What is claimed is:

1. A method of generating a synthetic image by transferring a feature onto an original image, the method comprising:

receiving or generating a feature matrix that is a numeric representation of a feature image depicting the feature, with each element of the feature matrix corresponding to a different pixel of the feature image;

receiving or generating a surrogate area matrix that is a numeric representation of an area, within the original image, to which the feature will be transferred, with each element of the surrogate area matrix corresponding to a different pixel of the original image;

normalizing the feature matrix relative to a portion of the feature matrix that does not represent the feature; and

generating the synthetic image based on (i) the surrogate area matrix and (ii) the normalized feature matrix.

2. The method of claim 1, wherein:

the original image is an image of a container; and

the feature is a defect associated with the container or contents of the container.

3. The method of claim 2, wherein:

the container is a syringe; and

the feature is a defect associated with a barrel of the syringe, a plunger of the syringe, a needle shield of the syringe, or a fluid within the syringe.

4. The method of claim 2, wherein:

the container is a vial; and

the feature is a defect associated with a wall of the vial, a cap of the vial, a crimp of the vial, or a fluid or lyophilized cake within the vial.

5. The method of any one of claims 1-4, wherein normalizing the feature matrix includes normalizing the feature matrix on a per-row or per-column basis.

6. The method of claim 5, wherein normalizing the feature matrix on a per-row or per-column basis includes, for each row or column of the feature matrix:

generating a feature row histogram of element values for the row or column of the feature matrix.

7. The method of claim 6, wherein normalizing the feature matrix on a per-row or per-column basis further includes, for each row or column of the feature matrix:

identifying a peak portion of the feature row histogram that corresponds to a portion of the row or column of the feature matrix that does not represent the feature; and

for each element of the row or column of the feature matrix, subtracting a center value of the peak portion from a value of the element.

8. The method of claim 7, wherein subtracting the center value of the peak portion from the value of the element includes subtracting (i) an average value of all values in the row or column that correspond to the peak portion from (ii) the value of the element.

9. The method of any one of claims 1-8, further comprising:

for each row or column of the surrogate area matrix,

generating a surrogate area row histogram,

identifying a peak portion of the surrogate area row histogram, and

determining a number range representative of a width of the peak portion of the surrogate area row histogram,

wherein generating the synthetic image includes generating the synthetic image based on (i) the number range for each row or column of the feature matrix, and (ii) the normalized feature matrix.

10. The method of claim 9, wherein generating the synthetic image includes, for each row or column of the feature matrix:

for each element of the row or column of the feature matrix, determining whether the element of the feature matrix has a value within the number range; and

modifying an original image matrix that is a numeric representation of the original image by either (i) when the element of the feature matrix has a value within the number range, retaining an original value of a corresponding element in the original image matrix, or (ii) when the value of the element of the feature matrix is not within the number range, setting the corresponding element in the original image matrix equal to a sum of the original value and the value of the element in the feature matrix.

11. The method of claim 10, wherein generating the synthetic image includes converting the modified original image matrix to a bitmap image.

12. The method of any one of claims 1-11, wherein receiving or generating the feature matrix includes rotating the feature matrix or the feature image.

13. The method of claim 12, wherein rotating the feature matrix or the feature image includes rotating the feature matrix or the feature image by an amount that is based on (i) a rotation of the feature depicted in the feature image and (ii) a desired rotation of the feature depicted in the feature image.

14. The method of claim 13, further comprising:

determining the desired rotation based on a position of the area to which the feature will be transferred.

15. The method of any one of claims 1-14, wherein receiving or generating the feature matrix includes resizing the feature matrix or the feature image.

16. The method of any one of claims 1-15, further comprising:

repeating the method for each of a plurality of features corresponding to different features in a feature library.

17. The method of any one of claims 1-16, further comprising:

generating a plurality of synthetic images by repeating the method for each of a plurality of original images.

18. The method of claim 17, further comprising:

training a neural network for automated visual inspection using the plurality of synthetic images and the plurality of original images.

19. The method of claim 18, further comprising:

inspecting a plurality of images for depicted defects using the trained neural network.

20. A system comprising:

one or more processors; and

one or more non-transitory, computer-readable media storing instructions that, when executed by the one or more processors, cause the system to perform the method of any one of claims 1 to 19.

