



Explainable AI-infused ultrasonic inspection for internal defect detection

Adithyaa Karthikeyan, Akash Tiwari, Yuhao Zhong, Satish T.S. Bukkapatnam (2)*

Department of Industrial and Systems Engineering, Texas A&M University, College Station, TX, United States

ARTICLE INFO

Article history:

Available online 20 May 2022

Keywords:

Artificial intelligence
Inspection
Quality assurance
Ultrasonic

ABSTRACT

While AI and imaging technologies are dramatically transforming the process and machine condition monitoring, product inspection remains confined to probing the geometry and surface morphology. Subsurface and bulk inspection remain prohibitively slow and imprecise. This paper presents an explainable AI (XAI)-infused ultrasound imaging approach for rapid detection of artifacts including product defects. The approach led to the discovery of correlated spatial patterns in the images located away from the artifacts. This discovery enabled accurate (> 80%) detection of artifacts that are not discernible with the current image segmentation methods, and it could profoundly impact product quality and (cyber)security assurance technologies.

© 2022 CIRP. Published by Elsevier Ltd. All rights reserved.

1. Introduction

Although the manufacturing community has been at the forefront of adopting the sensor and AI technologies for industrial quality assurance [1–3], inspection of bulk products to detect internal morphological defects remains slow, expensive, and imprecise, even with the latest non-destructive evaluation (NDE) techniques [4,5]. For example, X-ray computed tomography typically takes 5–15 hours to scan products as small as 100 mm³ in volume [4]. Other NDE techniques, such as conventional eddy current, thermography, ferromagnetic, and chemical profiling suffer from limited (a few mm) depth of penetration and poor detection-sensitivity for various materials [4,6]. Consequently, even the quality-critical microelectronics industry bypasses bulk product inspection, and instead it mostly tests the functional performance.

The limitations of the current bulk inspection technologies also impede the assurance of authenticity of products against counterfeiting and cyber-attacks [7]. Reflecting the emergence of a manufacturing-as-a-service (MaaS) paradigm, industrial product authentication technologies are growing at an annual rate of over 10% to exceed \$250B by 2026 [8]. The prominent technologies, employing QR codes, radiofrequency identification, etc., merely mark the product packaging or at best, a product surface [7]. These surface markers are vulnerable to tamper and wear. Attempts have been made to embed alternative geometric (e.g., 3D QR codes [9]) and intrinsic markers (e.g., morphological or microstructure features [10]). As with quality assurance, the limitations of the current bulk product inspection curtail the viability of these innovative authentication methods.

Interestingly, the rapidly increasing resolutions and cost reductions make imaging technologies attractive for bulk inspection to assure both the quality and authenticity of products [11,12]. While the manufactured parts are opaque to optical imaging, reverberation

and attenuation effects severely degrade the signals from various electromagnetic imaging methods, leading to poor quality and authenticity assurance. Ultrasound methods are considered the most promising for bulk inspection of manufactured products [4].

Ultrasound imaging instruments use piezoelectric elements to release ultrasound pressure waves into a product at frequencies of up to 100 MHz [13]. The boundaries of various internal artifacts (e.g., defects and authentication markers) of the product generate echoes of incident ultrasound. Images of the internal structure of the product are reconstructed based on the intensity of echo received from varying depths of the product. Ultrasound imaging is quick and can be performed in real time with 25–100 fps rate, in contrast to other inspection techniques [13].

Although fast and cost-effective, the sensitivity of ultrasound sensors tends to be poor, especially to discern the artifacts (e.g., defects and embedded codes) in many polymeric and composite materials [14]. Infusing the physics-based image reconstruction with recent AI methods can enhance the detection of the product artifacts one leap forward. Among the AI methods, several supervised machine learning methods have been applied to enhance the sensitivity and the prediction power of NDE and ultrasound methods (e.g. [15]). They predominantly used experimental data to learn purely “black box” models to segment or classify the artifacts in ultrasound images. The models are neither based on nor contribute to the understanding of the underlying physical phenomena. Hence, they cannot be generalized beyond the training scenarios.

In contrast, the new class of *explainable AI* (XAI) approaches would not just complement the predictions from powerful machine learning methods, but would help discover the underlying physical processes [16]. Also, different from prior applications, this work aims to employ machine learning methods to enhance the discrimination power of the ultrasound images beyond what purely physics-driven image renderings would achieve.

