Automated Testing of Interaction-Requiring Devices with a Robotic Arm and Object Detection

Levente Daroczi[1], Beniamin Demeter[1], Csaba Sulyok[1], Gyöngyi Katona[2], Dénes Mihály[2] and Arnold Szász[2]

[1] Department of Computer Science, Faculty of Mathematics and Computer Science, Babeș-Bolyai University, Strada Ploiești 22-25, Cluj-Napoca 400157, Romania, levente.daroczi@stud.ubbcluj.ro; beniamin.czompa@stud.ubbcluj.ro; csaba.sulyok@ubbcluj.ro

[2] Codespring, Strada Constantin Brâncuși 69-71, Cluj-Napoca 400347, Romania katona.gyongyi@codespring.ro; mihaly.denes@codespring.ro; szasz.arnold@codespring.ro

Abstract: Testing user interfaces is fundamental for ensuring reliability and proper functionality. However, for many software systems, there is no framework available that would support automated testing, especially when the system requires manual input controls such as physical buttons. Currently, testing of these devices is largely done manually, which is extremely time-consuming, requires a high degree of attention, and can increase the potential for errors. The project aims to create an automated testing system that uses a robotic arm and object detection technology. The system is capable of testing devices with physical buttons and mechanical controls, ensuring precise and repeatable interactions, as well as validating device feedback. This reduces the need for manual intervention while significantly increasing the efficiency of the testing process. The solution consists of a web application and a backend server. The web application provides the user interface that allows for creating tests and viewing results, while the backend server is responsible for executing the tests, controlling the robotic arm, and performing object detection processes.

Keywords: object detection; robotic arm; automation; testing; AI

1 Introduction

In the domain of software development, integration testing of a finished product is at least as important as the separate verification of individual components. While the latter ensures the correct operation of the building blocks, the former guarantees their proper cooperation. [1] Testing, however, becomes significantly more complex when the software is not running on a general-purpose computer but is

instead designed for dedicated hardware. [2] In such cases, standard testing infrastructure is often lacking, especially for microcontroller-based devices. [3] Without such support, it is difficult or even impossible to issue automatic test instructions or to retrieve their results from the target hardware. The challenge is further increased if the device is equipped with physical input controls that require human interaction, such as push buttons, touchscreens, or rotary switches. Testing these using conventional tools can be particularly cumbersome, especially once the product has completed its final manufacturing phase. [4, 5]

The current paper proposes a solution for automated testing of such unconventional systems, relying on the following tools: a robotic arm, a camera, and advanced object detection software. The arm simulates human interaction by handling physical inputs such as button presses, touchscreen navigation, or turning rotary switches. The recognition unit assists in controlling the robotic arm and enables result verification. Its primary task is to identify the components of the device undertest through the camera feed.

Unlike conventional testing methods, the use of a robotic arm is essential for devices equipped with physical input interfaces that require human interaction. While various testing frameworks exist--such as *Appium*[1], which communicates with Android devices using the *Android Debug Bridge (ADB)*[2] for Android devices or software specialized in UI/UX testing-these solutions are typically limited to a specific platform or input method. There is a lack of a universally applicable system capable of handling devices with diverse input interfaces, such as touchscreens, physical buttons, or rotary switches.

With an accurate description of components of the test device and deep learning-based recognition models, complete multi-step tests can be executed without human intervention. Introducing an automated system offers several advantages over traditional manual testing. It eliminates human errors and significantly frees up time for test engineers. [5] The system can independently perform testing, so the user only needs to review the results and the camera footage documenting the test environment's state at the end of execution.

Currently, there are few solutions on the market capable of software testing for devices requiring interaction by combining robotic arm and object detection technology. Among the best-known similar systems are *MATT*[3], developed by *Adapta Robotics*, and *QUACO Pro*[4] from the portfolio of *Sastra Robotics*. These solutions use proprietary devices that have been specifically designed for this purpose. As a result, in certain cases, they may be more efficient than systems based on general-purpose robotic arms.

1 Source: https://appium.io/docs/en/latest/
2 Source: https://developer.android.com/tools/adb/
3 Source: https://www.adaptarobotics.com/matt/
4 Source: https://sastrarobotics.com/products/quaco-pro/

2 Interface Element Recognition in Images

The current project aims to allow navigation on/interaction with any user interface, including touchscreens, physical buttons, knobs and other actuators. The current section focuses on the technology choices enabling this.

