Towards Autonomous Verification: Integrating Cognitive AI and Semantic Digital Twins in Medical Robotics

Patrick Mania[1], Michael Neumann[1], Franklin Kenghagho Kenfack[1] and Michael Beetz[1]

Abstract— **In medical laboratory environments, where precision and safety are critical, the deployment of autonomous robots requires not only accurate object manipulation but also the ability to verify task success to comply with regulatory requirements. This paper introduces a novel imagination-enabled perception framework that integrates cognitive AI with semantic digital twins to allow medical robots to simulate task outcomes, compare them with real-world results, and autonomously verify the success of their actions. Our approach addresses challenges related to handling small and transparent objects commonly found in sterility testing kits and other related consumables. By enhancing the RoboKudo perception system with parthood-based reasoning, we enable more accurate task verification through focused attention on object subparts. Experiments show that our system significantly improves performance compared to traditional object-centric methods, increasing accuracy in complex environments without the need for extensive retraining. This work demonstrates a novel concept in making robotic systems more adaptable and reliable for critical tasks in medical laboratories.**

I. Introduction

In recent years, the integration of robotic systems into medical laboratory settings has gained significant attention due to the potential for increased efficiency, precision, and safety in tasks that are often repetitive and prone to human error. However, the successful deployment of such systems in critical environments, such as sterility testing processes, necessitates not only the ability to perform complex manipulation tasks but also the capability to verify the success of these tasks autonomously. A key challenge in this domain is ensuring that robots can accurately assess whether their actions have achieved the desired effect, especially when dealing with delicate and transparent objects, such as those found in commonly used laboratory testing kits.

To overcome these challenges, we introduce a novel framework that extends imagination-enabled robot perception[10] specifically for the complex demands of medical robotics, enabling robots to autonomously hypothesize, simulate, and verify task outcomes in real-time. By simulating the post-action environment within a semantic digital twin—a highly detailed virtual replica of the real-world scenario equipped with semantic knowledge—the robot can compare this hypothesized state with the actual outcome observed through its sensors. This capability enables the system to autonomously verify task completion or identify the need for corrective actions, which is a crucial requirement for flexible robotic systems to meet regulatory standards where correct task execution must be proven.

Our work is inspired by the principles outlined in [17] which emphasizes the importance of cognitive AI systems that can reason about their actions and verify their outcomes. The work define five core domains: Functionality, Physics, Intent, Causality, and Utility (FPICU). These domains are key to enabling AI systems to reason autonomously and verify task success. Understanding the functional properties of objects helps the system predict and validate the outcomes of interactions, especially with transparent or small objects found in medical labs. Causal reasoning allows the AI to predict and verify the physical results of actions, ensuring tasks like placing or clamping are correctly performed. Additionally, intuitive physics helps the system to reason about the interactions of the robot and objects within the environment. Applying these principles to robotic systems in medical labs is crucial, as it enables creating robots that not only perform tasks but also understand their actions, hypothesize intended effects, and autonomously confirm successful task execution.

In the demanding environment of sterility testing, where tasks often involve transparent and small objects like those in the Steritest™ NEO kit(see Fig. 1), traditional vision systems face significant challenges. Accurately representing such objects in digital twin simulations is challenging, complicating real-world and hypothesized imagery comparisons. Our approach addresses this by incorporating geometric and material knowledge, improving similarity detection between expected and observed states.

Furthermore, we leverage the flexibility and modularity of the RoboKudo framework[9], which allows for the dynamic adaptation of perception pipelines to the specific demands of different tasks and the usage of multiple vision experts. By integrating object knowledge and task-specific features into this framework, we can simplify the perception process while maintaining high verification accuracy. This approach not only streamlines the process but also reduces the need for retraining vision models for each new task, making the system more versatile and easier to deploy in various laboratory scenarios.

Our key contributions are as follows:

1) We provide the concept for imagination-enabled medical robot perception, outlining the key mechanisms and models for task success verification.
2) To tackle the specific requirements of lab items, we

The research reported in this paper has been (partially) supported by the German Research Foundation DFG, as part of Collaborative Research Center (Sonderforschungsbereich) 1320 Project-ID 329551904 "EASE - Everyday Activity Science and Engineering", University of Bremen (http://ease-crc.org/). The research was conducted in subproject R02. This work was also supported by the European Union's Horizon 2020 research and innovation programme under grant agreement No 101017089 as part of the TraceBot project.