~~receive or generate a feature matrix that is a numeric representation of a feature image depicting a feature, with each element of the feature matrix corresponding to a different pixel of the feature image,~~

~~receive or generate a surrogate area matrix that is a numeric representation of an area, within an original image, to which the feature will be transferred, with each element of the surrogate area matrix corresponding to a different pixel of the original image,~~

~~normalize the feature matrix relative to a portion of the feature matrix that does not represent the feature, and~~

~~generate a synthetic image based on (i) the surrogate area matrix and (ii) the normalized feature matrix.~~

~~21. The system of claim 20, wherein normalizing the feature matrix includes normalizing the feature matrix on a per-row or per-column basis.~~

~~22. The system of claim 21, wherein normalizing the feature matrix on a per-row or per-column basis includes, for each row or column of the feature matrix:~~

~~generating a feature row histogram of element values for the row or column of the feature matrix.~~

~~23. The method of claim 22, wherein normalizing the feature matrix on a per-row or per-column basis further includes, for each row or column of the feature matrix:~~

~~identifying a peak portion of the feature row histogram that corresponds to a portion of the row or column of the feature matrix that does not represent the feature; and~~

~~for each element of the row or column of the feature matrix, subtracting a center value of the peak portion from a value of the element.~~

~~24. The system of claim 23, wherein subtracting the center value of the peak portion from the value of the element includes subtracting (i) an average value of all values in the row or column that correspond to the peak portion from (ii) the value of the element.~~

~~25. The system of any one of claims 20-24, wherein the instructions further cause the system to:~~

for each row or column of the surrogate area matrix,

generate a surrogate area row histogram,

identify a peak portion of the surrogate area row histogram, and

determine a number range representative of a width of the peak portion of the surrogate area row histogram,

wherein generating the synthetic image includes generating the synthetic image based on (i) the number range for each row or column of the feature matrix, and (ii) the normalized feature matrix.

26. The system of claim 25, wherein generating the synthetic image includes:

for each row or column of the feature matrix,

for each element of the row or column of the feature matrix, determining whether the element of the feature matrix has a value within the number range, and

modifying an original image matrix that is a numeric representation of the original image by either (i) when the element of the feature matrix has a value within the number range, retaining an original value of a corresponding element in the original image matrix, or (ii) when the value of the element of the feature matrix is not within the number range, setting the corresponding element in the original image matrix equal to a sum of the original value and the value of the element in the feature matrix; and

converting the modified original image matrix to a bitmap image.

27. The system of any one of claims 20-26, wherein receiving or generating the feature matrix includes rotating the feature matrix or the feature image.

28. The system of claim 27, wherein rotating the feature matrix or the feature image includes rotating the feature matrix or the feature image by an amount that is based on (i) a rotation of the feature depicted in the feature image and (ii) a desired rotation of the feature depicted in the feature image.

29. The system of claim 28, wherein the instructions further cause the system to:

determine the desired rotation based on a position of the area to which the feature will be transferred.

30. The system of any one of claims 20-29, wherein receiving or generating the feature matrix includes resizing the feature matrix or the feature image.

31. A method of generating a synthetic image by removing a defect depicted in an original image, the method comprising:

masking a portion of the original image that depicts the defect;

generating the synthetic image, at least in part by

calculating correspondence metrics indicative of correspondences between (i) portions of the original image adjacent to the masked portion and (ii) other portions of the original image, and

using the calculated correspondence metrics to fill the masked portion of the original image with a defect-free image portion; and

training a neural network for automated visual inspection using the synthetic image.

32. The method of claim 31, wherein calculating the correspondence metrics includes:

calculating first statistics for each of a plurality of adjacent portions of the original image that are adjacent to the masked portion;

calculating second statistics for each of a plurality of portions of the original image that are outside of both (i) the masked portion and (ii) the plurality of adjacent portions; and

calculating the correspondence metrics based on the first statistics and the second statistics.

33. The method of claim 31 or 32, wherein:

the original image is an image of a container; and

the defect is associated with the container or contents of the container.

34. The method of claim 33, wherein:

the container is a syringe; and

the defect is associated with a barrel of the syringe, a plunger of the syringe, a needle shield of the syringe, or a fluid within the syringe.

35. The method of claim 33, wherein:

the container is a vial; and

the defect is associated with a wall of the vial, a cap of the vial, a crimp of the vial, or a fluid or lyophilized cake within the vial.