2.1 Theoretical Background

Object detection algorithms are capable of recognizing different shapes, icons, and even shorter text. They can detect specific visual patterns and structures, which help identify buttons, surfaces, and other elements [6, 4]. The detection employed in this project is based on deep learning, and enables the identification and localization of various objects in images or videos. The goal of their usage is to detect UI widgets on screens, read label texts and identify interactive elements [7, 8].

Among the various technologies, the YOLO (You Only Look Once) algorithm family is chosen [9, 10], as it is capable of identifying multiple objects in a single image. This is particularly important on a user interface, where real-time and accurate detection plays a key role. YOLO is a single-stage detector that has undergone numerous developments over the years, continuously integrating the latest methods into its architecture. The current project uses YOLOv8 as it was the newest and most advanced version at the time of development. [11]

YOLO recognizes objects and their positions with "a single glance", that is, with a single image processing, hence the name *You Only Look Once*. In the background, YOLO uses a CNN (convolutional neural network) that helps predict the bounding box of different objects in an image and their associated probabilities. Convolutional neural networks are very effective in processing visual data, as features can efficiently pass from initial convolutional layers to later ones. [12]

2.2 Labeling and Training Process

The first step in model training and object detection is creating a well-structured and properly labelled dataset. In order for the models to accurately and reliably recognize the desired objects under all conditions, the dataset is created in a varied environment. Images are captured under different light conditions, continuously changing light sources, and the device to be tested is photographed from multiple angles and in different positions. Additionally, the camera position is regularly modified so that the model would be able to handle perspective differences (see Figure 1).

In the case of interaction elements of the mobile application (*FestivApp*), some similarities could be noticed, so the decision was made to group all buttons according to these similarities. From the main menu page, buttons were grouped into separate models for navigation buttons, tabs, and various action buttons, such as adding to favourites.

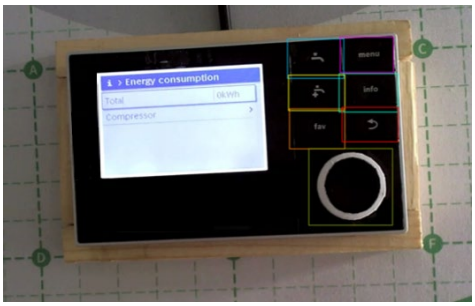

Figure 1
Buttons of a controller used for heating systems with labels

After the training process, the result contains images and files showing various statistics that help evaluate the performance of the model and the best-performing model.

2.3 Recognition Process

The project contains a dedicated recognition layer that is responsible for identifying different objects based on given parameters. In order for this layer not to be specific to a particular test device, it has been designed so that no modifications to this layer are necessary when introducing additional devices.

3 Robotic Arm

3.1 Dobot Magician Lite

The application uses the *Magician Lite* robotic arm developed by *Dobot Robotics* for controlling the test device (see Figure 2). This is a general-purpose robotic arm designed for educational purposes, which can be controlled via hardware, software, or Python programs. [13]

The robotic arm is capable of executing instructions with a repeatability of *0.2* mm, providing more than adequate precision for the use case of this project. [13] Multiple types of end effectors can be attached to the end of the arm, such as a pen holder, a soft gripper, or a rotatable suction cup capable of vacuum-based gripping of small objects (see Table 1).

From a hardware perspective, out of the three available end effectors, only two were actively used, as the gripper unit did not prove useful in this application environment. The pen holder unit, combined with a touch-sensitive stylus, is suitable for controlling touchscreen devices such as a smartphone.

Table 1
DOBOT Magician Lite end effectors [13]

End Effectors	
Pen holder	Pen diameter: 8-12 mm
Suction cup	Built-in air pump drive, operates under negative pressure, with suction cup diameters of mm and 20 mm
Soft gripper	Built-in air pump drive, operates under both positive and negative pressure, maximum opening and closing distance: 50 mm

In contrast, the suction cup unit can be effectively used to operate devices equipped with push buttons or rotary knobs.

The design of the system enables the automatic testing of devices with varying types, sizes, and input interfaces. However, integration into the application requires the following conditions to be met: (1) the surfaces of the device intended for control or observation must be located within the working area of the arm; (2) these components must be placed on the upper side of the device; (3) the input interfaces must be physically accessible and operable by the arm. Supported components include touchscreens, LCD or other displays, push buttons, rotary knobs, and other physical control elements.