[1]Institute for Artificial Intelligence, University of Bremen, Germany

propose the introduction of object knowledge and parthood relations to provide an attention mechanism of a visual similarity process.
3) We describe the required changes to a robot perception system to extend the typically object-centric analysis towards a more generic analysis scheme, which includes the analysis of subparts of objects in combination with semantic annotations.

II. Related Work

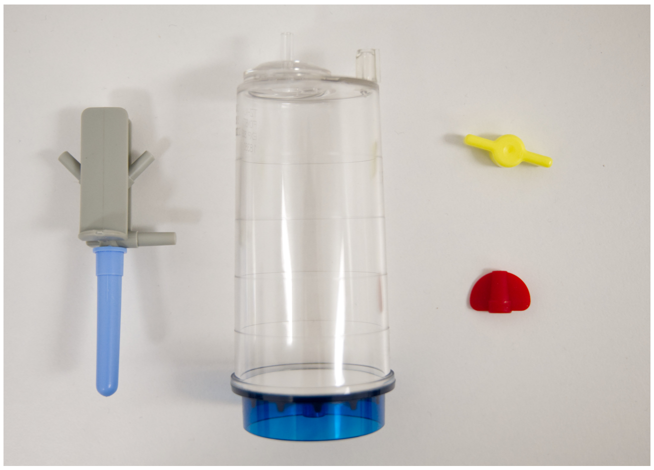

Fig. 1. Photo of key consumables in the Steritest™ NEO kit for sterility testing. It is used as the main example throughout our studies. The red and yellow cap can be attached to the canister (middle) to seal outlets on the top and the bottom. The grey needle on the left is shipped with a protective blue cap.

Robotics in medical and laboratory environments has seen significant growth in recent years, particularly in the automation of repetitive and precise tasks such as drug preparation, sample handling, and sterility testing. Notable advancements have been made in automating processes that traditionally required human intervention, with a focus on increasing efficiency and ensuring sterility in controlled environments.

Medical robots, such as those designed for the automation of repetitive tasks, often operate in highly structured environments where precision is crucial for tasks like sample pipetting, sorting, and sealing. Systems like Tecan's Freedom EVO [2] and Beckman Coulter's Biomek series [7] are widely used for automating laboratory workflows, yet they are limited to rigid, pre-programmed tasks that rely on specialized hardware and offer little flexibility for adaptation. Introducing more flexible robotic systems is a promising solution to solve more repetitive task and increase the overall level of automation in laboratories[15].

In life science laboratories, robotics have been widely deployed for repetitive tasks such as liquid handling, sample transport, and high-throughput screening. Wolf and Széll review existing technologies, such as XYZ gantry robots for liquid handling and selective compliance articulated robots for sample transportation, which are commonly employed in pharmaceutical research and development[14]. Although mature, these systems are rigid, excelling at specific tasks but lacking the flexibility to handle dynamic environments or adapt to new processes. Studies like [1] and [13] highlight how current lab automation boosts productivity but still requires significant human intervention due to specialized hardware and the need for robust task completion verification to meet regulatory standards.

In the specific context of sterility testing, the Steritest™ NEO kit (see Figure 1) represents an example, where manual steps, such as placing caps or securing seals, still demand human involvement due to the challenges associated with manipulating small, flexible and transparent components. Research in medical robotics has yet to fully address the need for systems capable of not only performing these manipulations but also autonomously verifying their successful execution to fulfill regulatory requirements.

The importance of precision and minimizing human error is further emphasized by Zaninotto and Plebani in [16], who explored the role of automation in hospital laboratories. They stress that automation not only ensures higher accuracy but also reduces human error in critical diagnostic processes. Mencacci et al. also noted that automation significantly decreases human involvement, thereby enhancing safety and efficiency in laboratory workflows[11]. However, the inherent limitations of these systems, primarily the lack of cognitive capabilities, still necessitate human intervention for complex or unpredictable tasks.

