

What is claimed is:

1. A method for reducing usage of processing resources when training a plurality of neural networks to perform automated visual inspection for a plurality of respective defect categories, with each of the plurality of respective defect categories being associated with a respective feature of containers or container contents, the method comprising:
  - obtaining, by one or more processors, a plurality of container images;
  - generating, by one or more processors processing the plurality of container images, a plurality of training image sets each corresponding to a different one of the plurality of container images, wherein for each training image set,
    - generating the training image set includes generating a different training image for each of the plurality of respective defect categories,
    - generating a different training image for each of the plurality of respective defect categories includes generating a first training image for a first defect category associated with a first feature, and
      - generating the first training image includes
        - identifying the first feature in the container image that corresponds to the training image set, and
        - generating the first training image such that the first training image (i) encompasses only a subset of the container image that corresponds to the training image set, and (ii) depicts the identified first feature;
    - and
    - training, by one or more processors and using the plurality of training image sets, the plurality of neural networks to perform automated visual inspection for the plurality of defect categories.
2. The method of claim 1, wherein training the plurality of neural networks includes training each of the plurality of neural networks to infer a presence or absence of defects in a different one of the plurality of respective defect categories.
3. The method of claim 2, wherein training each of the plurality of neural networks includes, for each of the plurality of training image sets, using a different training image to train a different one of the plurality of neural networks.
4. The method of any one of claims 1 through 3, wherein identifying the first feature includes identifying the first feature using template matching.
5. The method of any one of claims 1 through 3, wherein identifying the first feature includes identifying the first feature using blob analysis.
6. The method of any one of claims 1 through 5, wherein the first feature is a feature having a variable position relative to one or more container components.
7. The method of claim 6, wherein the plurality of container images is a plurality of syringe images, and wherein the first feature is a syringe plunger.
8. The method of claim 6, wherein the plurality of container images is a plurality of cartridge images, and wherein the first feature is a piston.

9. The method of claim 6, wherein the first feature is a meniscus of a fluid within a container, and wherein the first defect category is a presence of one or more particles in or near the meniscus.

10. The method of any one of claims 1 through 6, wherein the plurality of container images is a plurality of syringe images, and wherein either:

- the first feature is a syringe plunger and the first defect category is one or both of (i) a plunger defect or (ii) a presence of one or more particles on the syringe plunger;
- the first feature is a syringe barrel and the first defect category is one or both of (i) a barrel defect or (ii) a presence of one or more particles within the syringe barrel;
- the first feature is a syringe needle shield and the first defect category is one or both of (i) an absence of the syringe needle shield or (ii) misalignment of the syringe needle shield; or
- the first feature is a syringe flange and the first defect category is one or both of (i) a malformed flange or (ii) a defect on the syringe flange.

11. The method of any one of claims 1 through 6, wherein the plurality of container images is a plurality of cartridge images, and wherein either:

- the first feature is a cartridge piston and the first defect category is one or both of (i) a piston defect or (ii) a presence of one or more particles on the cartridge piston;
- the first feature is a cartridge barrel and the first defect category is one or both of (i) a barrel defect or (ii) a presence of one or more particles within the cartridge barrel; or
- the first feature is a cartridge flange and the first defect category is one or both of (i) a malformed flange or (ii) a defect on the cartridge flange.

12. The method of any one of claims 1 through 6, wherein the plurality of container images is a plurality of vial images, and wherein either:

- the first feature is a vial body and the first defect category is one or both of (i) a body defect or (ii) a presence of one or more particles within the vial body;
- the first feature is a vial crimp and the first defect category is a defective crimp; or
- the first feature is a lyophilized cake and the first defect category is a crack or other defect of the lyophilized cake.

13. The method of any one of claims 1 through 12, wherein:

- generating a different training image for each of the plurality of respective defect categories further includes generating a second training image for a second defect category associated with a second feature; and
- generating the second training image includes generating the second training image such that the second training image depicts the second feature.

14. The method of claim 13, wherein the second image encompasses an entirety of the container image that corresponds to the training image set.

15. The method of claim 13, wherein generating the second image includes generating the second image by down-sampling at least a portion of the container image that corresponds to the training image set to a lower resolution.

16. The method of claim 13, wherein generating the second training image includes identifying the second feature in the container image that corresponds to the training image set.

17. The method of any one of claims 1 through 16, wherein generating the plurality of training image sets each corresponding to a different one of the plurality of container images includes:

digitally aligning at least some of the plurality of container images, at least in part by resampling the at least some of the plurality of container images.