The robotic arm comes with a special mat that determines the position of the arm base and defines its workspace. While these markings are not strictly necessary for the operation of the robotic arm, the six reference points (*A–F*) on the mat play a key role in the system. On one hand, they visually designate the workspace defined by the application, but their most important function is to support calibration.

To align the camera image with the other device coordinate system, at least four points with known coordinates are needed both in the camera image and within the workspace. The reference points placed on the mat serve precisely this purpose: with human assistance, the arm can read their spatial coordinates, while they also appear as recognizable objects in the camera image. This ensures that visual perception and arm movement are precisely synchronized. (see Section 3.3).

3.2 Camera

The camera is an essential component of the application, as it enables the recognition of test device components and plays a crucial role during the calibration of the robotic arm. During a test run, whenever component detection is required, the control unit retrieves the current frame from the camera and runs the YOLO model on it. For calibration, the reference points on the mat are also identified using a pre-trained YOLOv8 model executed on the camera image.

Figure 2
The test surface as recorded from above by the stationary camera

The primary goal in programming the device is to enable the simplest possible usage during test execution. For this purpose, an abstraction layer is used to encapsulate the required methods.

3.3 Calibration

Precise calibration of the arm is essential for running tests. Early tests demonstrate that the pixel coordinates of components detected by object detection algorithms and the arm's own coordinate system are not compatible. This is resolved by applying a projective transformation, which enables the mapping between the two planes if at least 4 reference points are precisely known for both. [14] The more available points, the more accurate the fit. [15] The coordinates of the six reference points on the special mat can be retrieved from the camera image via the camera module, using the previously trained YOLOv8 model. In the coordinate system of the robotic arm, these must be registered manually, which requires human intervention. To ensure easy development and modularity, calibration is outlined and executed using a behavior tree (see Figure 3). [16]

Device Selection is an atomic element that waits for the user's input. Once the camera is selected, it initiates camera initialization. In case of error, this information propagates to the top level, and calibration becomes invalid. This reaction is similar in all cases, except for components capable of error handling.

Reference Point Detection is a selector-type component that also executes its descendants from left to right but stops at the first successful execution and passes the result up. According to the figure, detection is first attempted with object detection, and if unsuccessful, manual adjustment is required. If object detection is successful, no further intervention is needed.

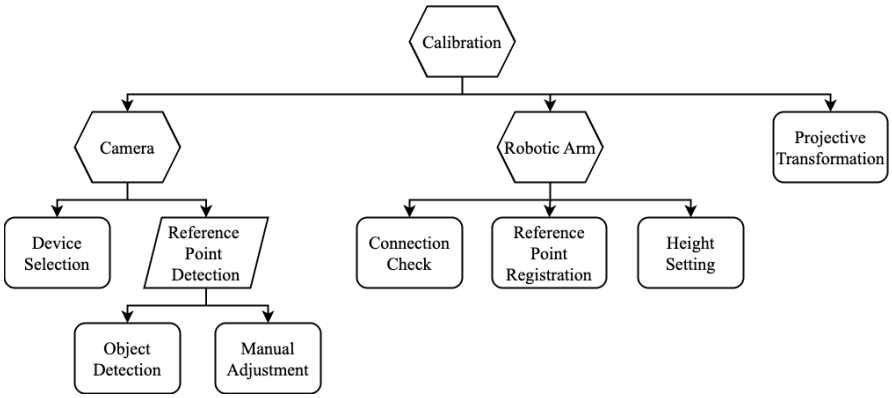

Figure 3
Behavior tree of the calibration process

During execution of the atomic *Object Detection* element, the previously mentioned YOLOv8 model for reference point detection is run. To reduce errors, the model is retried up to five times in case of unsuccessful detection. If no satisfactory result is achieved after this, the component signals an error. The behavior tree library allows information sharing via a key-value storage system called the b*lackboard*. Through this, behavior tree components can access the registered coordinates of the reference points.

The atomic *Manual Adjustment* element serves as a fallback to ensure that the robotic arm can always be calibrated. The user must adjust the camera so that the reference points projected onto the camera image coincide with the real points. In other words, the user must find a predefined camera angle and distance to ensure valid coordinates.

After the *Camera* component completes successfully, the *Robotic Arm* component is executed, which is also sequential. Its first step is running the atomic *Connection Check* element, which is considered successful if the program communicates properly with the robotic arm. If the robotic arm is not connected or is used by another program, an error is signaled.

