

*Neuro-Symbolic AI for Compliance Checking of Electrical Control Panels**

VITO BARBARA

University of Calabria, Rende, Italy
(e-mail: barbara.vito@unical.it)

MASSIMO GUARASCIO

ICAR-CNR, Arcavacata, Italy
(e-mail: massimo.guarascio@icar.cnr.it)

NICOLA LEONE

University of Calabria, Rende, Italy
(e-mail: nicola.leone@unical.it)

GIUSEPPE MANCO

ICAR-CNR, Arcavacata, Italy
(e-mail: giuseppe.manco@icar.cnr.it)

ALESSANDRO QUARTA

Sapienza University of Rome, Rome, Italy
(e-mail: alessandro.quarta@uniroma1.it)

FRANCESCO RICCA

University of Calabria, Rende, Italy
(e-mail: francesco.ricca@unical.it)

ETTORE RITACCO

University of Udine, Udine, Italy
(e-mail: ettore.ritacco@uniud.it)

submitted 24 May 2023; accepted 12 June 2023

Abstract

Artificial Intelligence plays a main role in supporting and improving smart manufacturing and Industry 4.0, by enabling the automation of different types of tasks manually performed by domain experts. In particular, assessing the compliance of a product with the relative schematic is a time-consuming and prone-to-error process. In this paper, we address this problem in a specific industrial scenario. In particular, we define a Neuro-Symbolic approach for automating the compliance verification of the electrical control panels. Our approach is based on the

* The authors express sincere gratitude to the anonymous referees for their valuable suggestions, which contributed to the improvement of our work. We would also like to acknowledge the exceptional support of Dimitri Buelli, Stefano Ierace, Salvatore Iiritano, Giovanni Labocchetta, and Valerio Pesenti elapsed during the development of the system presented in this paper. Their expertise and dedication have been instrumental in shaping the success of our research. Research partially supported by MISE (today MIMIT) under projects “MAP4ID – Multipurpose Analytics Platform 4 Industrial Data” Proj. N. F/190138/01-03/X44 and MUR under PNRR project PE0000013-FAIR, Spoke 9 – Green-aware AI – WP9.1. This work has been carried out while Alessandro Quarta was enrolled in the Italian National Doctorate on Artificial Intelligence run by Sapienza University of Rome with University of Calabria. A preliminary version appeared in CEUR Workshop Proceedings vol. 3204 pp. 247–253.

combination of Deep Learning techniques with Answer Set Programming (ASP), and allows for identifying possible anomalies and errors in the final product even when a very limited amount of training data is available. The experiments conducted on a real test case provided by an Italian Company operating in electrical control panel production demonstrate the effectiveness of the proposed approach.

KEYWORDS: automated quality control systems, answer set programming, computer vision, data scarcity

1 Introduction

With the rise of new technologies, industry moved a step forward to a new era in the field of manufacturing. This complex transformation, including the integration of emerging paradigms and solutions such as *Artificial Intelligence* (AI), *Human-Computer Interaction*, *Cloud Computing*, *Industrial Internet Of Things* (IIoT) and *Blockchain*, is referred as *Industry 4.0*. The impact of the field is witnessed by the effort to promote its development within several national economic policies. For example, the Italian Ministry of Development (nowadays called Ministry of Enterprise and Made in Italy, and identified by the MIMIT acronym) is funding the application of AI to the manufacturing processes to improve efficiency and push the development and modernization of Italian SMEs. In this evolving scenario, Quality Control (QC) is greatly benefiting from the adoption of advanced AI tools and techniques, that can allow for speeding up or automatizing processes of assessment about integrity, working capability, and durability of the products (Javaid *et al.* 2022). In particular, the automation of the compliance verification process for products is among the promising applications of AI for QC that poses a significant challenge for all manufacturing-related businesses, because it can make more efficient a necessary but costly and time-consuming operation in the supply chain.

Among the projects funded by MIMIT is the one titled “Multipurpose Analytics Platform 4 Industrial Data” (MAP4ID), where one of the main use cases is precisely the development of an AI capable of automating the compliance checking of Electrical Control Panels (ECPs).

Basically, an ECP is an enclosure, typically a metal or plastic box, which contains electrical components to control and monitor various mechanical processes, motors, sensors, and actuators. ECPs are employed to regulate a wide variety of components used in industry: for example, they allow to control of mechanical equipment, electrical devices, manufacturing machinery, etc.

One of the basic QC tasks in the manufacturing of ECPs requires checking the *compliance* of the produced control panels with their schematics. Automating this task is particularly relevant since it is currently manually performed by human experts, which makes the whole process inefficient, expensive, and prone to errors. The release of defective ECPs (due to poor quality control) can cause exposure to penalties by the customer and compromise the company’s reputation. The adoption of AI-based tools can greatly mitigate these risks by enabling continuous monitoring of the whole production chain and early detection of issues in each stage of the production process.

Main Problem. In this work, we devise an innovative approach combining Deep Learning (DL) (Goodfellow et al. 2016), and Answer Set Programming (ASP) (Brewka et al. 2011; Gelfond and Lifschitz 1991) to support the QC for the production of electrical control panels. Here, the main task consists in identifying anomalies in the final product, such as the lack, the misplacing, or the wrong connectivity of the electrical components in the cabinet of the ECP, by just analyzing an image of the assembled product. Important requirements are that the AI must be capable of producing the results of the compliance-checking task in a very short time (in the order of seconds) and with high accuracy (>90%), to enable the integration into a tool assisting human inspectors that delivers real-time and robust performance.

This problem is made very challenging for standard DL approaches by the following main issues:

1. *Data scarcity.* Although the companies can produce sufficient amounts of data, semantics, and labels are often missing from images. In particular, in our scenario, such a problem affects both the data representations that is, the pictures depicting the ECPs, and the correspondent schematics. Indeed, supervised information about the position, dimensions, and typology of the installed components is missing for the pictures. As regards the schematics, although they seem to provide a more detailed representation, the possibility of easily translating them into actionable constraints (in the form of grammar rules) strongly depends on the underlying software used to produce them.
2. *Custom Designs.* Despite ECP being made of standard components, there is no standard set of schematics for ECPs. Usually, the design of a solution is customized and very specific for the needs of a specific customer. Thus, the AI must be able to work with different schematics without requiring any additional training.

