WHAT IS CLAIMED:

1. A method of generating a high resolution image of a scene using an imager array including a plurality of imagers that each capture an image of the scene, and a forward imaging transformation for each imager, the method comprising:

obtaining input images using the plurality of imagers;

determining an initial estimate of at least a portion of a high resolution image using a plurality of pixels from the input images;

determining a high resolution image that when mapped through the forward imaging transformation matches the input images to within at least one predetermined criterion using the initial estimate of at least a portion of the high resolution image;

wherein each forward imaging transformation corresponds to the manner in which each imager in the imaging array generate the input images; and

wherein the high resolution image has a resolution that is greater than any of the input images.

2. The method of claim 1, wherein the forward imaging transformation comprises:

a blur function; and

decimation.

3. The method of claim 2, wherein the blur function further comprises:

a lens blur function for each imager; and

a sensor blur function for each imager.

4. The method of claim 2, wherein the forward imaging transformation further comprises applying geometric transformations related to the geometry of the imagers.

5. The method of claim 2, wherein the forward imaging transformation further comprises applying geometric transformations related to parallax.

6. The method of claim 2, wherein the forward imaging transformation further comprises applying photometric transformations related to the photometric characteristics of the imagers.

7. The method of claim 1, wherein the method uses an imaging prior including photometric calibration data and obtaining input images using the plurality of imagers further comprises:

capturing images using the plurality of imagers; and

photometrically normalizing each of the captured images using the photometric calibration data to obtain the input images.

8. The method of claim 1, wherein the method uses an imaging prior including geometric calibration data and obtaining images using the plurality of imagers further comprises:

capturing images using the plurality of imagers; and

applying scene independent geometric corrections to the captured images using the geometric calibration data to obtain the input images.

9. The method of claim 1, further comprising determining scene dependent parallax information with respect to the input images.

10. The method of claim 9, wherein the parallax information includes scene dependent geometric transformations.

11. The method of claim 10, wherein the parallax information also includes occlusion maps.

12. The method of claim 1, wherein determining an initial estimate of at least a portion of a high resolution image using a plurality of pixels from the captured images further comprises:

fusing at least portions of the input images to form the initial estimate of at least one portion of the high resolution image.

13. The method of claim 12, wherein fusing at least portions of the input images to form the initial estimate of at least one portion of the high resolution image comprises:

populating a high resolution grid corresponding to the pixel locations of the at least a portion of the initial estimate of the high resolution image with pixels from the input images using geometric correction information; and

interpolating the high resolution grid to obtain filtered pixel values for each pixel in the initial estimate of the high resolution image.

14. The method of claim 13, wherein interpolating the high resolution grid to obtain filtered pixel values for each pixel in the initial estimate of the high resolution image comprises interpolating pixel values at pixel locations on the high resolution grid on which no pixel from an input image is located.

15. The method of claim 13, wherein the geometric correction information is selected from the group consisting of scene independent geometric corrections, and scene dependent parallax corrections, and combinations thereof.

16. The method of claim 13, wherein fusing at least portions of the input images to form the initial estimate of at least one portion of the high resolution image further comprises:

assigning a depth value for each pixel on the high resolution grid; and

using the depth values to direct the interpolation of the high resolution grid.

17. The method of claim 16, wherein using the depth values to direct the interpolation of the high resolution grid comprises:

assigning relative weights to the pixels that are interpolated based upon their depth value; and

interpolating the pixels using their assigned weights.

18. The method of claim 16, further comprising:

determining a high resolution occlusion map;

wherein using the depth values to direct the interpolation of the high resolution grid comprises:

identifying a pixel within an occlusion zone using the high resolution occlusion map;

identifying a neighborhood of pixels around the identified pixel; and

performing interpolation using only those pixels whose depth is greater than a threshold.

19. The method of 18, wherein the neighborhood of pixels varies in size based upon the number of pixels populated onto the high resolution grid in the neighborhood of the identified pixel.

20. The method of claim 12, wherein fusing at least portions of the input resolution images to form the initial estimate of at least one portion of the high resolution image further comprises performing filtering to remove pixels that are outliers from the high resolution grid.

21. The method of claim 12, wherein fusing at least portions of the input images to form the initial estimate of at least one portion of the high resolution image comprises:

populating a high resolution grid corresponding to the pixel locations of the at least a portion of the initial estimate of the high resolution image with pixels from the input images using geometric correction information;

obtaining at least a portion of an image from another color channel, wherein the at least a portion of the image from the other color channel is at least as high resolution as the high resolution grid; and

interpolating the high resolution grid to obtain pixel values for each pixel in the initial estimate of the high resolution image using cross correlation between the pixels on the high resolution grid and the at least a portion of the image from the other color channel.

22. The method of claim 1, wherein determining a high resolution image that when mapped through the forward imaging transformation matches the input images to within at least one predetermined criterion using the initial estimate of at least a portion of the high resolution image comprises:

transforming the initial estimate of at least a portion of the high resolution image using at least one forward imaging transformation;

comparing the transformed initial estimate of at least a portion of the high resolution image to at least a portion of at least one input image; and

refining the estimate of the high resolution image based upon the comparison.

23. The method of claim 22, wherein determining a high resolution image that when mapped through the forward imaging transformation matches the input images to within at least one predetermined criterion using the initial estimate of at least a portion of the high resolution image further comprises:

transforming, comparing and refining estimates until the at least one predetermined criterion is satisfied.

24. The method of claim 22, wherein transforming an estimate of at least a portion of the high resolution image using at least one forward imaging transformation, comprises:

applying geometric transformations to the pixels of the estimate of at least a portion of the high resolution image;

applying a blur function to the pixels of the estimate of at least a portion of the high resolution image; and

decimating the warped and blurred pixels of the estimate of at least a portion of the high resolution image.

25. The method of claim 24, wherein the blur function comprises:

a lens blur function; and

a sensor blur function.

26. The method of claim 24, wherein the geometric transformations are selected from the group consisting of scene independent geometric transformations, scene dependent parallax transformations, and combinations thereof.

27. The method of claim 24, wherein comparing a transformed estimate of at least a portion of the high resolution image to at least a portion of at least one input image comprises:

using geometric transformations to identify pixels in at least a portion of at least one input image that correspond to pixels in the transformed estimate of at least a portion of the high resolution image; and

determining differences between pixels in the transformed estimate of at least a portion of the high resolution image and the identified corresponding pixels in at least a portion of at least one input image.

28. The method of claim 27, wherein using geometric transformations to identify pixels in at least a portion of at least one input image that correspond to pixels in the transformed estimate of at a least a portion of the high resolution image comprises:

identifying the pixel in the input image specified by the geometric transformation for at least a pixel from the transformed estimate of at least a portion of the high resolution image , when a geometric transformation exists for the pixel in the transformed estimate of at least a portion of the high resolution image; and

identifying a pixel in at least one input image based upon the geometric transformations of pixels in the neighborhood of a pixel from the transformed estimate of at least a portion of the high resolution image, when a geometric transformation does

not exist for the pixel in the transformed estimate of at least a portion of the high resolution image.