The need for cognitive robotics arises from the shortcomings of current systems, which are typically limited to task-specific applications. As noted in [15], many existing robotic solutions are highly focused on performing a singular task, lacking the flexibility and adaptability to manage diverse workflows autonomously. Cognitive robotics—systems that can reason, adapt, and make decisions based on the task environment—are essential to overcome this challenge. These systems can dynamically modify their behavior based on real-time data, which is especially crucial in medical or laboratory settings where variability is common.

An excellent example of this evolution towards cognitive robotics is the integration of AI into automated experimentation. Burger et al. [4] demonstrated a mobile robotic chemist that autonomously optimized photocatalyst formulations, conducting nearly 700 experiments in eight days. Using Bayesian optimization, the system efficiently selected chemical combinations for hydrogen production, showcasing AI and robotics' ability to handle complex decision-making tasks traditionally dependent on human expertise.

While the system presented in [4] is primarily focused on chemical research, the underlying principles of autonomous operation, adaptability, and hypothesis testing are applicable to the medical domain as well. Such advancements can be translated into healthcare and laboratory automation, enabling robots to perform more complex tasks, adapt to changing conditions, and reduce the burden on human operators. This shift from task-specific automation to cognitive, adaptable systems is critical for advancing the field.

In conclusion, the current generation of robotics has made significant strides in automating repetitive tasks in medical and laboratory settings, but limitations remain due to their lack of cognitive adaptability. The future lies in integrating AI and cognitive capabilities, as seen in the mobile robotic chemist, to enable robots to not only perform tasks but also autonomously hypothesize, adapt, and optimize their actions. The novel approach of integrating imagination-enabled robot perception, as proposed in this paper, represents an advancement in this direction, providing robots with the ability to predict outcomes and adjust autonomously, making them better suited for life-critical applications.

III. System overview

To create a system within medical robotics that has a „mind's eye" mechanism to imagine how task outcomes shall look like, a comprehensive architecture composed of multiple interacting components is necessary. This system (see Fig. 2) operates by simulating the target environment and allowing the robot to form an internal hypothesis of the expected outcome for manipulation tasks, which is then compared to the actual observed state in the real world (RW). This chapter outlines the essential components and mechanisms that must be developed and integrated to enable this process, as well as the challenges involved.

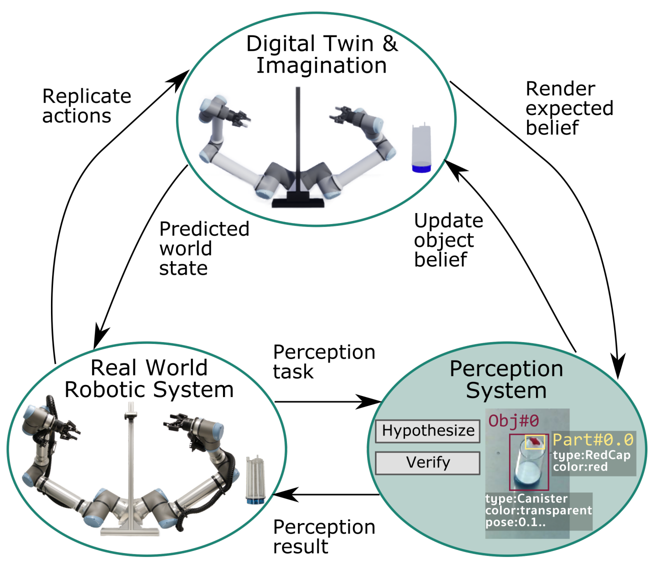

Fig. 2. High-level system overview depicting the interaction of the key components of an imagination-enabled robotic system. A digital twin is continuously maintained by the robot control program and the perception system of the robotic system to enable task success verification.

A. Semantic Digital Twin: A High-Fidelity Simulation Environment

At the core of this approach is the *semantic digital twin* of the selected use case. It is a high-fidelity simulation that mirrors the RW environment in terms of both objects and their physical properties. The digital twin is not limited to the robot itself but includes all relevant elements within the target environment, such as incubators, work surfaces, fixtures, and medical devices like pumps. The goal is to create a highly realistic simulation that can support imagination-enabled reasoning about the process steps the robot must perform. This simulation will be denoted as artificial world (AW).