Contribution. In this work, we define a solution approach composed of two main phases:

1. First, a Deep Learning based solution allows for recognizing the electrical components (object detection) from the images of the panels and reconstructing the scheme. In this phase, a number of data augmentation strategies are also exploited to cope with the lack of labeled data.
2. Then, an Answer Set Programming-based system is used to compare the scheme reconstructed from the picture with its original schematic in order to discover possible mismatches/errors.

The contribution of this paper can be located in the challenge of providing a suitable combination of learning and reasoning through the development of integrated components, which, nowadays, is identified by the buzzword *neuro-symbolic AI* (d'Avila Garcez et al. 2015). Actually, our system can be classified as a *Neural—Symbolic* architecture (or architecture of Type 3) according to Henry Kautz's taxonomy (Kautz 2022), where DL is used for sensing (detecting components) and a reasoner (ASP-based) is used for checking conformance and detecting issues.

Although based on a conceptually straight combination of DL and ASP, an experiment conducted on (scarce) data provided by an Italian SME leader in the production of ECPs

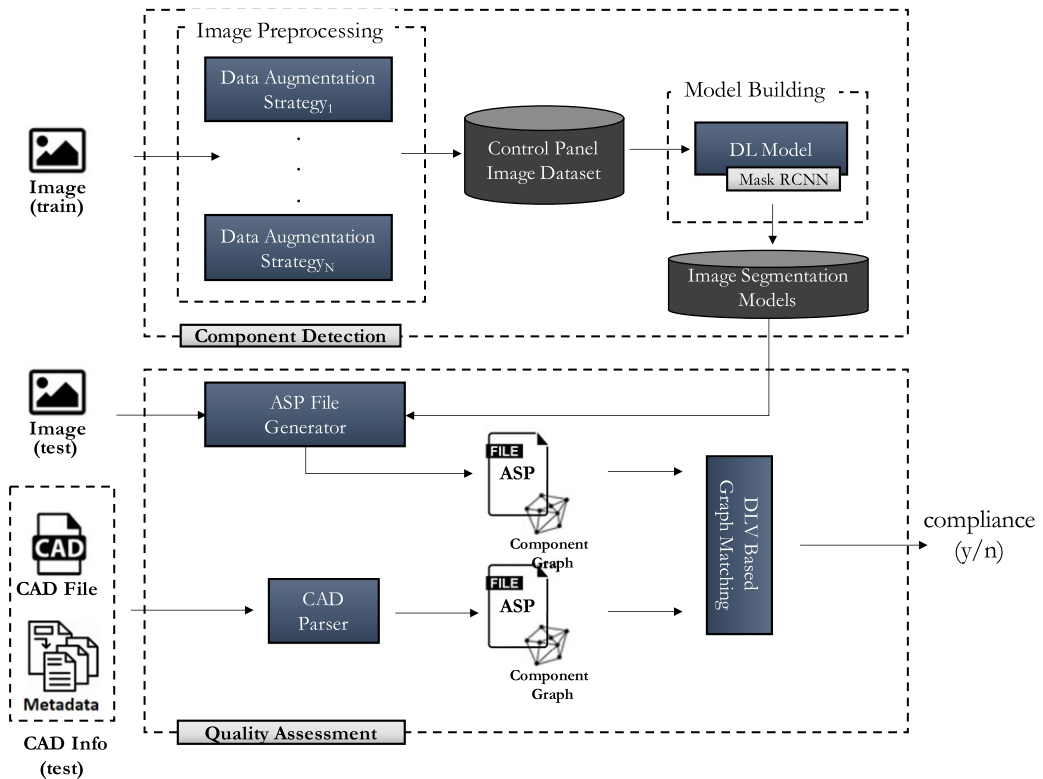


Fig. 1. Framework for automatic compliance verification.

confirms that *our neuro-symbolic system can deliver the expected performance*, which is the main acceptability criteria to be fulfilled by a successful real-world application.

2 Framework overview

In this section, we illustrate the solution approach devised to address the main problem of how to automate the compliance verification process of control panels. As highlighted in Section 1, an effective solution for this problem has to cope with the challenges of understanding image contents and extraction of the constraints encoded in schematics, while coping with the issues of lack of labeled data and unlabeled data distribution. To this aim, we defined the framework shown in Figure 1 that includes two main macro-modules, respectively named *Component Detection* and *Quality Assessment*.

The former is devoted to recognizing the electrical components assembled in the cabinet. Basically, it includes the modules characterizing the adopted machine learning methodology, whose main objective is to identify the components of the panel from a picture. Specifically, a set of *Data Augmentation and generation* techniques builds a suitable dataset (robust to overfitting), which feeds a Convolutional Neural Network (CNN) based model trained to perform the component detection.

The latter exploits ASP to tackle the task of compliance checking. It automatically compares the control panel scheme built starting from the neural network output and the corresponding schematic to highlight any anomalies.

3 Component detection via deep learning

The component detection is meant to recognize, given a picture representing a panel, the type and geometric position of each component within the panel. This is a preliminary and fundamental step since, in order to check the compliance of the cabinets with their schematics, we need first to understand their composition. The main problem in this step is given by the scarcity of data, as well as the lack of labeling annotations. This is a typical scenario characterizing industrial processes: the quality of a machine learning model relies on the data used to train it; however, the latter requires an accurate labeling process that is time and resource-consuming and hence difficult to obtain. In our framework, we address these issues by exploiting a synthetic data generation process that allows us to enrich the starting training set.

3.1 Data augmentation and generation

Basically, our synthetic data generation method is fed with three inputs: (i) a picture showing an empty cabinet, (ii) a catalog including all the available components that can be installed in a cabinet, and (iii) a limited number of real pictures that will be manipulated in order to add noisy elements in the generated data, by exploiting a suitable strategy described in the following.

The core idea is to enrich the ground truth (consisting of a limited number of images) with synthetic images, where the area of the empty cabinet is filled with random components picked from the catalog. Notice that, at this stage, we are not interested in generating compliant panels, since our only objective at this stage is to build a suitable object detector that is capable of recognizing both the component and its geometrical position and extension. The size of the catalog and the randomness of the composition allow us to generate a suitable number of images where each component can be included with a suitable frequency, thus making the result dataset robust to object detection and segmentation learning tasks.

