



iJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 10 Issue: VI Month of publication: June 2022

DOI: <https://doi.org/10.22214/ijraset.2022.44930>

www.ijraset.com

Call:  08813907089

E-mail ID: ijraset@gmail.com

Non-Destructive Defect Detection Using MASKR-CNN on Industrial Radiography of Machines

Aman Kumar Mahapatra¹, Ajith Haridasan², Danish Sharma³, Anuj Singh⁴, Dr. Sudhan M.B⁵

^{1, 2, 3, 4}Dept of Computer Science & Engg., MVJ College of Engineering, Bangalore, India

⁵Dept of AI & ML, MVJ College of Engineering, Bangalore, India

Abstract: Many industrial processes, particularly those requiring casting or welding, rely heavily on quality control. Manual quality control processes, on the other hand, are frequently time-consuming and error-prone. To address the increased demand for high-quality products, sophisticated visual inspection technologies are becoming increasingly important in manufacturing lines. Convolutional Neural Networks have recently demonstrated exceptional performance in image classification and localization tasks. Based on the Mask Region-based CNN architecture, this research proposes a solution for detecting casting errors in X-ray pictures. The suggested defect detection system conducts flaw identification and segmentation on input pictures at the same time, making it appropriate for a variety of defect detection jobs. It is demonstrated that training the network to conduct defect detection and defect instance segmentation at the same time leads in greater defect detection accuracy than training on defect detection alone. Transfer learning is used to minimize training data requirements while increasing the trained model's prediction accuracy. More precisely, the model is trained using two huge publically available picture datasets before being fine-tuned using a relatively modest metal casting X-ray dataset. The trained model's accuracy outperforms state-of-the-art performance on the GRIMA database of Xray images (GDxray) Castings dataset and is quick enough to be deployed in production. On the GDxray Welds dataset, the system likewise works well. A variety of in-depth research are being undertaken to investigate how transfer learning, multi-task learning, and multi-class learning affect the trained system's performance.

Index Terms: Non-Destructive Testing, Mask R-CNN, Image Segmentation and Annotation

I. INTRODUCTION

Quality control is an essential component of every production process [1]. Manufacturers must boost production rates while adhering to strict quality control constraints in order to satisfy growth objectives. According to a recent assessment, the most essential technological innovation for manufacturing company success is the creation of superior quality management systems [2]. To address the increased demand for high-quality products, sophisticated visual inspection technologies are becoming increasingly important in manufacturing lines. Casting and welding processes, for example, might generate faults in the product that are harmful to the final product's quality [3]. Air holes, foreign-particle inclusions, shrinkage voids, fractures, wrinkles, and casting fins are examples of common casting flaws [4]. If these casting flaws go unnoticed, they can lead to catastrophic failure of crucial mechanical components like turbine blades, brake calipers, or vehicle driveshafts. Early identification of these faults can allow defective goods to be recognised early in the production process, saving time and money [5]. In order to ensure consistent and cost-effective inspection, automated quality control might be implemented. The key motivations for automated inspection systems include quicker inspection rates, increased quality expectations, and the requirement for more quantitative product assessment that is not limited by the effects of human weariness. Nondestructive assessment techniques allow a product to be checked throughout the manufacturing process without endangering the product's quality. There are several nondestructive evaluation techniques available for creating two-dimensional and three-dimensional pictures of an object. Real-time X-ray imaging technology is frequently employed in defect detection systems in industry, such as on-line weld defect inspection [5]. Ultrasonic and magnetic particle inspection may also be utilized to determine the size and location of casting flaws in cast components [6, 7].

Many cutting-edge object identification algorithms have been built utilizing the region-based convolutional neural network (R-CNN) architecture [10]. R-CNN generates bounding boxes, or region suggestions, by a method known as selective search. To detect objects, selective search examines the picture via different-sized windows and attempts to group neighboring pixels by texture, color, or intensity for each size. Once the suggestions have been developed, R-CNN warps the region to a regular square size and runs it through a feature extractor. A support vector machine (SVM) classifier is then used to predict whether or not an item exists in the picture. Each component of the object identification network is replaced with a deep neural network in more modern object detection designs, such as region-based fully convolutional networks (R-FCN) [11].

Using current improvements in computer vision, this paper develops a quick and accurate flaw identification method. The suggested flaw detection method is built on the Mask R-CNN architecture [12]. This architecture does object identification and instance segmentation at the same time, making it helpful for a variety of automated inspection tasks. The suggested system is trained and tested using the GRIMA X-ray image dataset, which was released by Grupo de Inteligencia de Maquina (GRIMA) [13].

The rest of this article is structured as follows: The first part offers an overview of related works, while the second portion provides an overview of CNNs.

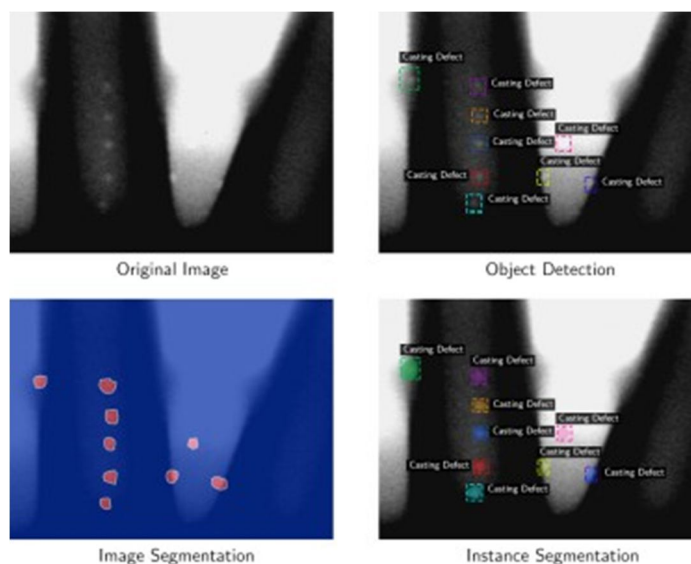


Figure 1 shows a variety of computer vision tasks for detecting casting flaws.

II. RELATED WORKS

The Defect Detection System section contains a full explanation of the proposed defect detection system. The section Implementation Details and Experimental Results shows how the system is trained to detect casting errors, as well as the key experimental results and a comparison with similar systems in the literature. The paper concludes with a number of detailed investigations, a detailed discussion of the findings, and a brief conclusion.