Reference Point Registration is an atomic element requiring manual intervention, during which the real coordinates of the reference points are recorded with the user's help. This is a simple process in which the end effector of the robotic arm must be placed on all six points from *A* to *F*. The height relative to the mat does not matter, as only the *X* and *Y* coordinates are relevant. Once the robotic arm is positioned on a reference point, the system queries and records the current coordinates of the end effector. At the end of the process, the scanned data are also stored in the shared *blackboard.*

The atomic *Height Setting* element is similar to the previous one, except here only the height of the end effector matters. The user must set this value by adjusting the

end effector, and the system records it. This is the default height level the robotic arm will use during the operations unless overridden. The set height value is communicated to the robotic arm. For both the *Reference Point Registration* and *Height Setting* components, error checking is performed to filter out values outside the workspace.

After successful execution of the *Camera* and *Robotic Arm* components, all conditions are met to run the *Projective Transformation* element. In this step, the mapping between the two planes is established using the method described above, and a transformation function is built that can convert coordinates between the two systems. Although at least four reference points are required for mapping, this does not guarantee that the transformation exists. The underlying system of equations may have no solution, for example, if the reference points are collinear. [15] In such cases, calibration must be repeated.

On the user interface, the calibration process includes an additional step for selecting the test device. However, this only sends feedback to the server and is not part of the behavior tree, as it is always considered a successful operation.

4 Tests

Testing plays an important role in the development of any software or hardware system, as it ensures that applications and devices function as expected. The goal of testing is to identify potential errors, thereby ensuring the error-free operation of the application. Software testing can rely on well-established methods and automated tools/frameworks, while hardware testing is based on a much more complex process. [17]

Software tests typically run in a virtual environment where all parameters can be precisely controlled. In contrast, hardware testing must account for physical factors such as mechanical elements of devices, which can present significant challenges. For all these tests, a structure had to be devised that can support testing for touchscreen or other hardware devices as well.

4.1 Graph of the Operation of Applications and Hardware Devices

Some interactive elements (such as buttons, scroll wheels, touchscreen surfaces) allow the user to move from one state of the system to another. These states can be considered different pages or user interfaces, between which the user can navigate through specific actions. For example, pressing the *Back* button returns the user from the current page to a previous state, while a *Like* button does not change the navigational state but performs a modification on the current page (for example, by updating the state of an element).

After studies, the idea emerged that the structure enabling navigation can be effectively represented in the form of a graph, regardless of whether a given action actually results in a state change or merely performs a modification within the current state.

This approach allows for the structured mapping of the operation of applications and hardware devices, and supports the definition and execution of automated tests. The states (pages) can be interpreted as the nodes of the graph, while the interactive elements that perform various actions – such as navigation, button presses, or settings adjustments – form the edges of the graph. With this method, it is possible to visually and logically model the options available to the user within a given system, and how they can interact with individual elements (see the navigation graph of the controller used for heating system control, Figure 4).

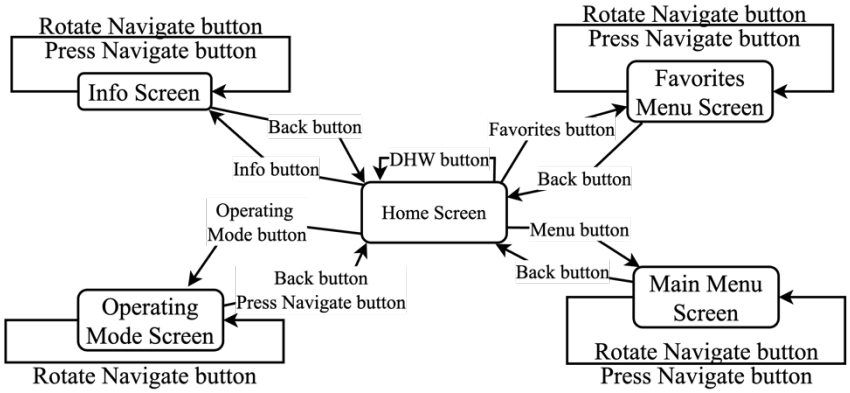

Figure 4

Navigation graph of the controller used for managing heating systems

The graph-based representation not only enables the structured mapping of the operation of applications and hardware devices, but also contributes to the efficient design and execution of automated tests. Since this type of representation allows for the precise definition of testing paths, potential errors can be identified more easily, and it can be ensured that every interaction is covered during the testing process.

From an implementation perspective, the navigation structure must be represented for every application, so a solution had to be found that is easily modifiable, flexible, and does not impose limitations when adding extensions, new features, or restructuring applications.