Unlike traditional low-poly or low-fidelity simulators typically used in robotics and automation, which prioritize simplicity over detail, our simulation environment must support realistic rendering of objects. This is achievable due to advancements in the availability of high-quality CAD-models and modern simulation technologies that allow detailed material properties, geometries and behaviors to be replicated in the virtual environment. For example, the system must accurately simulate both the visual appearance and the physical behavior of objects, such as the movement of liquid in a container or the attachment of a cap to a canister.

The challenge here lies in the balance between realism and computational efficiency. Photorealistic environments and highly detailed object simulations are computationally expensive, especially when these simulations must run in real-time and interact dynamically with the robot's actions.

B. Actable Simulations: Dynamic and Real-Time Interaction

For the imagination process to function, the digital twin must be *actable*, meaning it can replicate the dynamics of the robot's actions in the RW. The robot arm movements, gripper interactions, and physical manipulations of objects must be mirrored in the AW in real-time. This includes processes such as picking and placing objects or performing more complex manipulations like attaching caps to containers. To achieve this, the simulation must support physics simulation and real-time capabilities to ensure the robot's imagined actions in the AW can be kept synchronous with its real-world operations.

One of the primary challenges in achieving real-time synchronization between the RW and the AW is maintaining the fidelity of physics-based interactions while ensuring the simulation remains responsive. For example, when a robot attempts to grasp an object in the RW, the same action must be instantly replicated in the AW, with accurate physics governing the interaction between the gripper and the object. As consumable items in the targeted use case are also regulary attached and removed from each other (e.g. blue cap and needle; red cap and canister), the simulation must support the handling of object compositions efficiently.

C. Perception System: Bridging the Gap Between RW and AW

An essential component of the system is the perception system, which connects the robot's vision in the RW with its imagistic reasoning counterpart in the AW. The robot's perception system, using standard computer vision techniques such as object detection and pose estimation, identifies objects in the RW and then synchronizes the detections of the RW with the AW. This ensures that the digital twin remains an accurate reflection of the real-world environment at any given time.

Once the perception system has established the current state of the RW, it can request renderings from virtual sensors placed in identical positions within the AW. These renderings allow the system to *compare the expected visual appearance* of objects in the AW with their observed appearance in the RW. This comparison is critical for determining whether the robot's actions have produced the intended outcomes.

For instance, if the robot attempts to place a red cap on a canister but fails to grasp the cap, the system may not detect this failure due to noisy tactile feedback. In the AW, the system simulates the attachment of the red cap as intended, resulting in the cap being placed on the canister. In contrast, in the RW, the same motion is executed, but the initial failure to grasp the cap means that no cap is actually attached. When comparing the outcomes in both environments after the action is performed, the canister in the RW will lack the red cap, while in the AW, the cap is expected to be present. To identify this discrepancy, the system should report that the expected outcome was not achieved, signaling a failure in the grasping and placement of the red cap.

D. Key System Components

To achieve the aforementioned functionality, the system integrates several key components:

- Robot Operating System(ROS) plays a critical role in managing the robot's operations, enabling the communication of its state to various process nodes, including the simulation environment. ROS provides real-time data exchange and standardized description formats, allowing the AW to replicate the robot's structure and movements accurately. Its broad compatibility with different robot arms and effectors makes it a flexible platform for developing a wide range of complex robotic applications.
- A high-fidelity simulation environment which is based on Unreal Engine, a feature-rich game engine known for its ability to produce photorealistic renderings and simulate complex physics. While Unreal Engine alone lacks specific capabilities for robotics simulation, the use of URoboSim[12] extends its functionality. URoboSim allows the import and simulation of robots using standard ROS formats, enabling seamless integration between the virtual and real worlds. It also provides the ability to simulate the dynamics of non-robotic elements, such as lab equipment, ensuring that the digital twin is a comprehensive replica of the RW environment.
- The final component, is a perception system that provides a strong basis for real world computer vision tasks but is also able to exploit the imagery generated by the virtual sensor in the AW. For this part, we use the RoboKudo perception system. It was designed to support imagistic reasoning fed by game engine belief states. We extended the system to support symbolic and geometric knowledge about objects which can be used to visually analyze subparts of the objects depending on the task-context. In the example above where the robot fails to attach a red cap to a canister, we can for example ask RoboKudo to focus task success verification only on the region in the image where the red cap is supposed to be put. By instructing the system where the attention shall be directed to, we were able to increase the robustness of the visual comparison for task success estimation as demonstrated in our evaluation.

