming expertise.

At the core of our framework is the support for executing *iCub Applications*, reusable, high-level therapy programs defined as structured sequences of robot actions. Each action corresponds to a capability of the robot, such as gaze control, speech synthesis, gesture execution, or facial expression rendering. These applications are executed under the WoZ paradigm, which places a human operator in the control loop, enabling real-time guidance of robot behavior based on the patient's engagement and responses. This approach supports patient-specific adaptation while maintaining the safety and repeatability required in clinical settings.

To support the operator during the execution of each application, the framework incorporates three key design choices. First, each *iCub Application* is modeled as a Finite State Machine (FSM), where individual states correspond to specific robot actions and transitions represent the flow of interaction. According to the WoZ paradigm, these transitions are triggered by the therapist in real time, based on their observation of the patient's behavior. This approach ensures that clinicians retain full control over the session while allowing for dynamic, personalized adjustments. An example of an FSM representation of an *iCub Application* is illustrated in Figure 2.

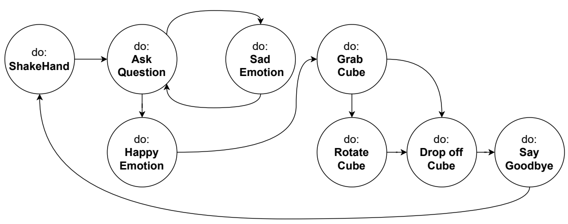

Figure 2: Example of a Finite State Machine representation of an iCub Application

Second, to facilitate therapist interaction with the application during execution, the system provides a web-based dashboard based on the selected *iCub Application*. This dashboard is composed of modular graphical widgets, as shown in Figure 3. The main widget consists of an FSM visualizer that can be used by the operator to track and control the execution of the application. Other dashboard widgets correspond to a specific robot function or protocol control, such as Text-to-Speech control, emotion display configuration, and video stream monitoring.

Third, the frontend dashboard is dynamically generated at runtime based on definitions retrieved from the backend via the REST API. This design allows the addition or customization of therapy applications solely through backend configuration, without requiring any changes to the frontend code. A key feature of this approach is the ability to define configurable session arguments—parameters that can be selected by the user at the start of the session (e.g., task difficulty, target emotion, or number of repetitions). The structure and allowed values of these arguments are fully defined in the backend and automatically reflected in the user interface at runtime. Moreover, the layout of the dashboard widgets is highly customizable and can be modified, imported, or exported as a JSON configuration file. This enables flexible adaptation of the interface to specific clinical needs or user preferences, supporting a variety of use cases and session types. Finally, custom Python methods defined in the main application backend are automatically exposed as REST endpoints and can be selectively included in the dashboard on request. This allows developers to quickly integrate new robot capabilities or protocol logic without additional frontend development, further enhancing the modularity and extensibility of the framework.