To store the data in memory, structured JSON files are used. In this format, nodes (pages) and edges (interaction elements) can be easily defined, which helps to accurately and comprehensibly represent the navigation tree of the application or system. This ensures that the system remains easily modifiable and expandable, while the data remain well-structured and simple to manage.

4.2 Test Structure

The tests that the user can create consist of steps. Each step represents edges in the graph, that is, operations that test the functionality of the given test device. For example, checking the correct operation of a button or examining the response of an interaction element may belong here. The sum of the individual steps constitutes the complete test. Since the steps of the tests are determined by the graph structure, they automatically fit into the navigation process of the system and provide an opportunity to test functionality.

The parameters stored in the JSON objects of the navigation graph contain not only the data that the user needs to know directly, but also other background information that is essential for the operation of the system but not necessary for the user.

Table 2
Description of the edge object fields

Field	Description
id	The unique identifier of the edge, which allows for unambiguous identification.
componentName	The name of the interactive component found on the application interface.
from	The node (page) from which the navigation starts.
to	The node (page) to which the navigation occurs.
static	Boolean value: for static (true) edge there is no navigation, while for dynamic (false) there is.
deltaHeight	A numeric value that modifies the default Y-coordinate (vertical position) of the touch event. Useful when the component must be touched slightly above or below its default center.
push	A boolean value indicating the nature of the interaction: true means the action requires a firm press (e.g., long press or physical button), while false means a simple touch (tap) is sufficient.
action	The name of the interaction operation.
direction	Specifies the direction of the gesture or interaction (e.g., "*up*", "*down*", "*left*", "*right*")—commonly used for scroll or rotate actions.
hold	The duration (in seconds) for which the touch interaction should be held. This is relevant for gestures like long press or press-and-hold operations.

There are two main object lists in the *JSON* file: *nodes* (nodes) and *edges* (edges). Within the *nodes* object list, other objects are stored that have an *id* and a *name* field, where the *id* represents a unique identifier, and the *name* is the name of the page within an application/interface. Within the *edges* object list, the edges are stored (see edge object fields in Table 2).

During test device integration, parameter values for each edge were determined by experimenting with different inputs and interpolating the optimal results.

4.3 Test Execution

The tests and their correctness conditions are received by the backend server in *JSON* format, which then checks and processes them.

Each test step has a behavior tree assigned to it, which is responsible for recognizing objects in the image provided by the camera and executing the desired operation (see Figure 5).

The use of a behavior tree is advantageous because if any part of a step's execution fails (for example, an object is not recognized or an interaction cannot be performed), then the entire step execution is considered unsuccessful. Consequently, the test cannot continue, as the step was not executed correctly. The system clearly indicates if any step did not execute as intended. Furthermore, if a step's behavior tree fails, the robot stops executing the remainder of the test. The user is then notified of how many steps were successfully executed before the failure occurred. At this stage, the system is not capable of recovering from the failure or returning to the initial state, and manual intervention may be required to reset or diagnose the issue.

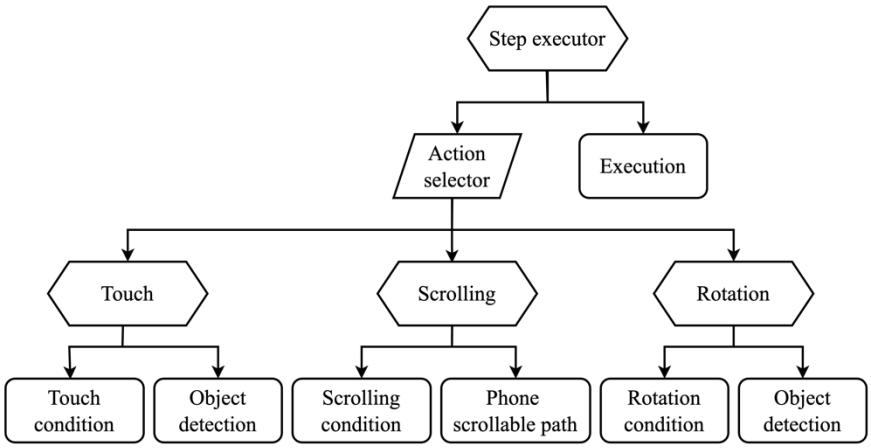

Figure 5
Behavior tree for the execution of steps

The *Step executor* is a sequential component that runs all of its direct subordinate elements, the *Action selector* and the *Execution*. If any of these components do not execute correctly, the entire step is considered unsuccessful.