While RoboKudo has been applied in mobile manipulation tasks previously, our work introduces several key ideas that elevate its applicability to the medical domain. The primary novelty lies in the integration of imagistic reasoning within a digital twin environment for task verification, especially in challenging scenarios with small, transparent objects. Unlike existing systems that rely on static object detection, our approach dynamically hypothesizes task outcomes and visually compares predicted and actual results in real time, leveraging detailed object knowledge and intuitive physics reasoning. This combination allows for improved detection of task success, addressing the limitations of earlier works(see Section V), which often fail to capture subtle, task-specific details or rely solely on predefined parameters. By incorporating parthood reasoning, our system surpasses standard object-centric methods, making it adaptable to complex, composite objects commonly found in medical labs.

IV. Imagination-enabled Perception process

After providing an high-level overview of the proposed system, this section details the practical implementation of our imagination-enabled medical robot perception system. By leveraging the RoboKudo perception framework and an Unreal Engine-based semantic digital twin, we integrate real-world actions with a high-fidelity simulation environment (i.e. AW) to enable autonomous verification of task success. Figure 3 illustrates the comparison process and the interplay of AW, RW and the perception system.

A. Artificial World

The first stage in implementing the system is the creation of the AW using Unreal Engine as the simulation platform. To achieve real-time synchronization with the robot operating in the RW, we employed the following components:

We begin by loading a URDF of the robot's kinematic structure into Unreal Engine, enabled by URoboSim. URoboSim interfaces with ROS to replicate the robot's joint movements in real-time. By reading the joint states from ROS topics, the robot's exact movements in the RW are mirrored in the AW, ensuring that manipulation tasks are synchronized between both environments. This ensures that every action performed by the robot in the RW has a corresponding simulation in the AW, allowing for precise hypothesis generation and outcome comparison.

Relevant objects used in the robot's tasks are integrated into the simulation as CAD models. Simple objects are loaded as standard meshes, while more complex objects, such as a needle with a detachable needle cap, are implemented using Unreal Engine Blueprints. This approach enables the AW to support dynamic object interactions during the task execution. Additionally, these object models are enhanced with physical properties, ensuring that simulated interactions, such as picking, placing, or assembling components, reflect their real-world counterparts accurately.