29. The method of claim 28, wherein determining differences between pixels in the transformed estimate of at least a portion of the high resolution image and the identified corresponding pixels in at least a portion of at least one input image comprises:

determining the difference in value between a pixel in the transformed estimate of at least a portion of the high resolution image and each of the identified corresponding pixels in the input images;

assigning weights to the determined differences in values; and

accumulating a weighted difference using the determined differences in value and the assigned weights.

30. The method of claim 29, wherein determining differences between pixels in the transformed estimate of at least a portion of the high resolution image and the identified corresponding pixels in at least a portion of at least one input image further comprises:

determining the difference in value between a pixel in the transformed estimate of at least a portion of the high resolution image and pixels within the neighborhood of each of the identified corresponding pixels in the input images.

31. The method of claim 29, wherein assigning a weight to the determined difference in values between a pixel in the transformed estimate of at least a portion of the high resolution image and a corresponding pixel in an input image further comprises:

computing a weight based upon a decimated neighborhood of pixels surrounding the pixel in the transformed estimate of at least a portion of the high resolution image and the neighborhood of pixels surrounding the corresponding pixel in the input image.

32. The method of claim 29, further comprising accumulating the weights used to accumulate the weighted difference.

33. The method of claim 22, wherein comparing a transformed estimate of at least a portion of the high resolution image to at least a portion of at least one input image comprises:

determining differences between pixels in the transformed estimate of at least a portion of the high resolution image and pixels in at least a portion of at least one input image.

34. The method of claim 33, wherein determining differences between pixels in the transformed estimate of at least a portion of the high resolution image and pixels in at least a portion of at least one input image comprises:

determining the difference in value between a pixel in the transformed estimate of at least a portion of the high resolution image and each corresponding pixel in the input images;

assigning weights to the determined differences in values; and

filtering the differences in values using the assigned weights.

35. The method of claim 34, wherein determining differences between pixels in the transformed estimate of at least a portion of the high resolution image and pixels in at least a portion of at least one input image further comprises:

determining the difference in value between a pixel in the transformed estimate of at least a portion of the high resolution image and pixels within the neighborhood of the corresponding pixels in the input images.

36. The method of claim 34, wherein assigning a weight to the determined difference in values between a pixel in the transformed estimate of at least a portion of the high resolution image and a corresponding pixel in an input image further comprises:

computing a weight based upon a decimated neighborhood of pixels surrounding the pixel in the transformed estimate of at least a portion of the high resolution image and the neighborhood of pixels surrounding the corresponding pixel in the input image.

37. The method of claim 34, further comprising accumulating the weights used to accumulate the weighted difference.

38. The method of claim 22, wherein refining the estimate of the high resolution image based upon the comparison comprises:

mapping the comparison of the transformed initial estimate of at least a portion of the high resolution image and the at least a portion of at least one input image through a backward imaging transformation, which is the reverse of the forward imaging transformation; and

updating the estimate using at least the transformed comparison.

39. The method of claim 38, wherein the comparison of the transformed initial estimate of at least a portion of the high resolution image and the at least a portion of at

least one input image includes weighted gradients for at least a portion of the initial estimate of the high resolution image and corresponding accumulated weights.

40. The method of claim 39, wherein the weights of the weighted gradients are all equal.

41. The method of claim 39, wherein mapping the comparison of the transformed initial estimate of at least a portion of the high resolution image and the at least a portion of at least one input image through a backward imaging transformation, which is the reverse of the forward imaging transformation, comprises:

upsampling the weighted gradients and the accumulated weights.

applying a blur function to the upsampled weighted gradients and the accumulated weights;

applying geometric corrections to the blurred and upsampled weighted gradients and the accumulated weights;

accumulating the geometrically corrected blurred and upsampled weighted gradients and accumulated weights; and

normalizing the accumulated geometrically corrected, blurred and upsampled weighted gradients using the accumulated weights.

42. The method of claim 41, wherein the blur function comprises:

the transpose of a lens blur function; and

the transpose of a sensor blur function.

43. The method of claim 41, wherein the geometric corrections are selected from the group consisting of the reverse of the scene independent geometric

transformations, the reverse of the scene dependent geometric transformations, and combinations thereof.

44. The method of claim 38, wherein updating the estimate using at least the transformed comparison comprises:

modifying the initial estimate by combining the initial estimate of at least a portion of the high resolution image with at least the backward transformed comparison.

45. The method of claim 38, further comprising:

generating an intra-channel prior gradient; and

updating the estimate using an intra-channel prior gradient;

wherein the intra-channel prior gradient term is determined so that updating the estimate using the intra-channel prior gradient enforces localized image constraints.

46 The method of claim 38, wherein the imager array is configured to capture images in multiple color channels, the method further comprising:

generating an inter-channel prior gradient; and

updating the estimate using the inter-channel prior gradient;

wherein the inter-channel prior gradient is determined so that updating the estimate using the inter-channel prior gradient enforces cross-channel image constraints.

47. The method of claim 1, wherein determining a high resolution image that when mapped through the forward imaging transformation matches the input images to

within at least one predetermined criterion using the initial estimate of at least a portion of the high resolution image, comprises:

identifying pixels in the initial estimate of at least a portion of the high resolution image corresponding to pixels in at least one input image using at least one forward imaging transformation;

comparing the corresponding pixels; and

refining the estimate of the high resolution image based upon the comparison.

48. The method of claim 47, wherein determining a high resolution image that when mapped through the forward imaging transformation matches the input images to within at least one predetermined criterion using the initial estimate of at least a portion of the high resolution image, further comprises applying a blur function to pixels in the initial estimate of at least a portion of the high resolution image.

49. The method of claim 48, wherein the blur function comprises:

a lens blur function; and

a sensor blur function.

50. The method of claim 47, wherein identifying pixels in the initial estimate of at least a portion of the high resolution image corresponding to pixels in at least one input image using at least one forward imaging transformation comprises:

selecting a pixel position in the initial estimate of at least a portion of the high resolution image; and

using geometric transformations to identify pixels in at least a portion of at least one input image.

51. The method of claim 50, wherein the geometric transformations are selected from the group consisting of scene independent geometric transformations, scene dependent parallax transformations, and combinations thereof.

52. The method of claim 50, wherein using geometric transformations to identify pixels in at least a portion of at least one input image comprises:

identifying at least one pixel in the input image specified by the geometric transformation for at least the selected pixel from the initial estimate of at least a portion of the high resolution image, when a geometric transformation exists for the pixel in the initial estimate of at least a portion of the high resolution image; and

identifying at least one pixel in at least one input image based upon the geometric transformations of pixels in the neighborhood of the selected pixel from the initial estimate of at least a portion of the high resolution image, when a geometric transformation does not exist for the pixel in the initial estimate of at least a portion of the high resolution image.