Traditional computer vision approaches have been pretty thoroughly explored in their ability to identify and segment casting flaws. Background subtraction is a prominent technique for removing the flaws and random noise from the preprocessed picture [14, 15]. With different degrees of effectiveness, background removal has been used in the welding flaw identification job. Background removal, on the other hand, tends to be very sensitive to the picture's location and random image noise.

One of the most prominent matched filtering approaches that has been presented is the modified median (MODAN) filtering method. Casting faults may be distinguished from the casting's structural outlines using the MODAN-Filter, a median filter with custom filter masks [20]. Wavelet-based approaches have been suggested by a number of researchers [4, 21]. Wavelet and frequency-based techniques often identify flaws as high-frequency portions of the picture, when contrasted to low-frequency background areas [22]. When identifying flaws, many of these systems do not include local and global picture information, making it difficult to distinguish between design elements like holes and edges and casting faults.

It is customary in many classic computer vision systems to manually define a number of characteristics that may be used to categorize individual pixels. In order to determine whether an individual picture pixel is a defect or not, a local neighborhood of characteristics surrounding the pixel is generated and used to determine the classification.

Localized wavelet decomposition and statistical descriptors such as mean and standard deviation are two of the most used features [4]. However, recent computer vision algorithms based on convolutional neural networks (CNNs) have essentially supplanted these earlier fuzzy logic approaches.

According to literature, automatic surface inspection (ASI) may also be performed. Local aberrations in homogenous textures are what ASI refers to as surface flaws. ASI techniques may be categorized into four categories based on the surface texture characteristics [24].

The texture primitives and displacements may be modeled using structural approaches. Primitive measurements [25], edge features [26], and morphological operations [27] are all popular structural techniques. Another strategy is to use statistical techniques to examine the distribution of pixel values. Stochastic textures like ceramic tiles, castings, and wood may benefit from the statistical method. Co-occurrence matrix [30] and histogram-based method (HBM) [28] are two of the most often used statistical approaches. Third, filter-based approaches use texture pictures to apply filter banks. It is possible to separate filter-based approaches into spatial (31), frequency (32), and spatial-frequency (33). To conclude, model-based techniques generate picture representations by simulating numerous defects' attributes [34].

As a result of well-archived experimental datasets, such as the GRIMA database of X-ray pictures (GDXray), the scientific community, including this study, is substantially enriched. [35] compares the performance of different basic approaches for flaw segmentation using the GDXray Welds series, but only qualitatively evaluates each method. Patches of varying sizes are used in [36] to examine casting fault identification using various computer vision algorithms.

GDXray Castings series images are reduced to 32x32 pixels and used to train and test a variety of classifiers. A basic LBP descriptor with a linear SVM classifier provides the best results [36]. An 86.4 percent patch classification accuracy was achieved using many deep learning algorithms. Pretrained neural networks may be fed with resized patches of the original size of 32 by 32 by 3 pixels when using deep learning methods [36, 37]. With 90.5 percent accuracy on the binary classification of 25 pixel patches, [38] uses deep CNNs for weld fault segmentation. Several machine learning algorithms have recently been effectively applied to the issue of object identification.

Faster RegionBased CNN (Faster R-CNN) and Single Shot Multibox Detector (SSD) are two significant neural network techniques [39, 40]. There are numerous parallels between the two techniques, although the latter places a higher value on speed than accuracy in the review process. [41] provides a comparison of several object detection networks. [12] Mask R-CNN is an extension of Faster R-CNN that concurrently does object identification and instance segmentation. According to earlier research, a defect detection system can be built using the Faster R-CNN framework. A defect detection system with object identification and instance segmentation capabilities is being developed in this new study to expand on previous success.

III. CONVOLUTIONAL NEURAL NETWORK

Computer vision has made great strides in picture categorization, object identification, and image segmentation during the last several years. In many image processing jobs, the introduction of deep CNNs has resulted in significant advances.

An introduction to CNNs is provided in this area of the site. The reader is directed to [43] for a more in-depth explanation. A CNN uses a series of mathematical procedures to transform each image's pixels into a feature-rich representation. It is possible to express images as a three-dimensional tensor with the following dimensions: height (H), width (W), color channels.

Layers are the terms used to describe the many processing processes that the input goes through when it is fed into the system. It is possible to think of each i -th layer as having inputs and outputs of a different kind, each of which may be thought of as a different kind of transformation. It is common to refer to a layer's outputs as a feature map. The development of a complicated nonlinear function that can transfer high-dimensional input (such as pictures) to usable outputs (such as classification labels) may be achieved by merging numerous layers [43]. $f(x) = f_N(f_2(f_1(x_1; 1); 2)); N, (1)$ where x_1 is the input to the CNN and $f(x)$ is the output may be thought of informally as the composition of a number of functions.

Layer types in current CNNs include convolutions, pools and batch normalizations. For each parameterized kernel, the convolution layer applies a function $f_i(x_i; i)$ on the input tensor, x_i . An order 3 Tensor with dimensions H_i , W_i and D_i may be used as an input to the algorithm. Tensor of order 3 convolution kernel is also $H/W/D_i$ in size. Convoluting the kernel and input is accomplished by taking the kernel's dot product over all of the input's spatial locations.

Image gradients may be obtained by converging specific kinds of kernel images with their respective input images. The first few convolutional layers in most current CNN designs extract characteristics such as edges and textures. Shapes of objects may be extracted from images using convolutional layers in the network's lower levels.

It is by design that deep neural networks are parameterized nonlinear functions.. This nonlinearity is introduced by applying an activation function to the output of a neural network layer. The sigmoid function has traditionally been utilized as a nonlinear activation function in neural networks, however this has changed recently. The Rectified Linear Unit (ReLU) is the most often utilized activation function in current architectures since it performs best in terms of runtime and generalization error. Each value, z , in the input tensor x_i , is represented by the formula $f(z) = \max(0, z)$.

Unless otherwise noted, the ReLU activation function is used as the activation function in the defect detection system described in this article. It is also typical in current CNN designs to use pooling layers. Pooling layers are used to gradually reduce the representation's spatial dimension.

As a result, overfitting may be controlled. The input tensor's spatial dimensions are generally averaged or multiplied by the pooling layer's maximum value. A 2×2 or 3×3 section of the input tensor is commonly used for the pooling procedure. Networks may be built that allow for an evolutionary and hierarchical growth of raw data into sophisticated feature representations by stacking pooling and convolutional layers.

