

SYSTEMS AND METHODS FOR SYNTHESIZING HIGH RESOLUTION IMAGES USING SUPER-RESOLUTION PROCESSES

FIELD OF THE INVENTION

[0001] The present invention generally relates to imaging and more specifically relates to array cameras that utilize super resolution processes to synthesize high resolution images from multiple low resolution images.

BACKGROUND

[0002] In a typical imaging device, light enters through an opening (aperture) at one end of the imaging device and is directed to an image sensor by one or more optical elements such as lenses. The image sensor consists of pixels that generate signals upon receiving light via the optical element. Commonly used image sensors include charge-coupled device image sensors (CCDs) and complementary metal-oxide semiconductor (CMOS) sensors.

[0003] Conventional digital cameras typically achieve color separation by performing color separation in the optical path and using a separate image sensor for the wavelengths of light corresponding to each of the primary colors (i.e. RGB), using an image sensor with color separation and multiple signal collection capability within each pixel, or by applying filters over a single sensor so that individual pixels detect wavelengths of light corresponding to one of the primary colors. Use of filters is particularly common in cameras that have a small form factor, such as cameras incorporated in mobile phone handsets and other consumer electronics devices including but not limited to, laptop computers and televisions. A common filter that is formed on image sensors is the Bayer filter, the pattern of which includes 50% green filters, 25% red filters, and 25% blue filters. The output of an image sensor to which a Bayer filter is applied can be reconstructed as a color image using interpolation techniques.

[0004] Image sensors are subject to various performance constraints including, among others, dynamic range, signal to noise (SNR) ratio and low light sensitivity. The dynamic range is defined as the ratio of the maximum possible signal that can be captured by a pixel to the total noise signal. The SNR of a captured image is, to a great extent, a measure of image quality. In general, as more light is captured by the pixel, the higher the SNR. The light sensitivity of an image sensor is typically determined by the intensity of light incident upon the sensor pixels. At low light levels, each pixel's light gathering capability is constrained by the low signal levels incident upon each pixel.

[0005] A challenge associated with increasing the number of pixels in an image sensor is that the lens system is dimensioned to span the image sensor. The problem is most acute with mobile cameras, such as those used in mobile phones and consumer electronics devices, where the form factor of the lens system can significantly impact the overall form factor of the mobile device.

[0006] In response to the constraints placed upon a traditional digital camera based upon the camera *obscura*, a new class of cameras that can be referred to as array cameras have been proposed. Array cameras are characterized in that they include multiple arrays of pixels, each having a separate lens system. Examples of 2, 3 and 4 array cameras in which each array of pixels captures light from a different band of the visible spectrum and the captured images are combined to create a full color image is disclosed in U.S. Patent 7,199,348 to Olsen et al., the disclosure of which is incorporated by reference herein in its entirety. U.S. Patent 7,262,799 to Suda, the disclosure of which is incorporated herein by reference in its entirety, discloses a 2 x 2 array camera including one sensor used to sense a red (R) image signal, one sensor used to sense a blue (B) image signal and, two sensors used to sense green (G) image signals.

SUMMARY OF THE INVENTION

[0007] Systems and methods are disclosed that use super-resolution (SR) processes to fuse information from a plurality of low resolution images captured by an imager array to synthesize a high resolution image. In many embodiments, the objective and subjective quality of the obtained super-resolution image is increased through signal restoration. In several embodiments, the SR process incorporates cross-channel fusion. In a number of embodiments, the imager array includes imagers having different fields of view. In many embodiments, aliasing is introduced into the low resolution images to enable improved recovery of high frequency information through SR processing.

[0008] One embodiment includes 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. In addition, 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.

[0009] In a further embodiment, the forward imaging transformation includes a blur function, and decimation.

[0010] In another embodiment, the blur function further includes a lens blur function for each imager, and a sensor blur function for each imager.

[0011] In a still further embodiment, the forward imaging transformation further includes applying geometric transformations related to the geometry of the imagers.

[0012] In still another embodiment, the forward imaging transformation further includes applying geometric transformations related to parallax.

[0013] In a yet further embodiment, the forward imaging transformation further includes applying photometric transformations related to the photometric characteristics of the imagers.

[0014] In yet another embodiment, the method uses an imaging prior including photometric calibration data and obtaining input images using the plurality of imagers further includes 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.

[0015] In a further embodiment again, the method uses an imaging prior including geometric calibration data and obtaining images using the plurality of imagers further includes 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.

[0016] Another embodiment again further includes determining scene dependent parallax information with respect to the input images.

[0017] In a further additional embodiment, the parallax information includes scene dependent geometric transformations.

[0018] In another additional embodiment, the parallax information also includes occlusion maps.

[0019] In a still yet further embodiment, determining an initial estimate of at least a portion of a high resolution image using a plurality of pixels from the captured images further includes fusing at least portions of the input images to form the initial estimate of at least one portion of the high resolution image.

[0020] In still yet another embodiment, fusing at least portions of the input images to form the initial estimate of at least one portion of the high resolution image includes 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.

[0021] In a still further embodiment again, interpolating the high resolution grid to obtain filtered pixel values for each pixel in the initial estimate of the high resolution image includes interpolating pixel values at pixel locations on the high resolution grid on which no pixel from an input image is located.

[0022] In still another embodiment again, the geometric correction information is selected from the group made up of scene independent geometric corrections, and scene dependent parallax corrections, and combinations thereof.

[0023] In a still further additional embodiment, fusing at least portions of the input images to form the initial estimate of at least one portion of the high resolution image further includes 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.

[0024] In still another additional embodiment, using the depth values to direct the interpolation of the high resolution grid includes assigning relative weights to the pixels that are interpolated based upon their depth value, and interpolating the pixels using their assigned weights.

[0025] A yet further embodiment again also includes determining a high resolution occlusion map. In addition, using the depth values to direct the interpolation of the high resolution grid includes 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.

[0026] In yet another embodiment again, 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.

[0027] In a yet further additional embodiment, 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 includes performing filtering to remove pixels that are outliers from the high resolution grid.

[0028] In yet another additional embodiment, fusing at least portions of the input images to form the initial estimate of at least one portion of the high resolution image includes 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.

[0029] In a further additional embodiment again, 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 includes 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.

[0030] In another additional embodiment again, 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 includes transforming, comparing and refining estimates until the at least one predetermined criterion is satisfied.

[0031] In a still yet further embodiment again, transforming an estimate of at least a portion of the high resolution image using at least one forward imaging transformation, includes 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.

[0032] In still yet another embodiment again, the blur function includes a lens blur function, and a sensor blur function.

[0033] In a still yet further additional embodiment, the geometric transformations are selected from the group made up of scene independent geometric transformations, scene dependent parallax transformations, and combinations thereof.

[0034] In still yet another additional embodiment, 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 includes 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.

[0035] In a still yet further additional embodiment again, 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 includes 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.

[0036] In still yet another additional embodiment again, 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 includes 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.

[0037] In an alternate embodiment, 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 includes 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.

[0038] In a further alternate embodiment, 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 includes 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.

[0039] Another alternate embodiment further includes accumulating the weights used to accumulate the weighted difference.

[0040] In a still further alternate embodiment, 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 includes 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.

[0041] In a yet further alternate embodiment, 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 includes 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.

[0042] In yet another alternate embodiment, 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 includes 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.

[0043] In a further alternate embodiment again, 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 includes 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.

[0044] Another alternate embodiment again also includes accumulating the weights used to accumulate the weighted difference.

[0045] In a further alternate additional embodiment, refining the estimate of the high resolution image based upon the comparison includes 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.

[0046] In another alternate additional embodiment, 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.

[0047] In a still yet further alternate embodiment, the weights of the weighted gradients are all equal.

[0048] In still yet another alternate embodiment, 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, includes upsampling the weighted gradients and the accumulated weights.

[0049] A still further alternate embodiment again also includes applying a blur function to the upsampled weighted gradients and the accumulated weights, and 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.

[0050] In still another alternate embodiment again, the blur function includes the transpose of a lens blur function, and the transpose of a sensor blur function.

[0051] In a still further alternate additional embodiment, the geometric corrections are selected from the group made up of the reverse of the scene independent geometric transformations, the reverse of the scene dependent geometric transformations, and combinations thereof.

[0052] In still another alternate additional embodiment, updating the estimate using at least the transformed comparison includes 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.

[0053] A yet further alternate embodiment again also includes generating an intra-channel prior gradient, and updating the estimate using an intra-channel prior gradient. In addition, the intra-channel prior gradient term is determined so that updating the estimate using the intra-channel prior gradient enforces localized image constraints.

[0054] In yet another alternate embodiment again, the imager array is configured to capture images in multiple color channels, the method further includes generating an inter-channel prior gradient, and updating the estimate using the inter-channel prior gradient. In addition, the inter-channel prior gradient is determined so that updating the estimate using the inter-channel prior gradient enforces cross-channel image constraints.

[0055] In a yet further alternate additional embodiment, 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, includes 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.

[0056] In yet another alternate additional embodiment, 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 includes applying a blur function to pixels in the initial estimate of at least a portion of the high resolution image.

[0057] In a further alternate additional embodiment again, the blur function includes a lens blur function, and a sensor blur function.

[0058] In another alternate additional embodiment again, 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 includes 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.

[0059] In a still yet further alternate embodiment again, the geometric transformations are selected from the group made up of scene independent geometric transformations, scene dependent parallax transformations, and combinations thereof.

[0060] In still yet another alternate embodiment again, using geometric transformations to identify pixels in at least a portion of at least one input image includes 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.

[0061] In a still yet further alternate additional embodiment, comparing corresponding pixels includes 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.

[0062] In still yet another alternate additional embodiment, 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 includes 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.

[0063] In a still yet further alternate additional embodiment again, 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 includes 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.

[0064] In still yet another alternate additional embodiment again, 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 includes 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.

[0065] Another further embodiment also includes accumulating the weights used to accumulate the weighted difference.

[0066] In still another further embodiment, refining the estimate of the high resolution image based upon the comparison includes 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.

[0067] In yet another further embodiment, the blur function includes the transpose of a lens blur function, and the transpose of a sensor blur function.

[0068] In another further embodiment again, updating the estimate using the blurred and normalized gradients includes 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.

[0069] Another further additional embodiment, further includes generating an intra-channel prior gradient, and updating the estimate using an intra-channel prior gradient. In addition, the intra-channel prior gradient term is determined so that updating the estimate using the intra-channel prior gradient enforces localized image constraints.

[0070] In still yet another further embodiment, the imager array is configured to capture images in multiple color channels, the method further includes generating an inter-channel prior gradient, and updating the estimate using the inter-channel prior gradient. In addition, the inter-channel prior gradient is determined so that updating the estimate using the inter-channel prior gradient enforces cross-channel image constraints.

[0071] In still another further embodiment again, 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 includes 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.

[0072] In still another further additional embodiment, 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.