The *Action selector* is a selector component that stops after running the first successful direct subordinate element. This allows the system to decide which operation the given step belongs to.

The *Action selector* component has three subordinate elements: *Touch*, *Scrolling*, and *Rotation*. All three are sequential components, from which the system selects

which operation to execute based on the supported operations, and also sets the necessary parameters.

During the execution of these operations, the *Touch*, *Scrolling*, and *Rotation condition* components can return success or error states, which decide which operation should continue.

For the *Touch* and *Rotation* operations, an interface component must be recognized on the test device. The *Object detection* node is responsible for this task, which sets the appropriate coordinates of the component using the so-called *blackboard* keys. These coordinates can be read by later nodes and used to execute the given operation.

In the case of *Scrolling*, it is not necessary to recognize a component, but rather to determine the position of the device, the so-called *Phone scrollable path*. This requires two coordinates that describe the path where the scrolling operation should be performed. Using computer vision, the frame of the phone screen can be outlined, and using distance analysis, the section where the user typically scrolls can be determined. To solve this task, the *fitLine* function of *OpenCV* is applied. [18] In this case as well, the two coordinates needed to determine the section are stored using the *blackboard* keys.

After the *Action selector* has successfully completed the tasks necessary for the operation, the *Execution* node reads from the *blackboard* the scrolling coordinates (if they exist), and the component centre point coordinates, which were set by the *Object detection* nodes. For each operation, the appropriate functions belonging to the robot module are called.

4.4 Test Validation

To check the correctness of tests, predefined conditions are needed that ensure the success of the test can be evaluated. These conditions allow the user to decide on a test result without reviewing the testing process on video or personally monitoring the execution.

The system provides the opportunity to create customized correctness conditions. The user can select those interface elements that must be present after the test execution, as well as those that should not appear. If the user does not define custom conditions, the default criterion for success is that at least 60% of the expected interaction elements should be recognizable after execution.

For other testing possibilities, *OCR* (optical character recognition) algorithms are applied, which are capable of extracting textual information based on images. [19] The user can specify an expected text for each step that should appear on the screen of the device being tested after the step is executed. This provides an opportunity for the system to automatically check whether the device has actually entered the desired state. The best performing solution is *EasyOCR*, an open-source *Python* library that supports character recognition in multiple languages.

To increase the accuracy of text recognition, the images go through multi-step preprocessing before the *OCR* algorithm runs. The input image is first converted to greyscale, then the contrast is increased using the *CLAHE (Contrast Limited Adaptive Histogram Equalization)* method to better distinguish the text from the background. After that, a sharpening filter is applied, which further highlights the characters, and finally, the image is enlarged so that the *OCR* can work from a higher resolution.

5 Environmental Requirements

To ensure that the test runs properly and object detection works accurately, certain environmental conditions must be met. For optimal recognition performance, it is important that the test device is evenly illuminated from all directions, minimizing shadows and reflections. There must be no direct light source above the test device, as reflected light can hinder object recognition. The screen brightness should also not be too high or too low; the optimal range is between 45% and 80%. The device should be positioned facing away from the robot arm, preferably in the center of the workspace, ensuring the camera has a clear and unobstructed view.

During testing, it was observed that if a light source is placed directly above the test device and causes strong reflections, it can significantly reduce the effectiveness of object detection. In such cases, the recognition of interaction elements may become inaccurate or completely fail. Therefore, it is important that light sources are positioned in a way that prevents disruptive glare on the surface of the test device. In addition, the focus of the camera and exposure settings also influence recognition performance. For example, if a mobile phone screen has a very high brightness in a poorly lit environment, the camera may struggle to focus properly, resulting in a blurred or overexposed image. In such cases, the objects may be difficult to identify even with the human eye, which also complicates automatic analysis.

6 Architecture

The application consists of two distinct layers: a processing layer and a presentation layer. The processing layer contains all functionalities related to object detection and device control, while the presentation layer communicates via *HTTP* requests to display relevant features to the user (see Figure 6). During the execution of the processing layer, it is crucial to consider hardware resources, as this layer requires greater computational capacity due to model execution and multithreading. Thanks to the clear separation of layers, the presentation layer does not impose significant hardware requirements and can be easily deployed elsewhere.

Communication between the layers is handled via a *REST API*, with the exception of video streaming, which uses *sockets* for increased speed.