Recent results in an XAI method called local interpretable model-agnostic explanations (LIME) [16] are adapted to discover the

* Corresponding author.

E-mail address: satish@tamu.edu (S.T.S. Bukkapatnam).

physical relationships captured in these machine learning models. This novel approach is applied to discern, using ultrasound imaging, the markers (artifacts) distributed in polymeric blocks to mimic pores and authentication codes. Even the state-of-the-art ultrasound image segmentation methods [12] could not detect these markers. The results suggest that the present approach can detect these markers to accuracies exceeding 80%. More importantly, this approach provides a robust, fast, and cost-effective solution to the quality and cybersecurity challenges in the emerging MaaS paradigm. The rest of this paper is organized as follows: The details of the explainable AI-infused ultrasound inspection principles are presented in Section 2, an experimental study and results are detailed in Section 3, and the conclusions are presented in Section 4.

2. XAI-infused ultrasound inspection

2.1. Principles of ultrasound inspection

Fig. 1 illustrates operating schematics of ultrasound imaging. It employs multiple piezoelectric elements within a transducer that generate and receive pressure waves. Pressure waves generated from an element travel downward into the product. Their propagation velocity $c(y)$ at depth y is determined by acoustic impedance $Z(y)$ and density $\rho(y)$ of the material at location y of the product as $c(y) = Z(y)/\rho(y)$. Whenever the pressure waves encounter an artifact, causing a major inhomogeneity in media (Fig. 1a), the waves are scattered at boundaries B within the product. A portion of the incident wave is reflected back to the piezoelectric element. Intensity of the reflected wave I_r in relation to incident wave I_i depends on the acoustic impedances Z_1 and Z_2 of the media on both sides of the boundary B as

$$I_r/I_i = (Z_1 - Z_2)^2 / (Z_1 + Z_2)^2 \quad (1)$$

The piezoelectric elements generate a voltage signal $v_r(t)$ at time t by aggregating the reflected wave(echo) it receives as

$$v_r(t) = v_{pe}(t) * (f_m(y) * h_{pe}(y, t)) \quad (2)$$

where ‘*’ denotes a convolution sum (aggregation), $f_m(y) = (\rho(y) - \rho_0) / \rho_0 - 2(c(y) - c_0)/c_0$ is the quantity of interest to capture in the ultrasound image and it captures inhomogeneity within the object [13], ρ_0 and c_0 are the nominal values of the density and velocity, respectively, $v_{pe}(t)$ is the excitation impulse response that aggregates $f_m(y)$ in time, $h_{pe}(y, t)$ is the spatial impulse response of the pulse echo relating the transducer geometry to the acoustic field that aggregates $f_m(y)$ in space.

The ultrasound image is in effect obtained by inverting Eq. (2) to retrieve a surrogate of $f_m(y)$, and this results in a single image line from the crystal element (Fig. 1(b)). The line reconstructions from $v_r(t)$ obtained from multiple piezoelectric elements are juxtaposed to obtain the signal s . The final rendering of the ultrasound image is given by $\xi = \Psi\{G[V\{R_0(s)\}]\}$ [17], where R_0 is a reconstruction operator (that performs the inversion operation outlined earlier), $V\{\cdot\}$ is a nonlinear operator that determines the intensities, G is used to achieve the correct image dimension and $\Psi\{\cdot\}$ is an optional operator for image enhancement and noise filtering.

In reality, the quality of the image ξ depends largely on the acoustic mismatch $Z_1 - Z_2$ at boundary B . A small acoustic mismatch results in




Fig. 1. Operating principle of ultrasound imaging (a) Acoustic mismatch ($Z_1 - Z_2$) at a boundary B (b) Signal received at piezoelectric element 1 and (c) A representative ultrasound image.

poor contrast due to a very small echo generated at this depth. If the impedance of the material is high (as with polymers), the propagating waves, and hence the signals are further attenuated. Also, large acoustic mismatches degrade the resolution of the images below this depth. Additionally, side lobes are produced in the ultrasound beam causing the signals to travel in alternate directions and positioning errors. These phenomena and their manifestation in the ultrasound images for various material systems are not well understood [17]. Recent advances in XAI can help gain a deeper understanding of how these complex phenomena ultrasound manifest in the ultrasound images.

