Deep Learning for Visual Inspection and Classification of Tire Defects

Deep learning is an artificial intelligence method by which a computer model can parse inputs and produce outputs in a way that is inspired by how neural networks in the human brain work. With the advent of self-driving vehicles, facial recognition, and surveillance cameras able to automatically detect suspicious behavior, computer vision is a quickly developing field within deep learning. Visual identification tasks that were once the sole domains of human inspectors are increasingly achievable by intelligent computer vision systems. Tire manufacturers have an opportunity to use this technology to advance their existing equipment investment to provide even better quality control.

Towards the beginning of this century, many tire manufacturers began investing in tire geometry systems using 3D laser profile sensors. These sensors scan each sidewall and the tread of an inflated and rotating tire to create a 3D image. The specialized tire geometry software examines each image to detect, measure, and grade any geometry defects along with radial and lateral runouts. Since these images contain most of the exterior of the inflated tire, why not use them as a second chance to check for visual defects?

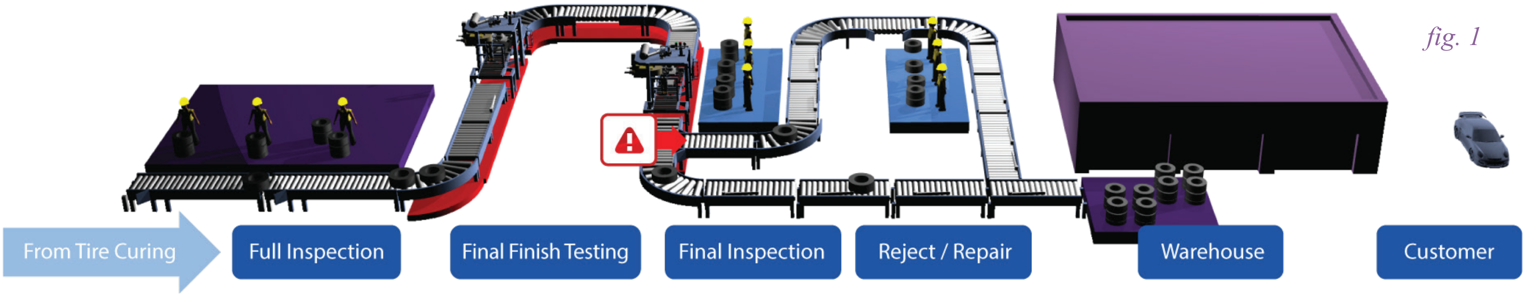

fig. 1

In a typical tire factory, tires are fully inspected, inside and outside, prior to final finish testing by a specialized workforce of tire inspectors. Downstream from this inspection, tires may be rejected or flagged for further inspection by any final finish test machine (e.g. tire uniformity, dynamic balancing, tire geometry, or X-ray machines); otherwise, they are sent to the warehouse without additional inspection. Therefore, it's possible that tires can suffer visual damage by their handling after their initial inspection but still end up in the warehouse (*fig. 1*). These tires eventually get installed on vehicles, where the customer can notice the visual defects, reflecting poorly on your quality control.

For a small investment, a tire factory could add visual inspection software to further examine its already-collected 3D laser profile images. This allows a tire factory to gain even more value from its prior large investment in tire geometry systems.

Visual Inspection from Tire Geometry Images

Tire geometry systems can be located on a tire uniformity machine, a dynamic balancing machine, or even on their own dedicated machine. On the majority of these machines, the tire is inflated and rotating at a constant 60 rpm (or one revolution per second), and typically three sensors—covering the top sidewall, bottom sidewall, and tread— collect measurement data in one revolution (one second). The tire geometry system provides its typical geometry measurements (bulge, depression, lateral runout, radial runout, wobble, etc.) to the machine's control system, which uses them to provide final grading results about each tire.