36. The method of any one of claims 31-35, further comprising:

generating a plurality of synthetic images by repeating the method for each of a plurality of original images,

wherein training the neural network for automated visual inspection includes training the neural network using the plurality of synthetic images.

37. The method of claim 36, wherein training the neural network for automated visual inspection includes training the neural network using the plurality of synthetic images and the plurality of original images.

38. The method of any one of claims 31-37, further comprising:

inspecting a plurality of images for defects using the trained neural network.

39. A system comprising:

one or more processors; and

one or more non-transitory, computer-readable media storing instructions that, when executed by the one or more processors, cause the system to

mask a portion of an original image that depicts a defect,

generate a synthetic image, at least in part by

calculating correspondence metrics indicative of correspondences between (i) portions of the original image adjacent to the masked portion and (ii) other portions of the original image, and

using the calculated correspondence metrics to fill the masked portion of the original image with a defect-free image portion, and

train a neural network for automated visual inspection using the synthetic image.

40. The system of claim 39, wherein calculating the correspondence metrics includes:

calculating first statistics for each of a plurality of adjacent portions of the original image that are adjacent to the masked portion;

calculating second statistics for each of a plurality of portions of the original image that are outside of both (i) the masked portion and (ii) the plurality of adjacent portions; and

calculating the correspondence metrics based on the first statistics and the second statistics.

41. The system of claim 39 or 40, wherein:

the original image is an image of a container; and

the defect is associated with the container or contents of the container.

42. The system of claim 41, wherein:

the container is a syringe; and

the defect is associated with a barrel of the syringe, a plunger of the syringe, a needle shield of the syringe, or a fluid within the syringe.

43. The system of claim 41, wherein:

the container is a vial; and

the defect is associated with a wall of the vial, a cap of the vial, a crimp of the vial, or a fluid or lyophilized cake within the vial.

44. A method of generating synthetic images by removing or modifying features depicted in original images, or by adding depicted features to the original images, the method comprising:

training a partial convolution model that includes (i) an encoder with a series of convolution layers and (ii) a decoder with a series of transpose convolution layers, at least by, for each of a plurality of training images, applying the training image and a corresponding training mask as separate inputs to the partial convolution model;

generating the synthetic images, at least by, for each of the original images, applying (i) either the original image or a modified version of the original image, and (ii) a corresponding mask, as separate inputs to the trained partial convolution model; and

training a neural network for automated visual inspection using the synthetic images.

45. The method of claim 44, wherein the method generates the synthetic images by removing defects depicted in the original images.

46. The method of claim 45, wherein:

generating the synthetic images includes, for each of the original images, applying the original image and the corresponding mask as separate inputs to the trained partial convolution model;

the original image is an image that depicts a corresponding defect; and

the corresponding mask, when applied to the original image, obscures the corresponding defect.

47. The method of claim 46, wherein the training images, the original images, and the synthetic images depict containers.

48. The method of claim 47, wherein the containers are syringes.

49. The method of claim 48, wherein the corresponding defect is a syringe barrel defect, a syringe plunger defect, or a defect associated with syringe contents.

50. The method of claim 47, wherein the containers are vials.

51. The method of claim 50, wherein the corresponding defect is a vial wall defect, a vial cap defect, a vial crimp defect, or a defect associated with vial contents.

52. The method of claim 44, wherein the method generates the synthetic images by adding depicted defects to the original images or modifying features depicted in the original images.

53. The method of claim 52, wherein:

generating the synthetic images includes, for each of the original images,

modifying the original image by replacing a particular portion of the original image with a feature image, and

applying the modified original image and the corresponding mask as separate inputs to the trained partial convolution model; and

the corresponding mask, when applied to the modified original image, obscures the particular portion of the original image.

54. The method of claim 53, wherein generating the synthetic images includes, for each of the original images, automatically generating the corresponding mask to correspond to the particular portion of the original image.

55. The method of claim 53 or 54, wherein the training images, the original images, and the synthetic images depict containers.

56. The method of claim 55, wherein the containers are syringes.

57. The method of claim 56, wherein:

the method generates the synthetic images by adding depicted defects to the original images; and

the feature image is an image of a syringe barrel defect, a syringe plunger defect, or a defect associated with syringe contents.