53. The method of claim 47, wherein comparing corresponding pixels comprises:

determining differences between pixels in the initial estimate of at least a portion of the high resolution image and the identified corresponding pixels in at least one input image.

54. The method of 53, wherein determining differences between pixels in the initial estimate of at least a portion of the high resolution image and the identified corresponding pixels in at least a portion of at least one input image comprises:

determining the difference in value between a pixel in the initial estimate of at least a portion of the high resolution image and each of the identified corresponding pixels in the input images;

assigning weights to the determined differences in values; and

accumulating a weighted difference using the determined differences in value and the assigned weights.

55. The method of claim 54, wherein determining differences between pixels in the initial estimate of at least a portion of the high resolution image and the identified corresponding pixels in at least a portion of at least one input image further comprises:

determining the difference in value between a pixel in the initial estimate of at least a portion of the high resolution image and pixels within the neighborhood of each of the identified corresponding pixels in the input images.

56. The method of claim 54, wherein assigning a weight to the determined difference in values between a pixel in the initial estimate of at least a portion of the high resolution image and a corresponding pixel in an input image further comprises:

computing a weight based upon a decimated neighborhood of pixels surrounding the pixel in the initial estimate of at least a portion of the high resolution image and the neighborhood of pixels surrounding the corresponding pixel in the input image.

57. The method of claim 54, further comprising accumulating the weights used to accumulate the weighted difference.

58. The method of claim 54, wherein refining the estimate of the high resolution image based upon the comparison comprises:

normalizing the accumulated weighted gradients using the accumulated weights;

applying a blur function to the normalized gradients; and

updating the estimate using the blurred and normalized gradients.

59. The method of claim 58, wherein the blur function comprises:

the transpose of a lens blur function; and

the transpose of a sensor blur function.

60. The method of claim 58, wherein updating the estimate using the blurred and normalized gradients comprises:

modifying the initial estimate by combining the initial estimate of at least a portion of the high resolution image with at least the blurred and normalized gradients.

61. The method of claim 58, further comprising:

generating an intra-channel prior gradient; and

updating the estimate using an intra-channel prior gradient;

wherein the intra-channel prior gradient term is determined so that updating the estimate using the intra-channel prior gradient enforces localized image constraints.

62. The method of claim 58, wherein the imager array is configured to capture images in multiple color channels, the method further comprising:

generating an inter-channel prior gradient; and

updating the estimate using the inter-channel prior gradient;

wherein the inter-channel prior gradient is determined so that updating the estimate using the inter-channel prior gradient enforces cross-channel image constraints.

63. The method of claim 1, wherein determining a high resolution image that when mapped through the forward imaging transformation matches the input images to within at least one predetermined criterion using the initial estimate of at least a portion of the high resolution image comprises:

generating an estimate of at least a portion of the high resolution image; and

applying an intra-channel prior filter to the estimate of at least a portion of the high resolution image, where the intra-channel prior filter is configured to preserve edges while removing noise.

64. The method of claim 63, wherein the intra-channel prior filter is configured to increase the sparseness of the coefficients of a transform, when the transform is applied to the estimate of at least a portion of the high resolution image.

65. The method of claim 64, wherein increasing the sparseness further comprises thresholding of the transform coefficients according to a predetermined criterion.

66 The method of claim 65, wherein the predetermined criterion is selected from the group consisting of hard thresholding, soft thresholding, and combinations thereof.

67. The method of claim 64, wherein the transform is selected from the group consisting of sparsifying transforms, wavelets, directional transforms, and combinations thereof.

68. The method of claim 1, wherein:

the imager array captures images in multiple color channels; and

the initial estimate of at least a portion of a high resolution image is an initial estimate of at least a portion of a high resolution image in a first color channel.

69. The method of claim 68, further comprising:

placing a plurality of pixels from input images in a second color channel on a high resolution grid; and

determining at least a portion of a high resolution image in the second color channel using at least the pixels in the second color channel placed on the high resolution grid and at least a portion of a high resolution image in another color channel.

70. The method of claim 69, wherein determining at least a portion of a high resolution image in the second color channel using at least the pixels in the second color channel placed on the high resolution grid and at least a portion of a high resolution image in another color channel comprises:

interpolating the pixels on the high resolution grid based upon their correlation with the pixels in the at least a portion of the high resolution image in the other color channel and the correlation between pixels in the high resolution image in the other color channel.

71. The method of claim 70, interpolating the pixels on the high resolution grid based upon their correlation with the pixels in the at least a portion of the high resolution image in the other color channel and the correlation between pixels in the high resolution image in the other color channel comprises interpolating pixel values at pixel locations on the high resolution grid on which no pixel from an input image is located.

72. The method of claim 70, wherein the high resolution image that is determined using the initial estimate of at a least a portion of the high resolution image in a first color channel that when mapped through the forward imaging transformation matches the input images to within at least one predetermined criterion is a high resolution image in the first color channel.

73. The method of claim 70, wherein the high resolution image that is determined using the initial estimate of at a least a portion of the high resolution image in a first color channel that when mapped through the forward imaging transformation matches the input images to within at least one predetermined criterion is a high resolution image in multiple color channels.

74. The method of claim 1, wherein determining a high resolution image that when mapped through the forward imaging transformation matches the input images to within at least one predetermined criterion using the initial estimate of at least a portion of the high resolution image, comprises:

transforming pixels from an estimate of at least a portion of the high resolution image using at least one forward imaging transformation;

comparing the transformed pixels to at least a portion of at least one input image; and

refining the estimate of the high resolution image based upon the comparison.

75 The method of claim 74, wherein the pixels from the estimate that are transformed using the forward imaging transformation are selected based upon an estimated high resolution occlusion map.

76. The method of claim 74, wherein the pixels from the estimate that are transformed using the forward imaging transformation are selected based upon an estimated high resolution focus map.

77 The method of claim 74, wherein the pixels from the estimate that are transformed using the forward imaging transformation are selected based upon a predetermined threshold with respect to SNR.

78. The method of claim 74, wherein at least one portion of the initial estimate that is transformed using the forward imaging transformation is selected based upon a comparison of a previous estimate and a portion of at least one input image.

79 The method of claim 1, further comprising generating a depth map for the high resolution image.

80. The method of claim 79, wherein generating the depth map further comprises:

determining depth information for pixels in the high resolution image based upon the input images, parallax information, and the characteristics of the imager array; and

interpolating the depth information to obtain depth information for every pixel in the high resolution image.

81. The method of claim 79, wherein the depth map is used to determine a focus map.

82. The method of claim 81, wherein the focus map identifies pixels having depths in the depth map that are within a specified depth of a defined focal plane.

83. The method of claim 82, further comprising rendering the high resolution image using the focus map.

84. The method of claim 83, further comprising:

rendering the high resolution image at full resolution having a depth within a specified range of the defined focal plane;

blurring the remaining pixels in the high resolution image; and

rendering the blurred pixels.