[0073] In yet another further embodiment again, increasing the sparseness further comprises thresholding of the transform coefficients according to a predetermined criterion.

[0074] In yet another further additional embodiment, the predetermined criterion is selected from the group made up of hard thresholding, soft thresholding, and combinations thereof.

[0075] In another further additional embodiment again, the transform is selected from the group made up of sparsifying transforms, wavelets, directional transforms, and combinations thereof.

[0076] In still yet another further embodiment again, 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.

[0077] Still yet another further additional embodiment also includes 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.

[0078] In still another further additional embodiment again, 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 includes 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.

[0079] In yet another further additional embodiment again, 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.

[0080] In still yet another further additional embodiment again, 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.

[0081] In another further alternate embodiment, 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.

[0082] In still another further alternate embodiment, 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, includes 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.

[0083] In yet another further alternate embodiment, the pixels from the estimate that are transformed using the forward imaging transformation are selected based upon an estimated high resolution occlusion map.

[0084] In another further alternate embodiment again, the pixels from the estimate that are transformed using the forward imaging transformation are selected based upon an estimated high resolution focus map.

[0085] In another further alternate additional embodiment, the pixels from the estimate that are transformed using the forward imaging transformation are selected based upon a predetermined threshold with respect to SNR.

[0086] In still yet another further alternate embodiment, 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.

[0087] Still another further alternate embodiment again also includes generating a depth map for the high resolution image.

[0088] In still another further alternate additional embodiment, generating the depth map further includes 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.

[0089] In yet another further alternate embodiment again, the depth map is used to determine a focus map.

[0090] In yet another further alternate additional embodiment, the focus map identifies pixels having depths in the depth map that are within a specified depth of a defined focal plane.

[0091] Another further alternate additional embodiment again also includes rendering the high resolution image using the focus map.

[0092] Still yet another further alternate embodiment again also includes 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.

[0093] Still yet another further alternate additional embodiment also includes 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.

[0094] In still yet another further alternate additional embodiment again, the depth map is used to perform depth metering.

[0095] In a specific embodiment, 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 includes determining a second high resolution image with respect to a second field of view. In addition, the first and second high resolution images form a stereo pair.

[0096] In another specific embodiment, determining the second high resolution image with respect to a second field of view further includes 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.

[0097] In still another specific embodiment, 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.

[0098] In yet another specific embodiment, 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.

[0099] In another specific embodiment again, the imager array includes a plurality of imagers with fields of view that capture different magnifications of the scene.

[00100] In another specific additional embodiment, obtaining input images using the plurality of imagers comprises only obtaining images from imagers having fields of view of the same magnification.

[00101] In still yet another specific embodiment, the forward imaging transformation includes filtering pixels based upon their magnification.

[00102] Still another specific embodiment again includes 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.

[00103] In still another specific additional embodiment, the geometric correction information is selected from the group made up of scene independent geometric corrections, and scene dependent parallax corrections, and combinations thereof.

[00104] In yet another specific embodiment again, using the depth values to direct the interpolation of the high resolution grid includes assigning relative weights to the pixels

that are interpolated based upon their depth value, and interpolating the pixels using their assigned weights.

[00105] Yet another specific additional embodiment, further includes determining a high resolution occlusion map. In addition, using the depth values to direct the interpolation of the high resolution grid includes 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.

[00106] In another specific additional embodiment again, 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.

[00107] In still yet another specific embodiment again, 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 includes interpolating pixel values at pixel locations on the high resolution grid on which no pixel from an input image is located.

[00108] Still yet another specific additional embodiment includes 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. In addition, the first and second synthesized images form a stereo pair, and the first and second high resolution images have higher resolution than the resolution of each of the images obtained from the imager array.

[00109] In still another specific additional embodiment again, 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 includes 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. In addition, each forward imaging transformation corresponds to the manner in which each imager in the imaging array generate the input images, and the high resolution image has a resolution that is greater than any of the input images.

[00110] In yet another specific additional embodiment again, 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 includes 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.

[00111] Another further specific embodiment includes 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. In addition, 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 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.

[00112] Still another further specific embodiment includes 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. In addition, 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.

BRIEF DESCRIPTION OF THE DRAWINGS

[00113] FIG. 1 is a block diagram of an array camera in accordance with an embodiment of the invention.

[00114] FIGS. 2A – 2D illustrate imager configurations of imager arrays in accordance with embodiments of the invention.

[00115] FIG. 3 illustrates an image processing pipeline in accordance with an embodiment of the invention.

[00116] FIG. 4 illustrates a SR module in accordance with an embodiment of the invention.

[00117] FIG. 5 is a flow chart illustrating a pilot fusion process in accordance with an embodiment of the invention.

[00118] FIG. 6 conceptually illustrates placement of LR image samples on a fusion grid in accordance with an embodiment of the invention.

[00119] FIG. 7 is a flow chart illustrating an SR process that obtains optimal SR image estimates from LR images in a manner that accounts for uncertainty in the upstream processing of the LR images in accordance with an embodiment of the invention.

[00120] FIG. 8 is a flow chart illustrating a process for obtaining an optimal SR image estimate using a gradient descent optimization technique in accordance with an embodiment of the invention.

[00121] FIG. 9 is a flow chart illustrating a process for refining an HR image estimate by computing a gradient including likelihood, intra-channel, and/or inter-channel gradient terms in accordance with an embodiment of the invention.

[00122] FIG. 10 illustrated a process for determining a likelihood gradient for an HR image estimate in accordance with an embodiment of the invention.

[00123] FIG. 11 conceptually illustrates the forward warping of pixel positions of an HR image estimate into the frame of coordinates of a target LR imager in accordance with an embodiment of the invention.

[00124] FIG. 12 conceptually illustrates the forward warping of pixel positions of an HR image estimate, where the pixels at the positions did not originate from a target LR imager, into the frame of coordinates of the target LR imager in accordance with an embodiment of the invention.

[00125] FIG. 13 is a flow chart showing a process for placing warped pixels on a decimation grid in accordance with an embodiment of the invention

[00126] FIG. 14A is a flow chart showing a process for formulating a likelihood gradient term in accordance with an alternative embodiment of the invention.

[00127] FIG. 14B illustrates the process of forming gradients between a current HR image estimate and input LR images in an SR process in accordance with embodiments of the invention.

[00128] FIGS. 15A and 15B illustrates imager arrays including different categories of luma imagers, where each category of luma imager shares a different field of view, in accordance with embodiments of the invention.

[00129] FIG. 16 is a flow chart showing a process for synthesizing a stereo pair of high resolution images using an array camera in accordance with an embodiment of the invention.

DETAILED DESCRIPTION

[00130] Turning now to the drawings, systems and methods in accordance with embodiments of the invention are disclosed that use super-resolution (SR) processes to use information from a plurality of low resolution (LR) images captured by an array camera to produce a synthesized higher resolution image. In a number of

embodiments, the SR processes include an initial pilot fusion that serves as an initial estimate for an optimization process that determines optimal SR image estimates. In several embodiments, the optimization process accommodates uncertainty in the upstream processing pipeline. In many embodiments, the optimization process involves using a forward imaging transformation to synthesize LR images in the frame of coordinates of each of the LR imagers in the array camera. In a number of embodiments, the simulated LR images are compared to the input LR images and the comparison used to improve a high resolution (HR) image estimate. In several embodiments, the array camera incorporates imagers having different fields of view. SR processing can be performed on the LR images captured by imagers having the same field of view. Array cameras that incorporate SR processes in accordance with embodiments of the invention are discussed further below.

ARRAY CAMERA ARCHITECTURE

[00131] An array camera architecture that can be used in a variety of array camera configurations in accordance with embodiments of the invention is illustrated in FIG. 1. The array camera 100 includes an imager array 110, which is connected to an image processing pipeline module 120 and to a controller 130.

[00132] The imager array includes two or more imagers, each of which receives light through a separate lens system. The imager array can also include other circuitry to control imaging parameters and sensors to sense physical parameters. The control circuitry can control imaging parameters such as exposure times, gain, and black level offset. In one embodiment, the circuitry for controlling imaging parameters may trigger each imager independently or in a synchronized manner. The sensor can include a variety of other sensors, including but not limited to, dark pixels to estimate dark current at the operating temperature. Imager arrays that can be utilized in array cameras in accordance with embodiments of the invention are disclosed in PCT

Publication WO 2009/151903 to Venkataraman et al., the disclosure of which is incorporated herein by reference in its entirety.

[00133] The image processing pipeline module 120 is hardware, firmware, software, or a combination for processing the images received from the imager array 110. The image processing pipeline module 120 processes the multiple LR images captured by the camera array and produces a synthesized higher resolution image in accordance with an embodiment of the invention. In a number of embodiments, the image processing pipeline module 120 provides the synthesized image data via an output 122.

[00134] The controller 130 is hardware, software, firmware, or a combination thereof for controlling various operation parameters of the imager array 110. The controller 130 receives inputs 132 from a user or other external components and sends operation signals to control the imager array 110. The controller 130 can also send information to the image processing pipeline module 120 to assist processing of the LR images captured by the imager array 110.

[00135] Although a specific array camera architecture is illustrated in FIG. 1, alternative architectures that enable the capturing of LR images and application of SR processes to produce a synthesized high resolution image can also be utilized in accordance with embodiments of the invention. Operation of array cameras, imager array configurations, and processing of LR images to produce a synthesized HR image in accordance with embodiments of the invention are discussed further below.

IMAGER ARRAY CONFIGURATIONS

[00136] As is disclosed in PCT Publication WO 2009/151903, an imager array can include any $N \times M$ array of imagers such as the imager array (200) illustrated in FIG. 2A. In several embodiments, the array is square. In other embodiments, the array is rectangular. In a number of embodiments, the array can also be irregular. Each of the imagers typically includes its own filter and/or optical elements and can image different wavelengths of light. In a number of embodiments, the imager array includes imagers

that sense red light (R), imagers that sense green light (G), and imagers that sense blue light (B). The human eye is more sensitive to green light than to red and blue light, therefore, an increase in the resolution of a synthesized image can be achieved using an array that includes more imagers that sense green light than imagers that sense red or blue light.

[00137] A 5 x 5 imager array (210) including 17 imagers that sense green light (G), four imagers that sense red light (R), and four imagers that sense blue light (B) is illustrated in FIG. 2B. A 5 x 5 imager array (220) including 13 imagers that sense green light (G), six imagers that sense red light (R), and six imagers that sense blue light (B) is illustrated in FIG. 2C. In several embodiments, the imager array also includes imagers that sense near-IR wavelengths, which can be used to improve the performance of the array camera in low light conditions. A 5 x 5 imager array including near-IR sensors is illustrated in FIG. 2D. The imager array 230 includes six imagers that sense red light (R), six imagers that sense blue light (B), nine imagers that sense green light and eight imagers that sense near-IR. The imagers indicated with an X can be either imagers that sense green light or imagers that sense near-IR wavelengths. In many embodiments, the distribution of imagers that sense green light and imagers that sense near-IR wavelengths is determined so that distribution is symmetrical. Although a specific configuration for imagers that sense blue light (B) and imagers that sense red light (R) is illustrated, these imagers can be switched with respect to each other. Additional imager array configurations are disclosed in U.S. Provisional Patent Application Serial No. 61/281,662 to Venkataraman et al., the disclosure of which is incorporated by reference herein in its entirety.