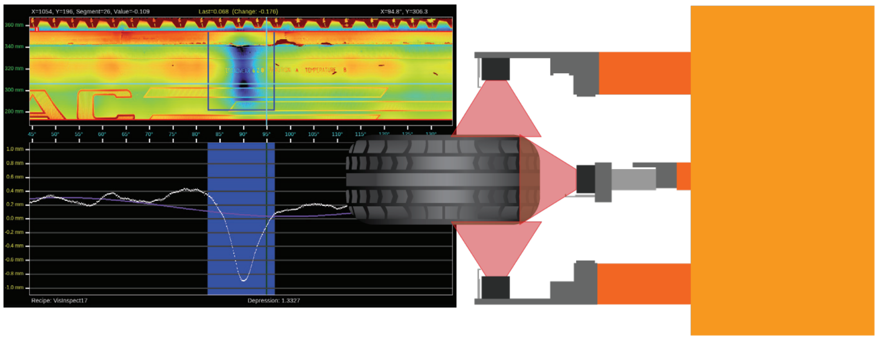

Laser Profile Sensors scan each sidewall and the tread of an inflated and rotating tire, creating a 3D image.

Many of these tire geometry systems exist today using 3D laser profile sensors, providing at least 1000 profiles per revolution and containing a field of view of at least 75mm of tire data per profile. Newer sensors, such as the Gocator 2530 from LMI Technologies Inc., can provide over 4000 profiles per revolution. Other sensors, used mostly for tread/RRO measurements, provide a field of view of up to 400mm per profile. Acquiring more profiles per revolution greatly increases the capability of the visual inspection software to detect even the smallest defects, but tire geometry data collected at 1000 profiles per revolution can still detect many of the larger ones. The image data from the 3D laser profile sensors is stored for each tire so that they could be reviewed later if the downstream final tire inspectors (or classifiers) find fault with the geometry measurements.

The only visual defects measured by the tire geometry systems are bulges and depressions in the sidewalls and bumps or dents in the tread (caused by a tread over-splice or under-splice). Other visual defects appearing in the tire geometry images go undetected. These could be caused by contaminated tire molds (*fig. 2*) or rubber not flowing correctly in the tire mold. Or it could be defects from damage due to tire handling equipment or even damage from improper processing by machines upstream to the tire geometry systems, such as white sidewall grinders (buffers), tire uniformity optimizers (grinding appearance issues), or trimming stations.

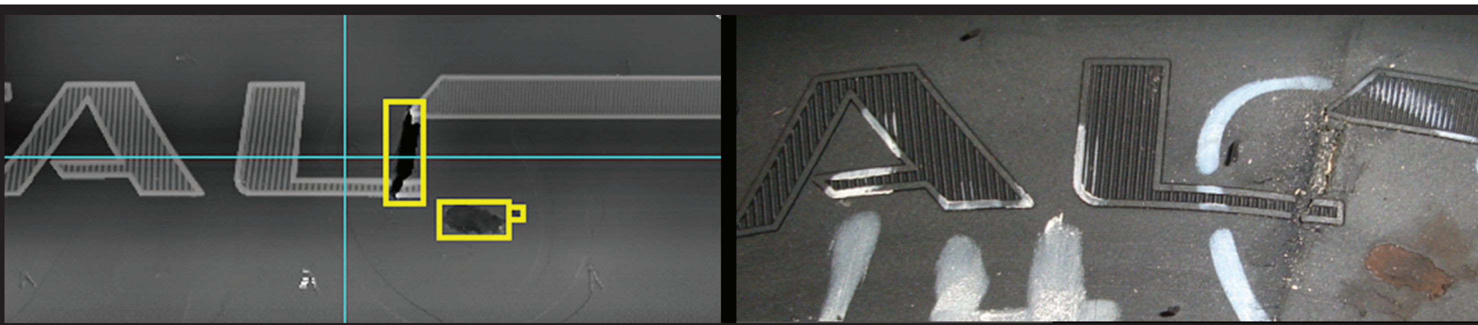

fig. 2

This is where visual inspection software can help. Visual inspection software processes the high-resolution laser images from the tire geometry systems in a completely different manner, focusing on sharpening and flattening the image to provide the clearest picture. It detects and classifies any visual defects found in the tire geometry images. It could also find objects, such as the DOT code (*fig. 3*) and the tread wear indicator bars (even measuring their heights). Such software could provide higher-quality assistance to a human inspector (or classifier), or even replace the most mundane aspects of their work, freeing them up to process many more tires.