A loss function is used to train a neural network [43]. The difference between the neural network's current output and the ground truth is often measured by the loss function. Any neural network with layers that are differentiable may have their gradient of the loss function calculated in terms of the parameters. The numerical gradients may be determined more quickly using the backpropagation technique [47]. Stochastic gradient descent (SGD) is a gradient-based optimization approach that may be used to determine the parameters that reduce the loss function.

A. Residual Networks

The architecture of a neural network is a key factor in determining the network's performance. More complicated characteristics may be calculated from a picture using deeper networks. The vanishing gradient issue [48] makes it increasingly difficult to train a neural network as its depth increases.

In order to circumvent many of the problems that beset deep neural networks, ResNet was built. Overcoming the gradient vanishing issue is made easier by using residual connections [48]. ResNet-101, a large ResNet variation with 101 trainable layers, is employed as the neural network's backbone in this study [48].

Aside from solving the picture classification challenge, ResNet may be utilized for a variety of other image processing applications as well. The intermediate layers' outputs may be utilized to represent the picture at a higher level. A feature extractor rather than a classification network, ResNet is employed in this manner.




Fig. 2 : A residual network cell layout

IV. DEFECT DETECTION SYSTEM

In this part, a technique for detecting casting faults in X-ray pictures is suggested. Automated inspections may benefit from the suggested method, which concurrently detects and segmented defects. At its core is the Mask R-CNN architecture, which serves as the foundation for the detection system [12]. Figure 3 shows the four components that makeup the fault detection system.

To begin, the first module builds a high-level feature representation of the input picture. With the second module, a CNN is used to identify areas of interest (RoIs) in the picture using the featured image as a starting point. CNNs are used in the third module to categorize the items in each region of interest (RoI). The picture segmentation performed in the fourth module is used to create a binary mask for each area. The rest of this section explains each module in great depth.

A. Feature Extraxtion

As the first step in the proposed defect identification method, pixels in a picture are transformed into a high-level feature representation. The VGG-16 architecture is used by several CNN-based object identification algorithms [10, 39, 49]. Recent research, on the other hand, has shown that more current feature extractors provide superior outcomes [41]. On the GDxray Castings dataset, we demonstrated that an object identification network with the ResNet-101 feature extractor outperformed a network with a VGG-16 feature extractor in terms of bounding-box prediction accuracy [42]. As a result, the feature extraction module uses the ResNet-101 architecture as its backbone.

There are around 27 million parameters and 101 trainable layers in the ResNet-101 feature extractor. The limited GDXray dataset makes it unlikely that the network can be taught to extract useful characteristics from input photos. CNN-based feature extractors have the unique ability of generating features that may be used to a variety of image processing applications.

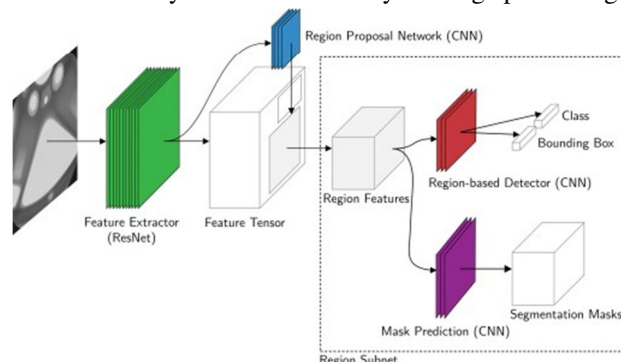


Fig. 3 depicts the suggested defect detection system's neural network design. Using ResNet-101, a network of convolutional neural networks, a detector based on regions and a network for region-based proposal, this system is comprised of four main components.

For example, the suggested casting defect detection system is trained using the massive ImageNet dataset [50] in order to take use of this characteristic. Many characteristics are extracted during training, but only a few are applicable to the more straightforward job of detecting casting defects. Training the object identification network using GDXray Castings dataset, the system learns which attributes are most predictive of casting faults and discards the rest. Because it's quicker to get rid of unnecessary features than it is to learn new ones, this method tends to perform effectively.

B. Region Proposal Network

The region suggestion network is the second module in the proposed defect detection system (RPN). Any size feature map may be sent into the RPN, and it will return a list of rectangular object recommendations, each with a score reflecting the chance that the area includes a certain kind of item. The ResNet-101 feature extractor output is convolved with a tiny CNN to provide region recommendations. $n \times n$ spatial windows of the ResNet-101 feature map are fed into this little CNN. At a high level, the RPN output is a vector representing the current sliding position's bounding box coordinates and likelihood of objects.

Boxes with anchors: Casting flaws may be found in a variety of sizes and shapes. In order to correctly detect casting faults, a variety of box shapes must be evaluated at various points in the picture. Anchor boxes are a typical name for these structures. It is possible to include every possible item in the picture by using anchors of varying aspect-ratios and scales. The RPN calculates the probability of each anchor box containing an item at each sliding point.

There are 15 anchors in each sliding position thanks to the usage of anchor boxes with three scales and five aspect ratios. A larger picture will have a greater number of anchors. Each of the 15WH anchors in the usual convolutional feature map (generally $W \times H$) is about 42,400 by 42,400 pixels.




Fig. 4 Maps from the "conv 4" ResNet feature extractor's final layer. the original Xray picture scaled and padding, (c) a feature map that purports to capture horizontal gradients features that seem to catch long straight vertical edges (c) and features that appear to capture hole-shaped objects (d).

When designing the anchor boxes, they are sized and scaled to fit the dataset's items. People and automobiles are often detected using anchor boxes with areas of 12,822, 25,622, and 5122 pixels with aspect ratios of 1:1, 1:2 and 2:1. [39] In contrast, many casting faults in the GDXray dataset are less than 20 20 pixels. Since the smallest anchor box is 16x16 pixels, it is decided that this is the size. The following ratios are used: 1:1, 1:2, and 2:1.

It is possible to employ factors of 1, 2, 4, 8, and 16. Because the majority of the flaws in the dataset are less than 64 64 pixels, it is reasonable to assume that employing scales 1, 2, and 4 will be enough for the defect identification challenge. Pretrained on a dataset with numerous huge items, the object identification network incorporates greater sizes in order to prevent constraining the system during this phase.