[00138] Each imager in the imager array captures a LR image. Image information captured by imagers, whose sensitivity is tuned by some means (e.g., color filters) to the same portion of the wavelength spectrum, can be considered as constituting a channel of imaging information. When different imagers are tuned to different portions of the

wavelength spectrum, then multiple channels of image information are produced that can be provided as inputs to a SR process in accordance with embodiments of the invention. In several embodiments, the imager array captures a single channel of information (B/W), or 3 channels of information (RGB). In many embodiments, the imager array captures the 3 RGB channels and a fourth near-IR channel that can also be used during SR processing to produce a synthesized higher resolution image.

[00139] Although specific imager array configurations are disclosed above, any of a variety of regular or irregular layouts of imagers including imagers that sense visible light, portions of the visible light spectrum, near-IR light, other portions of the spectrum and/or combinations of different portions of the spectrum can be utilized to capture LR images that provide one or more channels of information for use in SR processes in accordance with embodiments of the invention. The processing of captured LR images is discussed further below.

IMAGE PROCESSING PIPELINES

[00140] The processing of LR images to obtain an SR image in accordance with embodiments of the invention typically occurs in an array camera's image processing pipeline. In many embodiments, the image processing pipeline performs processes that register the LR images prior to performing SR processes on the LR images. In several embodiments, the image processing pipeline also performs processes that eliminate problem pixels and compensate for parallax.

[00141] An image processing pipeline incorporating a SR module for fusing information from LR images to obtain a synthesized HR image in accordance with an embodiment of the invention is illustrated in FIG. 3. In the illustrated image processing pipeline 120, pixel information is read out from imagers 110 and is provided to a photometric conversion module 304 for photometric normalization. The photometric conversion module can perform any of a variety of photometric image processing processes including but not limited to one or more of photometric normalization, Black

Level calculation and adjustments, vignetting correction, and lateral color correction. In several embodiments, the photometric conversion module also performs temperature normalization. In the illustrated embodiment, the inputs of the photometric normalization module are photometric calibration data and the captured LR images. The photometric calibration data is typically captured during an offline calibration process. The output of the photometric conversion module 304 is a set of photometrically normalized LR images. These photometrically normalized images are provided to a parallax detection module 308 and to a super-resolution module 314.

[00142] Prior to performing SR processing, the image processing pipeline detects parallax that becomes more apparent as objects in the scene captured by the imager array approach the imager array. In the illustrated embodiment, parallax (or disparity) detection is performed using the parallax detection module 308. In several embodiments, the parallax detection module 308 generates an occlusion map for the occlusion zones around foreground objects. In many embodiments, the occlusion maps are binary maps created for pairs of LR imagers. In many embodiments, occlusion maps are generated to illustrate whether a point in the scene is visible in the field of view of a reference LR imager and whether points in the scene visible within the field of view of the reference imager are visible in the field of view of other imagers. In order to determine parallax, the parallax detection module 308 performs scene independent geometric corrections to the photometrically normalized LR images using geometric calibration data 306 obtained via an address conversion module 302. The parallax detection module can then compare the geometrically and photometrically corrected LR images to detect the presence of scene dependent geometric displacements between LR images. Information concerning these scene dependent geometric displacements can be referred to as parallax information and can be provided to the super-resolution module 314 in the form of scene dependent parallax corrections and occlusion maps. Geometric calibration (or scene-independent geometric correction)

data 306 can be generated using an off line calibration process or a subsequent recalibration process. The scene-independent correction information, along with the scene-dependent geometric correction information (parallax) and occlusion maps, form the geometric correction information for the LR images.

[00143] Once the parallax information has been generated, the parallax information and the photometrically normalized LR images are provided to the super-resolution module 314 for use in the synthesis of one or more HR images 316. In many embodiments, the super-resolution module 314 performs scene independent and scene dependent geometric corrections (i.e. geometric corrections) using the parallax information and geometric calibration data 306 obtained via the address conversion module 302. The photometrically normalized and geometrically registered LR images are then utilized in the synthesis of an HR image. The synthesized HR image 316 may then be fed to a downstream color processing module 564, which can be implemented using any standard color processing module configured to perform color correction and/or chroma level adjustment. In several embodiments, the color processing module performs operations including but not limited to one or more of white balance, color correction, gamma correction, and RGB to YUV correction.

[00144] In a number of embodiments, image processing pipelines in accordance with embodiments of the invention include a dynamic refocus module. The dynamic refocus module enables the user to specify a focal plane within a scene for use when synthesizing an HR image. In several embodiments, the dynamic refocus module builds an estimated HR depth map for the scene. The dynamic refocus module can use the HR depth map to blur the synthesized image to make portions of the scene that do not lie on the focal plane to appear out of focus. In many embodiments, the SR processing is limited to pixels lying on the focal plane and within a specified Z-range around the focal plane.

[00145] In several embodiments, the synthesized high resolution image 316 is encoded using any of a variety of standards based or proprietary encoding processes including but not limited to encoding the image in accordance with the JPEG standard developed by the Joint Photographic Experts Group. The encoded image can then be stored in accordance with a file format appropriate to the encoding technique used including but not limited to the JPEG Interchange Format (JIF), the JPEG File Interchange Format (JFIF), or the Exchangeable image file format (Exif).

[00146] Processing pipelines similar to the processing pipeline illustrated in FIG. 3 that can also be utilized in an array camera in accordance with embodiments of the invention are described in PCT Publication WO 2009/151903, the disclosure of which is incorporated by reference above. Although a specific image processing pipeline is described above, super-resolution processes in accordance with embodiments of the invention can be used within any of a variety of image processing pipelines that register the LR images prior to super-resolution processing in accordance with embodiments of the invention. Super-resolution modules and super-resolution processes in accordance with embodiments of the invention are discussed further below.

SUPER-RESOLUTION

[00147] SR processes typically use LR input images forming one or more input spectral channels to generate a higher resolution image corresponding to some or all the input spectral channels. The LR images acquired by an imager array in accordance with an embodiment of the invention are typically afflicted by geometric distortions inherent to the optics and manufacturing processes used to fabricate the imager array, which are not scene dependent, and parallax distortions, which are scene dependent. Therefore, SR processes in accordance with embodiments of the invention can use as inputs information about initial scene-independent geometric corrections of the LR images, and about correcting the scene-dependent geometric distortions introduced by parallax in the LR images. The scene-independent geometric corrections and the

scene-dependent geometric corrections (i.e., parallax) are referred to as the geometric corrections for the LR images. Any of a variety of spatial registration approaches can be used to determine the information about the scene-independent geometric corrections for the LR imagers, and typically this portion of the spatial registration of the image content of the acquired LR images is generated using the output of well-known geometric calibrations techniques, which can be performed "off-line".

[00148] Many SR processes assume very accurate geometric correction (or registration) of LR images prior to SR processing. SR processes in accordance with embodiments of the invention assume uncertainty in the geometric correction of the LR images provided as inputs to the SR processes and estimate the most likely super-resolved image given the uncertainty introduced by the imager design and image processing operations, such as the geometric correction of the LR images, performed prior to the SR processing. In a number of embodiments, the most likely super-resolved image is determined using a Bayesian approach in which an objective function imposes constraints related to the desired characteristics of the synthesized image and fidelity of the synthesized image to the observed LR images. The SR processes can be applied on an entire HR image or on portions of the HR image (i.e. selective SR processing). In other embodiments, alternative optimizations can be performed in which information from the LR images is used to synthesize an HR image with desired characteristics. In several embodiments, the SR process is initialized by performing a pilot fusion process that joins (or places) and processes data from multiple input LR images onto a higher resolution output grid to produce a first fused image. The pilot fusion process can involve filling missing samples on the output grid and filtering potentially misplaced samples. The SR process then uses the pilot fused image along with the original LR images to perform an optimal image fusion and restoration according to a predetermined objective function. The objective function can be designed to enforce the fidelity of the synthesized HR image to the observed data. In several

embodiments, fidelity is enforced by passing the synthesized HR image through an image degradation model and can also involve imposing predetermined constraints on the images, including but not limited to intra-channel constraints and inter-channel consistency constraints. Although the SR processes are described above as being initialized by a pilot fusion process, other initial estimates of the HR image estimate appropriate to the requirements of specific applications can be utilized to initialize an SR process in accordance with embodiments of the invention. In addition, SR processing can be applied to one or more portions of the HR image and the portions combined.

[00149] In many embodiments, a forward imaging transformation (reflective of an image degradation model) is constructed incorporating an imaging prior that is provided as an input to the SR process. The imaging prior can include but is not limited to modeled or measured information about the optics of the cameras used to capture the LR images in any channel, and information concerning the LR imagers including but not limited to characterized information (e.g. data about noise characteristics), and dynamic sensor information or sensor-state information, such as the analog gain used to capture each of the LR images. A variety of SR processes and the usage of imaging priors in SR processing in accordance with embodiments of the invention are discussed in detail below.

SUPER-RESOLUTION PROCESSING

[00150] An SR module in accordance with an embodiment of the invention is conceptually illustrated in FIG. 4. The SR module 314 receives as inputs photometrically normalized image data, in possibly multiple spectral or color channels, and scene-dependent geometric correction information (parallax information) and scene-independent geometric correction information (geometric calibration data). Through a sequence of operations, the SR module generates a higher resolution image corresponding to some or all of the input spectral channels. In many embodiments, the

SR module is provided with photometrically normalized LR images and parallax information. In several embodiments, the SR module is provided photometrically normalized and geometrically corrected LR images, and parallax information. In instances in which either parallax or scene independent differences are small, the SR module need not utilize parallax information and/or scene independent geometric calibration data. In many instances, an SR process utilizes information from a channel that is not included in the synthesized HR image. For example, an input near-IR channel may be used by the SR module to generate higher resolution images, however, the SR module may not output a higher resolution near-IR channel. The SR module can also generate HR depth maps, which can be utilized for a variety of purposes, information to allow for dynamic refocusing of HR images, and information for generating LR or HR stereo pair images for use in 3D applications. Where multiple channels are provided to the SR module, each channel has a native resolution (i.e. the resolution in which the imaging information forming the channel was captured). The geometric correction information provided to the SR module is sufficient to register all, or subsets, of the captured LR images in a common system of coordinates and can use a reference system of coordinates, which can be that of an arbitrarily or purposefully-picked image from the LR image set of a channel. Alternatively, all LR input images can be spatially registered across the multiple input spectral channels. Any known spatial registration technique can be used for the purpose of providing registration information as input to an SR process in accordance with embodiments of the invention including ways differing from those outlined above with respect to the image processing pipeline illustrated in FIG. 3. In many embodiments, the SR module also receives as inputs sensor state information such as the analog gain of the imagers used to capture each of the LR images (not shown). In other embodiments, the SR module utilized in SR processing can utilize any of a variety of inputs associated with sensor state. An SR module in accordance with an embodiment of the invention is illustrated in FIG. 4. The

SR module 314 includes a Pilot Fusion module 402, a Dynamic Focus module 404, a Maximum A Posteriori Super Resolution (MAP SR) module 406, and a Cross Channel Processing module 408. Either photometrically normalized LR images and geometric calibration data, or photometrically normalized and geometrically registered LR images can be provided as inputs to the Pilot Fusion module 402 in conjunction with parallax information. The Pilot Fusion module 402 is hardware, software, firmware, or a combination thereof for performing a pilot fusion process. The term fusion can be used to refer to a process by which image data samples from multiple images having lower resolution are processed and placed appropriately onto a generally higher resolution grid. In the illustrated embodiment, the pilot fusion process joins (or places) and processes data from multiple input LR images of a channel onto a (typically) higher resolution output grid. This is done such that the LR image samples are placed appropriately onto the output grid (using the geometric corrections – i.e. scene independent geometric calibration data and scene dependent parallax information), and such that positions on the grid where there are missing samples get filled, and a filtering of potentially misplaced samples (outliers) occurs.