18. The method of claim 17, wherein digitally aligning at least some of the plurality of container images includes: detecting an edge within one or more of the plurality of container edges; and comparing a position of the detected edge to a position of a reference line.

19. The method of any one of claims 1 through 18, further comprising:  
prior to obtaining the plurality of container images, capturing, by an imaging system, the plurality of container images.

20. A system comprising one or more processors and one or more memories, the one or more memories storing instructions that, when executed by the one or more processors, cause the one or more processors to:

obtain a plurality of container images;

generate, by processing the plurality of container images, a plurality of training image sets each corresponding to a different one of the plurality of container images, wherein for each training image set,

generating the training image set includes generating a different training image for each of a plurality of respective defect categories,

generating a different training image for each of the plurality of respective defect categories includes generating a first training image for a first defect category associated with a first feature, and

generating the first training image includes

identifying the first feature in the container image that corresponds to the training image set, and

generating the first training image such that the first training image (i) encompasses only a subset of the container image that corresponds to the training image set, and (ii) depicts the identified first feature; and

train, using the plurality of training image sets, a plurality of neural networks to perform automated visual inspection for the plurality of defect categories, each of the plurality of defect categories being associated with a respective feature of containers or container contents.

21. The system of claim 20, wherein training the plurality of neural networks includes training each of the plurality of neural networks to infer a presence or absence of defects in a different one of the plurality of respective defect categories.

22. The system of claim 21, wherein training each of the plurality of neural networks includes, for each of the plurality of training image sets, using a different training image to train a different one of the plurality of neural networks.

23. The system of any one of claims 20 through 22, wherein identifying the first feature includes identifying the first feature using (i) template matching, or (ii) blob analysis.

24. The system of any one of claims 20 through 23, wherein the first feature is a feature having a variable position relative to one or more container components.

25. The system of claim 24, wherein the plurality of container images is a plurality of syringe images.

26. The system of claim 25, wherein the first feature is a syringe plunger, a meniscus of a fluid, a syringe barrel, a syringe needle shield, or a syringe flange.

27. The system of claim 24, wherein the plurality of container images is a plurality of cartridge images.

28. The system of claim 27, wherein the first feature is a cartridge piston, a meniscus of a fluid, a cartridge barrel, or a cartridge flange.

29. The system of claim 24, wherein the plurality of container images is a plurality of vial images.

30. The system of claim 29, wherein the first feature is a vial body, a vial crimp, a meniscus of a fluid, or a lyophilized cake.

31. The system of any one of claims 20 through 30, further comprising:  
an imaging system configured to capture the plurality of container images,  
wherein the instructions cause the one or more processors to obtain the plurality of container images from the imaging system.

32. A method of using an efficiently trained neural network to perform automated visual inspection for detection of defects in a plurality of defect categories, with each of the plurality of defect categories being associated with a respective feature of containers or container contents, the method comprising:

obtaining, by one or more processors, a plurality of neural networks each corresponding to a different one of the plurality of defect categories, the plurality of neural networks having been trained using a plurality of training image sets each corresponding to a different one of a plurality of container images, wherein for each training image set and corresponding container image,

the training image set includes a different training image for each of the plurality of respective defect categories, and

at least some of the different training images within the training image set consist of different portions of the corresponding container image, with the different portions depicting different features of the corresponding container image;

obtaining, by one or more processors, a plurality of additional container images; and

performing automated visual inspection on the plurality of additional container images using the plurality of neural networks.

33. The method of claim 32, wherein the plurality of container images is a plurality of syringe images, and wherein the different features include a syringe plunger, a syringe barrel, a syringe needle shield, and/or a syringe flange.

34. The method of claim 32, wherein the plurality of container images is a plurality of cartridge images, and wherein the different features include a cartridge piston, a cartridge barrel, a barrel defect, and/or a cartridge flange.

35. The method of claim 32, wherein the plurality of container images is a plurality of vial images, and wherein the different features include a vial body, a vial crimp, and/or a lyophilized cake.

36. The method of any one of claims 32 through 35, wherein at least one of the different training images includes an entirety of the corresponding container image.

37. The method of any one of claims 32 through 36, wherein the different training images include images down-sampled to different resolutions.