85. The method of claim 83, further comprising:

rendering the high resolution image at full resolution having a depth within a specified range of the defined focal plane;

blurring the pixels in the input images; and

rendering the remainder of the high resolution image using the blurred pixel information from the input images.

86. The method of claim 79, wherein the depth map is used to perform depth metering.

87. The method of claim 1, wherein the high resolution image that when mapped through the forward imaging transformation matches the input images to within at least one predetermined criterion is determined with respect to a first field of view, the method further comprising:

determining a second high resolution image with respect to a second field of view;

wherein the first and second high resolution images form a stereo pair.

88. The method of claim 87, wherein determining the second high resolution image with respect to a second field of view further comprises:

determining an initial estimate of at least a portion of the second high resolution image using a plurality of pixels from the input images; and

determining a high resolution image that when mapped through the forward imaging transformation matches the input images to within at least one predetermined criterion using the initial estimate of at least a portion of the second high resolution image.

89. The method of claim 1, wherein pixels in the input images are flagged and the flagged pixels are treated as missing values when determining a high resolution image that when mapped through the forward imaging transformation matches the input images to within at least one predetermined criterion using the initial estimate of at least a portion of the high resolution image.

90. The method of claim 89, wherein the flagged pixels are also treated as missing values when determining an initial estimate of at least a portion of a high resolution image using a plurality of pixels from the input images.

91. The method of claim 1, wherein the imager array includes a plurality of imagers with fields of view that capture different magnifications of the scene.

92. The method of claim 91, wherein obtaining input images using the plurality of imagers comprises only obtaining images from imagers having fields of view of the same magnification.

93. The method of claim 91, wherein the forward imaging transformation comprises filtering pixels based upon their magnification.

94. A method of fusing a plurality of input images, the method comprising:

populating a high resolution grid corresponding to the pixel locations of the at least a portion of a fused high resolution image with pixels from the input images using geometric correction information;

assigning a depth value for each pixel on the high resolution grid; and

interpolating the high resolution grid to obtain filtered pixel values for each pixel in the initial estimate of the high resolution image, where the depth values are used to direct the interpolation of the high resolution grid.

95. The method of claim 94, wherein the geometric correction information is selected from the group consisting of scene independent geometric corrections, and scene dependent parallax corrections, and combinations thereof.

96. The method of claim 94, wherein using the depth values to direct the interpolation of the high resolution grid comprises:

assigning relative weights to the pixels that are interpolated based upon their depth value; and

interpolating the pixels using their assigned weights.

97. The method of claim 94, further comprising:

determining a high resolution occlusion map;

wherein using the depth values to direct the interpolation of the high resolution grid comprises:

identifying a pixel within an occlusion zone using the occlusion map;

identifying a neighborhood of pixels around the identified pixel; and

performing interpolation using only those pixels whose depth is greater than a threshold.

98. The method of 97, wherein the neighborhood of pixels varies in size based upon the number of pixels populated onto the high resolution grid in the neighborhood of the identified pixel.

99. The method of claim 94, wherein interpolating the high resolution grid to obtain filtered pixel values for each pixel in the initial estimate of the high resolution image, where the depth values are used to direct the interpolation of the high resolution grid comprises interpolating pixel values at pixel locations on the high resolution grid on which no pixel from an input image is located.

100. A method of generating a stereo pair of images of a scene using an imager array, comprising:

obtaining a plurality of images from the imager array;

synthesizing a first high resolution image of the scene with respect to a first field of view using pixels from at least two of the plurality of images; and

synthesizing a second high resolution image of the scene with respect to a second field of view using pixels from at least two of the plurality of images;

wherein the first and second synthesized images form a stereo pair; and

wherein the first and second high resolution images have higher resolution than the resolution of each of the images obtained from the imager array.

101 The method of claim 100, wherein synthesizing an image of a scene with respect to a specific field of view using the pixels from at least two of the plurality of images comprises:

determining an initial estimate of at least a portion of a high resolution image with respect to the field of view using a plurality of pixels from the input images; and

determining a high resolution image with respect to the field of view that when mapped through the forward imaging transformation matches the input images to within at least one predetermined criterion using the initial estimate of at least a portion of the high resolution image;

wherein each forward imaging transformation corresponds to the manner in which each imager in the imaging array generate the input images; and

wherein the high resolution image has a resolution that is greater than any of the input images.

102. The method of claim 100, wherein synthesizing an image of a scene with respect to a specific field of view using the pixels from at least two of the plurality of images comprises:

populating a grid corresponding to the pixel locations of at least a portion of the synthesized image with pixels from the input images using geometric correction information with respect to the field of view; and

interpolating the high resolution grid to obtain filtered pixel values.

103. An array camera, comprising:

an imager array including a plurality of imagers;

memory containing parameters defining a forward imaging model for the imager array; and

a processor configured to obtain a plurality of input images using the imager array and store the input images in memory;

wherein the processor is configured to determine an initial estimate of at least a portion of a high resolution image using a plurality of pixels from the input images; and

wherein the processor is configured to determine a high resolution image that when mapped through the forward imaging transformation matches the input images to within at least one predetermined criterion using the initial estimate of at least a portion of the high resolution image.

104. An array camera configured to capture stereo pairs, comprising:

an imager array including a plurality of imagers;

memory containing parameters defining a forward imaging model for the imager array; and

a processor configured to obtain a plurality of input images using the imager array and store the input images in memory;

wherein the processor is configured to synthesizing a stereo pair of high resolution images using the plurality of input images, where each image in the stereo

pair is synthesized with respect to a different field of view using pixels from at least two of the plurality of images, and is higher resolution than each of the plurality of input images.

CLAIMS:

1. A method of generating a high resolution image of a scene using an imager array including a plurality of imagers that each capture an image of the scene, and a forward imaging transformation for each imager, the method comprising:

obtaining input images captured by a plurality of imagers using a processor configured by image processing pipeline software, where the input images capture a scene in which depths of points in the imaged scene vary and each of the input images differs from the other input images due to:

scene independent geometric distortions inherent to the optics and manufacturing processes used to fabricate each of the plurality of imagers; and

scene dependent geometric displacements due to parallax experienced by each of the plurality of imagers based upon the different depths of the points in the imaged scene; and

determining scene dependent parallax information with respect to the input images based upon disparity relative to a reference point of view resulting from the different depths of points in the imaged scene using the processor configured by the image processing pipeline software, where the scene dependent parallax information comprises scene dependent geometric transformations;

determining a total shift for each of a plurality of pixels relative to the reference point of view, where the total shift of a given pixel location is the combination of a scene independent geometric correction determined for the given pixel location using geometric calibration data and the scene dependent geometric transformation determined for the given pixel location;

performing a super-resolution process utilizing at least a portion of plurality of input images and the total shift for each of the plurality of pixels relative to the reference point of view as inputs, where the super-resolution process comprises:

determining an initial estimate of at least a portion of a high resolution image from a plurality of pixels from the input images using the processor configured by the image processing pipeline software based upon a total shift for each of the plurality of pixels relative to the reference point of view;

determining a high resolution image that when mapped through a forward imaging transformation matches the input images to within at least one predetermined criterion using the processor configured using the image processing pipeline software based upon the initial estimate of at least a portion of the high resolution image;

wherein each forward imaging transformation corresponds to the manner in which each imager in the plurality of imagers captures the input images and comprises applying geometric transformations based upon the total shift at a given pixel location, which is the combination of a scene independent geometric correction determined for the given pixel location using geometric calibration data and the scene dependent geometric transformation determined for the given pixel location; and

wherein the high resolution image has a resolution that is greater than any of the input images.