[00151] When initially placing the LR image samples onto a different (higher resolution) grid (the fusion grid), the LR samples may fall irregularly on the grid such that there may be empty grid positions, or grid positions having multiple samples. The fusion grid may have the same precision as the precision of the geometric corrections, or a higher or lower precision than the precision of the geometric corrections, in which case the transferred sampled onto the grid will fall arbitrarily onto the grid (i.e., including in-between grid nodes). The fusion described below may be preceded by a nearest neighbor procedure that would place the geometrically registered LR samples onto the fusion grid nodes. In any case, the fusion described in this disclosure can be used to interpolate the transferred LR samples onto the fusion grid. However, in many embodiments interpolation is only used to fill empty grid positions and not to interpolate

pixel values at locations occupied by pixels from the input images. Irrespective of the approach, it is the task of the pilot fusion process to generate a first, fused image, onto the fusion (output) grid. Where the SR processing is performed on portions of the image, then the first fused image or images obtained by the pilot fusion are obtained using pixels from one or more portions in corresponding locations in each of the LR images.

[00152] The Dynamic Focus module 404 may determine a high resolution depth map, or an equivalent representation of it. Using the LR occlusion maps provided at the input of the SR processing as part of the geometric correction information, an HR occlusion map may also be determined. Alternatively, the high resolution depth map, or a representation of it, and the HR occlusion maps, may be determined prior to the Dynamic Focus module, and this module can simply use the information to perform processes related to the dynamic focusing of the synthesized HR image. Similarly, an HR occlusion map can also be used in relation to subsequent SR processing including Pilot Fusion and the MAP module operation. The Dynamic Focus module uses the HR depth map to determine the set of pixels on which subsequent blocks in the SR processing pipeline operate (i.e., the pixels containing information concerning the "in-focus" areas of the HR image estimate). In many embodiments, the modules subsequent to the Dynamic Focus module operate only at the positions on the HR grid that are identified by the Dynamic Focus module based on the determined HR depth map, and either pre-determined, or interactive specifications of the desired pixels to be in focus in the HR image estimate. The rest of the pixels may be allowed to remain in the processing state attained prior to the Dynamic Focus module, or may be further blurred to form the "out-of-focus" areas in the generated HR image estimate. In several embodiments, the Dynamic Focus module is placed before the Pilot Fusion module. In that case, the out-of-focus area on the HR grid may be obtained directly by blurring LR image samples that correspond to those areas (i.e., a pilot fusion in those

areas may not occur). Some padding of the in-focus and out-of-focus areas may be used to ensure the proper processing (e.g., blurring) at the boundaries between in-focus and out-of-focus areas. The high-resolution depth map (or a representation of it) can be determined using the parallax information of the LR samples transferred onto a fusion grid, along with information about the geometry of the imager (e.g., inter-LR image baselines). In several embodiments, a high resolution depth map is generated by fusing onto the high-resolution fusion grid the parallax information corresponding to each LR sample available from all the LR imagers (regardless of channel), using the baseline information to determine the HR depth or an equivalent representation of it, and interpolating the depth at fusion grid positions where no LR samples are present. In a number of embodiments, a Dynamic Focus module is provided separately from the SR module. The SR module can generate a high resolution image and a high resolution depth map, and the Dynamic Focus module can be used to render pixels within a specified depth range from a defined focal plane at full resolution, and to render blurred pixel information for pixels outside of the specified depth range. In several embodiments, the blurred pixel information is obtained by blurring the high resolution image. In many embodiments, the blurred pixel information is obtained by blurring the input images.

[00153] The MAP SR module 406 is hardware, software, firmware, or a combination thereof that uses a maximum a posteriori (MAP) estimate to obtain good estimates for pixels of a synthesized HR image based upon at least the observed LR images and, in many instances, other observed information and/or prior constraint information. In a number of embodiments, the SR process estimates an HR image for a single channel such as the Green channel based upon at least the fused image output by the Pilot Fusion module 402, the LR image data (raw, photometrically normalized or photometrically normalized and geometrically registered LR images), geometric correction information including an occlusion map, and an imaging degradation model

[discussed below] that can include an imaging prior (where available). As is discussed further below, the HR image estimate for the first channel (e.g. the Green channel) can then be used in cross channel processes to estimate SR images for the other channels and to synthesize a full color SR image. In other embodiments, the SR process is a multi-channel statistical estimation process for determining the SR images for each channel. In a number of embodiments, the SR process initializes the SR image estimate(s) with the pilot fusion estimate, and then proceeds iteratively to refine the estimate(s) using the input information and an adaption to the evolution of the iteration process. In several embodiments, SR processes in accordance with embodiments of the invention attempt to synthesize the SR image(s) most likely to yield the captured LR images when a forward imaging transformation is applied to the synthesized HR image(s), where the forward imaging transformation is based upon what is known about the characteristics of the imagers in the array camera.

[00154] The MAP SR module 406 can utilize information, where available, concerning the characteristics of the imagers. In many embodiments, the MAP SR module has access to the forward imaging transformation specification for each imager and a noise model(s). In several embodiments, the MAP SR module accesses an imaging prior, whose elements are used to define the forward imaging transformation. The imaging prior can be defined as: modeled or measured information about the optics of the cameras used to capture the LR images in any channel, and information about the imagers. The information about the imagers may consist of both characterized information, such as data about noise characteristics, sensor resolution, sensor blur, and dynamic sensor operation information (or sensor-state information) such as the analog gain used to capture certain LR images. In a number of embodiments, the imaging prior includes a blur function, H_p^c , which can be obtained by modeling or measuring the lens characteristics for each LR camera. The blur function is discussed further below, where p indexes the LR camera, and c indexes the color channel for a

camera. The noise statistics can also be determined using any of a number of well-known models of imaging noise. Thus, the read noise can be modeled as an additive signal-independent noise, while the shot noise can be modeled as a signal dependent noise, whose variance is modeled (using Poisson distribution assumptions) or measured as a function of signal level. The total noise variance is taken as the sum of the individual noise variance components (where other components such as fixed-pattern noise are assumed to have been compensated for previously).

[00155] In the illustrated embodiment, the Cross Channel Processing module 408 estimates HR images for the Blue and Red color channels using the HR Green channel estimate generated by the MAP SR module 406. The output of the Cross Channel Fusion module is a synthesized HR image including color information from multiple channels. In other embodiments, the cross channel processing combines multiple-channel HR images output by the MAP SR module. In many embodiments, the cross channel processing utilizes additional channels of information, including but not limited to a near-IR channel, that are utilized to enhance the generation of SR images for one or more color channels (e.g. R, G, B).

[00156] In a number of embodiments, pilot fusion and SR processes can be applied to an entire HR image, selectively applied to portions of an HR image, and/or applied to blocks of an HR image that are combined. When the SR processes are applied to subsets of pixel positions in the HR image, the subsets of pixel positions can be determined based on the computation of per-pixel or per-region SNR values, occlusion maps, HR depth maps (or equivalent quantities such relative depth maps), and resulting gradient values from the output of an SR iteration. Pilot fusion, SR processes, and the application of SR processes to subsets of pixel positions in an HR image in accordance with embodiments of the invention are discussed further below.

GEOMETRIC CALIBRATION AND REGISTRATION

[00157] Types of scene independent geometric distortion degradations that can be accounted for in the calibration phase include but are not limited to, so called “barrel” distortions (for example the absolute geometric distortion of a camera chosen as the reference camera) and inter-camera geometric distortions relating the reference camera and each of the other cameras in the array. Effective calibration can correct captured LR images of an arbitrary scene using scene independent geometric correction information in such a way that the remaining geometric warping among the corrected LR images primarily consists of the parallax shift. The combination of the geometric calibration data for a pixel position (scene-independent correction information) and the parallax shift information (scene-dependent correction information) for a pixel position can also be referred to as the total shift experienced by the pixel at that pixel position. In reality, the corrected images will also include some positional noise associated with imprecision in the geometric correction information. The reference system of coordinates for geometric correction operations can be that of an arbitrarily or purposefully-picked image from the LR image set in a channel. Alternatively, all LR input images can be spatially registered using geometric corrections across the multiple input spectral channels. Yet another possibility is to geometrically correct the images spatially with respect to a camera position, which is not in the set of captured images (i.e. a “virtual camera” position). Indeed, any coordinate system appropriate to a specific application can be utilized in accordance with an embodiment of the invention.

[00158] Depending on the severity of the maximum parallax allowed for a specific imager before scene-dependent geometric correction (i.e., parallax correction) is required to remove parallax distortions, inter-camera parallax detection may or may not be necessary prior to the SR processing. Indeed, an SR process in accordance with embodiments of the invention can mitigate parallax to some degree without specifically

compensating for parallax using a form of uncertainty processing in accordance with embodiments of the invention. Where a parallax detection module is present, a parallax detection process can be performed to detect the remaining warping among the LR cameras in the array, after accounting for the scene-independent geometric correction (using calibration information) to facilitate the SR processing. In several embodiments, the parallax detection module utilizes at run time the information about the scene independent geometric corrections that are to be performed among the LR images to “dynamically” establish a spatial correspondence among the LR images for the same observed object in the scene being imaged for the purpose of parallax detection. The SR processing also utilizes the scene-independent and scene-dependent portions of the geometric correction information. Pilot fusion processes and SR processes in accordance with embodiments of the invention are discussed further below.