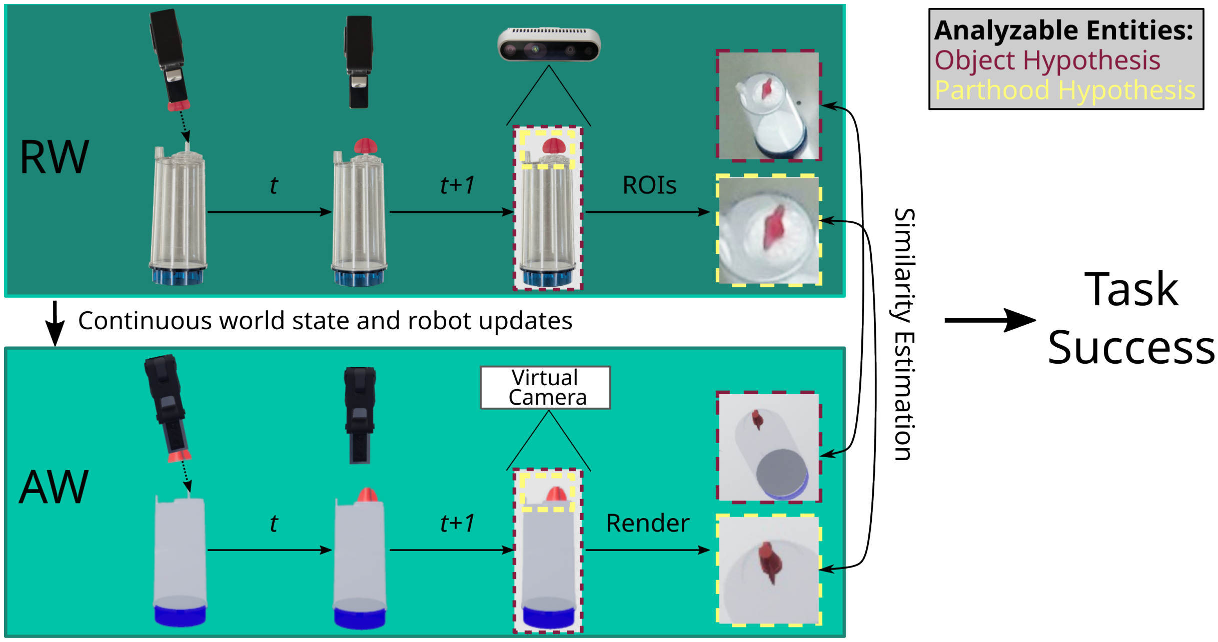

Fig. 3. Illustration of the imagination-enabled task success verification. Operations performed in the Real World (RW) are continuously replicated into a simulation environment called Artificial World(AW). It acts as a digital twin which allows the robot to create a visual expectation of the current world state. Actions of the robot as well as perception beliefs (e.g. object detections) are synchronized during task exection to maintain a world model which can be rendered by a virtual sensor. These renderings as well as the images from the RW camera can be compared for visual similarity before or after manipulation actions. The system can also guide the attention if necessary, for example when expected changes in the image (e.g. „Is the red cap on the canister?") only effect a small region in the image space of the affected object. Ultimately, the visual comparison between the RW camera images and the rendered expectation can be used to verify the success of the intended action.

By combining real-time synchronization of robot movements with detailed object modeling, the AW provides a fully actable simulation that allows the robot to generate hypotheses about the expected outcomes of its actions.

B. Perception system integration

The next critical step in the system implementation is integrating the perception pipeline with RoboKudo. It realizes the interpretation of real-world data and its synchronization with the AW. RoboKudo is responsible for two key functions:

1) Processing Real-World Perception: RoboKudo accepts task-specific perception queries, such as identifying canisters or verifying the attachment of a cap. It processes real-world camera images through a set of Annotators—modular experts within the framework. These Annotators generate ObjectHypotheses, which identify potential objects in the scene, and additional object-related information (called Annotations) such as shape, color, and class. This perception data is continuously updated in the system's belief state, forming a real-time understanding of the objects present and their attributes. By tracking changes in object states across frames, RoboKudo ensures that the RW representation remains consistent with the robot's ongoing actions.

2) Synchronizing with the AW: The perception belief state generated in the RW is used to keep the AW representation up-to-date. Whenever new objects are detected or object properties change (e.g., position, orientation), RoboKudo sends the necessary updates to the AW. Following these updates, the system requests a rendering of the AW to generate an image of the hypothesized outcome of the robot's actions. This image serves as the expected result of the task, which will be compared against RW observations. Internally, RoboKudo represents objects in the AW similar to the ones in the RW, making it possible to apply the same Annotators to both types of imagery to get semantic descriptions.

C. Generalized Visual Analysis Using Analyzables

During the system's early testing, we identified limitations in visual comparisons between the RW and AW for tasks involving small-scale visual changes. For example, attaching a small red cap to a canister often led to minimal changes in pixel space, making it difficult to verify task success through standard object-level analysis. To overcome this, we extended the RoboKudo framework with the concept of *Analyzables*.

Analyzables generalize the system's analysis capabilities by allowing it to focus on both whole objects and their subparts. Using the Parthood concept from the SOMA Ontology [3], the system can now use background knowledge to break down objects into smaller, meaningful components or features. For instance, in the red cap attachment scenario, the system isolates the region of interest (ROI) where the red cap is expected to be placed on the canister. This provides a more focused basis for visual analysis, enabling more accurate verification of task success.

RoboKudo's Annotators were extended to operate on the new Analyzable superclass, which comprises *ObjectHypotheses* (representing entire objects) and *ParthoodHypotheses* (representing parts of objects). For example, an Annotator tasked with verifying the placement of a red cap now only analyzes the specific region where the cap should be, rather than attempting a global image comparison. This more granular approach simplifies the task of verifying subtle changes in the scene, such as detecting whether the red cap is properly attached to the canister.