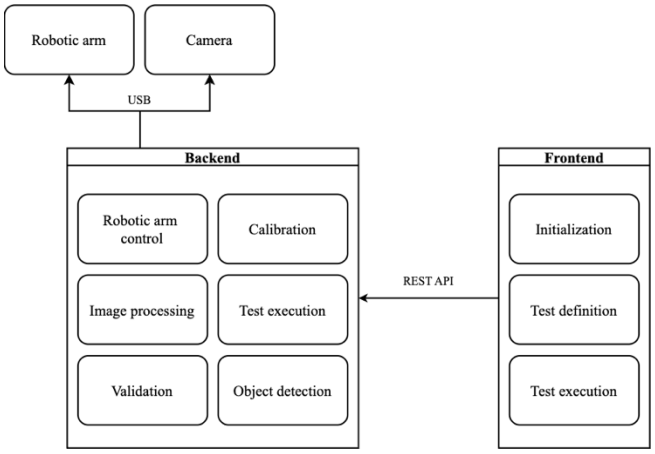

Figure 6
System architecture

The processing layer (backend) is implemented in Python, as the *Dobot Magician Lite* must be programmed in this language, and state-of-the-art AI tools such as YOLOv8 are also available in Python. This layer also provides an *API* for the presentation layer, functioning as a backend server.

The presentation layer is a web interface served by a *Flask* server because it offers all needed functionalities without unnecessarily burdening application performance. The robotic arm and camera are connected to the server exclusively via USB cable, which significantly complicates remote deployment.

The backend server provides numerous functionalities for the presentation layer, which are offered through an API following *REST* conventions, a logical and well-organized framework for endpoints. All endpoints use the convenient and natural JSON format for payload transmission. [20]

Persistent data storage is handled by a *PostgreSQL* database, chosen for its efficient handling of *JSON* data and excellent compatibility with Python.

The database includes tables such as *Tests*, *Configs* (loadable calibrations), and *Subjects* (test devices). The server can connect to a local or remote database, the location of which is configurable via environment variables.

The web application is written in *JavaScript* using the *React* library, resulting in an intuitive and clean interface. Since the application is not intended solely for experts, this is taken into account during interface design.

A designer is involved in the interface development, resulting in a design prototype created in the *Figma* design program, which serves as the basis for the website.

The required *React* components are built by customizing elements provided by the *Bootstrap* library.

All functionalities of the web application are accessible exclusively to authenticated users. User management is provided by a private *Keycloak* server, made available for the project by *Codespring*.

Conclusions

The system presented in the thesis provides the user with the ability to test interaction-requiring devices using a robotic arm based on object recognition. The testing process is accessible through a web application, which ensures the use of the application's functions. After calibrating the robotic arm, the user is able to perform a multi-step test. The system offers unique verification conditions for each test and step, which automatically evaluate the success of the test.

To support the evaluation of the system's practical effectiveness, a series of test executions was conducted under controlled conditions. A total of 100 test executions were performed, with the goal of evaluating the system's practical effectiveness. The test case consisted of 11 steps, including rotation and press actions, aimed at increasing the temperature on the controller used in heating systems. Out of the 100 tests, 88 were fully successful, while 11 resulted in application errors due to hardware-related limitations inherent to the educational robot, and one test failed. Additionally, 5 tests were identified as false positives, where the robot executed all steps, but one or more actions did not perform as expected, yet the test was still marked as successful. The execution time across all successful tests remained consistent, ranging between 1 minute and 14 seconds to 1 minute and 16 seconds. These initial results suggest a functional baseline but highlight the need for further development cycles focused on quantitative performance metrics. Incorporating measures such as detection accuracy, detailed error analysis, and step-level execution timing would strengthen the assessment of reliability and enable more targeted improvements.

The foundations of the system have proven to be quite usable; however, further developments are needed to increase reliability. Additionally, several functionalities can be integrated that enhance the user experience and bring the system closer to a market-ready state. The following development opportunities are highlighted:

- Importing and exporting test configurations from the database, as well as creating an overview dashboard where the user can view previously run tests and their results.
- Generating statistics based on executed tests and their outcomes, such as success rates, common errors, average run times, etc.
- Providing the ability to run multiple tests sequentially in an automated manner.

- Implement robust error recovery strategies in the event of test execution failure.
- Developing a system that allows users to introduce new test devices. This would include generating the necessary configuration files and training new models, preferably with minimal human intervention.
- Implementing backend server deployability and decoupling communication with the robotic arm from the wired connection.