Architecture: If an item is within a box at any given position, the RPN can forecast the location of this box and its likelihood of being contained. To begin with, the $n \times n$ input from the feature extractor is transformed into a 512-d feature vector by means of a fully connected neural network layer. Box-regression (loc) and box-classification (cls) layers both receive this feature vector as input (cls). For each anchor box, the class layer generates 2k scores estimating the likelihood of object and non-object. It contains 4k outputs, which represent the coordinate adjustments for each of the k boxes.

An in-depth explanation of the neural network's architecture may be found in [39]. An anchor box's objectness score is a measure of how likely it is that the box contains an object. It is possible to think of this objectness score as a technique to separate the foreground from the background of a picture. Proposals for regions are picked from the top n anchor boxes based on their objectness score at the conclusion of this round. Training: Developing the RPN includes minimizing a combined loss in classification and regression, as has been detailed. Using the intersection over union (IoU) metric, the best-matching defect bounding box b is picked for each anchor, a. a fault is presumed to be present, and the ground-truth class label $p = 1$ is applied. A vector encoding of the box b with regard to the anchor a is constructed and designated " in this example (b; a). Assuming a is free of defects, then the class label is assigned to $p = 0$. A bounding box's distance from its proposed area is captured during training by the location loss function L_{loc} [39].

With regard to an object's center coordinates, a is encoded as (x_c, y_c) , $\log w$ and $\log h$, where the box's width and height are defined as x_c and y_c respectively. The width and height of the anchor are w_a and h_a . According to Figure 9, a predicted bounding box and the actual ground truth are shown. The cross-entropy loss function, LCE, is used to represent the classification loss as a function of the predicted class $f_{cls}(I; a,)$ and p_a . The weighted sum of the location-based loss and the categorization loss is used to indicate the overall loss for a [41]: To put it more succinctly: $L(I; 5)$ is equal to the sum of the values of L_{loc} and L_{cls} . are weights selected to balance the loss of localisation and categorization. (5) is averaged across the collection of anchors and minimized with respect to parameters to train the object identification model.




Fig. 5 An analysis of the Microsoft Common Objects in Context dataset yielded these top 50 locations of interest. Blue indicates the projected areas of interest, whereas red indicates the places of ground-truth casting fault.

It's important to be able to apply what you learn. In order to use transfer learning, the RPN is a good choice since it detects "regions of interest" (RoIs) in pictures rather than recognising specific categories of objects. Using information gained in one context to increase generalization in another context is known as transfer learning. For domain-specific tasks with minimal training data, transfer learning has been shown to be especially useful [51, 52].

When using a large dataset to train an object detection network, the RPN learns to identify portions of the picture that are most likely to contain an item without making a distinction based on the object class. First, the object identification system is pre-trained on a big dataset with many different types of objects, such as the COCO dataset from Microsoft [53]. Casting flaws are among the interesting sections of an X-ray picture that are promptly detected using the RPN from the trained object detection system. This is the RPN's output after it was trained only on the COCO dataset (Figure 5).

C. Detectors Based on Regions

Only a limited number of area ideas are currently being selected by the flaw detection system. To categorize casting flaws, a region-based detector (RBD) is employed to fine-tune the bounding box coordinates for each area. Based on the Faster R-CNN object detection network, the RBD is the regressed bounding box shape is used to trim the output of the ResNet-101 feature extractor. The problem is that the bounding box determines the size of the input. RoIAlign is used to transform the input to a fixed-length feature vector to overcome this problem [12]. Because it uses a H:W grid of smaller subwindows, RoIAlign can align the large h:w RoI window. The precise values of the input features are computed via bilinear interpolation [54] at four regularly sampled positions inside each sub-window. RoIAlign is further described in [49] for the benefit of the reader. Regardless of the input size, the generated feature vector has spatial dimensions H W.

Using the RoIAlign layer as a starting point, a series of convolutional and fully-connected layers is applied. Two convolutional layers and two fully linked layers make up the RBD in the proposed defect detection system. Two output vectors are generated by the last fully linked layer: In the first vector, each of the K object classes and a background class are represented by a probability estimate. For one of the K classes, the second vector encodes revised bounding-box coordinates. Similarly to the RPN, the RBD is trained by minimizing a combined regression and classification loss function. Detailed information on the loss function and training procedure is provided in [49].

Segmentation of Defects: By anticipating a segmentation mask for each RoI, instance segmentation is carried out. Another CNN, known as the instance segmentation network, is used to predict segmentation masks. The output of the feature extractor is used to crop a block of features for input into the segmentation network. To encode M binary masks of resolution 28x28, one for each of the K classes, the instance segmentation network produces a 28x28xK-dimensional output for each RoI.

The output of the instance segmentation network is fed into a per-pixel sigmoid function during training. The average binary cross-entropy loss is defined as Mask. Mask is only specified on the i-th mask for a RoI associated with ground-truth class I (other mask outputs do not contribute to the loss). This definition of Mask enables the network to create masks for every class without rivalry amongst the various classes. This definition of Lmask. As a result, minimization of the combined RBD and mask loss may be used to train the instance segmentation network. Each class is projected to have one mask at the time of testing (K masks in total). It is, however, just the i-th mask that is employed, where I is the projected class by the RBD classification branch. It is then reduced in size to the RoI size and binarized at a threshold of 0.5, resizing the 28 x 28 floating-number mask.

Figure 6 shows a few examples of masks.

V. IMPLEMENTATION DETAILS AND EXPERIMENTAL RESULTS

As previously discussed, casting defects may be detected using a defect detection system. The GDXray dataset is the primary source of training and evaluation for the model. Car components, such as aluminum wheels and knuckles, make up the majority of this dataset's Castings series. Bounding Boxes that fit snugly around the casting faults are used to identify them. From 256 x 256 pixels up to 768 pixels wide, the photos in the collection are of varying sizes. According to [42], training and testing data are split in the same manner as indicated in this paper.

A. Training

Similar to several recent object identification networks, such as Faster R-CNN and Mask R-CNN, the model is taught 12 and 39 respectively. However, the GDXray dataset contains just a few photos of casting faults of varying sizes, which necessitates a lot of changes.

The longest edge of an image must exceed 768 pixels. 768 x 768 pixels are the final image dimensions after padding with black pixels. In addition, the visuals are flipped horizontally and vertically throughout training. During training and testing, no extra preprocessing is done to the photos.




Fig.6 Floating point masks in action. Predicted bounding boxes and segmentation masks are shown in the top and bottom rows, respectively. As expected by the instance segmentation module, masks are shown here at a resolution of 28 28 pixels. The segmentation masks in the proposed defect detection method are scaled to the projected bounding box's form and size.