D. Visual Similarity and Task Success Verification

The final step in the system's operation is the actual task verification process, which involves comparing the rendered image from the AW with the RW camera image. These are the key steps in this process (as illustrated in Fig 3):

1) Rendering the expected outcome: After updating the AW to reflect the robot's manipulation actions, a simulated image of the expected post-task environment is generated. This image is based on the updated state of the objects and robot in the AW.

2) ROI extraction: The system extracts ROIs from both the RW and AW images, focusing on areas relevant to the task, such as the area where the red cap is expected to be attached. This targeted comparison allows the system to ignore irrelevant parts of the image and concentrate on verifying the specific elements affected by the robot's action.

3) Verifying task success: The ROIs from the RW and AW are then compared using both pixel-wise and semantic analysis. If the expected outcome, such as the red cap being correctly placed, matches the observed result in the RW, the system concludes that the task was successful. If discrepancies are detected, the system identifies the task as incomplete or erroneous and triggers corrective actions.

V. Experiments

Existing lab automation systems do not integrate cognitive reasoning or imagistic perception. Our approach uniquely enables robots to conduct imagistic reasoning about task outcomes to verify task success. In our experiments, we studied different approaches for visual task success estimation. We were investigating which methods can deal best with the challenges that lie in the simulation of transparent objects such as the canister, where accurate renderings are the hardest to achieve. In contrast to the work in [10] where only big, textured objects have been used, we will provide results towards the imagination-based comparison of transparent and small objects.

In our study, we investigated the performance of three categories of methods (abbreviations for the methods as shown in Fig.4 are put in parentheses): 1) Neural-Network-based Appearance Similarity based on several state of the art networks: ConvNeXt(CNX)[8], Resnet50(Res)[6], DINO-ViT(ViT)[5]. 2) A color histogram similarity estimation (His). 3) Our semantic analysis method called SemanticComparisonEquality(SCE) which compares Annotations provided by the RoboKudo system for the targeted subparts of the comparison process. The results are further divided into comparison of the ROIs showing the full object (F) and only the targeted subpart (PH) based on ParthoodHypotheses. All methods have been implemented as individual Annotators in the RoboKudo framework.

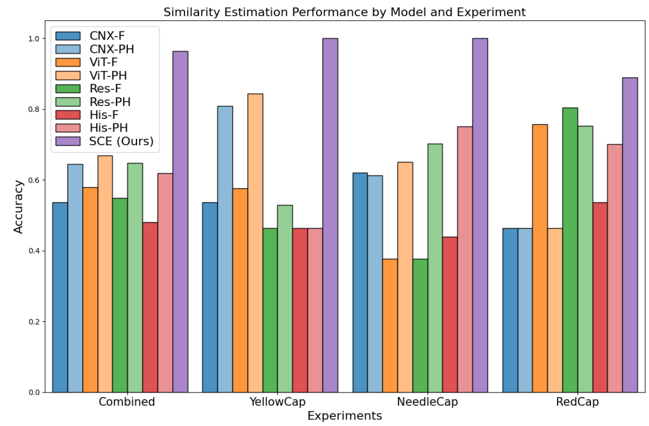

Fig. 4. Results of our similarity estimation experiments which verify the visual presence of indicators for task success. We have conducted two experiments on Canister-related experiments for the succcessful/unsuccessful attachment of the red cap and yellow cap as well as successful/unsuccesful removal of the blue needle cap on the needle (ref. to Figure 1 for a picture of these objects).

The experiment was done on three typical manipulation actions for the targeted sterility testing process: Attaching a red cap to the canister on the top, attaching a yellow cap to the canister bottom and removing the blue needle cap from the needle. 1713 images with varying conditions have been analyzed. We can observe in the combined results (see Fig.4) that on average, the NN-based similarity estimation and SCE method performed better than the histogram-based approach. When further dividing the results, we can observe that the addition of the ParthoodHypothesis as discussed in Section IV led to an improved performance in the *-PH* Results because of the focus on the relevant regions. The best performing method was SCE, which compares the semantic annotations provided by RoboKudo on the ParthoodHypotheses of AW and RW. It was able to generalize well between the different use cases by benefiting from the provided knowledge about the objects and the perception task.