38. A method of training an automated visual inspection (AVI) neural network to more accurately detect defects, the method comprising:

obtaining, by one or more processors, a plurality of container images;

for each container image of the plurality of container images, generating, by one or more processors, a corresponding set of new images, wherein generating the corresponding set of new images includes, for each new image of the corresponding set of new images, moving a portion of the container image that depicts a particular feature to a different new position; and

training, by one or more processors, the AVI neural network using the sets of new images corresponding to the plurality of container images.

39. The method of claim 38, wherein the plurality of container images is a plurality of syringe images, and wherein the particular feature is one of:

a syringe plunger;

a syringe needle shield; or

a wall of a syringe barrel.

40. The method of claim 38, wherein the plurality of container images is a plurality of cartridge images, and wherein the particular feature is a cartridge piston or a wall of a cartridge barrel.

41. The method of claim 38, wherein the plurality of container images is a plurality of vial images, and wherein the particular feature is a wall of a vial body or a crimp.

42. The method of claim 38, wherein the particular feature is a meniscus of a fluid within a container.

43. The method of claim 38, wherein the particular feature is a top of a lyophilized cake within a container.

44. The method of any one of claims 38 through 43, wherein moving the portion of the container image to the different new position includes shifting the portion of the container image along an axis of a substantially cylindrical portion of a container depicted in the container image.

45. The method of any one of claims 38 through 44, wherein generating the corresponding set of new images further includes, for each new image of the corresponding set of new images, low-pass filtering the new image after moving the portion of the container image to the different new position.

46. The method of any one of claims 38 through 45, wherein training the AVI neural network includes training the AVI neural network using (i) the sets of new images corresponding to the plurality of container images and (ii) the plurality of container images.

47. The method of any one of claims 38 through 46, further comprising:  
prior to obtaining the plurality of container images, capturing, by an imaging system, the plurality of container images.

48. A system comprising one or more processors and one or more memories, the one or more memories storing instructions that, when executed by the one or more processors, cause the one or more processors to:  
obtain a plurality of container images;  
for each container image of the plurality of container images, generate a corresponding set of new images, wherein generating the corresponding set of new images includes, for each new image of the corresponding set of new images, moving a portion of the container image that depicts a particular feature to a different new position; and  
train an automated visual inspection (AVI) neural network using the sets of new images corresponding to the plurality of container images.

49. The system of claim 48, wherein the plurality of container images is a plurality of syringe images, and wherein the particular feature is one of:

a syringe plunger;  
a syringe needle shield; or  
a wall of a syringe barrel.

50. The system of claim 48, wherein the plurality of container images is a plurality of cartridge images, and wherein the particular feature is a cartridge piston or a wall of a cartridge barrel.

51. The system of claim 48, wherein the plurality of container images is a plurality of vial images, and wherein the particular feature is a wall of a vial body or a crimp.

52. The system of claim 48, wherein the particular feature is a meniscus of a fluid within a container.

53. The system of claim 48, wherein the particular feature is a top of a lyophilized cake within a container.

54. The system of any one of claims 48 through 53, wherein moving the portion of the container image to the different new position includes shifting the portion of the container image along an axis of a substantially cylindrical portion of a container depicted in the container image.

55. The system of any one of claims 48 through 54, wherein generating the corresponding set of new images further includes, for each new image of the corresponding set of new images, low-pass filtering the new image after moving the portion of the container image to the different new position.

56. The system of any one of claims 48 through 55, further comprising:  
an imaging system configured to capture the plurality of container images,  
wherein the instructions cause the one or more processors to obtain the plurality of container images from the imaging system.

57. A method of performing automated visual inspection for detection of defects, the method comprising:  
obtaining, by one or more processors, a neural network trained using a plurality of container images and, for each container image, a plurality of augmented container images, wherein for each of the augmented container images, a portion of the container image that depicts a particular feature was moved to a different new position;  
obtaining, by one or more processors, a plurality of additional container images; and  
performing automated visual inspection on the plurality of additional container images using the neural network.

58. The method of claim 57, wherein for each of the augmented container images, the portion of the container image that depicts the particular feature was moved by shifting the portion of the container image along an axis of a substantially cylindrical portion of a container depicted in the container image.

59. The method of claim 57 or 58, wherein each of the augmented container images was low-pass filtered (i) after the portion of the container image that depicts the particular feature was moved to the different new position, and (ii) before training the neural network using the augmented container image.

60. A method of training an automated visual inspection (AVI) neural network to more accurately detect defects, the method comprising:  
obtaining, by one or more processors, a plurality of images depicting real containers;  
training, by one or more processors and using the plurality of images depicting real containers, a deep generative model to generate synthetic container images;  
generating, by one or more processors and using the deep generative model, a plurality of synthetic container images;  
and  
training, by one or more processors and using the plurality of synthetic container images, the AVI neural network.