Figure 6 shows how transfer learning is used to minimize training time and enhance model accuracy. The weights of a ResNet-101 network trained on the ImageNet dataset are used to initialize the ResNet-101 feature extractor. The COCO dataset is then used to train the fault detection algorithm [53]. As explained in [41], the model's learning rates are changed during pre training. This huge dataset guarantees that each model is properly seeded to locate common items before it is used for fault localization training. NVIDIA K80 GPUs are used to train the COCO dataset. There are two pictures per GPU and 100 sampled RoIs per image, with a 1:3 ratio of positive to negative. It is assumed that a ROI is positive in *Faster R-CNN* and negative in all other cases when it has IoU with a ground-truth box at least 0. After that, the GDXray dataset is used to fine-tune the flaw detection system: The RBD and instance segmentation output layers have been enlarged to include predictions for the dataset's 80 object classes.

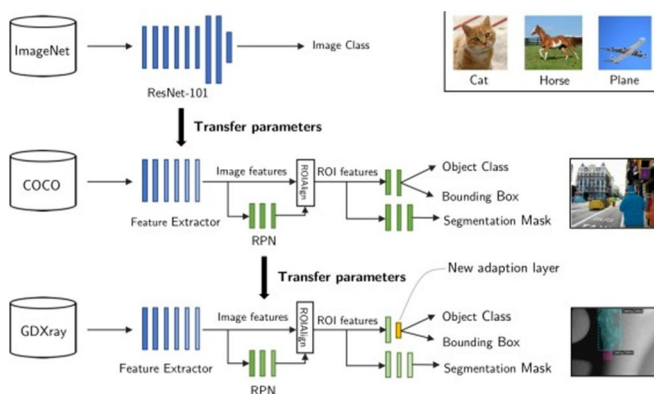


Fig.7 Using GDXray with transfer learning to train a suggested flaw detection system.

Defect and Background. Casting. Defect and Background Casting. Defect and Background Casting. Defect we start using a Gaussian distribution with a zero mean and a 0.01 standard deviation to generate the weights for the resized layers. While all parameters are held constant except for output layer settings, the GDXray dataset is used to train the defect detection system over an 80-epoch period. An additional 80 epochs of training are performed without any weights being held constant.

B. Inference

To test the flaw detection system, we used an Intel Xeon E5 desktop computer with eight cores running at 3.6 GHz, 32 GB of RAM, and a single NVIDIA GTX 1080 Ti GPU (GPU). The models are compared with and without the GPU. The RBD evaluates the top 600 area ideas for each picture based on objectivity scores from the RPN. Only the BARD's top 100 bounding boxes are projected to have masks. As a part of the proposed defect detection system, a bounding box prediction accuracy test is conducted with and without an instance-segmentation module. The GDXray Castings dataset is used to test the system's accuracy.

Each picture in the test set has been separately processed (no batching). The mean of average precision (mAP) is a statistic used to assess each model's accuracy [55]. Using the IoU metric, it is possible to verify whether a bounding box prediction is accurate. According to the formula of $\text{IoU} = \frac{\text{area}(\text{Bp} \cap \text{Bgt})}{\text{area}(\text{Bp} \cup \text{Bgt})}$ (6), the area of overlap between the predicted and ground truth bounding boxes Bp and Bgt must be more than 0.5 in order to be regarded as a valid detection. Predictions made for both the bounding box and segmentation mask are presented in the average precision (mAPbox and mAPseg) (mAPmask).

C. Main Results

Casting defects may be detected with a mAPbox of 0.957, which is higher than the previous state-of-the-art performance. The trained defect detection system's outputs are displayed in Figure 8. Compared to the Faster R-CNN model from [42], the suggested defect detection method has better accuracy and evaluation time. The increased accuracy is attributed in large part to the advantages of jointly predicting bounding boxes and segmentation masks, according to this theory. It takes the same amount of time to assess both systems on the CPU, while the suggested system is quicker when examined on a GPU. Mask R-CNN's advantage over Faster R-CNN's processing capabilities is most likely to blame for the discrepancy between the two implementations. [42] SSD ResNet-101, for example, has a substantially shorter assessment time than the flaw detection system presented in this work [43, 44]. [42]

The suggested defect detection system only achieves a mAPbox of 0.931 when trained without the segmentation module. When the proposed defect detection system is trained concurrently on casting defect detection and casting defect instance segmentation tasks, the bounding-box prediction accuracy of the system is greater.

In the literature, this is a well-known advantage of multitasking [12, 39, 49]. A common representation of the input picture (from the feature extractor) improves the accuracy of the bounding box and segmentation modules [56]. In addition, the suggested method is roughly 12 percent slower while concurrently detecting objects and segmenting images.

When object identification and instance segmentation are done concurrently, the memory requirements for training and testing are also larger than for pure object detection. 9 percent more RAM is needed for object identification and instance segmentation when using a GPU for inference than for object detection alone.

On the GDXray Castings test dataset, the suggested method produces just a few misclassifications. Two examples of misclassification are given in this section. An X-ray picture shows a possible fault in the X-ray equipment, which is detected by the suggested defect detection system. This flaw is not part of the GDXray castings dataset, hence it is classified as a mistake. By deleting bounding box predictions that are outside of the object being photographed. A misclassification using the IoU measure, where the bounding box coordinates are wrongly anticipated. In this scenario, the label is extremely subjective; the underlying truth may be labeled as two tiny flaws rather than one huge flaw, for example.

VI. DISCUSSION

Many tests were carried out throughout the development of the suggested casting fault detecting system to better understand the system. Here are the findings of the trials, as well as a discussion on how the suggested system works.

A. Accuracy vs. Speed Tradeoff

In most current object detecting systems, there is an inherent tradeoff between speed and accuracy [41]. Speed and accuracy are known to be influenced by the number of RBD area suggestions that are picked [12, 39, 49]. [12, 39, 49]. Reduced chances of missing an item are achieved by increasing the number of area recommendations, however this raises processing demands while assessing networks. 3000 region ideas are often used by researchers to successfully identify complicated objects. For the defect identification job, a variety of experiments were undertaken. Accuracy, assessment duration and number of area ideas. These findings suggest that the usage of 600 region suggestions provides a reasonable compromise between speed and accuracy.

B. Information Needed

To train an accurate classifier, like with many other deep learning tasks, a substantial quantity of labeled data is required. Each time the defect detection system is trained, it is given a varied set quantity of training data to see how it affects model accuracy. Each trained system's mAPbox and mAPmask performance is evaluated.