PILOT FUSION PROCESSES

[00159] A pilot fusion process in accordance with an embodiment of the invention is illustrated in FIG. 5. The process 500 includes placing (502) the LR image samples from each input LR image on a fusion grid in accordance with geometric correction information (i.e. scene independent geometric calibration data and scene dependent parallax information forming the total pixel shift) associated with the input LR images. Outlier image samples are then identified (504) and removed. Processes for identifying outliers in accordance with embodiments of the invention are discussed below and typically involve identifying pixels that possess very low levels of correspondence with neighboring pixels. Once outliers are removed from the fusion grid, the remaining image samples can be interpolated (506) to determine values at all fusion grid positions. In many embodiments, the pilot fusion generates a regularly sampled, initial estimate of the higher resolution image sought for by the SR processing, onto the fusion grid. This estimate represents the input to the next SR Processing module 406, which performs the super-resolution restoration processing. Specific pilot fusion processes

incorporating various techniques for placing LR image samples on a fusion grid, identifying and discarding outliers, and interpolating the remaining LR image samples are discussed below.

[00160] In many embodiments, the pilot fusion process is performed using input LR image data for a channel, the related geometric correction information, and imaging prior information (where available). The imaging prior information typically includes data concerning the estimated sensor noise statistics of the imagers that captured the LR images. However, any of a variety of imaging prior information can be utilized during the pilot fusion. A different resolution grid (typically higher resolution) from that of the input LR images is used in the pilot fusion process, which can be referred to as the fusion grid. As is discussed further below, the input LR images used in the pilot fusion can be image data originating from a single channel. In many embodiments, however, the pilot fusion for one channel can utilize image data from both the channel itself, and another channel that provides extra information that can be used in fusion. For example, the pilot fusion of image data from a specific visible channel (e.g. G) can include image data from both the specific visible channel and from a non-visible channel (e.g., infrared) or from a channel that combines the wavelengths of the specific visible channel with wavelengths from non-visible portions of the wavelength spectrum.

[00161] Placement of LR image samples on a fusion grid in accordance with an embodiment of the invention is illustrated in FIG. 6. In the illustrated embodiment, the pilot fusion fuses three LR images 604 intra-channel (i.e., using the image data only from the fused channel) and the first step in the pilot fusion process to place the LR image samples 606 onto the fusion grid 602 using the geometrical correction information provided at its input. The LR image samples 606 from each of the first, second, and third LR images are indicated on the fusion grid 602 as circles, triangles, and squares respectively. The first LR image (i.e. "Frame 1") is taken as the reference image (geometrically), and the other LR images are transferred onto the fusion grid

602. The result of the placement is typically an irregular placement of the image samples on the fusion grid 602. The placement outcome is typically a function of the difference between the precision of the geometric correction information and the resolution of the fusion grid. In many embodiments, the two resolutions are different, and samples may be placed arbitrarily on the fusion grid, including in between grid nodes. In several embodiments, a nearest neighbor process is used to place samples onto the fusion grid positions. The pilot fusion operates in an unified manner with respect to these cases, through the maintaining of lists of neighbors of a particular grid node position, which are used to facilitate the pilot fusion process. At each grid position on the fusion grid 602 there exist the extreme possibilities of having no samples there (i.e. an empty grid point), or of having multiple samples at a grid point. Although three LR images are shown in FIG. 6, pilot fusion processes in accordance with embodiments of the invention can be performed with respect to two or more LR images. Techniques for removing outliers and interpolating the image samples placed on a fusion grid are discussed further below.

[00162] When the LR image samples 606 are placed onto the fusion grid 602, the pilot fusion performs a filtering process to remove outlier samples and interpolate values at each grid position. In many embodiments, the filtering process is implemented using a locally adaptive interpolation process to: remove “outliers”, which can be defined as image samples that do not belong at their transferred positions for various reasons including those outlined above; fill-in missing samples; and interpolate a sample at a grid position using the neighboring available samples. The pilot fusion may interpolate all or a subset of grid positions. For example, the pilot fusion may interpolate only at missing fusion grid positions (i.e., where no samples exist).

[00163] Outlier removal is typically the first operation in each local neighborhood on the fusion grid 602, since it is undesirable to use outlier values in the processing of other samples in the local neighborhood. In a number of embodiments, outlier removal

is performed using a filtering operation that determines whether an LR image sample is an outlier using all neighboring LR image samples within a predetermined window, including stacks of samples, if contained in that window. In several embodiments, lists of neighboring samples are created within a certain Euclidean radius distance from the LR sample being tested for outlier status to accommodate the fact that the filtering operation takes place on the irregularly transferred LR samples.

[00164] In a number of embodiments, outliers are identified by constructing a measure based upon bilateral filter weights calculated using the tested LR image sample and its list of neighbors on the HR grid. Bilateral filters are typically used to perform edge-preserving smoothing by filtering images based upon both range and domain (i.e., spatial distance and photometric difference). A typical bilateral filter extends the concept of Gaussian smoothing by also weighting the filter coefficients based upon the corresponding relative pixel intensities. Pixels that are very different photometrically from the central pixel are weighted less even though they may be in close proximity to the central pixel. In the illustrated embodiment, a measure of the consistency of an LR image sample with its neighbors is obtained utilizing the bilateral filter weights that are computed based upon the LR image sample and its neighboring samples on the HR grid. The weights are summed and the sum normalized by the number of neighboring LR image samples to produce a value indicative of the similarity between the LR image sample point and its neighbors. A determination is then made that the LR image sample is an outlier (i.e. lacking consistency with the local LR transferred samples on the fusion grid) if the normalized sum of the weights is below a threshold. Where the imaging prior incorporates a pre-determined noise model, the bilateral filter weights described above utilize a range kernel that is a function (e.g., a scaled value) of the standard deviation of the noise estimate at the position and the spatial kernel is typically a predetermined Gaussian that decays with distance. In the absence of information from the imaging prior, predetermined values can be utilized to

determine the bilateral filter weights. Although specific approaches for identifying outliers based upon calculated bilateral filter weights are described above, outliers can be identified using a variety of techniques that identify LR image samples that have a low degree of correspondence with neighboring LR image samples in accordance with embodiments of the invention.

[00165] Once the outliers are identified, the next step in the pilot fusion process involves locally, and adaptively filtering the available samples in the neighborhood of each grid position to determine the value of an image sample at that position. For this process, a structured kernel local filtering can be used (see for example the process described in Trevor J. Hastie, et al. *The elements of statistical learning, data mining, inference, and prediction*, [2nd ed. 2003]), which obtains a value at each position on the fusion grid 602 by fitting a surface to the available image samples in the neighborhood of the grid position, and then re-sampling the fitted surface at the desired grid position. In a number of embodiments the fitted surface is obtained by using a weighted least squares technique to fit a given-order surface parameterization to the available image samples in the local window (i.e. list of neighborhood image samples) of the grid position. During this process, the degree of interpolation (or smoothing) can be controlled by the estimated signal to noise value estimated at each position on the fusion grid. Specifically, for low signal to noise at a given position, a higher degree of smoothing is performed. The above process produces an interpolated image sample at each grid position, thereby transforming the initially irregularly-placed image samples into interpolated, regular samples at all fusion grid positions. Although specific interpolation processes are described above, alternative interpolation processes can be utilized in accordance with embodiments of the invention. Alternatively, the interpolation process is only utilized to determine pixel values in empty grid positions and is not used to interpolate pixel values at grid positions in which pixels have already been placed.

[00166] In the process of filtering local LR image samples for interpolating a value at each fusion grid position, LR image samples are placed on the grid based upon a determination that some information carried by these LR image samples is useful. However, many of the LR image samples may be very close to the “outlier” threshold. Accordingly, many SR processes in accordance with embodiments of the invention weigh the contribution of LR image samples placed on the grid in any local interpolation that uses them. The weighting function, which assigns a scalar weight value to each LR sample transferred onto the fusion grid, can be provided to the SR processing by previous modules that influence the quality of the geometric correction information including but not limited to the parallax detection module. In a number of embodiments, the size of the local window used for the pilot fusion, centered at the current grid position being processed, can be made variable, based on a measure of local density of available samples. For example, for a grid position whose neighborhood contains a higher density of available image samples, the local window can be made smaller to allow for better detail preservation. However, for pilot fusion grid positions, which have a sparse-sample neighborhood, the local window size can be extended to process sufficient samples for a reliable interpolation. The order of the model used in the structured kernel filtering can vary depending on the size of the local window and the number of samples in the window.

[00167] In a preferred embodiment, the pilot fusion generates a regularly sampled, initial estimate of the higher resolution image sought for by the SR processing, onto the fusion grid 602. This estimate represents an input to the next SR processing stage, which performs the super-resolution restoration processing.

DYNAMIC FOCUS PROCESSING

[00168] As illustrated in FIG. 4, a dynamic focus module can be used to generate high resolution estimates of the depth of points in the imaged scene from the point of view of a reference imager, other imagers in the array, or virtual points of view not coincident

with any of the imagers in an imager array. Similarly, by using low-resolution pair-wise camera occlusion maps (e.g., generated by a parallax detection module), a high resolution occlusion map can be generated from the point of view of the camera chosen as reference or any other camera in the array. Uses of the obtained HR depth map can include but are not limited to depth metering in the scene, creating an SR image with particular characteristics in terms of what image content is in-focus and which is out-of-focus, and/or enabling a dynamic re-focusing of the rendered SR image, even after the image data has been captured by the imager. The details of these processing approaches are given below.

[00169] An HR depth map can be determined using the point of view (POV) of a reference imager or from other points of view. In order to determine a HR depth map of the scene observed from the POV of the reference camera, a cross-channel fusion is performed, but instead of fusing actual pixel values, the parallax information corresponding to all the available pixels from all LR imagers regardless of color channel, are fused onto the common HR depth fusion grid. Also available at each position on the HR depth fusion grid where a pixel would land, are the inter-imager baselines relating the source LR imager (i.e. the LR imager that captured the pixel located at the pixel position), and the reference LR imager. With this information, one can use the relationship between depth and parallax to determine an HR depth map of the scene visible from the POV of the reference camera. This can be done in absolute terms by using the formula:

$$depth[m] = \frac{baseline[m] * focalLength[m]}{disparity[m]}$$

[00170] The disparity in the above equation is alternatively referred to as parallax, and the focalLength is the focal length for the lens of a the originating LR imager (i.e. the LR

imager from which the pixel originated), and the baseline is the inter-imager baseline. Further, since the parallax is referred to in this disclosures in units of HR pixels, the further definition of the disparity (or parallax) in pixel units is given by:

$$disparity[pixels] = \frac{disparity[m]}{pixelSize[m]}$$

[00171] As discussed above, there are positions on the HR depth map grid that will remain empty following the fusion (fewer in this case since, as compared to the pilot fusion process in one channel described above, we are fusing together onto the HR depth fusion grid the parallax information of all pixels from all LR cameras regardless of channel). The depth values at all or a subset of grid positions (e.g., missing depth positions) may be interpolated using the fused HR depth map. A number of interpolation techniques can be used, as known in the art. Additionally, when computing the HR depth map, simple constraints can be imposed such that the HR depth map is validated against the geometrical characteristics of the imager (e.g., a filtered depth such as for a missing pixel, should not result in a conflict with other pixels when geometric transformations or corrections are applied to the pixel).