61. The method of claim 60, wherein training the deep generative model includes training a generative adversarial network (GAN).

62. The method of claim 61, wherein:  
training the GAN includes applying, as inputs to a discriminator neural network, (i) the plurality of images depicting real containers and corresponding real image labels, and (ii) synthetic images generated by a generator neural network and corresponding fake image labels; and  
generating the plurality of synthetic container images is performed by the trained generator neural network.

63. The method of claim 61, wherein:

training the GAN includes training a cycle GAN; and  
generating the plurality of synthetic container images includes transforming images of real containers that are not associated with any defects to images that exhibit a defect class.

64. The method of claim 61, wherein generating the plurality of synthetic container images includes seeding a respective particle location for each of the plurality of synthetic container images.

65. The method of claim 64, wherein seeding the respective particle location for each of the plurality of synthetic container images includes randomly seeding the respective particle location for each of the plurality of synthetic container images.

66. The method of claim 60, wherein training the deep generative model includes training a variational autoencoder (VAE).

67. The method of claim 66, wherein:  
training the VAE includes encoding each of the plurality of images into a latent space; and  
generating the plurality of synthetic container images includes randomly sampling the latent space.

68. A system comprising one or more processors and one or more memories, the one or more memories storing instructions that, when executed by the one or more processors, cause the one or more processors to:  
obtain a plurality of images depicting real containers;  
train, using the plurality of images depicting real containers, a deep generative model to generate synthetic container images;  
generate, using the deep generative model, a plurality of synthetic container images; and  
train, using the plurality of synthetic container images, the AVI neural network.

69. The system of claim 68, wherein the deep generative model is a generative adversarial network (GAN).

70. The system of claim 69, wherein:  
training the GAN includes applying, as inputs to a discriminator neural network, (i) the plurality of images depicting real containers and corresponding real image labels, and (ii) synthetic images generated by a generator neural network and corresponding fake image labels; and  
generating the plurality of synthetic container images is performed by the trained generator neural network.

71. The system of claim 69, wherein:  
the GAN is a cycle GAN; and  
generating the plurality of synthetic container images includes transforming images of real containers that are not associated with any defects to images that exhibit a defect class.

72. The system of claim 69, wherein generating the plurality of synthetic container images includes seeding a respective particle location for each of the plurality of synthetic container images.

73. The system of claim 72, wherein seeding the respective particle location for each of the plurality of synthetic container images includes randomly seeding the respective particle location for each of the plurality of synthetic container images.

74. The system of claim 68, wherein the deep generative model is a variational autoencoder (VAE).

75. The system of claim 74, wherein:  
training the VAE includes encoding each of the plurality of images into a latent space; and  
generating the plurality of synthetic container images includes randomly sampling the latent space.

76. A method of performing automated visual inspection for detection of defects, the method comprising:  
obtaining, by one or more processors, a neural network trained using synthetic container images generated by a deep generative model;  
obtaining, by one or more processors, a plurality of additional container images; and  
performing automated visual inspection on the plurality of additional container images using the neural network.

77. The method of claim 76, wherein obtaining the neural network includes obtaining a neural network trained using a generative adversarial network (GAN).

78. The method of claim 76, wherein obtaining the neural network includes obtaining a neural network trained using a cycle generative adversarial network (GAN).

79. The method of claim 76, wherein obtaining the neural network includes obtaining a variational autoencoder (VAE).

80. A method for evaluating reliability of a trained automated visual inspection (AVI) neural network, the method comprising:  
obtaining, by one or more processors, a container image;  
generating or receiving, by one or more processors, a heatmap of the container image, wherein the heatmap indicates which portions of the container image contribute most to an inference made by the trained AVI neural network, and wherein the inference is an inference of whether the container image depicts a defect;  
determining, by one or more processors, whether the trained AVI neural network made the inference for the correct reason, at least in part by analyzing the heatmap; and  
generating, by one or more processors, an indicator of whether the trained AVI neural network made the inference for the correct reason.

81. The method of claim 80, wherein:  
the trained AVI neural network is trained to infer a presence or absence of defects associated with a particular container zone; and  
analyzing the heatmap includes generating a metric indicative of a level of heatmap activity in a region of the heatmap that corresponds to the particular container zone.

82. The method of claim 81, wherein determining whether the trained AVI neural network made the inference for the correct reason includes comparing the metric to a threshold value.