Increasing the size of the training dataset from 1100 to 2308 photos improves the accuracy of object recognition (mAPbox) and segmentation (mAPmask). Comparing the performance of instance segmentation and defect identification, it seems that a substantial quantity of training data is needed.

C. Supplementing the Training Equipment

In certain circumstances, training data augmentation has been shown to improve prediction accuracy by increasing the amount of training datasets [12, 49]. This section examines the impact of several commonly used picture enhancement methods on the testing accuracy.

At training time, photos are flipped horizontally in a random fashion. When training CNNs, the label of an object is agnostic to horizontal flipping, making this strategy advantageous. However, vertical flipping is less prevalent since numerous items like vehicles and trains seldom appear upside down. Using Gaussian blur to decrease random noise generated by the camera or picture compression method is a widespread practice in image processing.

In this work, the Gaussian blur augmentation approach used a Gaussian kernel with a standard deviation of 1.0 pixels for each training picture. The trained model's resistance to noise in the input pictures may also be improved by adding Gaussian noise to the training images [58]. Images are manipulated to include 0.05 of the picture dynamic range's standard deviation as zero-mean Gaussian noise in this investigation. The dynamic range of an image is defined as the range between the darkest pixel and the brightest pixel in the picture. This approach is utilized exclusively for instructional purposes, while the original photographs are used for testing purposes.




Fig. 8 Examples of casting faults detected by the proposed technique.

D. Learning Transfer

According to this research, transfer learning is primarily responsible for the system's great prediction accuracy. Before training on the GDXray Casting dataset, the system generates relevant picture characteristics and excellent area recommendations for GDXray casting images. As a result, the ResNet feature extractor is initialized using weights that have been pre-trained on the ImageNet dataset and then trained on the COCO dataset for defect identification. The impact of transfer learning is assessed using three different training methods: Scheme (a) uses the GDXray Castings dataset to train the proposed defect detection system without using the ImageNet or COCO datasets for pretraining. Random weights are assigned to the feature extraction layers via Xavier initialization [59]. While the identical training procedure is performed in Training Scheme (b), this time the weights for the feature extractor are pre-trained using weights from the ImageNet dataset. The "Defect Detection System" section describes how training scheme (c) leverages pre trained ImageNet weights COCO pretraining.

The GDXray Castings test dataset is used to assess each trained system in Table 1. The resultant system has a low mAPbox of 0.651 on the GDXray Castings test dataset as a consequence of training method (a). Pretrained ImageNet is used to initialize the feature extractor in training scheme (b), resulting in a higher mAPbox value of 0.874 on the same dataset. The mAPbox of 0.957 is achieved by completely exploiting transfer learning in training scheme (c). On the GDXray Castings training dataset, the trained systems are also given in. As can be observed from the results, transfer learning does not alter the system's capacity to fit to the training dataset, but rather its ability to generalize predictions to unseen pictures.

E. Segmentation of weld defects using multi-class learning

The capacity to apply a model to a wide range of activities is very useful in a variety of contexts. Both the GDXray Castings dataset and the GDXray Welds dataset were used to train the proposed defect detection method. Welds in this collection range in width from 3176 to 4998 pixels and are annotated using high-resolution X-ray pictures. A total of 704 pictures are generated by horizontally dividing each high-resolution image into 8 smaller photos for testing and training. For the training and testing sets, 80 and 20 percent of photos are randomly given. Instead of castings, this dataset has just been marked with masks for welds. Closed forms in the masks are identified using a binary border-following method [61] and the bounding boxes are then wrapped around each shape.

Welding and Casting defect detection training sets are used to teach the defect detection system. On the GDXray Welds test dataset, the defect detection system achieved a segmentation accuracy mAPmask of 0.850 for casting flaws and welding faults. Figure 9 illustrates some of the predictions that may be made. Even with just 88 high-resolution photos in the GDXray Welds dataset, it is possible to say that the identification and segmentation of welding flaws is very accurate. The combined training on both datasets does not increase the accuracy of casting flaw identification.




Fig. 9 An picture from the GDXray Welds series was used to compare weld flaw detections to actual welds to provide the ground truth. There are no ground truth bounding boxes shown since this is essentially an instance segmentation operation.

F. The Zero-Shot Learning Approach to Defect Detection in Other Datasets

Defects in a broad variety of items should be able to be classified by a competent defect detection system. Because it can identify problems in objects that are not in the training dataset, the defect detection system is considered to be generalising effectively. A zero-shot transfer is a technique in machine learning that involves utilising a previously trained model to generate predictions on a new dataset without having to retrain it. The suggested flaw identification system is put to the test on a variety of X-ray pictures from different sources to see whether it can generalise its results




Fig. 10 An X-ray picture of a jet turbine blade with defect detection and segmentation findings. There were no turbine blade photos in the training set. Four of the five flaws in the picture are appropriately identified by the defect detection system. Top right defects are wrongly classed as both Casting and Welding defects.

Training Scheme	Feature Extractor Initialization	Pretraining on MS COCO Dataset	GDxRay Castings Training Set		GDxRay Castings Test Set	
			mAP_{bbox}	mAP_{mask}	mAP_{bbox}	mAP_{mask}
a	Xavier Initialization [60] (Random)	No	0.970	0.960	0.651	0.420
b	Pretrained ImageNet Weights	No	1.00	0.981	0.874	0.721
c	Pretrained ImageNet Weights	Yes	1.00	0.991	0.957	0.930

TABLE I

Quantitative results indicating the influence of transfer learning on the accuracy of the trained defect detection system. The bounding box prediction accuracy mAP_{bbox} and instance segmentation accuracy mAP_{mask} are reported on the GDxRay Castings training dataset and GDxRay Castings test dataset.

As illustrated in Figure 10, the system accurately detects a number of faults in an X-ray picture of a jet turbine blade. In all, there are five casting flaws in the jet turbine blade, and four of them have been appropriately recognised. There are no turbine blades in the GDxRay dataset, therefore the system's failure to find a casting fault is not unusual. However, the system's capacity to detect faults in photos from various datasets shows that it is generalizable and resilient.