[00172] While the processing above can be used to generate a HR depth map where the obtained values are actual depth values, a more economical but similar process can be used to only determine a relative HR depth map in the scene, which can be used for dynamic focus purposes. For this purpose, only the parallax and the baselines information is needed, such that their ratios (reflective of relative depth) can be used to determine the relative HR depth map of points in the scene viewed by the reference camera.

[00173] The absolute or relative HR depth map determined as described above, can be used for dynamic focus purposes at some point in the rest of the SR processing. For example, the user of the system in which the current processing is embedded, may

want to only render in focus only a particular depth plane (plus or minus a depth range). In other words, a shallower depth of field is required. By the processing described above, this can be easily done by either automatic selection, or by acting to SR-process and render only the depth range selected. In many embodiments, the user is presented with an image of the scene (at possibly lower resolution) where that selection can be made. In this way, an interactive selection of a plane of focus and a range of focus when generating an SR image can be performed. In many embodiments, since an HR depth map is estimated and can be saved to storage by the processing in conjunction with the HR image estimation, the desired (and varying) focus selection can be done at capture time, or just as well, after the images have been captured (i.e., post-capture focusing).

[00174] The HR depth map and the HR parallax information (which is used in the SR processing described below) can be equivalent quantities as shown by the above equations, and thus, any operations described in one domain can equivalently be performed in the equivalent domain. Thus, the constrained HR depth map described above can also be obtained as a byproduct of the SR processing described below (where forward warping operations are described).

[00175] Similarly, the use of the absolute or relative HR depth map can be extended to benefit the efficient computation of the HR image. Thus, if only a limited depth of focus is required, the SR processing can concentrate only on the estimation of the pixels that are representative of points in the scene that fall within the desired depth of focus, since they are the ones that will have to be rendered in focus and with maximal resolution. The rest of the pixels in the HR image may be left in the state attained at the point before the Dynamic Focus module, or can even be blurred to further effect their "out-of-focus" appearance.

[00176] The HR depth information can also be used for an enhanced method of pilot fusion. Thus, instead of performing the full pilot fusion as described above as a first step, the HR depth information is first estimated by the procedure described above.

Then, this HR depth information is input to the Pilot Fusion module. The HR depth information is used such that, when operating the interpolating kernel (e.g., the structured kernel filter) in a local neighborhood, samples now also carry with them estimated HR depth information. In one embodiment this information is used for generating the interpolated pixel value at the current position of the HR fusion grid, in addition to using the actual values of the pixels in the local window centered at the pixel being interpolated, their confidence values, and their distance from the current position being interpolated. For example, in occlusion zones estimated for the point of view of the reference imager (where the current HR estimate is sought), samples are much sparser and typically include pixels belonging to the foreground object (which caused the occlusion) and the background object(s). In this case, the corresponding local HR depth map can be used to eliminate pixels belonging to the foreground object from the interpolation, when interpolating the value of a pixel in the occlusion zone (which is known to belong to the background object given the mechanics of the occlusion). Thus, a higher-quality interpolation can be performed. The same procedure can be extended to other scenarios for the HR fusion process.

SR PROCESSING

[00177] The MAP SR module 406 uses a statistical estimation approach to determine the super-resolved images for each output spectral channel. An SR process in accordance with an embodiment of the invention is illustrated in FIG. 7. The process 700 obtains (702) as its inputs the pilot fused image, the photometrically normalized LR image data for each input spectral channel, geometric correction information, and an imaging prior if available. The pilot fused image is selected (704) as an initial estimate in an optimization process that seeks to obtain (706) optimal HR image estimates given the input LR images captured by the imager array and a forward imaging transformation (built based on an image degradation model). In a number of embodiments, the SR process utilizes a proprietary form of a maximum a posteriori

(MAP) objective criterion formulation, which is then optimized with respect to the SR image to be estimated for at least one output spectral channel. In many embodiments, the objective criterion formulation utilizes a specific forward imaging transformation (reflective of an image degradation model) that describes the image degradation that occurs within each imager during the capturing of an LR image. In several embodiments, the criterion used in the optimization includes two main terms: a data fidelity measure of the estimate; and an intra-channel prior term, which imposes intra-channel constraints on the estimated SR image solution. In many embodiments, the estimation process is a multi-channel estimation process and a third inter-channel prior term is used in the optimization criterion, which places constraints on the relationship among the SR image solutions for the multiple output channels. The third term is obviously not required when dealing with a single channel and in many embodiments is also not required where dealing with multiple channels as other processes can be used pre- or post-SR processing to constrain and/or register the SR image solutions with respect to each other. In several embodiments, terms associated with any of a number of criteria appropriate to the application can be used in formulation of an appropriate objective function for use in the estimation of optimal SR images.

DEGRADATION MODELS

[00178] In a number of embodiments, one of the captured LR images in the color channel that is the subject of the SR processing is taken as the reference image, and all geometrical transformations of the other LR images are taken with respect to this reference image (i.e., any geometric correction information provided at the input of the SR processing, which may include scene-independent geometric correction and parallax information). In which case, an estimated high resolution image x^c generated by an SR process in accordance with an embodiment of the invention represents the

high resolution image that the reference imager would see, were it a high resolution camera [at the same resolution as the desired SR output resolution that is being sought in a particular channel c]. In other embodiments, the estimate x^c is from the perspective of a virtual camera. As noted above, the MAP process used to estimate x^c involves applying forward imaging transformations to the estimated HR image to obtain synthesized LR images that can be compared to the observed LR images. The forward imaging transformations (possibly specific to each LR camera) applied to the estimate x^c can include but are not limited to scene dependent geometric displacements (i.e. parallax), scene independent geometric distortions, lens and sensor blur, sensor degradation, and decimation. Application of forward imaging transformations to a high resolution image estimate x^c are discussed below.

[00179] An LR image captured by imager p in channel c can be denoted as y_p^c . In many embodiments, the forward imaging transformation involves applying a first matrix operand $W_{p,parallax}^c$ to x^c . $W_{p,parallax}^c$ represents a geometric transformation describing the parallax view of the p^{th} imager before its lens acts upon it, compared to the estimated high resolution image x^c of the scene as viewed by the reference imager. In several embodiments, a second matrix operand $W_{p,dist}^c$ is also applied, which describes the scene independent geometric warp for each LR imager (if the LR imagers have not been corrected for scene independent differences, see discussion above). Furthermore, a lens blur operator $H_{p,lens}^c$ can also be applied, which produces the optical image for LR imager p . The image degradation can be further modeled using the operator $H_{p,sensor}^c$ that describes the extent to which the sensor blurs the image formed on the sensor by the lens, and a decimation operator D . Finally, imaging noise n_p^c can be added to obtain the p^{th} simulated LR image y_p^c .

[00180] Without loss of generality, and for simplicity of description, many embodiments of the invention incorporate a forward imaging transformation in which

the decimation operator D is assumed to be the same for all imagers in a channel c , and the blur operators $H_{p,lens}^c$ and $H_{p,sensor}^c$ are assumed to be invariant across the images of a given channel. A number of embodiments, however, utilize a spatially-variant lens blur $H_{p,lens}^c$ to reflect the manner in which the imager optics behave in reality.

[00181] The joint spatially invariant lens-sensor blur (also referred to as total blur) operator (H^c) can be denoted as follows:

$$H^c = H_{sensor}^c H_{lens}^c.$$

[00182] The observed LR image y_p^c is thus the result of degrading (by the above operators) the original high-resolution image x^c , with the addition of noise n_p^c . On the latter variable, it can be assumed that the noise statistics are drawn from the same distribution for all images in all channels. Also, the noise statistics (i.e., variance) can be modeled or measured as part of the formation of the imaging prior..

[00183] Although a specific forward imaging transformation is discussed above, alternative forward imaging transformations can be utilized in accordance with embodiments of the invention and the formulation of the SR objective criterion modified accordingly.

ESTIMATING AN SR IMAGE

[00184] A MAP process in accordance with many embodiments of the invention can involve iteratively refining an estimate of an HR image using an objective function. In several embodiments, the iterations utilize a gradient descent technique as follows, where the current HR image estimate is expressed as a vector $x^c(n)$:

$$x^c(n+1) = x^c(n) - \beta^c \nabla_{x^c(n)} J(x^c(n)) \quad [1]$$

where

n - indexes the iteration number

β^n - is the step-size for the iteration

$\nabla_x(\cdot)$ - is the gradient operator

$J(x^c(n))$ - is the MAP objective function parametrized by the HR estimate

[00185] For simplicity, the vector $x^c(n)$ can be denoted as x^c during a specific iteration. The gradient descent described above is a process that can be performed by the SR processing module of an array camera in accordance with embodiments of the invention to obtain an optimal SR image estimate(s). In various embodiments, other optimization techniques can also be used to obtain SR image estimate(s). A process for using gradient descent to obtain optimal SR image estimates in accordance with an embodiment of the invention is illustrated in FIG. 8. The process 800 includes initializing (802) the first estimate, x^0 , with the output of the pilot fusion process described previously. The gradient is determined (804) and a new estimate obtained (806). The process can either iterate until a predetermined optimality criterion is satisfied (808) or for a finite number of iterations.

[00186] When the process illustrated in FIG. 8 is applied with respect to the objective function referred to in Eq. (1), the gradient $\nabla_{x^c} J(x^c)$ has the form:

$$\nabla_{x^c} J(x^c) = VT(x^c) + \psi_{\Lambda_1}^{\text{intra}}(x^c) + \Phi_{\Lambda_2}^{\text{inter}}(x^c, \hat{x}^{c'}), \quad [2]$$

where

"c'" - signifies a channel different than c,

" Λ " - signifies that some of the SR channels will have constraints imposed using already finalized SR estimates of other channels, as mentioned above.

[00187] The gradient vector in Eq. (2) (having a value corresponding to each pixel in the current HR estimate) may arise from the use of different norm types (e.g., l1-, l2-,

Lorentz], in the objective function. The choice of norms does not limit the generality of the approach.

[00188] The matrices/vectors in Eq. (1) have appropriate sizes once the size of the vector x^c is given. Also,

c, c' - index the channel; the current channel being processed, and another channel, respectively.

x^c - the SR image for the current channel in the frame of coordinates of the reference camera.

$\Gamma(.)$ - is a gradient vector of a norm of the difference between the current HR estimate and all, or subsets of, the input LR images, that was used in the formulation of the objective function (i.e., the "likelihood-term" gradient).

$\hat{x}^{c'}$ - the super-resolved image for another channel c' , in the geometric reference frame of the reference camera (its use will be detailed further below)

$\psi_{\Lambda_1}^{\text{intra}}(.)$ - is a gradient vector of an intra-prior term formulation in the objective function using a particular norm, [locally-]weighted by diagonal matrix Λ_1 , which can be the identity matrix. This referred to as the intra-prior term gradient.