83. The method of claim 81, wherein:  
analyzing the heatmap further includes generating one or more other metrics indicative of levels of heatmap activity in one or more other regions of the heatmap that correspond to the one or more other container zones; and  
determining whether the trained AVI neural network made the inference for the correct reason is based on the metric and the one or more other metrics.

84. The method of claim 80, wherein analyzing the heatmap includes comparing the heatmap to a reference heatmap.

85. The method of claim 84, wherein the reference heatmap is a composite heatmap derived from a plurality of heatmaps.

86. The method of claim 84 or 85, wherein the inference is an inference that the container image depicts a defect, and wherein the reference heatmap is a reference heatmap generated based on one or more container images known to not depict defects.

87. The method of claim 84 or 85, wherein the inference is an inference that the container image does not depict a defect, and wherein the reference heatmap is a reference heatmap generated based on one or more container images known to depict defects.

88. The method of any one of claims 80 through 87, wherein generating or receiving the heatmap includes generating the heatmap at least in part by:

sequentially, for each image portion of a plurality of different image portions in the container image, (i) masking the image portion, (ii) generating a replacement image portion for the masked image portion by resampling based on pixels surrounding the masked image portion, and (iii) generating a respective inference confidence score by applying the container image, with the replacement image portion, to the trained AVI neural network,

wherein the heatmap includes indications of the respective inference confidence scores for the plurality of different image portions.

89. The method of any one of claims 80 through 87, wherein generating or receiving the heatmap includes generating the heatmap using gradient-weighted class activation mapping (grad-CAM).

90. A system comprising one or more processors and one or more memories, the one or more memories storing instructions that, when executed by the one or more processors, cause the one or more processors to:

obtain a container image;

generate or receive a heatmap of the container image, wherein the heatmap indicates which portions of the container image contribute most to an inference made by the trained AVI neural network, and wherein the inference is an inference of whether the container image depicts a defect;

determine whether the trained AVI neural network made the inference for the correct reason, at least in part by analyzing the heatmap; and

generate an indicator of whether the trained AVI neural network made the inference for the correct reason.

91. The system of claim 90, wherein:

the trained AVI neural network is trained to infer a presence or absence of defects associated with a particular container zone; and

analyzing the heatmap includes generating a metric indicative of a level of heatmap activity in a region of the heatmap that corresponds to the particular container zone.

92. The system of claim 91, wherein determining whether the trained AVI neural network made the inference for the correct reason includes comparing the metric to a threshold value.

93. The system of claim 91 or 92, wherein:

analyzing the heatmap further includes generating one or more other metrics indicative of levels of heatmap activity in one or more other regions of the heatmap that correspond to the one or more other container zones; and

determining whether the trained AVI neural network made the inference for the correct reason is based on the metric and the one or more other metrics.

94. The system of claim 90, wherein analyzing the heatmap includes comparing the heatmap to a reference heatmap.

95. The system of any one of claims 90 through 94, wherein generating or receiving the heatmap includes generating the heatmap at least in part by:

sequentially, for each image portion of a plurality of different image portions in the container image, (i) masking the image portion, (ii) generating a replacement image portion for the masked image portion by resampling based on pixels surrounding the masked image portion, and (iii) generating a respective inference confidence score; and

applying the container image, with the replacement image portion, to the trained AVI neural network,

wherein the heatmap includes the respective inference confidence scores for the plurality of different image portions.

96. The system of any one of claims 90 through 94, wherein generating or receiving the heatmap includes generating the heatmap using gradient-weighted class activation mapping (grad-CAM).

97. A method for evaluating reliability of a trained automated visual inspection (AVI) neural network, the method comprising:

obtaining, by one or more processors, a container image;

generating or receiving, by one or more processors, data indicative of a particular area within the container image, wherein the particular area indicates a portion of the container image in which the trained AVI neural network has detected an object corresponding to a defect category;

determining, by one or more processors, whether the trained AVI neural network correctly identified the object, at least in part by comparing a position of the particular area to a position of a user-identified area within the container image; and

generating, by one or more processors, an indicator of whether the trained AVI neural network correctly identified the object.

98. The method of claim 97, wherein the data indicative of the particular area within the container image is data indicative of a boundary within the container image.

99. The method of claim 98, wherein the data indicative of the boundary within the container image is data indicative of a bounding box within the container image.