2. The method of claim 1, wherein the forward imaging transformation comprises:
a blur function; and
decimation.

3. The method of claim 2, wherein the blur function further comprises:
a lens blur function for each imager; and
a sensor blur function for each imager.

4. The method of claim 2, wherein the forward imaging transformation further comprises applying scene independent geometric transformations related to the different geometries of each of the imagers in the plurality of imagers.

5. The method of claim 2, wherein the forward imaging transformation further comprises applying photometric transformations related to the different photometric characteristics of each of the imagers in the plurality of imagers.

6. The method of claim 1, wherein the method uses an imaging prior including photometric calibration data, and obtaining input images comprises photometrically normalizing each of the captured images using the photometric calibration data to obtain the input images.

7. The method of claim 1, wherein:
determining the initial estimate of at least a portion of the high resolution image comprises using an imaging prior including the geometric calibration data and applying scene independent geometric corrections to the input images using the geometric calibration data to obtain geometrically registered input images; and
determining the high resolution image that when mapped through the forward imaging transformation matches the input images to at least one predetermined criterion comprises determining the high resolution image that when mapped through the forward imaging transformation matches the geometrically registered input images to within at least one predetermined criterion.

8. The method of claim 1, wherein:
the depths of points in the imaged scene relative to a reference viewpoint vary due to the presence of foreground and background objects;
each of the input images also differ from the other input images due to occlusion zones surrounding foreground objects; and
the scene dependent parallax information also includes occlusion maps.

9. The method of claim 1, wherein determining an initial estimate of at least a portion of a high resolution image from a plurality of pixels from the captured images further comprises: fusing at least portions of the input images to form the initial estimate of at least one portion of the high resolution image.

10. The method of claim 9, wherein fusing at least portions of the input images to form the initial estimate of at least one portion of the high resolution image comprises:

populating a high resolution grid corresponding to the pixel locations of the at least a portion of the initial estimate of the high resolution image with pixels from the input images using the total shift for the pixels; and

interpolating the high resolution grid to obtain filtered pixel values for each pixel in the initial estimate of the high resolution image.

11. The method of claim 10, wherein interpolating the high resolution grid to obtain filtered pixel values for each pixel in the initial estimate of the high resolution image comprises interpolating pixel values at pixel locations on the high resolution grid on which no pixel from an input image is located.

12. The method of claim 10, wherein fusing at least portions of the input images to form the initial estimate of at least one portion of the high resolution image further comprises:

assigning a depth value for each pixel on the high resolution grid corresponding to the depth of a point in the imaged scene relative to a point of view; and

using the depth values to direct the interpolation of the high resolution grid.

13. The method of claim 12, wherein using the depth values to direct the interpolation of the high resolution grid comprises:

assigning relative weights to the pixels that are interpolated based upon their depth value; and

interpolating the pixels using their assigned weights.

14. The method of claim 12, further comprising:

determining a high resolution occlusion map;

wherein using the depth values to direct the interpolation of the high resolution grid comprises:

identifying a pixel within an occlusion zone using the high resolution occlusion map;

identifying a neighborhood of pixels around the identified pixel; and performing interpolation using only those pixels whose depth is greater than a threshold.

15. The method of 14, wherein the neighborhood of pixels varies in size based upon the number of pixels populated onto the high resolution grid in the neighborhood of the identified pixel.

16. The method of claim 9, wherein fusing at least portions of the input resolution images to form the initial estimate of at least one portion of the high resolution image further comprises performing filtering to remove pixels that are outliers from the high resolution grid.

17. The method of claim 9, wherein fusing at least portions of the input images to form the initial estimate of at least one portion of the high resolution image comprises:

populating a high resolution grid corresponding to the pixel locations of the at least a portion of the initial estimate of the high resolution image with pixels from the input images using geometric correction information;

obtaining at least a portion of an image from another color channel, wherein the at least a portion of the image from the other color channel is at least as high resolution as the high resolution grid; and

interpolating the high resolution grid to obtain pixel values for each pixel in the initial estimate of the high resolution image using cross correlation between the pixels on the high resolution grid and the at least a portion of the image from the other color channel.

18. The method of claim 1, wherein determining a high resolution image that when mapped through the forward imaging transformation matches the input images to within at least one predetermined criterion using the initial estimate of at least a portion of the high resolution image comprises:

transforming the initial estimate of at least a portion of the high resolution image using at least one forward imaging transformation;

comparing the transformed initial estimate of at least a portion of the high resolution image to at least a portion of at least one input image; and

refining the estimate of the high resolution image based upon the comparison.

19. The method of claim 18, wherein determining a high resolution image that when mapped through the forward imaging transformation matches the input images to within at least one predetermined criterion using the initial estimate of at least a portion of the high resolution image further comprises: transforming, comparing and refining estimates until the at least one predetermined criterion is satisfied.

20. The method of claim 18, wherein transforming the initial estimate of at least a portion of the high resolution image using at least one forward imaging transformation, comprises:

applying geometric transformations including geometric transformations related to parallax observed due to the different depths of points in the imaged scene to the pixels of the estimate of at least a portion of the high resolution image;

applying a blur function to the pixels of the estimate of at least a portion of the high resolution image; and

decimating the warped and blurred pixels of the estimate of at least a portion of the high resolution image.

21. The method of claim 20, wherein the blur function comprises:

a lens blur function; and

a sensor blur function.

22. The method of claim 20, wherein the geometric transformations further comprise scene independent geometric transformations.

23. The method of claim 20, wherein comparing the transformed estimate of at least a portion of the high resolution image to at least a portion of at least one input image comprises:

using geometric transformations including geometric transformations related to parallax observed due to the different depths of points in the imaged scene to identify pixels in at least a portion of at least one input image that correspond to pixels in the transformed estimate of at least a portion of the high resolution image; and

determining differences between pixels in the transformed estimate of at least a portion of the high resolution image and the identified corresponding pixels in at least a portion of at least one input image.