VI. Conclusion

This paper presents a novel framework for imagination-enabled medical robotic perception, enhancing task verification in laboratories. By integrating cognitive AI with a semantic digital twin, the system autonomously simulates and verifies tasks, effectively handling the challenges posed by transparent and flexible objects. Experiments demonstrated that focusing on object subparts significantly improves verification accuracy, particularly in complex lab environments.

The system is designed to be modular, scalable, and adaptable to various applications beyond sterility testing. Through the RoboKudo framework, new tasks and object classes can be incorporated by updating models and task-specific annotators, allowing adaptation to different laboratory processes. Future work will focus on extending the system to more complex tasks like liquid handling and drug preparation. Its scalability ensures it can be deployed across a wide range of robotic arms, lab equipment, and environments, making it a flexible solution for diverse laboratory automation needs.

REFERENCES

[1] "Automation Concepts for Life Sciences". In: *Automation Solutions for Analytical Measurements*. John Wiley & Sons, Ltd, 2017. Chap. 3, pp. 73–91. ISBN: 9783527805297. DOI: https://doi.org/10.1002/9783527805297.ch3.

[2] Laurent Bessemans et al. "Automated gravimetric calibration to optimize the accuracy and precision of TECAN freedom EVO liquid handler". In: *Journal of Laboratory Automation* 21.5 (2016), pp. 693–705.

[3] Daniel Beßler et al. "Foundations of the Socio-Physical Model of Activities (SOMA) for Autonomous Robotic Agents 1". In: *Formal ontology in information systems*. IOS Press. 2021, pp. 159–174.

[4] Benjamin Burger et al. "A mobile robotic chemist". In: *Nature* 583.7815 (2020), pp. 237–241.

[5] Mathilde Caron et al. "Emerging properties in self-supervised vision transformers". In: *Proceedings of the IEEE/CVF international conference on computer vision*. 2021, pp. 9650–9660.

[6] Kaiming He et al. "Deep residual learning for image recognition". In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2016, pp. 770–778.

[7] R Lehmann et al. "Biomek cell workstation: a variable system for automated cell cultivation". In: *Journal of laboratory automation* 21.3 (2016), pp. 439–450.

[8] Zhuang Liu et al. "A convnet for the 2020s". In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2022, pp. 11976–11986.

[9] Patrick Mania et al. "An Open and Flexible Robot Perception Framework for Mobile Manipulation Tasks". In: *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2024, pp. 17445–17451.

[10] Patrick Mania et al. "Imagination-enabled Robot Perception". In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2021.

[11] Antonella Mencacci et al. "Laboratory automation, informatics, and artificial intelligence: current and future perspectives in clinical microbiology". In: *Frontiers in Cellular and Infection Microbiology* 13 (2023), p. 1188684.

[12] Michael Neumann, Andrei Haidu, and Michael Beetz. "URoboSim — a simulation-based predictive modelling engine forcognition-enabled robot manipulation". In: *2nd Embodied AI Workshop (CVPR)*. 2021. URL: https://embodied-ai.org/papers/URoboSim.pdf.

[13] Anthony Remazeilles et al. "Robotizing the sterility testing process: scientific challenges for bringing agile robots into the laboratory". In: *Iberian Robotics conference*. Springer. 2022, pp. 223–234.

[14] A Wolf and K Széll. "A review on robotics in life science automation". In: *Proceedings of the AIS 2019 14th international symposium on applied informatics and related areas organized in the frame of Hungarian science festival*. 2019, pp. 106–111.

[15] Ádám Wolf et al. "Towards robotic laboratory automation Plug & Play: The "LAPP" framework". In: *SLAS technology* 27.1 (2022), pp. 18–25.

[16] Martina Zaninotto and Mario Plebani. "The "hospital central laboratory": automation, integration and clinical usefulness". In: *Clinical chemistry and laboratory medicine* 48.7 (2010), pp. 911–917.

[17] Yixin Zhu et al. "Dark, Beyond Deep: A Paradigm Shift to Cognitive AI with Humanlike Common Sense". In: *Engineering* 6.3 (2020), pp. 310–345.