58. The method of claim 55, wherein the containers are vials.

59. The method of claim 58, wherein:

the method generates the synthetic images by adding depicted defects to the original images; and

the feature image is an image of a vial wall defect, a vial cap defect, a vial crimp defect, or a defect associated with vial contents.

60. The method of any one of claims 53-55, wherein:

the method generates the synthetic images by modifying features depicted in the original images; and

the feature image is an image of a meniscus.

61. The method of claim 52, wherein:

the corresponding training mask, when applied to the training image, obscures a corresponding feature depicted in the training image; and

generating the synthetic images includes, for each of the original images, applying the original image and the corresponding mask as separate inputs to the trained partial convolution model.

62. The method of claim 61, wherein the training images, the original images, and the synthetic images depict containers.

63. The method of claim 62, wherein the containers are syringes.

64. The method of claim 63, wherein:

the method generates the synthetic images by adding depicted defects to the original images; and

the corresponding feature is a syringe barrel defect, a syringe plunger defect, or a defect associated with syringe contents.

65. The method of claim 62, wherein the containers are vials.

66. The method of claim 63, wherein:

the method generates the synthetic images by adding depicted defects to the original images; and

the corresponding feature is a vial wall defect, a vial cap defect, a vial crimp defect, or a defect associated with vial contents.

67. The method of any one of claims 61-66, wherein:

training the partial convolution model includes training the partial convolution model to add a specific type of defect; and

the corresponding feature corresponds to the specific type of defect.

68. The method of claim 61, wherein:

the method generates the synthetic images by modifying features depicted in the original images; and

the corresponding feature is a meniscus.

69. The method of any one of claims 44-68, wherein training the partial convolution model includes minimizing a plurality of losses, the plurality of losses including at least a valid loss, a hole loss, a perceptual loss, a style loss, and a variation loss.

70. A system comprising:

one or more processors; and

one or more non-transitory, computer-readable media storing instructions that, when executed by the one or more processors, cause the system to

train a partial convolution model that includes (i) an encoder with a series of convolution layers and (ii) a decoder with a series of transpose convolution layers, at least by, for each of a plurality of training images, applying the training image and a corresponding training mask as separate inputs to the partial convolution model,

generate a plurality of synthetic images, at least by, for each of a plurality of original images, applying (i) either the original image or a modified version of the original image, and (ii) a corresponding mask, as separate inputs to the trained partial convolution model, and

train a neural network for automated visual inspection using the synthetic images.

71. The system of claim 70, wherein the system generates the synthetic images by removing defects depicted in the original images.

72. The system of claim 71, wherein:

generating the synthetic images includes, for each of the original images, applying the original image and the corresponding mask as separate inputs to the trained partial convolution model;

the original image is an image that depicts a corresponding defect; and

the corresponding mask, when applied to the original image, obscures the corresponding defect.

73. The system of claim 72, wherein the training images, the original images, and the synthetic images depict containers.

74. The system of claim 70, wherein the system generates the synthetic images by adding depicted defects to the original images or modifying features depicted in the original images.

75. The system of claim 74, wherein:

generating the synthetic images includes, for each of the original images,

modifying the original image by replacing a particular portion of the original image with a feature image, and

applying the modified original image and the corresponding mask as separate inputs to the trained partial convolution model; and

the corresponding mask, when applied to the modified original image, obscures the particular portion of the original image.

76. The system of claim 75, wherein the training images, the original images, and the synthetic images depict containers.

77. The system of claim 75 or 76, wherein:

the system generates the synthetic images by adding depicted defects to the original images; and

the feature image is an image of a defect.

78. The system of claim 75 or 76, wherein:

the system generates the synthetic images by modifying features depicted in the original images; and

the feature image is an image of a meniscus.

79. The system of claim 74, wherein:

the corresponding training mask, when applied to the training image, obscures a corresponding feature depicted in the training image; and

generating the synthetic images includes, for each of the original images, applying the original image and the corresponding mask as separate inputs to the trained partial convolution model.

80. The system of claim 79, wherein the training images, the original images, and the synthetic images depict containers.

81. The system of claim 79 or 80, wherein:

the system generates the synthetic images by adding depicted defects to the original images; and

the corresponding feature is a defect depicted in the training image.

82. The system of claim 81, wherein:

training the partial convolution model includes training the partial convolution model to add a specific type of defect; and

the corresponding feature corresponds to the specific type of defect.