This simple approach can be further combined with other image augmentation strategies (*Image Processing* module in Figure 1) with the aim of yielding a training set that includes a sufficient and diversified number of examples for learning the model. In particular, our framework also includes traditional data augmentation strategies that is, *Gaussian Blur* and *PerspectiveTransform*. As regards the former, the idea consists in introducing imperfections into data so as to make component detection more resilient to data changes, it is obtained by averaging contiguous pixel values. The last one allows for applying random four-point perspective transformations to images.

The resulting dataset from this process will include all the necessary features for the training phase: (i) a large number of different pictures, (ii) the position of each component, (iii) the type of each component. Notably, since each component depicted in the synthetic pictures is randomly placed, the detection model will be forced to learn the intrinsic features of each component, instead of considering positional features, that may

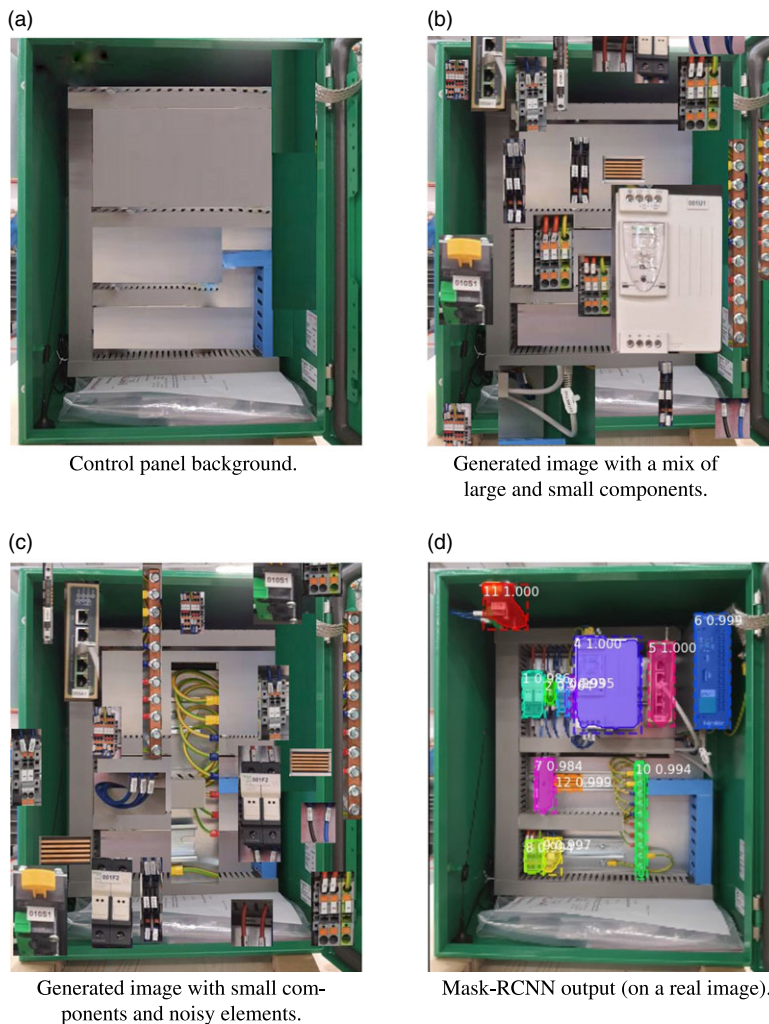


Fig. 2. Input and output of the component detection approach.

vary in the different schematics. Within a cabinet, there are other “auxiliary” elements that are simplified in a schematic, mainly separation boxes, metal runners, and cables. For simplicity, we call them noise to randomly add to the generated images in order to make the object detection model able to distinguish and ignore these elements.

Figure 2, show some examples of the data generation process. We can observe the empty cabinet (Figure 2a), and two instances where it is filled with random components (Figure 2b and 2c). Notice that the synthetic data does not necessarily represent a realistic situation. As already mentioned, this is not an issue since our purpose here is to strengthen the object detection and segmentation phase, which is discussed below.

3.2 Component detection

The *Model Building* module in Figure 1 allows for training the deep architecture used to perform the component detection. For this, we adopted the *Mask R-CNN* convolutional

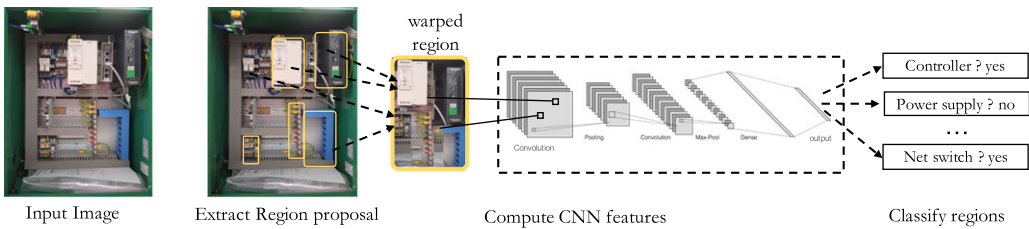


Fig. 3. R-CNN working flow.

neural architecture proposed by He *et al.* (2017). In general, R-CNN (Region based CNN) refers to a family of neural architectures adopting a *Multi-shot* approach. The underlying idea is a two-step process: first, different bounding boxes across possible regions of interest (RoIs) are extracted; then, such regions are independently evaluated through a CNN architecture in order to map them to any of the proposed classes (see Figure 3).

Mask R-CNN extends a specific architecture named *Faster R-CNN* (Ren *et al.* 2015) that includes two main components: (i) *Region Proposal Network* (RPN), a deep neural network aimed at extracting RoIs from the picture, and (ii) *Fast R-CNN*, a neural architecture that performs classification, by scaling a region to a predefined size thus enabling the computation of a set of CNN feature maps. The main advantage of the Faster R-CNN architecture is a suitable trade-off between competitive accuracy in terms of object recognition, and relative speed in the recognition phase. By contrast, other approaches based on Single-Shot architectures such as YOLO (Redmon *et al.* 2016) or SSD (Liu *et al.* 2016) focus on fast recognition, at the cost of recognition accuracy. This is clearly not acceptable in our scenario, where we aim at checking compliance, and missing a component in the picture would result in a failure in the check. Mask R-CNN further improves Faster R-CNN by introducing a further branch for predicting segmentation masks on each Region of Interest. The recognition of the mask is crucial in our scenario since it allows to precisely identify the geometrical position of the component within the panel. Technically, Mask R-CNN rebuilds the mask by resorting to an alignment component and a mask head, composed of two convolutional layers and capable of generating a mask for each RoI in order to segment the picture in a pixel-to-pixel fashion.