2.2. XAI and local interpretable model-agnostic explanation (LIME)

Despite decades of application of the ultrasound techniques, their physical understandings are not mature yet to correctly predict the various spatial patterns in the reconstructed ultrasound images, even for simple realistic products. The growing suite of sophisticated machine learning methods, such as deep convolution neural networks (CNNs) trained with labelled experimental ultrasound images can detect the internal defects and markers in manufactured products at orders of magnitude higher resolutions than what is possible with the current ultrasound reconstruction methods [15]. However, the high predictive power of these advanced machine learning methods comes at the expense of poor interpretability. While these methods can accurately detect the presence of artifacts, it is not straightforward to locate them. Explainable AI (XAI) and local Interpretation model-agnostic explanation (LIME) approaches [18] are garnering significant interest to explain the complex “black box” machine learning models. They can be adapted to further analyse black box models to locate the internal artifacts of a product from an ultrasound image, as well as to discover the dominant physical mechanisms that cause certain unusual ultrasound features to appear whenever a specific internal artifact is present.

Instead of considering the complex functional relationship that CNN and other black box models capture, LIME aims to explain the relationship local to the various neighbourhoods of the input space (here, a neighbourhood consists of ultrasound images that are similar to each other). This is achieved by generating synthetic perturbed image samples μ_i (see Fig. 2) within the neighbourhood of an input ξ and constructing a highly interpretable linear model using these samples that will hold only for that neighbourhood.




Fig. 2. An illustration of the LIME scheme.

The coefficients β of the resulting linear model can suggest the importance of the various components (here, a group of pixels) of an input (i.e., an image ξ). To explain the prediction of an already trained CNN classifier g around ξ , LIME first partitions ξ into M segments. Then, a binary vector $\xi' \in \{0, 1\}^M$ is used to alternatively represent the original image ξ , where each element of ξ' indicates the presence of a segment of ξ . By randomly setting some elements in ξ' to 0, LIME further generates n different binary vectors, $\mu'_i \in \{0, 1\}^M$, $i = 1, 2, \dots, n$. Each of these generated vectors represents a perturbed image μ_i ($i = 1, 2, \dots, n$) of ξ . Each perturbed image μ_i is weighted by a similarity index $\pi_\xi(\mu_i)$ as

$$\pi_\xi(\mu_i) = \exp\left(-D(\xi', \mu'_i)^2 / \delta^2\right), \quad (3)$$

where δ is the kernel width. The distance $D(\xi', \mu'_i)$ is the cosine of the angle between ξ' and μ'_i . Next, using μ'_i as inputs and the

corresponding outputs from g as the ground truth, LIME trains a linear model, defined by its coefficient vector β^* as:

$$\beta^* = \underset{\beta}{\operatorname{argmin}} \sum_{i=0}^n \pi_{\xi}(\mu_i) [g(\mu_i) - \langle \beta, \mu_i' \rangle]^2 + \lambda \langle \beta, \beta \rangle, \quad (4)$$

where $\mu_0 = \xi$, $\mu'_0 = \xi'$, $g(\mu_i)$ is the output of the CNN model, and λ is the ridge regularization term. The sign and magnitude of each element of β^* indicate positive/negative importance of the presence of the corresponding segment in ξ . Thus, a segment with a positive element implies that without the segment the image is less likely to be classified to have markers, and a negative coefficient implies the opposite.

Additionally, a total of $2^M - 1$ of possible perturbations exist to estimate β^* about every image ξ . The number of segments M is typically of the order of 10^2 , and employing all possible permutations is computationally intractable. Therefore, a much smaller perturbation sample size n is usually chosen [18]. This sampling introduces uncertainty in the resulting feature importance β^* . The following theoretical result was used to guide the selection of n .