VII. SUMMARY AND CONCLUSION

Metal casting faults may be detected and segmented simultaneously with this technology. The suggested system's ability to concurrently identify and segment defects makes it appropriate for a wide variety of automated quality control applications. For instance : on the GDxRay Castings dataset, the suggested defect detection method has a mean average precision (mAP_{bbox}) of 0.957%, which outperforms the state-of-the-art performance for flaw identification. Transfer learning, dataset augmentation, and multi-task learning are all used in the development of this highly accurate system. Through thorough ablation testing, we were able to quantify the advantages of each of these paradigms.

It is possible to accurately identify casting and welding faults using the technique proposed in this paper. Using the same network to train on other materials like wood or glass might be a future goal. Defects in several materials might be detected using the suggested defect detection system since it was built for multiclass detection.

Additive manufacturing applications might potentially benefit from the flaw detection technique proposed in this article. In a real-world production environment, the suggested defect detection system is precise and fast enough to be helpful. It is, however, difficult and time consuming to develop the system from scratch. Working on standardizing the representation of these models in the future will make it simpler to disseminate the trained models. On Other Datasets, Using Zero-Shot Learning for Defect Detection.

Classifying flaws for various types of objects is essential for an effective defect detection system. Generalization is achieved when a flaw detection system can identify problems in items that were not included as part of the training dataset. A zero-shot transfer is a technique in machine learning that involves utilizing a previously trained model to generate predictions on a new dataset without having to retrain it. The trained system is tested on a variety of X-ray pictures from various sources in order to evaluate the generalization features of the proposed fault detection method.

As illustrated in Figure 10, the system accurately detects a number of faults in an X-ray picture of a jet turbine blade. Four of the five casting flaws in the jet turbine blade have been accurately detected. The lack of jet engine turbine blades in the GDxray dataset explains why the algorithm failed to detect one of the image's casting faults. In any case, the system's generalizability and resilience are shown by its ability to detect faults in photos from various datasets.

VIII. ACKNOWLEDGEMENTS

NIST's Smart Manufacturing Systems Design and Analysis Program at NIST, the US Department of Commerce's national institute for standards and technology, provided assistance to the authors. Stanford University received financial support for this project under the NIST Cooperative Agreement 70NANB17H031. This article points out a few commercially available options. However, this does not always mean that these items are the best available for the task at hand or that NIST recommends or endorses them. As a further point of clarification, this information does not necessarily represent NIST or any other supporting U.S. government or business entities' views on any subject matter discussed in this material.

REFERENCES

- [1] The Metal Casting: Principles and Practice by T. R. Rao (1981). New Age International is a worldwide network of independent spiritual centres.
- [2] 2K, "The Future of Manufacturing: 2020 and Beyond," p. 375–383 in International Journal of Research in Advent Technology, eds. R. Rajkolhe and J. Khan, "Defects, causes and their cures in casting process: A review," 2014.
- [3] There are several ways to improve the identification of casting flaws using wavelet analysis; for example, a wavelet algorithm may be used to improve the detection of casting faults by applying the wavelet algorithm to a set of castings that have been inspected.
- [4] IEEE Transactions on Instrumentation and Measurement, vol. 62, no. 3, pp. 612–621, 2013, S. Ghorai, A. Mukherjee, M. Gangadharan, and P. K. Dutta. "Automatic flaw identification on hot-rolled flat steel products." IEEE Transactions.
- [5] "Implementing an ultrasonic inspection system to discover surface and interior faults in hot, moving steel using EMATs," vol. 49, no. 2, pp. 87–92. [6].
- [6] M. Lovejoy's Magnetic particle inspection: a practical guide
- [7] Science and business publishing house, Springer.
- [8] "Characterization of air void distribution in asphalt mixtures using x-ray computed tomography," vol. 14, no. 2, pp. 122–129, E. Masad, V. Jandhyala, N. Dasgupta, N. Somadevan, and N. Shashidhar. Why is real-world visual object identification difficult?
- [9] N. Pinto, et al. "Why is real-world visual object recognition difficult?," vol. 4, p. e27.
- [10] 'Rich feature hierarchies for accurate object identification and semantic segmentation,' in IEEE conference on computer vision and pattern recognition, pp. 580–587, R. Girshick, J. Donahue, T. Darrell, and J. Malik.
- [11] J. Dai, Y. Li, K. He, and J. Sun, "R-FCN: Object identification through region-based fully convolutional networks," in Advances in neural information processing systems (NIPS 2016), p. 379–387.
- [12] IEEE International Conference on Computer Vision (ICCV), pp. 2980–2988, 2017.
- [13] D. Mery, V. Riffó, U. Zscherpel, G. Mondragón, I. Lillo,
- [14] H. Lobel, and M. Carrasco, "GDxray: The database of x-ray pictures for nondestructive testing," vol. 34, no. 4, p. 42.
- [15] IEEE, "Background subtraction techniques: a review," in 2004 IEEE international conference on Systems, man and cybernetics (SMC), vol. 4, pp. 3099–3104, IEEE.
- [16] Digital radiography defect detection technique for porosity in magnesium castings [15]: V. Rebuffel, S. Sood, and B. Blakeley
- [17] 145–148, "Weld flaw identification in industrial radiography based digital image processing," by Nacereddine, Zelmat, Belaïfa, and Tridi, et al.
- [18] Automatic diagnosis of various welding faults in radiographic pictures, G. Wang and T. W. Liao, "Volume 35, No. 8, pages 519–528."
- [19] V. Kaftandjian, A. Joly, T. Odievre, Courbiere, C. and Hantrais, C., "Automatic identification and characterisation of aluminium weld defects: comparison between radiography, radioscopy and human interpretation," Society of Manufacturing Engineers, pp. 1179–1186, 2004.
- [20] D. Mery, T. Jaeger, and D. Filbert, "A review of approaches for automated detection of casting flaws," vol. 44, no. 7, pp. 428–436. [19].