Mask R-CNN relies on a backbone convolutional architecture. In our framework, we used ResNet (Residual Network) (He *et al.* 2016), a very deep CNN architecture characterized by *residual blocks* and *skip connections*. These two features guarantee both, fast convergence in the training stage and expressiveness/accuracy in the recognition phase. We further strengthened the training phase by exploiting *Transfer Learning*. In particular, we used a ResNet pre-trained on *COCO* dataset,¹ which was also fine-tuned for our specific scenario, by exploiting the generated dataset with the artificially generated labeled components.

Figure 2d shows the output of the recognition phase, on a real picture representing a true panel. We can see that the model successfully recognizes all available components, and additionally devises a contour of their geometric extension. These contours represent one of the inputs to the reasoning component.

¹ Available online at: <https://cocodataset.org/#home> [Last Accessed: June 2022].

4 Compliance checking in answer set programming

In this section, we describe the reasoning component of our architecture for compliance checking. In particular, this component has been implemented by resorting to Answer Set Programming (ASP). ASP is a well-established paradigm for declarative programming and non-monotonic reasoning developed in the area of Knowledge Representation and Reasoning (Baral 2003; Brewka *et al.* 2011; Bonatti *et al.* 2010; Gelfond and Lifschitz 1991). ASP has been employed to develop many academic and industrial applications of AI (Erdem *et al.* 2016; Gebser *et al.* 2020; Calimeri *et al.* 2016; Grasso *et al.* 2009; 2011; Dodaro *et al.* 2016). ASP is based on logic programming and non-monotonic reasoning, and it allows for flexible declarative modeling of search problems, by means of logic programs (collection of rules), whose intended models (answer sets) encode solutions (Baral 2003; Brewka *et al.* 2011). The specification (logic program) described in the following can be fed to an ASP system to actually compute the solutions to the modeled program (Lierler *et al.* 2016).

The reasoning module is fed by two handlers, named *ASP File Generator* and *CAD Parser*. The former component is devoted to translating the objects recognized by the neural model in ASP facts (a file containing a list of facts concerning coordinates and membership of the electrical component), similarly, the second one yields a list of facts from the input CAD image.

In the following, we focus on the core parts of our solution and simplify some technical aspects that do not impact the comprehension of the working principle of our solution. This is done with the aim of making the presentation more accessible and meeting space requirements. Hereafter, we assume the reader to be familiar with ASP. For details please refer to Brewka *et al.* (2011), Baral (2003), Gelfond and Lifschitz (1991).

4.1 Input specification

In ASP the input specification is made by a set of “facts,” which are assertions that model true sentences. Thus, the labeled schematic of the circuit (we informally refer to it as *cad*), and the output of the Mask R-CNN net (exemplified in Figure 2a) are converted in a set of ASP facts of the following form:

```
object(LABEL, ID, X_TOP_L, Y_TOP_L, X_BOT_R, Y_BOT_R, MEMBERSHIP).
```

These facts provide information about the components like their label, id, top-left and bottom-right coordinates, and membership. In particular, the membership is valued with “*cad*” if the object modeled is part of the schematic of the panel, and “*net*” if it is recognized by the neural network in the actual picture we are comparing to the schematic. Moreover, we also compute a graph of topological relations among objects, providing information on relative position and distance among objects. The relative position and the distance among components are actually calculated by our ASP program, but for simplicity, we assume here they are given in input as facts of the form:

```
between(ID, START_ID, END_ID, DIR, MEM).  
manhattan(ID1, ID2, DIST, MEM1, MEM2).
```

The `between` predicate denotes the neighbors for the component `ID` along the direction `DIR` having `MEM` as membership; additionally, the `manhattan` predicate specifies the

manhattan distance between the two components ID1 and ID2, where the terms MEM1 and MEM2 stand for their membership.

4.2 ASP program

We now present ASP program (see Encoding 1) that encodes in a uniform way (w.r.t. the input instance provided as a set of facts) the compliance problem. First, the graph is preprocessed (lines 2-3), by calculating useful information about the relative positions of the objects. Next, according to the “guess-and-check” programming methodology, a disjunctive rule guesses the mapping between “cad” components of the schematic and “net” components predicted by the neural network (see lines 6-7).

The disjunctive rule can be read as follows: “Given a cad component and a net component of the same type, the two can be mapped, or not.” The candidate solutions are filtered out by the constraints in lines 9-13, ensuring that the same element of the cad is not mapped twice, and the same element of the net is not mapped twice.

The optimal mapping is obtained by weak constraints in lines 15-35. In detail, the program first minimizes the cad elements without a mapping (lines 15-16), then (also in order of priority) the weak constraints in lines 18-31 ensure that “If a cad component ID1 is mapped to a net component ID2, ID1 neighbors should be mapped to ID2 neighbors.”

The mapping is further optimized considering the distance (lines 33-35) between cad components and net components. The distance is optimal when the elements are in the same position in “net” and “cad.” Finally, the program identifies components that are absent or in excess w.r.t. the schematic by rules in lines 37-40.

5 Evaluation

This section describes a suite of experiments we conducted, devoted to demonstrating the effectiveness of our approach and its suitability for the industrial scenario. Specifically, we are interested in evaluating the capability of DL-based approach in recognizing the cabinet components when no training data are available, and the scalability of the ASP-based technique in verifying the conformance of the image with the schematics.

5.1 Experimental setup

We set up the experimentation by considering the extreme scenario where no labeled examples are available. Therefore, our training set used includes only the synthetically generated images (by using the data augmentation techniques described in Section 3), while the real pictures of the EPCs are used to evaluate the predictive performances. The final result of this process is a training set composed of ~10,000 colored images synthetically generated with size (320 × 320) and a test set of 32 images depicting real control panels with the same size as the training ones.

The model discussed in Section 3 has been implemented in the form of a python prototype based on TensorFlow² library. The experiments were performed on an NVidia

² TensorFlow machine learning library: <https://www.tensorflow.org/> [Last Accessed: June 2022].