24. The method of claim 23, wherein using geometric transformations to identify pixels in at least a portion of at least one input image that correspond to pixels in the transformed estimate of at a least a portion of the high resolution image comprises:

identifying the pixel in the input image specified by the geometric transformation for at least a pixel from the transformed estimate of at least a portion of the high resolution image, when a geometric transformation exists for the pixel in the transformed estimate of at least a portion of the high resolution image; and

identifying a pixel in at least one input image based upon the geometric transformations of pixels in the neighborhood of a pixel from the transformed estimate of at least a portion of the high resolution image, when a geometric transformation does not exist for the pixel in the transformed estimate of at least a portion of the high resolution image.

25. The method of claim 24, wherein determining differences between pixels in the transformed estimate of at least a portion of the high resolution image and the identified corresponding pixels in at least a portion of at least one input image comprises:

determining the difference in value between a pixel in the transformed estimate of at least a portion of the high resolution image and each of the identified corresponding pixels in the input images;

assigning weights to the determined differences in values; and

accumulating a weighted difference using the determined differences in value and the assigned weights.

26. The method of claim 25, wherein determining differences between pixels in the transformed estimate of at least a portion of the high resolution image and the identified corresponding pixels in at least a portion of at least one input image further comprises:

determining the difference in value between a pixel in the transformed estimate of at least a portion of the high resolution image and pixels within the neighborhood of each of the identified corresponding pixels in the input images.

27. The method of claim 25, wherein assigning a weight to the determined difference in values between a pixel in the transformed estimate of at least a portion of the high resolution image and a corresponding pixel in an input image further comprises:

computing a weight based upon a decimated neighborhood of pixels surrounding the pixel in the transformed estimate of at least a portion of the high resolution image and the neighborhood of pixels surrounding the corresponding pixel in the input image.

28. The method of claim 25, further comprising accumulating the weights used to accumulate the weighted difference.

29. The method of claim 18, wherein comparing the transformed estimate of at least a portion of the high resolution image to at least a portion of at least one input image comprises:

determining differences between pixels in the transformed estimate of at least a portion of the high resolution image and pixels in at least a portion of at least one input image.

30. The method of claim 29, wherein determining differences between pixels in the transformed estimate of at least a portion of the high resolution image and pixels in at least a portion of at least one input image comprises:

determining the difference in value between a pixel in the transformed estimate of at least a portion of the high resolution image and each corresponding pixel in the input images;

assigning weights to the determined differences in values; and

filtering the differences in values using the assigned weights.

31. The method of claim 30, wherein determining differences between pixels in the transformed estimate of at least a portion of the high resolution image and pixels in at least a portion of at least one input image further comprises:

determining the difference in value between a pixel in the transformed estimate of at least a portion of the high resolution image and pixels within the neighborhood of the corresponding pixels in the input images.

32. The method of claim 30, wherein assigning a weight to the determined difference in values between a pixel in the transformed estimate of at least a portion of the high resolution image and a corresponding pixel in an input image further comprises:

computing a weight based upon a decimated neighborhood of pixels surrounding the pixel in the transformed estimate of at least a portion of the high resolution image and the neighborhood of pixels surrounding the corresponding pixel in the input image.

33. The method of claim 30, further comprising accumulating the weights used to accumulate the weighted difference.

34. The method of claim 18, wherein refining the estimate of the high resolution image based upon the comparison comprises:

mapping the comparison of the transformed initial estimate of at least a portion of the high resolution image and the at least a portion of at least one input image through a backward imaging transformation, which is the reverse of the forward imaging transformation; and

updating the estimate using at least the transformed comparison.

35. The method of claim 34, wherein the comparison of the transformed initial estimate of at least a portion of the high resolution image and the at least a portion of at least one input image includes weighted gradients for at least a portion of the initial estimate of the high resolution image and corresponding accumulated weights.

36. The method of claim 35, wherein the weights of the weighted gradients are all equal.

37. The method of claim 35, wherein mapping the comparison of the transformed initial estimate of at least a portion of the high resolution image and the at least a portion of at least one input image through a backward imaging transformation, which is the reverse of the forward imaging transformation, comprises:

upsampling the weighted gradients and the accumulated weights; applying a blur function to the upsampled weighted gradients and the accumulated weights;

applying geometric corrections including geometric transformations related to parallax observed due to the different depths of points in the imaged scene to the blurred and upsampled weighted gradients and the accumulated weights;

accumulating the geometrically corrected blurred and upsampled weighted gradients and accumulated weights; and

normalizing the accumulated geometrically corrected, blurred and upsampled weighted gradients using the accumulated weights.

38. The method of claim 37, wherein the blur function comprises:

the transpose of a lens blur function; and

the transpose of a sensor blur function.

39. The method of claim 37, wherein the geometric transformations further comprise scene independent geometric transformations.

40. The method of claim 34, wherein updating the estimate using at least the transformed comparison comprises:

modifying the initial estimate by combining the initial estimate of at least a portion of the high resolution image with at least the backward transformed comparison.

41. The method of claim 34, further comprising:

generating a local intra-channel prior gradient for individual pixel locations in the estimate of at least a portion of the high resolution image; and

updating the pixels in the estimate using the intra-channel prior gradient determined for the individual pixel locations;

wherein the intra-channel prior gradient term is determined so that updating the estimate using the intra-channel prior gradient enforces localized image constraints.

42. The method of claim 34, wherein the plurality of imagers is configured to capture images in multiple color channels, the method further comprising:

generating a local inter-channel prior gradient for individual pixel locations in the estimate of at least a portion of the high resolution image; and

updating the pixels in the estimate using the inter-channel prior gradient determined for the individual pixel locations;

wherein the inter-channel prior gradient is determined so that updating the estimate using the inter-channel prior gradient enforces cross-channel image constraints.

43. The method of claim 1, wherein determining a high resolution image that when mapped through the forward imaging transformation matches the input images to within at least one

predetermined criterion using the initial estimate of at least a portion of the high resolution image, comprises:

identifying pixels in the initial estimate of at least a portion of the high resolution image corresponding to pixels in at least one input image using at least one forward imaging transformation;

comparing the corresponding pixels; and

refining the estimate of the high resolution image based upon the comparison.

44. The method of claim 43, wherein determining a high resolution image that when mapped through the forward imaging transformation matches the input images to within at least one predetermined criterion using the initial estimate of at least a portion of the high resolution image, further comprises applying a blur function to pixels in the initial estimate of at least a portion of the high resolution image.

45. The method of claim 44, wherein the blur function comprises:

a lens blur function; and

a sensor blur function.

46. The method of claim 43, wherein identifying pixels in the initial estimate of at least a portion of the high resolution image corresponding to pixels in at least one input image using at least one forward imaging transformation comprises:

selecting a pixel position in the initial estimate of at least a portion of the high resolution image; and

using geometric transformations including geometric transformations related to parallax observed due to the different depths of points in the imaged scene to identify pixels in at least a portion of at least one input image.

47. The method of claim 46, wherein the geometric transformations further comprise scene independent geometric transformations.