Theorem 1. (Perturbed sample size) [18]: *The number of perturbed samples n required to achieve an uncertainty interval width w of feature importance at a user-specified confidence level α can be calculated as*

$$n = 4\varepsilon_j^2 / \left\{ \bar{\pi}_j \times \left[w / \Phi^{-1}(\alpha) \right]^2 \right\}, \quad (5)$$

where $\bar{\pi}_j = \sum \pi_{\xi}(\mu_i) / J$ is the average weight for the perturbations estimated from an initial J samples, ε_j^2 is the empirical sum of squared errors between the LIME linear model and g , weighted by $\pi_{\xi}(\mu_i)$ for $i = 1, \dots, J$, and $\Phi^{-1}(\alpha)$ is the two-tailed inverse normal cumulative distribution function at confidence level α .

It can be observed that, to keep the uncertainty interval w small, n needs to be large. On the other hand, when the error ε_j^2 is large, i.e., when the LIME model cannot accurately approximate g locally and thus cannot provide high-quality explanations, we also need a proportionally larger number of perturbed sample size. Therefore, it is recommended to only consider the explanations obtained from the LIME models that have a sufficient accuracy. Here, criteria such as R^2 (coefficient of determination) can be used to evaluate the performance of the LIME model.

3. Case study and results

To assess the performance of the XAI-infused ultrasound inspection, we considered 2cm cubic test products made of a transparent polymer (Veroclear). The products are embedded with spherical markers of diameter 0.48mm (see Fig 3(a)). The markers are dispersed at both shallow (1–3 mm depth) and deeper (9–12 mm depth) locations, emulating the occurrence of defects (e.g., internal pores and voids) and embedded authentication codes in a product. The products were printed using a Stratasys J750 Polyjet printer and inspected using an ultrasound imaging system from Aixplorer Ultimate with an SL 15–4 transducer. Although the difference in the acoustic impedance between Veroclear ($\sim 3 \times 10^6 \text{ Pa s m}^{-1}$) and air (415 Pa s m^{-1}) assures sufficient contrast to identify the markers in the ultrasound images, the large acoustic attenuation coefficient of Veroclear ($\sim 170 \text{ dB m}^{-1}$ compared to $< 10 \text{ dB m}^{-1}$ for metallic materials at 2MHz) severely impedes the discernibility of markers located more than 5 mm below a surface [14].




Fig. 3. (a) Isometric, front, and right-side views of the 2cm cubic test product embedded with a distribution of spherical markers. (b) Photograph of a 3D printed cube made of Veroclear material.

The experiments consisted of moving the components at a speed of $\sim 1\text{mm/sec}$ relative to the ultrasound instrument to collect image-frames at 18 different settings of frequency (5, 7.5, and 12 MHz), input power (0dB, -5dB and -10dB) and focal position (Top and bottom)

(Fig. 4a). Each image frame, shown as a light blue cross-section plane Fig 4(b), captures the markers located underneath the instrument at a specified time during a scan. The numbers indicated in red at the bottom of Fig. 4(b) represents the time (in seconds) at which the instrument traverses the cross section. Also, the higher the frequency setting, the higher are the resolution of the scan and signal attenuation. The higher the power, the greater is the scan illuminance (and scattering effects). The focal position indicates the depth that requires the best resolution. Also, based on the marker patterns, the ultrasound image-frames can be generally labelled into three classes, namely, those that capture the shallow markers (Region A), deeper markers (Region C) and no markers (Region B). The markers are distributed over the depths of 0.5 – 3mm from the top section for Region A, and 9–12 mm for Region C.




Fig. 4. (a) A schematic of the ultrasound scanning experimental setup. (b) A representation of the product scanning showing image frames and different regions of marker distribution, and the marker coordinates.

To apply the novel XAI-infused ultrasound method, CNN models are employed to classify these ultrasound image frames depending on whether a frame contains embedded markers. A CNN model extracts a compact set of features from an image via convolution. Then, it associates these features with the prespecified class of the image. In this way, it learns the feature pattern of each class and can classify a new image into the most likely class. A good classification accuracy may indicate that the CNN has learned to distinguish the ultrasound image frames that contain the markers.

One of the challenges encountered here is that the presence of markers is visible, even partially, only up to a depth of 3mm from the top of the cube. The markers embedded at depths greater than 3mm from the top of the cube such as those present in Region C and bottom sections of Region A, are not observable even with advanced image segmentation methods [12]. However, the presence of markers can modify the ultrasound image intensity patterns at locations far away from where the markers are placed. These correlated patterns across the image, caused by the presence of the marker, serve as a signature that can be learned by a CNN model.