Encoding 1 ASP program modeling compliance

```

1: ▷ Calculate auxiliary information
2:   previous(ID, Start_ID, D, M):- between(ID, Start_ID, _, D, M).
3:   after(ID, End_ID, D, M):- between(ID, _, End_ID, D, M).
4: ▷ Guess mapping between cad components and net components
5:   simpObject(C1, ID1, M) :- object(C1, ID1, _, _, _, M).
6:   mapped(ID1, ID2) || noMapped(ID1, ID2)
7:       :- simpObject(C1, ID1, "cad"), simpObject(C1, ID2, "net").
8: ▷ No element from the cad is mapped twice
9:   :- mapped(Cad_ID, Net_ID1), mapped(Cad_ID, Net_ID2),
10:      Net_ID1 != Net_ID2.
11: ▷ No element from the net is mapped twice
12:   :- mapped(Cad_ID1, Net_ID), mapped(Cad_ID2, Net_ID),
13:      Cad_ID1 != Cad_ID2.
14: ▷ Minimize the cad elements without a mapping
15:   atLeastOne(Cad_ID) :- mapped(Cad_ID, _).
16:   :~ simpObject(C1, ID1, "cad"), not atLeastOne(ID1). [1@3, C1, ID1]
17: ▷ Optimize mapping by relative position
18:   :~ mapped(Cad_ID1, Net_ID1), mapped(Cad_ID2, Net_ID2),
19:      previous(Cad_ID1, Cad_ID2, DIR, "cad"),
20:      not previous(Net_ID1, Net_ID2, DIR, "net").
21:   [1@2, Cad_ID1, Net_ID1, Cad_ID2, Net_ID2, DIR]
22:   :~ mapped(Cad_ID1, Net_ID1), mapped(Cad_ID2, Net_ID2),
23:      after(Cad_ID1, Cad_ID2, DIR, "cad"),
24:      not after(Net_ID1, Net_ID2, DIR, "net").
25:   [1@2, Cad_ID1, Net_ID1, Cad_ID2, Net_ID2, DIR]
26:   :~ mapped(Cad_ID1, Net_ID1),
27:      previous(Cad_ID1, Cad_ID2, DIR, "cad"),
28:      absent(_, Cad_ID2). [1@2, Cad_ID1, Net_ID1, Cad_ID2, DIR]
29:   :~ mapped(Cad_ID1, Net_ID1),
30:      after(Cad_ID1, Cad_ID2, DIR, "cad"),
31:      absent(_, Cad_ID2). [1@2, Cad_ID1, Net_ID1, Cad_ID2, DIR]
32: ▷ Optimize mapping by distance
33:   :~ mapped(Cad_ID, Net_ID),
34:      manhattan(Cad_ID, Net_ID, Dis, "cad", "net").
35:   [Dis@1, Cad_ID, Net_ID, Dis]
36: ▷ Identify absent and in excess components
37:   mappedCad(ID1):- mapped(ID1, _).
38:   mappedNet(ID1):- mapped(_, ID1).
39:   absent(C1, ID1):- simpObject(C1, ID1, "cad"), not mappedCad(ID1).
40:   excess(C1, ID1):- simpObject(C1, ID1, "net"), not mappedNet(ID1).

```

DGX Station equipped with 4 GPU V100 32GB. As described in Section 3, a ResNet instance (including 101 layers) is used as the backbone of the component detection model, the Mask R-CNN, that is trained over 200 epochs with *batch_size* = 2, while *Adam* is adopted as optimizer with learning rate $lr = 10^{-4}$.

To assess the capability of the proposed approach in detecting the components installed in the ECPs, a number of traditional measures and well-known metrics for the Object Recognition tasks have been used. In this sub-section, we briefly introduce and define such measures. The first measures we consider are the standard Precision and Recall metrics, defined as $p = \frac{TP}{TP+FP}$ and $r = \frac{TP}{TP+FN}$. Here, TP , FP , FN , and TN denote respectively the number of cases that are: positive and correctly classified, positive and incorrectly classified, negative and incorrectly classified, and negative and correctly classified. Hence, a Precision-Recall Curve can be obtained by computing and plotting the precision against the recall for different threshold values (i.e. the detection probabilities of the model).

In an object detection scenario, precision and recall represent the capability of the prediction model to identify the boxes that contain the target objects. In particular, for a given object the focus is on comparing the true bounding box with the predicted bounding box, and the TP , FP , FN , and TN values depend on the degree of overlap between these two boxes. Given two boxes, the *Intersection Over Union* (IoU) is defined as the fraction of the overlapping area between the ground truth b and the predicted bounding box \hat{b} :

$$IoU(b, \hat{b}) = \frac{b \cap \hat{b}}{b \cup \hat{b}} \quad (1)$$

Then, given a threshold θ , an object with a true bounding box b and a predicted bounding box \hat{b} is *positive* if $IoU(b, \hat{b}) > \theta$, and *negative* otherwise. For a given θ , it is possible to devise a precision-recall curve by plotting all p/r values relative to all objects and interpolating the resulting curve (He et al. 2020; Ren et al. 2015).

Since the values of precision and recall are defined on a given θ threshold, we can define (He et al. 2020; Ren et al. 2015) *Average Precision* and *Recall* as the area represented by integrating over all possible thresholds:

$$AP = \int_0^1 p(\theta) d\theta \quad (\text{and resp.}) \quad AR = \int_0^1 r(\theta) d\theta. \quad (2)$$

Finally, by averaging AP (resp. AR) over all class components we can finally obtain the *mean average precision* (mAP) and *mean average recall* (mAR) measures.

5.2 Evaluation results

Here, we discuss the results in terms of the effectiveness of the DL-Based detection model and the scalability of the ASP module. For the first aspect, the detection model exhibits optimal performances for both the quality measures, exhibiting values of $mAP = 0.954$ and $mAR = 0.935$. In order to evaluate the operational applicability of our approach in a real scenario, we conducted a further analysis by considering the values of precision and recall on a fixed IoU threshold $\theta = 0.5$. Basically, in this test case, precision and recall, respectively, represent the capability of the model to correctly recognize the components depicted in the picture and the percentage of recognized components w.r.t. the ground truth.