$\Phi_{\Lambda_2}^{\text{inter}}(.)$ - gradient vector of an inter-prior term formulation in the objective function using a particular norm, [locally-]weighted by diagonal matrix Λ_2 , which can be the identity matrix. This is referred to as the inter-prior term gradient.

V - is a diagonal matrix weighing the first gradient term in Eq. (2). It could be the identity matrix--if not, the entries on the diagonal of the weight matrix V can be determined as follows:

$$\text{diag}[V](k,l) = f(R(x^c), \text{SNR}(k,l)), \quad [3]$$

$f(.)$ is a function of a per-pixel local certainty function $R(x^c)$, and may also be dependent on the estimated SNR, at each $[k,l]$ position.

[00189] Function $f(\cdot)$ is such that for lower local SNRs estimated at position (k,l) , or for lower local confidence values (received as input by the SR processing) the weight ν at that position (k,l) is lower such that the first term in Eq. (2) is de-emphasized with respect to the second term (intra-prior gradient) in the overall gradient computation. This means that in low SNR areas more smoothing is encouraged rather than the enforcement of observed-data fidelity.

[00190] The last two terms in Eq. (2) measure a behavior of the estimated SR image with respect to a prior model of a natural image. The third term in Eq. (2), which enforces cross-channel constraints among the estimated SR channel images, imposes an edge location and orientation consistency among the output-channel SR images that are estimated. This term can use, but is not limited to using, a cross-channel Gibbs prior distribution formulation, or a co-directionality constraint. In many embodiments, the cross-channel prior is active only when estimating some channels c' based on a HR estimate already obtained for a channel c (e.g., the Red and Blue channels are estimated by transferring and filtering edge location and orientation as determined by the super-resolution estimate of the Green channel). Obviously, there are no cross-channel constraints where the camera array captures a monochrome image and the third term in Eq. (2) can be ignored. In many embodiments that involve multiple color channels, the third term in Eq. (2) is also ignored in favor of an alternate process for enforcing cross-channel constraints between the channels. When the third term in Eq. (2) is ignored, the iterative estimation proceeds using the first two terms, i.e., the likelihood gradient term, and the intra-prior gradient term. In several embodiments, different weighting functions are applied to each of the terms to obtain any of a variety of desired SR image characteristics.

ITERATIVE SR RESORATION PROCESSING

[00191] A process for iteratively refining an estimate of an SR image using a gradient descent approach is illustrated in FIG. 9. The process 900 is initialized (902) using an

HR estimate generated using a pilot fusion process. The process can also utilize (904) input information such as a degradation model and/or an imaging prior. The process then commences iterating (906). In each iteration, the likelihood gradient term is computed (908), the intra channel prior gradient term is computed (910), and the inter-channel prior gradient term is computed (912), and then a new estimate for the HR image is generated (914) using the gradient terms. The estimate of the HR image can be updated by adding the weighted gradient terms to the previous HR estimate. The process continues to iterate (916) until a predetermined number of iterations is reached (918). At which point, the SR image estimate is returned. Although the process illustrated in FIG. 9 terminates after a predetermined number of iterations, alternative termination conditions can be utilized in accordance with embodiments of the invention. The computation of the gradient terms in accordance with embodiments of the invention is discussed further below.

COMPUTING THE LIKELIHOOD GRADIENT TERM

[00192] In order to calculate the likelihood gradient term in accordance with Eq. (2), each pixel at position (k,l) in a current SR image estimate is tested for fidelity when forward-simulated through a forward imaging transformation, to pixel positions (i,j) in multiple LR input images, using a norm. The pixel positions (k,l) and (i,j) are related by the geometric correction information, including the scene-independent correction information (also referred to as geometric calibration), and the scene-dependent correction information (also referred to as parallax). It is possible that not every pixel position on the HR grid has scene-independent (calibration) data since there may not have been any pixel from an LR imager placed there in the fusion. The norm may be for example the l_2 -, l_1 -, or Lorentz norm. Forward-simulation means here that a current HR estimate is passed through a forward imaging transformation, which typically involves degradation by the operators corresponding to the imaging degradation model described above (i.e., including geometric warps, blurs, and decimation).

[00193] A process for generating a likelihood gradient term for a given color channel using the forward imaging transformation in accordance with an embodiment of the invention is illustrated in FIG. 10. The process 1000 utilizes the current HR estimate for the channel (1002) and the input information utilized in the generation of the HR estimate (1004) and then iterates over each (or a subset) of the input LR images (1006). For each input LR image considered, the HR estimate is geometrically forward warped (1008) in high resolution by the total warp (comprising the geometric correction information) relating the target LR imager and the reference LR imager, and the warped HR estimate is blurred (1010) by the total blur. The gradients between the warped and blurred HR image estimate at the decimation positions and the corresponding input LR images are then computed and accumulated (1012). Weight-related information for each of the gradients is computed (1014), decimation is performed (1016) at the computed gradient positions and upsampling of the resulting gradient data (1018) is performed. The upsampled gradient data is then blurred (1020) by the transpose of the total blur, and the blurred gradient data is backward-warped (1022) using the total warp information into the reference frame of the HR estimate. The weighted gradients are then accumulated (1024) in the HR reference frame of coordinates of the reference imager and the next iteration (1026) performed. The process iterates across all (or a subset) of the input LR images (1028) and outputs (1030) a final gradient vector.

[00194] The input information utilized in the calculation of a likelihood gradient term in accordance with an embodiment of the invention typically includes information such as geometric correction information (including parallax information and scene independent geometric correction information), occlusion information, and imaging prior information (such as per-pixel SNR estimates in the HR frame). Not all of the input information is necessary for calculating the likelihood gradient term. If the parallax information is not available and/or the parallax is small, uncertainty

processing can compensate for a lack of parallax information. This is also true with respect to small scene-independent geometric corrections, which may not need to be performed prior to the SR processing. Incorporation of uncertainty processing into SR processes in accordance with embodiments of the invention is discussed further below. The following discussion of calculating the likelihood gradient term assumes both portions of the geometric correction information, i.e., the scene dependent parallax information and the scene independent geometric calibration data are available during the SR processing.

WARPING AND BLURRING

[00195] The geometric warping and blurring of an HR image estimate in accordance with embodiments of the invention can be performed in a number of different ways, including warping and blurring the entire HR image estimate into the frame of coordinates of all, or a subset, of LR imagers, followed by per-pixel gradient computations with respect to each input LR image being considered. The warping and blurring of the HR image estimate with respect to the frame of coordinates of an LR imager involves forward-warping the HR image estimate using the available geometric correction information (also referred to as warping information) for the LR imager, which may include scene dependent parallax information and scene independent geometric calibration information, into the geometrical frame of coordinates of the LR imager. In many embodiments, the warping process simply involves a spatial shift (or translation) of a pixel. Once the HR image estimate has been warped onto the decimation grid of a target LR imager, the warped HR image estimate can be blurred by a total blur including both lens blur and sensor blur. The decimation grid of a target LR imager is a high resolution grid, and decimation can be performed on the warped and blurred HR image to simulate an LR image in the frame of coordinates of the target LR imager.

[00196] In other embodiments, the gradients are formed sequentially as the processing progresses through the pixels of the current HR image estimate. When gradients are formed sequentially, sufficiently large blocks of data (patches) in the images involved can be used so that the warping and blurring operations can function properly, and the current pixel being processed in the HR image has a fully-populated, large-enough, spatial neighborhood.

[00197] Warping of the entire HR image estimate with respect to the frame of coordinates of an LR imager is described in detail below.

WARPING AN ENTIRE HR IMAGE ESTIMATE

[00198] When a pixel in an HR image estimate originated from a specific pixel in an LR image, then “origination” information can be associated with the pixel position on the HR grid. The value of the pixel at that grid position may change through processing, but the origination information at that position remains and is available for use in SR processing. The origination information can identify the input LR image that the pixel is from (i.e. the source LR image) and the coordinates (i,j) of the pixel within the source LR image. When pixels are forward warped onto the decimation grid of a specific LR image, then all of the pixel positions with origination information from that LR image fall at regularly spaced intervals (i.e. the super-resolution factor/decimation interval) on the high resolution decimation grid starting at the (1,1) position (i.e. the position at which the grid starts). Typically, only a subset of pixel positions within an HR image estimate have origination information. Therefore, not all of the pixels in the HR image estimate can be directly warped back into an LR imager’s frame of coordinates. Accordingly, an adaptive warping process can be used that performs warping based upon conditions including whether a specific pixel has origination information.

[00199] In embodiments where the array camera captures multiple color channels, the origination information available for a pixel may relate the pixel to a source LR imager from a different color channel to the color channel of the HR image estimate.

Many adaptive warping processes in accordance with embodiments of the invention treat pixels that originated from a different color channel as a separate class of pixel positions and utilize only the spatial registration (geometric correction) information of the pixels originating from LR imagers from all color channels during the warping of the current HR image estimate. Although, in many embodiments only the geometric correction information of the pixels in the LR images in the same color channel as the reference LR camera are used for the purpose of warping and forming gradients at each point in the current HR estimate, the warping process need not be limited to utilizing just the spatial registration information for the current color channel LR images. Adaptive warping processes and the treatment of different classes of pixel positions are described below.

ADAPTIVE WARPING PROCESSES

[00200] As discussed above, adaptive warping processes apply different warping processes to various classes of pixel positions when warping an HR image estimate from the frame of coordinates of the reference imager into the frame of coordinates of a specific LR imager, which can be referred to as the “target” LR imager. The term pixel position refers to a coordinate on the HR grid to which one or more pixels from LR images map. The pixels can be from the same color channel as the HR image estimate or from a different color channel. In instances where the only pixels that map to a pixel position are from a different color channel, then the geometric correction information can be utilized in estimation of the pixel value at the pixel position. Stated another way, the actual pixel value from the different color channel is not used, however, geometric correction information can be used to adaptively warp such positions for the purpose of determining the likelihood gradient term (see description below). In a number of embodiments, the adaptive warping process applies different warping processes to the following classes of pixel positions:

- **Type 1** pixel positions are positions on the HR grid that originated from the target LR imager (irrespective of the color channel of the target LR imager).
- **Type 2** pixel positions are positions on the HR grid that originated from an LR imager other than the target LR imager (irrespective of the color channel). An LR imager other than the target LR imager, from which a **Type 2** pixel position originated, can be referred to as a “source” LR imager. In several embodiments, **Type 2** pixel positions can originate from a source LR imager within any color channel (i.e. **Type 2** pixel positions are not necessarily limited to the color channel which is same as that of the HR image estimate).
- **Type 3** pixel positions are synthetic pixel positions (i.e. positions within the HR grid to which none of the pixels in the LR images can map) that have neighboring **Type 1** and/or **Type 2** pixels. Such pixels are typically the result of a fusion process used to generate the original HR image estimate.
- **Type 4** pixel positions are positions that do not fall into any of the above categories. Such pixel positions are synthetic pixel positions that do not have any neighboring **Type 1** or **Type 2** pixel positions. **Type 4** pixel positions are not warped with respect to the target LR imager. It is likely that at least some of them will be able to be warped into some other target LR imager (than the current one) in turn. Any empty positions created when the HR image estimate is warped into the frame of coordinates of the target LR imager due to **Type 4** pixel positions can be locally interpolated. In other embodiments, the neighborhood for **Type 4** pixel positions is increased such that there are no **Type 4** pixel positions (i.e., they become Type 3 pixel positions).