83. The system of claim 79 or 80, wherein:

the system generates the synthetic images by modifying features depicted in the original images; and

the corresponding feature is a meniscus.

84. The system of any one of claims 70-83, wherein training the partial convolution model includes minimizing a plurality of losses, the plurality of losses including at least a valid loss, a hole loss, a perceptual loss, a style loss, and a variation loss.

85. A method of assessing synthetic images for use in a training image library, the method comprising:

calculating, based on pixel values of a set of images, metrics indicative of differences between (i) each image in the set of images and (ii) each other image in the set of images;

generating, based on the calculated metrics, a threshold difference value; and

for each of a plurality of synthetic images,

calculating a synthetic image metric based on pixel values of the synthetic image, and

determining acceptability of the synthetic image based on the synthetic image metric and the threshold difference value.

86. The method of claim 85, further comprising:

for each of the plurality of synthetic images, either adding the synthetic image to the training image library, or omitting the synthetic image from the training image library, based on the acceptability of the synthetic image.

87. The method of claim 86, further comprising:

after adding or omitting each of the synthetic images, training a neural network for automated visual inspection using the training image library.

88. The method of claim 87, further comprising:

inspecting a plurality of container images for depicted defects using the trained neural network.

89. The method of any one of claims 85-88, wherein calculating the metrics indicative of differences between each image and each other image includes:

calculating a mean-squared-error between (i) the pixel values of each image in the set of images and (ii) the pixel values of each other image in the set of images.

90. The method of claim 89, wherein generating the threshold difference value includes:

determining a maximum mean-squared-error from among the mean-squared-errors calculated for image pairs formed from the set of images.

91. The method of claim 89 or 90, wherein calculating the synthetic image metric includes:

calculating a plurality of synthetic image mean-squared-errors, each synthetic image mean-squared-error being a mean-square error between (i) the pixel values of the synthetic image and (ii) the pixel values of a respective image in the set of images.

92. The method of claim 91, wherein determining acceptability of the synthetic image is based on (i) a maximum of the synthetic image mean-squared-errors and (ii) the maximum mean-squared-error.

93. The method of claim 91, wherein determining acceptability of the synthetic image is based on (i) an average of the synthetic image mean-squared-errors and (ii) the maximum mean-squared-error.

94. The method of any one of claims 91-93, wherein determining acceptability of the synthetic image includes:

presenting, on a display, a histogram representing the synthetic image mean-squared-errors for the synthetic images.

95. The method of any one of claims 85-94, wherein the set of images is a set of container images that have not been digitally altered and the synthetic images are digitally-altered container images.

96. A system comprising:

one or more processors; and

one or more non-transitory, computer-readable media storing instructions that, when executed by the one or more processors, cause the system to

calculate, based on pixel values of a set of images, metrics indicative of differences between (i) each image in the set of images and (ii) each other image in the set of images,

generate, based on the calculated metrics, a threshold difference value, and

for each of a plurality of synthetic images,

calculate a synthetic image metric based on pixel values of the synthetic image, and

determine acceptability of the synthetic image based on the synthetic image metric and the threshold difference value.

97. The system of claim 96, wherein calculating the metrics indicative of differences between each image and each other image includes:

calculating a mean-squared-error between (i) the pixel values of each image in the set of images and (ii) the pixel values of each other image in the set of images.

98. The system of claim 97, wherein generating the threshold difference value includes:

determining a maximum mean-squared-error from among the mean-squared-errors calculated for image pairs formed from the set of images.

99. The system of claim 97 or 98, wherein calculating the synthetic image metric includes:

calculating a plurality of synthetic image mean-squared-errors, each synthetic image mean-squared-error being a mean-square error between (i) the pixel values of the synthetic image and (ii) the pixel values of a respective image in the set of images.

100. The system of claim 99, wherein determining acceptability of the synthetic image is based on:

a maximum of the synthetic image mean-squared-errors and the maximum mean-squared-error; or

an average of the synthetic image mean-squared-errors and the maximum mean-squared-error.

101. The system of claim 99 or 100, wherein the instructions further cause the system to:

present, on a display, a histogram representing the synthetic image mean-squared-errors for the synthetic images.

102. The system of any one of claims 96-101, wherein the set of images is a set of container images that have not been digitally altered and the synthetic images are digitally-altered container images.