Figure 4a reports the resulting precision-recall curve. The resulting area is 0.947, which denotes a good performance of the detection model also considering the operational case. Figure 4b shows a more detailed picture of the model performances by plotting the pr-curve for each instance. As expected, for almost all instances, the yielded curves

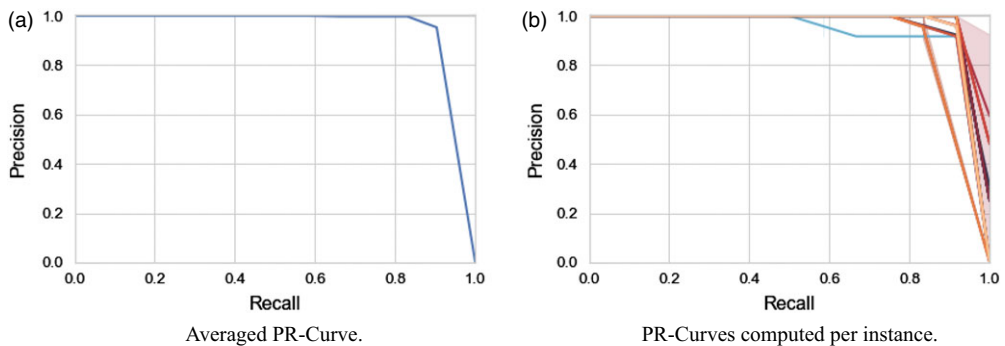


Fig. 4. Precision-Recall Curves on test set.

highlight the good predictive accuracy of the model in recognizing the different types of components, except for one case in which the quality is slightly lower. The above evaluation shows that the component detection module is effective in recognizing the components of a panel: in particular, prediction errors can occur in rare cases with inaccurate image acquisition (e.g. non-frontal framing or inclusion of elements external to the cabinet) as the catalog provides only a limited number of component perspectives. An example of such behavior is depicted in Figure 5, where accuracy is affected by the wrong perspective of the image. Since the ASP program performs the compliance task with optimal accuracy in our benchmark images, the accuracy of the system corresponds with the one of the neural model.

One might wonder whether the ASP component is efficient thus in a further experiment the execution time of the ASP-based component was measured. We generated instances of compliance testing in a range of 6–50 labels (types of components), and of 12–75 components and averaged over 500 samples the execution time needed by our ASP engine DLV2 (Alviano *et al.* 2017) to solve the instances. The results reported in Figure 6 show that our system provides answers in a short time, in the order of seconds for instances sized as real-world ones, and performance is acceptable (avg. 1.93 s, max about 18 s) also for instances of 75 components.

6 Related works

In this section, we survey some relevant works that try to address the product quality assurance problem by leveraging AI-based strategies, then we review some preliminary works proposing solutions to integrate ML techniques with logic programming.

Compliance Checking through Machine Learning. To the best of our knowledge, the problem of assessing the compliance of a product with its schematic through Artificial Intelligence techniques is new and quite unexplored however, some recent works tried to tackle similar tasks, in particular within *Predictive Maintenance* field. For instance, Tanuska *et al.* (2021) propose a comprehensive framework integrating Industrial Internet of Things (IIoT) devices, neural networks, and sound analysis for detecting anomalies in the production chain. Schmitt *et al.* (2020) define a holistic solution for quality inspection based on merging Machine Learning techniques and Edge Cloud Computing

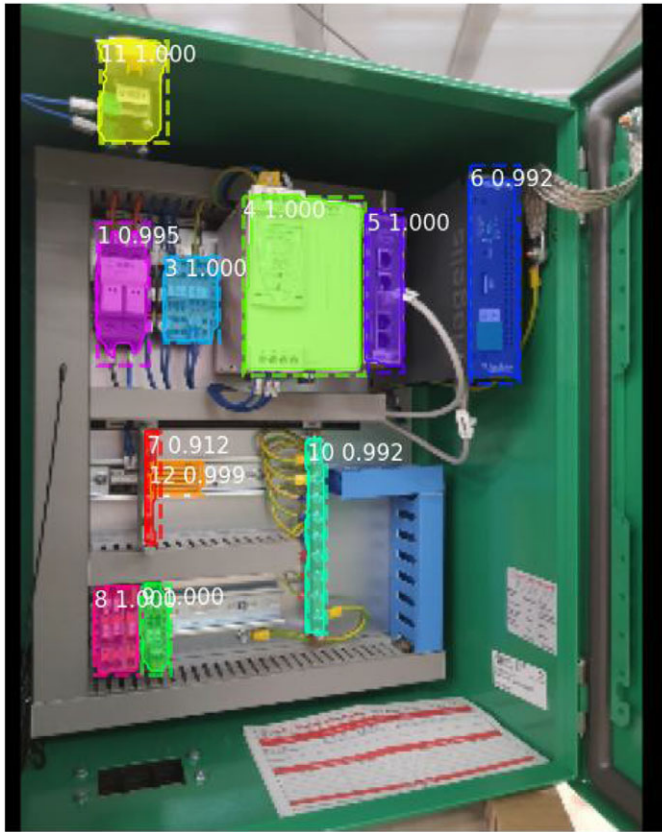


Fig. 5. An example of inaccurate image acquisition. The side perspective of the image does not match the component images in the catalog leading to inaccurate predictions.

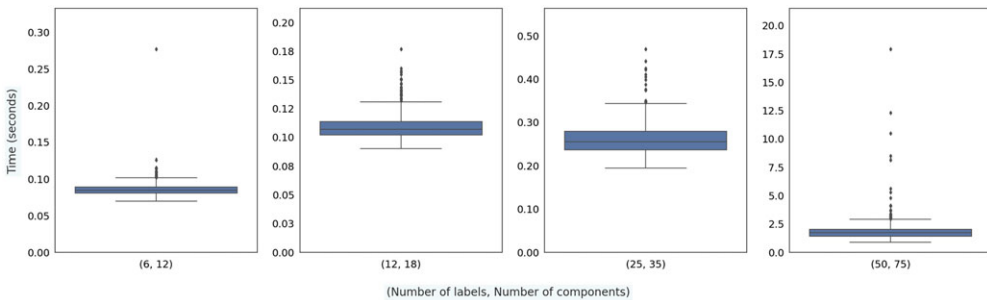


Fig. 6. Performance of the ASP-Based component (Execution time).