In this study, two different CNN models are considered: CNN1 to detect the shallow markers, and CNN2 to detect the markers at greater depths. The models take an image-frame as input and outputs 0 or 1 indicating the absence or presence of markers. The architecture of both CNNs comprises of three convolution layers having 32, 64 and 128 filters, respectively, followed by a max pooling layer, and finally a fully connected dense layer with 128 neurons. A stochastic gradient-descent optimizer is used to minimize loss and a binary cross entropy quantifier is used to measure the accuracy.

CNN1 used the dataset from Regions A and B. For this analysis, a total of 123 frames were collected, of which 69 frames were sampled from region A, and 54 frames (that are devoid of any markers) were sampled from Region B of the cube (refer Fig 4(b)). This data was then split into training and testing sets in the ratio of 0.8, and the model was trained for 50 epochs. A 10-fold cross validation on the training set was performed to mitigate bias in data splitting, and to better evaluate the model's performance over the testing dataset. The CNN1 model consistently achieved an accuracy of 94.9% on the test dataset as illustrated by the box plot in Fig 5(a).

CNN2 was used to detect the presence of markers at greater depths (i.e., in Region C). It used 62 frames from Region B and 50 frames from Region C. This data was split into training and testing sets in the ratio 0.8 and a 10-fold cross validation was performed as before. The model




Fig. 5. Box plot summary of the prediction accuracies for training, validation, and testing datasets over the 10-fold cross validation models corresponding to (a) CNN1 for the detection of markers at shallow depths (b) CNN2 for the detection of markers at greater depths

can achieve a prediction accuracy of 90% on the validation set and 78% with the testing set (Fig. 5(b)). Given the sparsity of data there is a chance that the models might have been overtrained. However, they possess considerable explanatory power as conveyed by the LIME Analysis. The sensitivity and specificity of detecting the markers can be improved further by fusing inspections made along different directions and combining multiple redundant marker measurements recorded in the vicinity of the marker region [7].

Fig. 6 summarizes a representative result from LIME analysis of CNN1. The markers present along a section (Fig. 6(a)) are hardly noticeable in the raw ultrasound images (Fig. 6(b)). The two bright spots near the top right corner of the image are the only possible indicators for the presence of markers. In contrast, the red segments from the LIME analysis (Fig. 6(c)) strongly indicate the presence of markers in the image-frame at these locations. Similarly, Fig 6(e) depicts the ultrasound scan pertaining to a frame without any markers. The green segments from LIME analysis (Fig 6(f)) contribute to strongly indicating the absence of embedded markers in the frame.




Fig. 6. (a) Representative view of the cube's actual design on its top section in Region A. (b) Ultrasound image of a frame captured in Region A. (c) The corresponding LIME explanation highlighting (in red) the important segments that determine the presence of artifacts. (d) Side view of the cube's design in Region B. (e) Ultrasound image of a frame captured in Region B. (f) The corresponding LIME interpretation highlighting (in green) the important segments that determine the absence of artifacts.

Similarly, for the more challenging case of CNN2 for detecting markers in Region C, the red segments near the middle as depicted in Fig 7(b) indicate the presence of embedded markers in the image-frame. The important segments near the top corners of the frame and, interestingly, far away from marker locations, are likely formed due to the scattering and reverberations of sound waves in the presence of the deeper markers. This is somewhat like a “butterfly effect”, where one needs to look for intensity variations elsewhere to locate a




Fig. 7. (a) Ultrasound image of a correctly classified frame captured in Region C of the scan containing markers at a depth of 9-12 mm near the centre of the cube. (b) The corresponding LIME explanation highlighting (in red) the important segments that indicate that the frame is from Region C.

defect at a particular location. Pertinently, the XAI-infused approach helps in determining where to look for these crucial patterns, and

these discovered patterns—largely ignored in the ultrasound literature—help with detecting artifacts located 3-4 times deeper compared to conventional ultrasonic inspections.