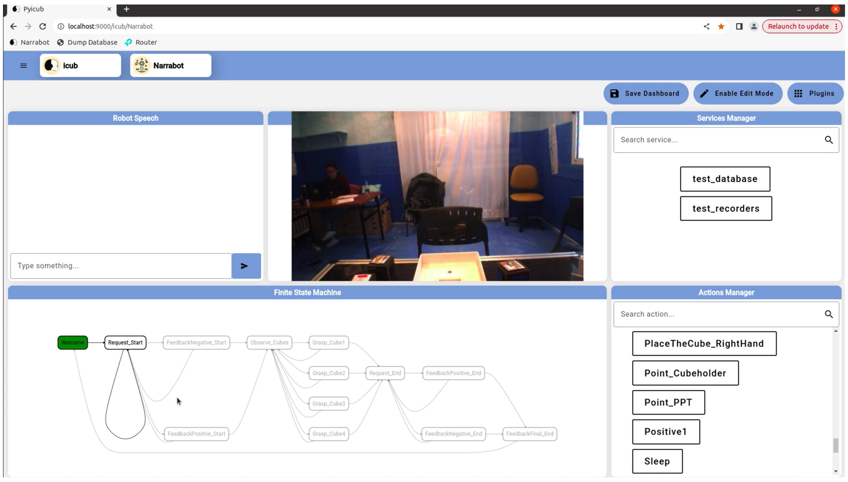

Figure 3: An example of configurable Dashboard offered by the framework

At the start of a rehabilitation session, the operator can select one of the available *iCub Applications* and, therefore, the frontend automatically configures the dashboard to reflect the structure and parameters of the selected FSM. Therapists can further personalize the interface by rearranging, resizing, or disabling widgets based on the therapeutic goals and interaction context. Through this configurable dashboard, the therapist remains central to the interaction, manually triggering state transitions in response to the patient's real-time behavior. This human-in-the-loop approach ensures that sessions are not only responsive and personalized, but also maintain the consistency and structure necessary for clinical reliability and experimental reproducibility.

The overall software architecture adopts a client–server model, as shown in Figure 4. The backend component, named `Pyicub`, handles robot coordination, FSM-driven session execution, and service orchestration. It exposes REST APIs to the frontend, enabling seamless, asynchronous communication between user-facing components and the robot's control logic.

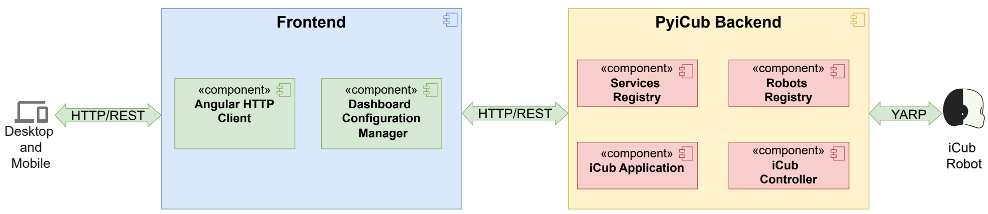

Figure 4: High-Level View of the Full-Stack REST Architecture for iCub Applications execution

The backend is structured into four main modules that handle robot communication and therapeutic session execution. The `iCub Controller` abstracts low-level interactions with the robot through the YARP middleware [5], a soft real-time, distributed communication framework based on named Ports. It exposes robot functionalities (e.g., gaze control, arm movement, speech output) as RESTful services. This abstraction allows the backend to trigger complex robot behaviors through simple HTTP calls, enabling hardware-independent control across deployments. The `iCub Application` module manages the logic of each therapeutic session, orchestrating robot actions by invoking the REST endpoints according to the protocol's flow. By composing sessions from modular, atomic actions, this component facilitates both reuse and rapid prototyping of new *iCub Applications*, aligning with the system's goal of flexibility and adaptability. The `RobotsRegistry` automatically detects available iCub units on the network at runtime, while the `ServiceRegistry` maintains a dynamic mapping between robot instances and their supported *iCub Applications*. Together, these components ensure correct routing of requests and enable scalable, distributed execution across multiple robots without requiring manual reconfiguration.

The frontend is a responsive Single Page Application (SPA) built with Angular. At its core, the `HTTP Client` handles all asynchronous REST communications with the backend and user-triggered events, keeping the interface synchronized with system logic. The `Dashboard Configuration Manager` controls the layout and widget setup, allowing the interface to adapt flexibly to different therapeutic sessions while ensuring an intuitive user experience.

In summary, the proposed system delivers a cohesive full-stack REST architecture that bridges the gap between iCub's technical capabilities and the real-world needs of clinicians and researchers. By modularizing robot functionality, supporting dynamic session configuration, and providing a user-friendly frontend, the system enables accessible, repeatable, and adaptive robot-assisted therapy—advancing the practical integration of socially assistive robotics in clinical practice.

4. Conclusion and Future Work

This work introduced a modular, full-stack REST framework that supports the configuration and execution of structured, adaptable therapeutic sessions using the iCub robot. By abstracting low-level control through a unified interface and adopting the Wizard-of-Oz paradigm, the system allows therapists to configure and execute *iCub Applications*—reusable therapeutic protocols composed of modular robot actions—directly via a web-based frontend. This approach enables real-time personalization based on patient needs, without requiring any programming skills, while preserving the structure of clinical workflows. The backend, built on the YARP middleware, ensures scalability, reusability, and hardware abstraction via RESTful APIs. Meanwhile, the frontend offers a dynamic, customizable dashboard that supports human-in-the-loop interaction. Overall, the proposed system advances the integration of socially assistive robotics by aligning technical flexibility with the practical needs of adaptive, human-centered care.