technology. A Deep Learning based approach for monitoring the process of sealing and closure of matrix-shaped thermoforming food packages is proposed by [Banus Paradell et al. \(2021\)](#). Specifically, Computer Vision techniques are exploited to process the images and perform quality checking. A comparison analysis performed by ranging different Convolutional Neural Network architectures (e.g. ResNet50, VGG19, ImageNet, etc.) highlights the best solutions to address this task. [Villalba-Diez et al. \(2019\)](#) propose a deep neural network (DNN) soft sensor enabling fast quality control for the Printing

Industry. Basically, the solution allows for comparing the scanned surface of the print with the correspondent file that generated it and performs an automatic quality control process by learning features through exposure to training data. Subakti and Jiang (2018) define and develop a deep learning-based framework to detect/recognize different machines and portions of machines for smart factories. MobileNets is used as the backbone for the machine recognition model, and it is deployed on mobile devices to support the operators in performing the machine classification through an augmented reality system. Experimental results on a real scenario show the capability of the approach in recognizing different machines and providing intuitive visualizations.

In Table 1, we compare the main approaches proposed in the literature and highlight the differences w.r.t. our solution. The main advantage of our approach (the only one based on a neuro-symbolic architecture) stays in the nature of the symbolic component that does not require additional training to deal with new (unseen) schematics. Another distinguishing feature is the ability to work with data scarcity (i.e. small training sets).

ML and ASP integration. The integration of inductive with deductive reasoning is an emerging problem in Artificial Intelligence (AI). Several proposals were made to implement the reasoning process in complex deep neural network (DNN) architectures (Kathryn and Mazaitis 2018; Rocktäschel and Riedel 2017; Yang *et al.* 2020; Lin *et al.* 2019; Donadello *et al.* 2017). The integration of deductive logical reasoning with the Deep Learning paradigm is a novel and quite unexplored research topic, although some recent works introduced interesting preliminary solutions (Ebrahimi *et al.* 2021). Concerning ASP, we recall that it is a declarative rule-based programming paradigm for knowledge representation and declarative problem-solving, that is known to be appropriate for executing complex knowledge-based applications (Erdem *et al.* 2016). One of the main issues is to incorporate high-dimensional vector space and pre-trained models for perception tasks as handled in deep learning, which limits the applicability of ASP in many practical applications involving data and uncertainty. Nonetheless, to overcome this issue a blending ASP with DL has been recently studied (Yang *et al.* 2020).

7 Conclusions and future works

Quality Control is a manually performed and prone-to-error task crucial for each company, indeed the release of defective products can damage the company's reputation and lead to the payment of penalties to the customer.

This paper describes a Neuro-symbolic approach to checking the compliance of electrical control panels with their schematics. A picture of a control panel is fed as input to a neural network to recognize the installed components and their locations, then an ASP-based module is used to compare the scheme reconstructed from the picture with its original blueprint and detect possible mismatches/errors.

The system can handle the lack of labeled data and is resilient to noise and variety in the specifications of schematics (no additional training is required, just an updated logical representation of the schematic). The overall system has been exploited in a practical use case provided by an Italian SME leader in the production of ECPs, where it has been shown to fulfill the requirements both in terms of accuracy and evaluation time.

Table 1. Comparison of the ML/DL-based approaches

Article	Application scenario	Solution	ML/DL model(s)	Neur. symb.	Labeled data
Tanuska et al. (2021)	IIoT, Sensors data	Combining Neural Networks and Sound Analysis	Feedforward Neural Networks	×	✓
Schmitt et al. (2020)	Surface Mount technology manufacturing	Merging ML with Edge Cloud Computing	Traditional ML Supervised techniques	×	✓
Banus Paradell et al. (2021)	Quality Control (Food Packages)	Computer Vision	ResNet, VGGNet, DenseNet	×	✓
Villalba-Diez et al. (2019)	Quality Control (Printing Industry)	Computer Vision	CNN (loosely inspired to AlexNet)	×	✓
Subakti and Jiang (2018)	Smart Factory (Machine Recognition)	Combination of Computer Vision and Augmented Reality	MobileNet	×	✓
<i>Our Solution</i>	ECP Compliance Verification	A framework based on DL and Answer Set Programming	Mask R-CNN	✓	×

Despite its practical utility, there is still room for improving the proposed framework. In fact, we plan to extend it along two research directions. Concerning the model, we can improve the learning phase by adopting a *Triplet Loss* (Kaya and Bilge 2019) architecture and by changing the model backbone (e.g. by resorting to *Vision Transformers* (Dosovitskiy *et al.* 2020)). Another potential issue is that the proposed model disregards the depth of the cabinet. In practice, we only consider a two-dimensional model where each component is placed on a plane. There are situations, however, where components partially overlap frontally but occupy different positions in depth. For these situations, a more accurate model that also addresses depth estimation should be considered.

The second line for possible is represented by the reasoning modules, where the logic programs can be calibrated to compute suggestions for the user, as well as suggest alternate schematic plans. Finally, one could study whether a tighter integration of the neural and logic-based components can enhance the results provided by the vision procedure.