48. The method of claim 46, wherein using geometric transformations to identify pixels in at least a portion of at least one input image comprises:

identifying at least one pixel in the input image specified by the geometric transformation for at least the selected pixel from the initial estimate of at least a portion of the high resolution image, when a geometric transformation exists for the pixel in the initial estimate of at least a portion of the high resolution image; and

identifying at least one pixel in at least one input image based upon the geometric transformations of pixels in the neighborhood of the selected pixel from the initial estimate of at least a portion of the high resolution image, when a geometric transformation does not exist for the pixel in the initial estimate of at least a portion of the high resolution image.

49. The method of claim 43, wherein comparing corresponding pixels comprises:

determining differences between pixels in the initial estimate of at least a portion of the high resolution image and the identified corresponding pixels in at least one input image.

50. The method of 49, wherein determining differences between pixels in the initial estimate of at least a portion of the high resolution image and the identified corresponding pixels in at least a portion of at least one input image comprises:

determining the difference in value between a pixel in the initial estimate of at least a portion of the high resolution image and each of the identified corresponding pixels in the input images;

assigning weights to the determined differences in values; and

accumulating a weighted difference for the pixel in the initial estimate of at least a portion of the high resolution image using the determined differences in value and the assigned weights.

51. The method of claim 50, wherein determining differences between pixels in the initial estimate of at least a portion of the high resolution image and the identified corresponding pixels in at least a portion of at least one input image further comprises:

determining the difference in value between a pixel in the initial estimate of at least a portion of the high resolution image and pixels within the neighborhood of each of the identified corresponding pixels in the input images.

52. The method of claim 50, wherein assigning a weight to the determined difference in values between a pixel in the initial estimate of at least a portion of the high resolution image and a corresponding pixel in an input image further comprises:

computing a weight based upon a decimated neighborhood of pixels surrounding the pixel in the initial estimate of at least a portion of the high resolution image and the neighborhood of pixels surrounding the corresponding pixel in the input image.

53. The method of claim 50, further comprising accumulating the weights used to accumulate the weighted difference for the pixel in the initial estimate of at least a portion of the high resolution image.

54. The method of claim 50, wherein refining the estimate of the high resolution image based upon the comparison comprises:

normalizing the accumulated weighted gradients for the pixel in the initial estimate of at least a portion of the high resolution image using the accumulated weights;

applying a blur function to the normalized gradients; and

updating the estimate using the blurred and normalized gradients.

55. The method of claim 54, wherein the blur function comprises:

the transpose of a lens blur function; and

the transpose of a sensor blur function.

56. The method of claim 54, wherein updating the estimate using the blurred and normalized gradients comprises:

modifying the initial estimate by combining the initial estimate of at least a portion of the high resolution image with at least the blurred and normalized gradients.

57. The method of claim 54, further comprising:

generating a local intra-channel prior gradient for individual pixel locations in the estimate of at least a portion of the high resolution image; and

updating the pixels in the estimate using the intra-channel prior gradient determined for the individual pixel locations;

wherein the intra-channel prior gradient term is determined so that updating the estimate using the intra-channel prior gradient enforces localized image constraints.

58. The method of claim 54, wherein the plurality of imagers is configured to capture images in multiple color channels, the method further comprising:

generating a local inter-channel prior gradient for individual pixel locations in the estimate of at least a portion of the high resolution image; and

updating the pixels in the estimate using the inter-channel prior gradient for individual pixel locations in the estimate of at least a portion of the high resolution image;

wherein the inter-channel prior gradient is determined so that updating the estimate using the inter-channel prior gradient enforces cross-channel image constraints.

59. The method of claim 1, wherein determining a high resolution image that when mapped through the forward imaging transformation matches the input images to within at least one predetermined criterion using the initial estimate of at least a portion of the high resolution image comprises:

generating an estimate of at least a portion of the high resolution image; and

applying an intra-channel prior filter to the estimate of at least a portion of the high resolution image, where the intra-channel prior filter is configured to preserve edges while removing noise.

60. The method of claim 59, wherein the intra-channel prior filter is configured to increase the sparseness of the coefficients of a transform, when the transform is applied to the estimate of at least a portion of the high resolution image.

61. The method of claim 60, wherein increasing the sparseness further comprises thresholding of the transform coefficients according to a predetermined criterion.

62. The method of claim 61, wherein the predetermined criterion is selected from the group consisting of hard thresholding, soft thresholding, and combinations thereof.

63. The method of claim 60, wherein the transform is selected from the group consisting of sparsifying transforms, wavelets, directional transforms, and combinations thereof.

64. The method of claim 1, wherein:

the imager array captures images in multiple color channels; and

the initial estimate of at least a portion of a high resolution image is an initial estimate of at least a portion of a high resolution image in a first color channel.

65. The method of claim 64, further comprising:

placing a plurality of pixels from input images in a second color channel on a high resolution grid; and

determining at least a portion of a high resolution image in the second color channel using at least the pixels in the second color channel placed on the high resolution grid and at least a portion of a high resolution image in another color channel.

66. The method of claim 65, wherein determining at least a portion of a high resolution image in the second color channel using at least the pixels in the second color channel placed on the high resolution grid and at least a portion of a high resolution image in another color channel comprises:

interpolating the pixels on the high resolution grid based upon their correlation with the pixels in the at least a portion of the high resolution image in the other color channel and the correlation between pixels in the high resolution image in the other color channel.

67. The method of claim 66, interpolating the pixels on the high resolution grid based upon their correlation with the pixels in the at least a portion of the high resolution image in the other color channel and the correlation between pixels in the high resolution image in the other color channel comprises interpolating pixel values at pixel locations on the high resolution grid on which no pixel from an input image is located.

68. The method of claim 66, wherein the high resolution image that is determined using the initial estimate of at a least a portion of the high resolution image in a first color channel that when mapped through the forward imaging transformation matches the input images to within at least one predetermined criterion is a high resolution image in the first color channel.

69. The method of claim 66, wherein the high resolution image that is determined using the initial estimate of at a least a portion of the high resolution image in a first color channel that when mapped through the forward imaging transformation matches the input images to within at least one predetermined criterion is a high resolution image in multiple color channels.

70. The method of claim 1, wherein determining a high resolution image that when mapped through the forward imaging transformation matches the input images to within at least one predetermined criterion using the initial estimate of at least a portion of the high resolution image, comprises:

transforming pixels from an estimate of at least a portion of the high resolution image using at least one forward imaging transformation;

comparing the transformed pixels to at least a portion of at least one input image; and

refining the estimate of the high resolution image based upon the comparison.

71. The method of claim 70, wherein the pixels from the estimate that are transformed using the forward imaging transformation are selected based upon an estimated high resolution occlusion map.

72. The method of claim 70, wherein the pixels from the estimate that are transformed using the forward imaging transformation are selected based upon an estimated high resolution focus map.

73. The method of claim 70, wherein the pixels from the estimate that are transformed using the forward imaging transformation are selected based upon a predetermined threshold with respect to SNR.