4. Conclusions

A novel explainable AI (XAI) - infused ultrasound imaging principle that enables a fast, holistic inspection of products manufactured from diverse materials, and detection of internal artifacts such as voids, pores and other defects is presented. Based on LIME analysis, it was discovered that distinct correlated spatial patterns are formed in the ultrasound images at locations that are far away from the artifact in the polymer cubic component with embedded markers. The results suggest that CNN models can detect internal artifacts that could not be discerned using any existing image segmentation method to accuracies exceeding 80% by effectively capturing the discovered spatial patterns. This result can profoundly impact the assurance of not just product quality, but also authentication and cybersecurity in the emerging manufacturing paradigm.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- [1] Peters J, Bryan J, Estler W, Evans C, Kunzmann H, Lucca D, Sartori S, Sato H, Thwaite E, Vanherck P (2001) Contribution of CIRP to the development of metrology and surface quality evaluation during the last fifty years. *CIRP Ann* : 471–488. 50/2.
- [2] Wang Z, Bukkapatnam ST, Kumara SR, Kong Z, Katz Z (2014) Change detection in precision manufacturing processes under transient conditions. *CIRP Ann* : 449–452. 63/1.
- [3] Caggiano A, Zhang J, Alfieri V, Caiazzo F, Gao R, Teti R (2019) Machine learning-based image processing for on-line defect recognition in additive manufacturing. *CIRP Ann* : 451–454. 68/1.
- [4] Brinksmeier E, Schneider E, Theiner W, Tönshoff H (1984) Nondestructive testing for evaluating surface integrity. *CIRP Ann* : 489–509. 33/2.
- [5] Wits WW, Carmignato S, Zanini F, Vaneker TH (2016) Porosity testing methods for the quality assessment of selective laser melted parts. *CIRP Ann* : 201–204. 65/1.
- [6] Ciampa F, Mahmoodi P, Pinto F, Meo M (2018) Recent advances in active infrared thermography for non-destructive testing of aerospace components. *Sensors* : 609.. 18/2.
- [7] Tiwari A, Villaseñor EJ, Gupta N, Reddy N, Karri R, Bukkapatnam ST (2021) Protection against counterfeiting attacks in 3D printing by streaming signature-embedded manufacturing process instructions. In: *Proceedings of the 2021 Workshop on Additive Manufacturing (3D Printing) Security*, 11–21.
- [8] Research & Markets, Electronic manufacturing services market research report by type, by application, by region - global forecast to 2027 - cumulative impact of COVID-19. <https://www.researchandmarkets.com/reports/5336821/> (accessed 2022).
- [9] Chen F, Yu JH, Gupta N (2019) Obfuscation of embedded codes in additive manufactured components for product authentication. *Adv Eng Mater* : 1900146. 21/8.
- [10] ElSayed KA, Dachowicz A, Panchal JH (2021) Information embedding in additive manufacturing through printing speed control. In: *Proceedings of the 2021 Workshop on Additive Manufacturing (3D Printing) Security*, 31–37.
- [11] Bukkapatnam ST, Iqbal AS, Kumara SR (2018) Planar random graph representations of spatiotemporal surface morphology: Application to finishing of 3-D printed components. *CIRP Ann* 67/1:495–498.
- [12] Iqbal AS, Bukkapatnam ST (2020) Consistent estimation of the max-flow problem: Towards unsupervised image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- [13] Jensen JRA (1991) A model for the propagation and scattering of ultrasound in tissue. *J Acoust Soc Am* : 182–190. 89/1.
- [14] Bakarić M, Miloro P, Javaherian A, Cox BT, Treeby BE, Brown MD (2021) Measurement of the ultrasound attenuation and dispersion in 3D-printed photopolymer materials from 1 to 3.5 MHz. *J Acoust Soc Am* : 2798–2805. 150/4.
- [15] Ye J, Ito S, Toyama N (2018) Computerized ultrasonic imaging inspection: From shallow to deep learning. *Sensors* : 3820.. 18/11.
- [16] Ribeiro MT, Singh S, Guestrin C (2016) Why should i trust you? Explaining the predictions of any classifier. In: *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 1135–1144.
- [17] Madore B, Meral FC (2012) Reconstruction algorithm for improved ultrasound image quality. *IEEE Trans Ultrason Ferroelectr Freq Control* : 217–230. 59/2.
- [18] Slack, D., Hilgard, S., Singh, S., Lakkaraju, H., 2020, How much should I trust you? modeling uncertainty of black box explanations, arXiv preprint arXiv:2008.05030.