- [21] Analysis of the properties of aluminium, steel, and superalloys by D. S. MacKenzie and G. E. Totten CRC Press.
- [22] Application of a novel image segmentation approach for the identification of flaws in castings is discussed in the article, "Application of a new image segmentation method to casting defect detection," pp. 431–439
- [23] "A visual inspection system for the surface flaws of highly reflected metal based on multi-class SVM," pp. 5930–5939, in
- [24] X-W. Zang, Y-Q. Ding, Y-Y. Lv, A.-Y. Shi, and R-Y. Liang For example, [23] V. Lashkia, "Defect Detection in X-Ray Images Using Fuzzy Reasoning," p. 261. According to [24] Xie Xie, "A review of current advancements in surface defect identification using texture analysis approaches," volume 7, issue 3, p. 324
- [25] Defect detection on colour texture surfaces: a review by J Kittler, R Marik, M Mirmehdi, M Petrou and J Song, published in the IAPR Workshop on Machine Vision Applications (MVA), pp. 558-567
- [26] Verifying edges for visual examination," vol. 20, no. 3, pp. 315–328, Wen and Xia [26].
- [27] Defect detection using lasers-based morphological image processing has been studied by B. Mallik Goswami and A. K. Datta in the 70th issue of the Journal of the American Society for Testing and Materials.
- [28] "Hierarchical categorization of surface flaws on dusty wood boards," by C.-W. Kim and A. J. Koivo, in the 15th issue of the Journal of Wood Science, pp. 713–721.
- [29] A colour and texture-based wood examination with non-supervised clustering was presented at the SCIA 2001 conference by M. Niskanen, O. Silvén, and H. Kauppinen.
- [30] There are a number of researchers that have worked together on a project called "Identifying and Locating Surface Defects in Wood: Part of an Automated Lumber Processing System," which has the working title "Identifying and Locating Surface Defects in Wood.
- [31] Comparing alternative filter sets for flaw identification in textiles, F. Ade, N Lins and M Unser, 14th International Conference on Pattern Recognition (ICPR), vol. 1, pp. 428–431, ICPR 2001, pp. 428–431.
- [32] Toriumi, K. and S. A. Hosseini Ravandi, S. A., "Fourier transform study of plain weave cloth look," vol. 65, no.11, p. 676–683.
- [33] Research frontier: a spatio-temporal model of memory formulation in the brain's spatio-temporal model, pp. 56–68.
- [34] A. Concise and C. B. Proena, "A fractal image analysis system for fabric inspection based on a box-counting algorithm," p. 1887–1895
- [35] Using canny, sobel, and gaussian filter edge detectors to identify weldment and casting defects: a comparative study.
- [36] F. Mirzaei, M. Faridafshin, A. Movafeghi, and R. Faghihi "Automatic fault detection in x-ray testing using computer vision" was presented by D. Mery and C. Arteta at the 2017 IEEE Winter Conference on Applications of Computer Vision (WACV).
- [37] "Going deeper with convolutions" was presented at the 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) at 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (CVPR 2015).
- [38] A steel arch bridge was studied in detail in "Experimental and Analytical Modal Analysis of Steel Arch Bridge," by W-X Ren, T Zhao, and I E Harik in "Volume 130, Number 7, Pages 1022–1031.
- [39] For real time object recognition, "Faster R-CNN: Towards Region Proposal Networks," by S. Ren, K. He, R. Girshick and J. Sun, was presented at NIPS 2015.
- [40] Single shot multibox detector (SSD): W. Liu, et al. "SSD: Single shot multibox detector," in 14th European Conference on Computer Vision (ECCV 2016), pp. 21–37, 2016.
- [41] "Speed/accuracy trade-offs for current convolutional object detectors," in IEEE Conference on Computer Vision and Pattern Recognition, p. 41-42 (CVPR 2017).
- [42] IEEE, "Automatic localisation of casting faults with convolutional neural networks," in 2017 IEEE International Conference on Big Data (Big Data 2017), pp. 1726–1735 (Measuring casting defects with convolutional neural networks), IEEE.
- [43] Convolutional neural networks, J. Wu, et al. An English translation is available at <https://wujx/teaching/15 CNN.pdf>.
- [44] There are 44 citations for I. Sobel's "An isotropic 3x3 image gradient operator" on pgs. 376–379.
- [45] In the 27th International Conference on Machine Learning (ICML), V. Nair and G. E. Hinton, "Rectified linear units enhance limited boltzmann machines," p. 807–814.
- [46] "Deep reinforcement learning with successor characteristics for navigation across comparable environments," by J. Zhang, J. T. Springenberg, J. Boedecker, and W. Burgard, pp. 2371–2378.
- [47] "Backpropagation over time: what it accomplishes and how to achieve it," vol. 78, pp. 1550–1560, by P. J. Werbos [47].
- [48] This work was presented at the IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2016) by KHe, X Zhang, S Ren, and J Sun as "Deep residual learning for picture recognition."
- [49] In 2015 IEEE International Conference on Computer Vision (ICCV), pages 1440–1448, R. Girshick presents "Fast r-CNN."
- [50] Imagenet big scale visual recognition task, O. Russakovsky et al., vol. 115, no. 3, pp. 211–252, with J Deng and others, "Imagenet large scale visual recognition challenge," pp. 211–252.
- [51] When it comes to the identification of safety guardrails in two-dimensional pictures, transfer learning and deep convolutional neural networks may help.
- [52] "Deep Transfer Learning for Image-Based Structural Damage Recognition," Computer-Aided Civil and Infrastructure Engineering, vol. 33, pp. 748–768. [52]
- [53] Microsoft COCO: Common objects in context is a paper published at the European Conference on Computer Vision (ECCV 2014) in Springer. [53] T-Y. Lin and his colleagues presented their findings at the European Conference on Computer Vision (ECCV 2014) in Springer.
- [54] On the other hand, [54] "Spatial transformer networks" in Advances in neural information processing systems paves the way for the development of more efficient and effective neural information processing systems in the near future.
- [55] C. Manning, P. Raghavan, and H. Schütze, "Introduction to Information Retrieval," Cambridge University Press, vol. 39, p. 39-55.[...]
- [56] R. Caruana, "Multitask learning," Journal of Experimental Psychology: General, pp. 41–75.
- [57] In "Kernel regression for image processing and reconstruction," by H. Takeda, S. Farsiu, and P. Milanfar, in Image Processing and Reconstruction, vol. 16, no. 2, 2003, pp. 349–366,
- [58] Stability training improves deep neural networks' resilience, as described in "Improving the robustness of deep neural networks by stability training," published in IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2016), (Las Vegas, United States), pp. 4480–4488 in 2016.



- [59] "Deep sparse rectifier neural networks," by X. Glorot, A. Bordes, and Y. Bengio, in the proceedings of the fourteenth international conference on artificial intelligence and statistics, p. 315–323.
- [60] A study by X. Glorot and Y. Bengio, "Understanding the difficulties of training deep feedforward neural networks," in Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, pages 249–256.
- [61] In "Topological structure analysis of digitised binary pictures by border following," by S. Suzuki and K. Abe, vol. 30, no. 1, pp



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)