74. The method of claim 70, wherein at least one portion of the initial estimate that is transformed using the forward imaging transformation is selected based upon a comparison of a previous estimate and a portion of at least one input image.

75. The method of claim 1, further comprising generating a depth map for the high resolution image.

76. The method of claim 75, wherein generating the depth map further comprises:

determining depth information for pixels in the high resolution image based upon the input images, scene dependent parallax information with respect to the input images based upon disparity relative to the reference viewpoint resulting from the different depths of points in the imaged scene, and the characteristics of the imager array; and

interpolating the depth information to obtain depth information for every pixel in the high resolution image.

77. The method of claim 75, wherein the depth map is used to determine a focus map.

78. The method of claim 77, wherein the focus map identifies pixels having depths in the depth map that are within a specified depth of a defined focal plane.

79. The method of claim 78, further comprising rendering the high resolution image using the focus map.

80. The method of claim 79, further comprising:

rendering the high resolution image at full resolution with respect to pixels having a depth indicated as being within a specified range of the defined focal plane by the depth map;

blurring the remaining pixels in the high resolution image; and rendering the blurred pixels.

81. The method of claim 79, further comprising:

rendering the high resolution image at full resolution with respect to pixels having a depth indicated as being within a specified range of the defined focal plane by the depth map;

blurring the pixels in the input images; and

rendering the remainder of the high resolution image using the blurred pixel information from the input images.

82. The method of claim 75, wherein the depth map is used to perform depth metering.

83. The method of claim 1, wherein the high resolution image that when mapped through the forward imaging transformation matches the input images to within at least one predetermined criterion is determined with respect to a first field of view, the method further comprising:

determining a second high resolution image with respect to a second field of view;

wherein the first and second high resolution images form a stereo pair.

84. The method of claim 83, wherein determining the second high resolution image with respect to a second field of view further comprises:

determining an initial estimate of at least a portion of the second high resolution image using a plurality of pixels from the input images; and

determining a high resolution image that when mapped through the forward imaging transformation matches the input images to within at least one predetermined criterion using the initial estimate of at least a portion of the second high resolution image.

85. The method of claim 1, wherein pixels in the input images are flagged and the flagged pixels are treated as missing values when determining a high resolution image that when mapped through

the forward imaging transformation matches the input images to within at least one predetermined criterion using the initial estimate of at least a portion of the high resolution image.

86. The method of claim 85, wherein the flagged pixels are also treated as missing values when determining an initial estimate of at least a portion of a high resolution image using a plurality of pixels from the input images.

87. The method of claim 1, wherein the imager array includes a plurality of imagers with fields of view that capture different magnifications of the scene.

88. The method of claim 87, wherein obtaining input images using the plurality of imagers comprises only obtaining images from imagers having fields of view of the same magnification.

89. The method of claim 87, wherein the forward imaging transformation comprises filtering pixels based upon their magnification.

90. A method of fusing a plurality of input images, the method comprising:

populating a high resolution grid corresponding to the pixel locations of the at least a portion of a fused high resolution image with pixels from the input images using geometric correction information using a processor configured by image processing pipeline software, where:

the input images capture a scene in which depths of points in the imaged scene vary and each of the input images differs from the other input images due to:

scene independent geometric distortions inherent to the optics and manufacturing processes used to fabricate each of the plurality of imagers; and

scene dependent geometric displacements due to parallax experienced by each of the plurality of imagers based upon the different depths of the points in the imaged scene; and

the geometric correction information comprises scene dependent geometric transformations determined with respect to the input images relative to a reference point of view based upon disparity relative to the reference point of view resulting from the different depths of points in the imaged scene;

assigning a depth value for each pixel on the high resolution grid, using the processor configured by the image processing pipeline software based upon the scene dependent geometric transformations determined with respect to the input images, based upon disparity relative to the reference point of view resulting from the different depths of points in the imaged scene; and

interpolating the high resolution grid to obtain filtered pixel values for each pixel in the initial estimate of the high resolution image, where the depth values are used to direct the interpolation of the high resolution grid.

91. The method of claim 90, wherein the geometric transformations further comprise scene independent geometric transformations.

92. The method of claim 90, wherein using the depth values to direct the interpolation of the high resolution grid comprises:

assigning relative weights to the pixels that are interpolated based upon their depth value; and

interpolating the pixels using their assigned weights.

93. The method of claim 90, further comprising:

determining a high resolution occlusion map;

wherein using the depth values to direct the interpolation of the high resolution grid comprises:

identifying a pixel within an occlusion zone using the occlusion map;

identifying a neighborhood of pixels around the identified pixel; and

performing interpolation using only those pixels whose depth is greater than a threshold.

94. The method of 93, wherein the neighborhood of pixels varies in size based upon the number of pixels populated onto the high resolution grid in the neighborhood of the identified pixel.

95. The method of claim 90, wherein interpolating the high resolution grid to obtain filtered pixel values for each pixel in the initial estimate of the high resolution image, where the depth values are used to direct the interpolation of the high resolution grid comprises interpolating pixel values at pixel locations on the high resolution grid on which no pixel from an input image is located.

96. An array camera, comprising:

an imager array including a plurality of imagers;

memory containing parameters defining a forward imaging transformation for the imager array; and

a processor configured by image processing pipeline software to obtain a plurality of input images using the imager array and store the input images in memory, where the input images capture a scene in which depths of points in the imaged scene vary and each of the input images differs from the other input images due to:

scene independent geometric distortions inherent to the optics and manufacturing processes used to fabricate each of the plurality of imagers; and

scene dependent geometric displacements due to parallax experienced by each of the plurality of imagers based upon the different depths of the points in the imaged scene; and

wherein the processor is configured by image processing pipeline software to:

determining scene dependent parallax information with respect to the input images based upon disparity relative to a reference point of view resulting from the different depths of points in the imaged scene using the processor configured by the image processing pipeline

software, where the scene dependent parallax information comprises scene dependent geometric transformations;

determining a total shift for each of a plurality of pixels relative to the reference point of view, where the total shift of a given pixel location is the combination of a scene independent geometric correction determined for the given pixel using geometric calibration data and the scene dependent geometric transformation determined for the given pixel location;

determine an initial estimate of at least a portion of a high resolution image from a plurality of pixels from the input images based upon the total shift for each of the plurality of pixels relative to the reference point of view; and

determine a high resolution image that when mapped through the forward imaging transformation matches the input images to within at least one predetermined criterion using the initial estimate of at least a portion of the high resolution image;

wherein the forward imaging transformation corresponds to the manner in which each imager in the plurality of images captures the input images and comprises applying geometric transformations based upon the total shift at a given pixel location, which is the combination of a scene independent geometric correction determined for the given pixel location using geometric calibration data and the scene dependent geometric transformation determined for the given pixel location; and

wherein the high resolution image has a resolution that is greater than any of the input images.