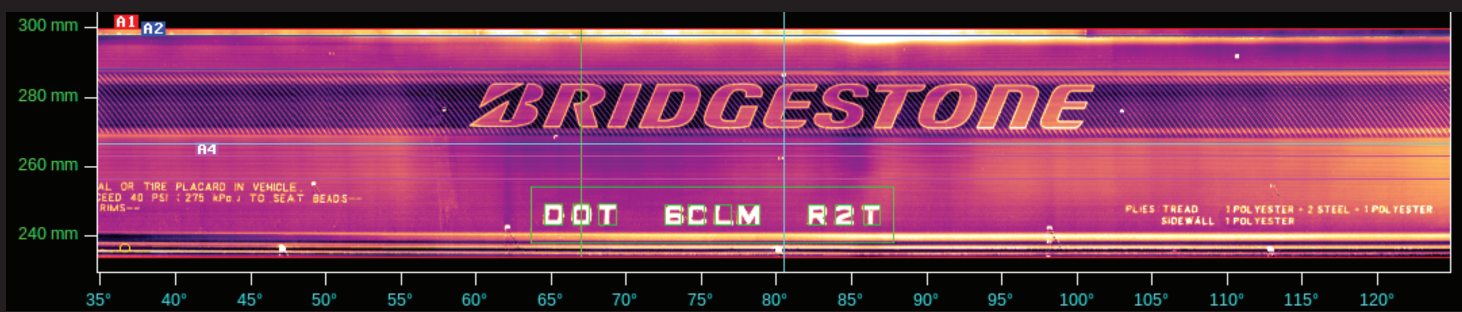

fig. 3

Final Inspection Station

At the beginning of the proposed solution, the visual inspection software provides its results and high-resolution images to the existing human-operated final inspection (or classifier) stations. Each station makes use of a large, widescreen touch panel to display a tire's visual inspection images with all potential defects identified. With the ability to pan and zoom within the images and the actual tire available for further examination, the inspector can manually classify each defect by selecting the appropriate classification from a dropdown list, including the ability to dismiss a defect misidentified by the software. The inspector can also manually identify any other defects, missed by the software, and classify them in the same manner. This feedback serves two purposes: (1) for flagging tires for rejection or repair, as part of improved quality control in the manufacturing process; (2) for training a deep learning model. The model would learn to detect and classify visual tire defects with high precision, becoming adept at the first purpose of flagging tires for rejection or repair and removing that task from the inspector.