Acknowledgement

The authors would like to express their sincere gratitude to Codespring for the opportunity to participate in this internship project as part of their mentoring program. We are especially thankful to our mentors, Gyöngyi Katona, Arnold Szász, and Dénes Mihály, for their invaluable guidance, continuous support, and expert advice throughout the development of the project. We also extend our appreciation to Csaba Sulyok for his coordination and oversight, which greatly contributed to the success of our work. Finally, we acknowledge the support of Babeș-Bolyai University, Cluj-Napoca, for fostering an environment that encourages practical and research-based learning experiences.

References

[1] B. J. Dilworth, A. Karlicek and L. Thibault, “An Approach to Component Testing: An Analytical Study,” pp. 341-345, 2019

[2] M. Pernice, D. Piumatti, E. Sánchez, P. Bernardi and R. Martorana, “An efficient strategy for the development of software test libraries for an automotive microcontroller family,” *Microelectronics Reliability,* Vol. 115, p. 113962, 2020

[3] J. Shi, W. Li, W. Wang and L. Guan, Facilitating Non-Intrusive In-Vivo Firmware Testing with Stateless Instrumentation, 2024

[4] K. Kato, S. Takekoshi and T. Shinagawa, Testing device drivers against hardware failures in real environments, 2016, pp. 1858-1864

[5] A. B. Ahmed, O. Mosbahi, M. Khalgui and Z. Li, “Toward a New Methodology for an Efficient Test of Reconfigurable Hardware Systems,” *IEEE Transactions on Automation Science and Engineering,* Vol. 15, No. 4, pp. 1864-1882, 2018

[6] L. Jiao, F. Zhang, F. Liu, S. Yang, L. Li, Z. Feng and R. Qu, “A Survey of Deep Learning-Based Object Detection,” 2019

[7] R. Kaur and S. Singh, “A comprehensive review of object detection with deep learning,” *Digital Signal Processing,* Vol. 132, p. 103812, 2023

[8] P. Kumar and E. Manash, “Deep learning: a branch of machine learning,” *Journal of Physics: Conference Series,* Vol. 1228, p. 012045, 2019

[9] M. Hussain, "YOLO-v1 to YOLO-v8, the Rise of YOLO and Its Complementary Nature toward Digital Manufacturing and Industrial Defect Detection," *Machines,* Vol. 11, 2023

[10] T. Diwan, G. Anirudh and J. V. Tembhurne, "Object detection using YOLO: challenges, architectural successors, datasets and applications," *Multimedia Tools and Applications,* Vol. 82, pp. 9243-9275, 2023

[11] U. Sirisha, P. S. Praveen, P. N. Srinivasu, P. Barsocchi and A. K. Bhoi, "Statistical Analysis of Design Aspects of Various YOLO-Based Deep Learning Models for Object Detection," *International Journal of Computational Intelligence System,* Vol. 16, p. 126, 2023

[12] U. Inc., "Ultralytics YOLO Docs," 2025 [Online] Available: https://docs.ultralytics.com/ [Accessed 3 4 2025]

[13] Shenzhen Yuejiang Technology Co. Ltd, "Dobot Magician Lite User Guide (DobotLab-based)," 6 12 2022 [Online] Available: https://dobot.com.mx/assets/download/magician-lite/Dobot%20Magician%20Lite%20User%20Guide%20(DobotLab-based).pdf

[14] R. Hartley and A. Zisserman, Multiple View Geometry, Cambridge University Press, 1999

[15] E. Dubrofsky, Homography estimation, 2009

[16] I. Shames, O. Biggar and M. Zamani, "An Expressiveness Hierarchy of Behavior Trees and Related Architectures," 2021

[17] P. E. Strandberg, "Automated System-Level Software Testing of Industrial Networked Embedded Systems," *arXiv,* vol. abs/2111.08312, 2021

[18] S. Gollapudi, Learn Computer Vision Using OpenCV: With Deep Learning CNNs and RNNs, Apress Berkeley, CA, 2019, pp. XX, 15

[19] A. Chaudhuri, K. Mandaviya, P. Badelia and S. K. Ghosh, Optical Character Recognition Systems, Springer International Publishing, 2017, pp. 9-41

[20] C. Rodriguez, M. Baez, F. Daniel, F. Casati, J. Trabucco, L. Canali and G. Percannella, REST APIs: A Large-Scale Analysis of Compliance with Principles and Best Practices, Springer International Publishing, 2016, pp. 21-39