The proposed solution was preliminarily validated through a case study involving technical operators and therapists who used the web application to conduct therapy sessions with actual patients and a single robot. Their feedback provided valuable insights into the usability and clinical applicability of the system, confirming its potential to support personalized, structured robot-assisted interventions in everyday therapeutic practice.

As future work, we plan to extend the experimental validation by involving a broader range of clinicians and therapists to evaluate the system's effectiveness in real-world therapeutic workflows. Furthermore, we intend to develop a graphical user interface that will enable users to define new *iCub Applications* directly from the browser. This tool will guide the composition of new FSM-based protocols by selecting and chaining available iCub actions—exposed as REST services through the YARP middleware—thus making application creation more accessible to non-technical users. Finally, we are exploring the integration of semi-autonomous behavior layers to complement the Wizard-of-Oz approach, enabling a more balanced interaction between clinician control and robot initiative in long-term interventions.

5. Declaration on Generative AI

The author(s) has not employed any Generative AI tool during the preparation of this work.

Acknowledgments

This work was supported by the Italian Ministry of Research, under the complementary actions to the NRRP "Fit4MedRob - Fit for Medical Robotics" Grant (# PNC0000007).

References

[1] A. Triantafyllidis, A. Alexiadis, K. Votis, D. Tzovaras, Social robot interventions for child healthcare: A systematic review of the literature, Computer Methods and Programs in Biomedicine Update 3 (2023) 100108. URL: doi.org/10.1016/j.cmpbup.2023.100108. doi:https://doi.org/10.1016/j.cmpbup.2023.100108.

[2] M. Kyrarini, F. Lygerakis, A. Rajavenkatanarayanan, C. Sevastopoulos, H. R. Nambiappan, K. K. Chaitanya, A. R. Babu, J. Mathew, F. Makedon, A survey of robots in healthcare, Technologies 9 (2021). URL: 10.3390/technologies9010008. doi:10.3390/technologies9010008.

[3] J. D. Kumar, N. Bansal, A. Kaushik, A. Sethi, Human-robot interaction: Designing effective interfaces for collaborative tasks, in: 2024 1st International Conference on Advances in Computing, Communication and Networking (ICAC2N), 2024, pp. 356–360. URL: 10.1109/ICAC2N63387.2024.10895441. doi:10.1109/ICAC2N63387.2024.10895441.

[4] G. Metta, L. Natale, F. Nori, G. Sandini, D. Vernon, L. Fadiga, C. von Hofsten, K. Rosander, M. Lopes, J. Santos-Victor, A. Bernardino, L. Montesano, The icub humanoid robot: An open-systems platform for research in cognitive development, Neural Networks 23 (2010) 1125–1134. URL: doi.org/10.1016/j.neunet.2010.08.010. doi:https://doi.org/10.1016/j.neunet.2010.08.010, social Cognition: From Babies to Robots.

[5] G. Metta, P. Fitzpatrick, L. Natale, Yarp: Yet another robot platform, International Journal of Advanced Robotic Systems 3 (2006) 8. URL: 10.5772/5761. doi:10.5772/5761.

[6] J. F. Kelley, An iterative design methodology for user-friendly natural language office information applications, ACM Trans. Inf. Syst. 2 (1984) 26–41. URL: 10.1145/357417.357420. doi:10.1145/357417.357420.

[7] A. Steinfeld, O. C. Jenkins, B. Scassellati, The oz of wizard: simulating the human for interaction research, in: Proceedings of the 4th ACM/IEEE International Conference on Human Robot Interaction, HRI '09, Association for Computing Machinery, New York, NY, USA, 2009, p. 101–108. URL: 10.1145/1514095.1514115. doi:10.1145/1514095.1514115.

[8] L. D. Riek, Wizard of oz studies in hri: a systematic review and new reporting guidelines, J. Hum.-Robot Interact. 1 (2012) 119–136. URL: 10.5898/JHRI.1.1.Riek. doi:10.5898/JHRI.1.1.Riek.