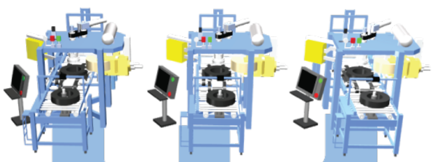

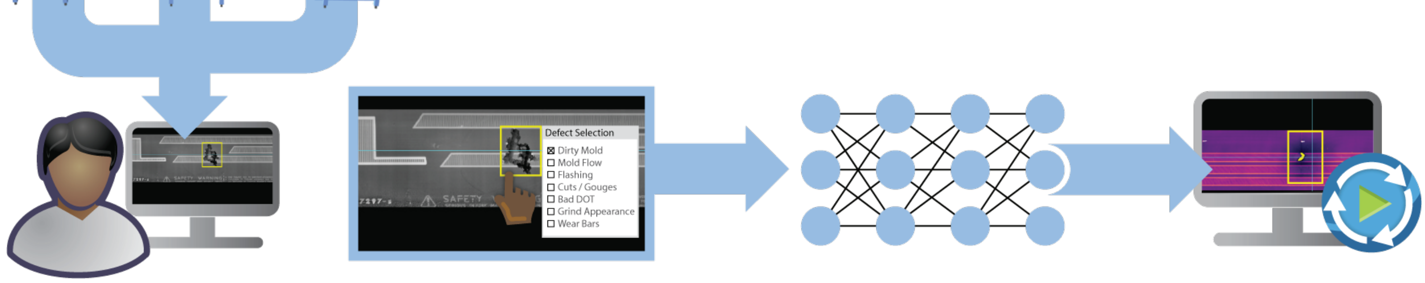

Deep Learning for Visual Inspection

One natural application of computer vision within deep learning is image processing. Within image processing, two common techniques are used: object detection and image classification. Object detection allows computer models to separate parts of images into background and foreground. The foreground parts are then analyzed, allowing for automatic detection, location, and classification of objects within an image. Image classification is when a computer model analyzes a whole image and then makes a decision—for instance, differentiating a photo of a cat from a person. These two techniques can be applied to visual tire inspection, whereby a computer model can input images from tires and output a list of classified defects from these images.

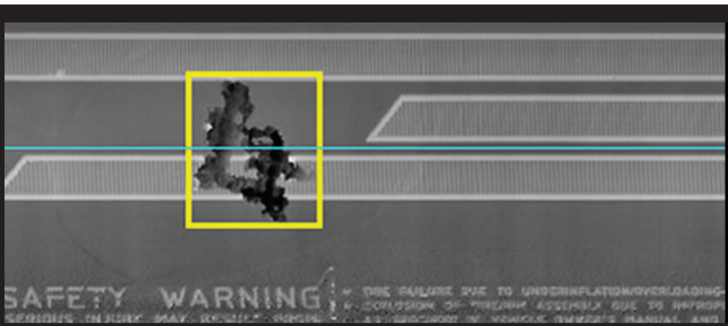

Image Classification : Do images contain defects to be flagged?

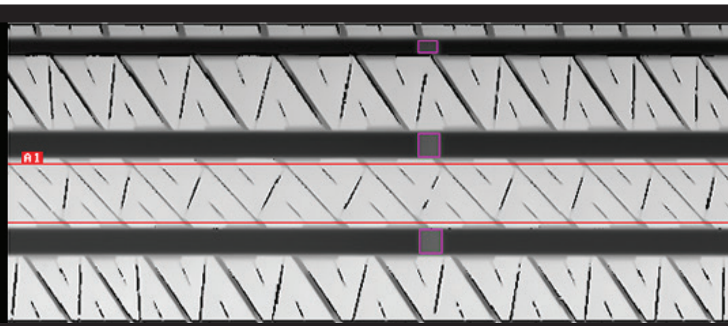

Object Detection : Find objects within tire image (tread wear bars)

Before the deep learning model is trained, visual inspection software uses rule-based image classification algorithms that are specific to finding defects that meet its preprogrammed criteria. Any defects beyond a predefined surface area or volume threshold are highlighted in its resultant images. Defects such as pits or blemishes from curing mold contamination, gouges or similar damage from tire handling equipment, unwanted pin vents or flashing, and damage from improper processing by upstream equipment, such as grind appearance issues, can all be detected using smart algorithms employing statistics and thresholds to differentiate pass from fail. Rule-based software algorithms can provide good results for identifying many visual defects, but these have an accuracy rate only high enough to assist with their classification by the human inspector.

In contrast, a deep learning model would train itself from the defects confirmed in the visual inspection images and classified by the final inspector. Once the accuracy of the model is validated, the visual inspection software is updated to begin using it to detect and classify defects without the assistance of the final inspector (fig. 4). It is common to report the accuracy of a model compared to the accuracy of a human doing the same task. It is not expected that either will be 100% accurate, but these models often outperform their human counterparts.