References

- ALVIANO, M., CALIMERI, F., DODARO, C., FUSCÀ, D., LEONE, N., PERRI, S., RICCA, F., VELTRI, P. AND ZANGARI, J. 2017. The ASP system DLV2. In *Proceedings of LPNMR 2017*. LNCS, vol. 10377. Springer, 215–221.
- BANUS PARADELL, N., BOADA, I., XIBERTA, P., TOLDRA, P. AND BUSTINS, N. 2021. Deep learning for the quality control of thermoforming food packages. *Scientific Reports* 11, 21887.
- BARAL, C. 2003. *Knowledge Representation, Reasoning and Declarative Problem Solving*. Cambridge University Press.
- BONATTI, P. A., CALIMERI, F., LEONE, N. AND RICCA, F. 2010. Answer set programming. In *25 Years GULP*. Lecture Notes in Computer Science, vol. 6125. Springer, 159–182.
- BREWKA, G., EITER, T. AND TRUSZCZYNSKI, M. 2011. Answer set programming at a glance. *Communications of the ACM* 54, 12, 92–103.
- CALIMERI, F., GEBSER, M., MARATEA, M. AND RICCA, F. 2016. Design and results of the fifth answer set programming competition. *Artificial Intelligence* 231, 151–181.
- D’AVILA GARCEZ, A. S., BESOLD, T. R., RAEDT, L. D., FÖLDIÁK, P., HITZLER, P., ICARD, T., KÜHNBERGER, K., LAMB, L. C., MIKKULAINEN, R. AND SILVER, D. L. 2015. Neural-symbolic learning and reasoning: Contributions and challenges. In *Proceedings of 2015 AAAI Spring Symposia*. AAAI Press.
- DODARO, C., GASTEIGER, P., LEONE, N., MUSITSCH, B., RICCA, F. AND SCHEKOTIHIN, K. 2016. Combining answer set programming and domain heuristics for solving hard industrial problems (application paper). *Theory and Practice of Logic Programming* 16, 5–6, 653–669.
- DONADELLO, I., SERAFINI, L. AND GARCEZ, A. D. 2017. Logic tensor networks for semantic image interpretation. In *Proceedings of IJCAI’17*. AAAI Press, 1596–1602.
- DOSOVITSKIY, A., BEYER, L., KOLESNIKOV, A., WEISSENBORN, D., ZHAI, X., UNTERTHINER, T., DEGHANI, M., MINDERER, M., HEIGOLD, G., GELLY, S., USZKOREIT, J. AND HOULSBY, N. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*.
- EBRAHIMI, M., EBERHART, A., BIANCHI, F. AND HITZLER, P. 2021. Towards bridging the neuro-symbolic gap: Deep deductive reasoners. *Applied Intelligence* 51, 9, 6326–6348.
- ERDEM, E., GELFOND, M. AND LEONE, N. 2016. Applications of answer set programming. *AI Magazine* 37, 3, 53–68.
- GEBSER, M., MARATEA, M. AND RICCA, F. 2020. The seventh answer set programming competition: Design and results. *Theory and Practice of Logic Programming* 20, 2, 176–204.
- GELFOND, M. AND LIFSCHITZ, V. 1991. Classical negation in logic programs and disjunctive databases. *New Generation Computing* 9, 3/4, 365–386.

- GOODFELLOW, I. J., BENGIO, Y. AND COURVILLE, A. C. 2016. *Deep Learning*. Adaptive Computation and Machine Learning. MIT Press.
- GRASSO, G., IIRITANO, S., LEONE, N. AND RICCA, F. 2009. Some DLV applications for knowledge management. In *LPNMR*. Lecture Notes in Computer Science, vol. 5753. Springer, 591–597.
- GRASSO, G., LEONE, N., MANNA, M. AND RICCA, F. 2011. ASP at work: Spin-off and applications of the DLV system. In *Logic Programming, Knowledge Representation, and Nonmonotonic Reasoning*. Lecture Notes in Computer Science, vol. 6565. Springer, 432–451.
- HE, K., GKIOXARI, G., DOLLÁR, P. AND GIRSHICK, R. 2020. Mask R-CNN. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 42, 2, 386–397.
- HE, K., GKIOXARI, G., DOLLÁR, P. AND GIRSHICK, R. 2017. Mask R-CNN. In *2017 IEEE International Conference on Computer Vision (ICCV)*, 2980–2988.
- HE, K., ZHANG, X., REN, S. AND SUN, J. 2016. Deep residual learning for image recognition. In *Proceedings of IEEE CVPR 2016*, 770–778.
- JAVOID, M., HALEEM, A., SINGH, R. P. AND SUMAN, R. 2022. Artificial intelligence applications for industry 4.0: A literature-based study. *Journal of Industrial Integration and Management* 7, 1, 83–111.
- KATHRYN, W. W. C. F. Y. AND MAZAITIS, R. 2018. Tensorlog: Deep learning meets probabilistic databases. *Journal of Artificial Intelligence Research* 1, 1–15.
- KAUTZ, H. A. 2022. The third AI summer: AAAI robert s. Englemore memorial lecture. *AI Magazine* 43, 1, 93–104.
- KAYA, M. AND BILGE, H. S. 2019. Deep metric learning: A survey. *Symmetry* 11, 9.
- LIERLER, Y., MARATEA, M. AND RICCA, F. 2016. Systems, engineering environments, and competitions. *AI Magazine* 37, 3, 45–52.
- LIN, B. Y., CHEN, X., CHEN, J. AND REN, X. 2019. KagNet: Knowledge-aware graph networks for commonsense reasoning. In *Proceedings of EMNLP-IJCNLP*. Association for Computational Linguistics, 2829–2839.
- LIU, W., ANGUELOV, D., ERHAN, D., SZEGEDY, C., REED, S., FU, C.-Y. AND BERG, A. C. 2016. Ssd: Single shot multibox detector. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- REDMON, J., DIVVALA, S., GIRSHICK, R. AND FARHADI, A. 2016. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- REN, S., HE, K., GIRSHICK, R. AND SUN, J. 2015. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Proceedings of NIPS - Volume 1*. MIT Press, 91–99.
- ROCKTÄSCHEL, T. AND RIEDEL, S. 2017. End-to-end differentiable proving. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017*, 3788–3800.
- SCHMITT, J., BÖNIG, J., BORGRÄFE, T., BEITINGER, G. AND DEUSE, J. 2020. Predictive model-based quality inspection using machine learning and edge cloud computing. *Advanced Engineering Informatics* 45, 101101.
- SUBAKTI, H. AND JIANG, J.-R. 2018. Indoor augmented reality using deep learning for industry 4.0 smart factories. In *2018 IEEE COMPSAC*, vol. 02, 63–68.
- TANUSKA, P., SPENDLA, L., KEBISEK, M., DURIS, R. AND STREMY, M. 2021. Smart anomaly detection and prediction for assembly process maintenance in compliance with industry 4.0. *Sensors* 21, 7, 2376.
- VILLALBA-DIEZ, J., SCHMIDT, D., GEVERS, R., MERÉ, J. B. O., BUCHWITZ, M. AND WELLBROCK, W. 2019. Deep learning for industrial computer vision quality control in the printing industry 4.0. *Sensors* 19, 18, 3987.
- YANG, Z., ISHAY, A. AND LEE, J. 2020. Neurasp: Embracing neural networks into answer set programming. In *Proceedings of IJCAI*, 1755–1762.