100. The method of claim 97, wherein:  
the trained AVI neural network is a segmentation neural network;  
the data indicative of the particular area within the container image is data that classifies each pixel of the particular area as a particular object; and  
the user-identified area within the container image is a user-selected set of pixels that consists of all pixels within the user-identified area.

101. The method of any one of claims 97 through 100, wherein comparing the position of the particular area to the position of the user-identified area includes either:

determining whether a center of the particular area falls within the user-identified area; or  
determining whether a center of the user-identified area falls within the particular area.

102. The method of claim 101, comprising:  
determining that the trained AVI neural network correctly identified the object at least by determining either (i) that the center of the particular area falls within the user-identified area, or (ii) that the center of the user-identified area falls within the particular area.

103. The method of claim 101, comprising:  
determining that the trained AVI neural network incorrectly identified the object at least by determining either (i) that the center of the particular area falls outside the user-identified area, or (ii) that the center of the user-identified area falls outside the particular area.

104. The method of any one of claims 97 through 100, wherein comparing the position of the particular area to the position of the user-generated bounding box includes either:

determining whether at least a threshold percentage of the particular area overlaps the user-identified area; or  
determining whether at least a threshold percentage of the user-identified area overlaps the particular area.

105. The method of claim 104, comprising:  
determining that the trained AVI neural network correctly identified the object at least by determining either (i) that at least the threshold percentage of the particular area overlaps the user-identified area, or (ii) that at least the threshold percentage of the user-identified area overlaps the particular area.

106. The method of claim 105, comprising:

determining that the trained AVI neural network incorrectly identified the object at least by determining either (i) that less than the threshold percentage of the particular area overlaps the user-identified area, or (ii) that less than the threshold percentage of the user-identified area overlaps the particular area.

107. The method of any one of claims 97 through 106, wherein the object is a particle.

108. The method of claim 107, wherein the AVI neural network is further trained to detect bubbles.

109. A system comprising one or more processors and one or more memories, the one or more memories storing instructions that, when executed by the one or more processors, cause the one or more processors to  
obtain a container image;  
generate or receive data indicative of a particular area within the container image, wherein the particular area indicates a portion of the container image in which the trained AVI neural network has detected an object corresponding to a defect category;  
determine whether the trained AVI neural network correctly identified the object, at least in part by comparing a position of the particular area to a position of a user-identified area within the container image; and  
generate an indicator of whether the trained AVI neural network correctly identified the object.

110. The system of claim 109, wherein the data indicative of the particular area within the container image is data indicative of a boundary within the container image.

111. The system of claim 110, wherein the data indicative of the boundary within the container image is data indicative of a bounding box within the container image.

112. The system of claim 109, wherein:  
the trained AVI neural network is a segmentation neural network;  
the data indicative of the particular area within the container image is data that classifies each pixel of the particular area as a particular object; and  
the user-identified area within the container image is a user-selected set of pixels that consists of all pixels within the user-identified area.

113. The system of any one of claims 109 through 112, wherein comparing the position of the particular area to the position of the user-generated bounding box includes either:  
determining whether a center of the particular area falls within the user-identified area; or  
determining whether a center of the user-identified area falls within the particular area.

114. The system of claim 113, comprising:  
determining that the trained AVI neural network correctly identified the object at least by determining either (i) that the center of the particular area falls within the user-identified area, or (ii) that the center of the user-identified area falls within the particular area.

115. The system of claim 113, comprising:

determining that the trained AVI neural network incorrectly identified the object at least by determining either (i) that the center of the particular area falls outside the user-identified area, or (ii) that the center of the user-identified area falls outside the particular area.

116. The system of any one of claims 109 through 112, wherein comparing the position of the particular area to the position of the user-generated bounding box includes either:

- determining whether at least a threshold percentage of the particular area overlaps the user-identified area; or
- determining whether at least a threshold percentage of the user-identified area overlaps the particular area.

117. The system of claim 116, comprising:

determining that the trained AVI neural network correctly identified the object at least by determining either (i) that at least the threshold percentage of the particular area overlaps the user-identified area, or (ii) that at least the threshold percentage of the user-identified area overlaps the particular area.

118. The system of claim 116, comprising:

determining that the trained AVI neural network incorrectly identified the object at least by determining either (i) that less than the threshold percentage of the particular area overlaps the user-identified area, or (ii) that less than the threshold percentage of the user-identified area overlaps the particular area.

119. The system of any one of claims 109 through 118, wherein the object is a particle.

120. The system of claim 119, wherein the AVI neural network is further trained to detect bubbles.