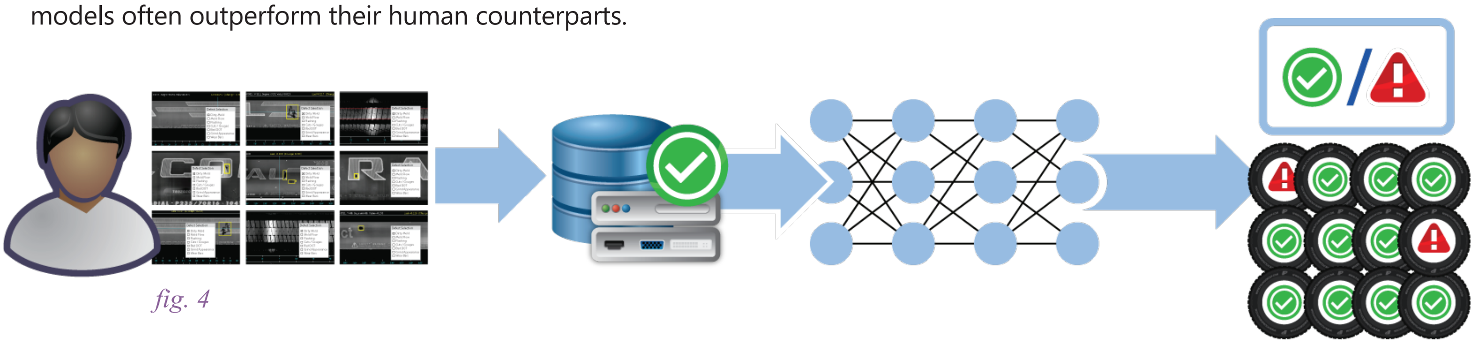

fig. 4

The availability of labeled data (i.e. classified defects) to train these models is one of the biggest limitations industries face when looking to implement deep learning. Input data labeling is a human task and can take armies of labelers dedicated to the task. If labeled data is not being produced as part of a normal production process, it can be prohibitively expensive to generate. That is why tire manufacturers have a unique opportunity to bring this transformative technology within reach by simply committing to capture data as part of an existing inspection process. Multiple tire factories within the same organization can share their captured data, training the deep learning model even faster.

Outside of detecting and classifying visual defects, the visual inspection software has additional image classification and object detection algorithms adding even more value to the tire geometry images. Such algorithms locate and read the DOT code on the sidewall, and, when using tread lasers with enough range, they measure the amount of rubber removed from force or runout grinding (optimization) and even measure the height of the tread wear indicator bars deep in the tread grooves.

Once the deep learning model is trained, it could run directly on the tire geometry systems, classifying any defects, reading the DOT code, measuring the height of the tread wear indicator bars, and even checking grind appearance. It could make those tasks part of the final geometry grading results. It could even call for extra 'clean up' grinding from the machine controller if grind appearance shows it is needed.

Since this is software, it can be updated and re-trained to find new problematic visual defects or to support new images from updated 3D laser profile sensors. Think of it like Tesla's Autopilot feature. Even the older model Tesla's can be updated to the latest Autopilot features, including their latest deep learning models trained on the feedback from the ever-increasing amount of Tesla's on the road, with the future hope that any Tesla will be capable of full self-driving.

Conclusion

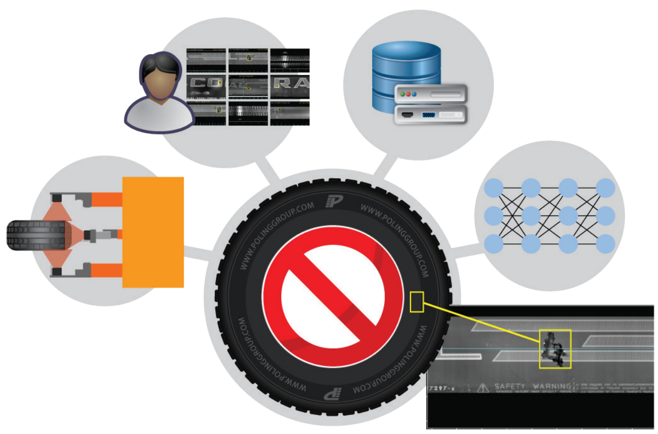

Using images provided by the existing tire geometry systems, tire manufacturers can prevent tires with visual defects from reaching their customers. Adding tire geometry testing using 3D laser profile sensors was a big investment. However, only a small investment in visual inspection software is needed to detect and classify visual tire defects with high precision. All of this is done by using those tire geometry images from your prior large investment. Rule-based software algorithms can provide good results for identifying many visual defects, but with confirmation and classification of those defects by the existing final inspectors, an expansive dataset can be created for training a deep learning model to detect and classify defects without the aid of final inspectors. This dataset guarantees increased quality from the plant floor to the tires on the road.