[00201] **Type 1** and **Type 2** pixel positions are pixels in an HR image estimate that have geometric correction (or total shift, or warp) information, because they originate from an LR imager. **Type 3** and **Type 4** pixel positions do not have associated geometric correction (or total shift, or warp) information, because they are synthesized pixel

positions. The warping of **Types 1, 2, and 3** pixel positions in accordance with embodiments of the invention is discussed below.

WARPING TYPE 1 PIXEL POSITIONS

[00202] The warping of **Type 1** pixel positions involves directly warping the pixels from the HR image estimate to the frame of coordinates of the LR imager using the total geometric transformation information from the HR image estimate to the frame of coordinates of the target LR imager. The warping process is illustrated in FIG. 11. As noted above, a separate warping is performed with respect to the HR image estimate and each input LR image. When the HR image estimate is being warped with respect to the frame of coordinates of a specific target LR image, the HR image estimate grid 1100 includes a number of pixel positions 1102 with origination information that indicates the pixels were originally from the target LR image (i.e. **Type 1** pixels). For these pixel positions the forward warping is direct, in that the total warp information for the target LR imager can be used to take the pixel values in the **Type 1** pixel positions from the current HR image estimate and place those pixel values in the proper positions 1104 on a decimation grid 1106 for the target LR imager.

WARPING TYPE 2 PIXEL POSITIONS

[00203] When an HR image estimate is being forward warped with respect to a specific target LR imager, many of the pixel positions of the HR image estimate do not originate from the target LR imager. These **Type 2** pixel positions can be warped to the frame of coordinates of the target LR imager by mapping the geometric transformation information (total shift) for the pixel positions through the reference LR imager. The total shift information is available with respect to the LR imager from which the **Type 2** pixel position originated, which as noted above can be referred to as the “source” LR imager, and the reference LR imager. The total shift information is also available with respect to the target LR imager and the reference LR imager. The combination of these

two sets of warping information can be used to generate total warping information to warp the **Type 2** pixel positions between the various LR imagers from which they originated (i.e., source LR imagers), and the frame of coordinates of the target LR imager.

[00204] Various approaches can be taken to determine the warping to apply to a **Type 2** pixel position. The parallax information can be mapped through the reference imager to determine the parallax shift to apply to the **Type 2** pixel position. The total shift to apply to the **Type 2** pixel position to shift it into the frame of coordinates of the target LR imager can then be obtained by adding the parallax shift to the geometric calibration data (scene-independent geometric transformation information) between the reference LR imager and the target LR imager. Processes for warping a **Type 2** pixel position into the frame of coordinates of the target LR imager by determining parallax shift information in accordance with embodiments of the invention are discussed further below.

[00205] In order to determine the appropriate parallax shift for **Type 2** pixel positions, the different parallax experienced by each LR imager can be considered. A **Type 2** pixel position and the warping of the **Type 2** pixel position into the frame of coordinates of a target LR imager in accordance with an embodiment of the invention are shown in FIG. 12. The HR image estimate grid 1200 includes a "square" **Type 2** pixel position 1202 that has parallax information associated with it. The parallax information, however, relates this pixel position 1202 to an LR imager other than the target LR imager. As such, the parallax information for the pixel 1202 cannot be used to directly warp the pixel onto the decimation grid for the target LR imager 1204. However, the transformed parallax information can be computed using inter-LR imager calibration information (i.e., baselines) provided as an input to the SR process.

[00206] In a number of embodiments, the inter-LR imager calibration information utilized in computing the transformed parallax information includes the baseline

distance (i.e. the distance between the center of each LR image) between the reference LR imager and the source LR imager, which can be denoted as *b_{source_x}* (vertical direction), *b_{source_y}* (horizontal direction), and the baseline distance with respect to the reference LR imager and the target LR imager, which can be denoted as *b_{target_x}* (vertical direction), *b_{target_y}* (horizontal direction). The parallax information associated with the source LR imager and the **Type 2** pixel position can be denoted by *parax* and *paray*, and the parallax information that needs to be determined to be applied to warp the **Type 2** pixel position to the frame of coordinates of the target LR imager can be denoted by *target_parax* and *target_paray*. The parallax information to be applied to warp the **Type 2** pixel to the frame of coordinates of the target LR imager can be computed in accordance with the following process:

```

if(parax ~= 0)
    target_parax = btarget_x/bsource_x*parax;
    target_paray = btarget_y/bsource_x*parax;

elseif(paray ~= 0)
    target_parax = btarget_x/bsource_y*paray;
    target_paray = btarget_y/bsource_y*paray;

else
    target_parax = 0;
    target_paray = 0;
end

```

[00207] Once the parallax information *target_parax* and *target_paray* is computed, the computed parallax information can be added to the geometric calibration data between the reference LR imager and the target LR imager for the current pixel position to obtain the total shift to be applied for that pixel position. The total shift can then be used to shift the pixel in the **Type 2** pixel position to place it on the decimation grid of the target LR imager in the appropriate pixel position 1206.

[00208] Although a specific process is discussed above, any of a variety of processes for determining the appropriate warping to apply to a pixel that does not originate from a target LR imager can be utilized to place the pixel on a decimation grid for the target LR imager in accordance with embodiments of the invention. Accordingly, any other appropriate technique can be utilized to warp **Type 2** pixel positions onto the decimation grid of a target LR imager in accordance with an embodiment of the invention.

WARPING OF TYPE 3 PIXEL POSITIONS

[00209] Many pixel positions in an HR image estimate do not have origination or associated geometric transformation information (or warping information), i.e., **Type 3** pixel positions. **Type 3** pixel positions are illustrated as the "X" pixel positions 1212 in FIG. 12. In order for a pixel position to be a **Type 3** pixel position, a determination has been made that there are pixel positions, within a given neighborhood of the pixel position, that are **Type 1** or **Type 2** pixel positions (i.e. originating from an LR imager). The **Type 1** and/or **Type 2** pixel positions can be used to warp the **Type 3** pixel positions onto the decimation grid of a target LR imager. Referring again to FIG. 12, both "X" pixel positions 1212 have as neighbors **Type 1** or **Type 2** pixel positions. The SR processing has access to warping information for these neighboring pixel positions, albeit corresponding to various LR imagers that are not necessarily the target LR imager.

[00210] In several embodiments, the total shift for each **Type 3** pixel position can be determined by:

a) determining the warping information for each of the neighboring **Type 1** and/or **Type 2** pixel positions to shift the pixels onto the decimation grid of the target LR imager in a manner similar to the manner outlined above with respect to **Type 1** and **Type 2** pixel positions, and

b) determining the shift to apply to the **Type 3** pixel position based upon the shift information for the **Type 1** and/or **Type 2** pixel positions according to some criterion.

[00211] In several embodiments, determining the warping information for a **Type 3** pixel position based upon the warping information of neighboring **Type 1** and/or **Type 2** pixel positions involves computing a set of weights that relate the value of the pixel at the **Type 3** pixel position to the values of the pixels in neighboring **Type 1** and/or **Type 2** pixel positions. This relation can be expressed using bilateral (or non-local mean) weights, which measure a bilateral proximity (e.g., spatial and range-wise) of the **Type 3** pixel to its **Type 1** and/or **Type 2** pixel neighbors.

[00212] In a number of embodiments, the number of **Type 1** and/or **Type 2** pixel position neighbors that the **Type 3** pixel position can be denoted by *count_valid_neighbors*, and the parallax information for each of the **Type 1** or **Type 2** pixel positions for warping the pixels to the frame of coordinates of the target LR imager are denoted by *target_parax_cand(u)* and *target_paray_cand(u)*, $u=1\dots\text{count_valid_neighbors}$. In addition, the weights applied to the parallax of each neighboring **Type 1** or **Type 2** pixel position are denoted by $w(u)$, $u=1\dots\text{count_valid_neighbors}$. In these embodiments, the parallax information (*target_parax*, *target_paray*) for each **Type 3** pixel position is obtained by applying the weights $w(u)$ to the parallax information for the neighboring **Type 1** and/or **Type 2** pixel positions in accordance with the following process:

```
target_parax=0;
target_paray=0;
```

```
for u=1:count_valid_neighbors
```

```
    target_parax=target_parax+(w(u)/wsum*target_parax_cand(u));
    target_paray=target_paray+(w(u)/wsum*target_paray_cand(u));
```

end

where

$$wsum = \sum_{u=1}^{count_valid_neighbors} w(u) .$$

[00213] The obtained parallax information can then be added to the geometric calibration data to obtain a total shift for the **Type 3** pixel position and the pixel value can be shifted onto the decimation grid of the target LR imager. Although a specific process for shifting a **Type 3** pixel position is outlined above, other processes for filtering warping information of neighboring pixel positions to obtain a total shift used to warp the **Type 3** pixel position onto the decimation grid of a target imager can be utilized in accordance with embodiments of the invention.

STACKS OF PIXELS

[00214] During the process of warping pixels from the HR image estimate onto a decimation grid for a target LR imager, it is possible that multiple pixels are warped to the same position on the decimation grid. In this case, "collision stacks" are formed, which are resolved to at most one pixel per position that can then be processed by the subsequent operations (e.g., blur, decimation operators). For this purpose, the type of pixels that form a collision stack at a given position (k,l) on a decimation grid (see for example the 2-stack 1214 shown in FIG. 12), are accommodated. The collision stacks are formed from pixels that originated from **Type 1**, **2** and **3** pixel positions. In several embodiments, the following two stage process is utilized for selecting a single pixel from a collision stack:

- 1) apply rules for the formation of the stacks; and
- 2) select a pixel from the collision stack based upon on the types of pixels in the stack.

[00215] In many embodiments, the following rules are applied during the warping process to control the formation of a collision stack at a given position $\langle k, l \rangle$:

i) if the occlusion map for position $\langle k, l \rangle$ on the decimation grid for a target LR imager indicates that the pixel is occluded in the reference LR camera, then the pixel is discarded. When a pixel position is occluded to the reference LR imager, the position is not used in the subsequent formation of a gradient and can therefore be discarded;

ii) if there are no pixels in position $\langle k, l \rangle$, place the pixel in the position;

iii) if pixels from **Type 1** and/or **Type 2** pixel positions have already been placed at position $\langle k, l \rangle$, and another pixel from a **Type 1** or **Type 2** pixel position is warped there, add the pixel to the collision stack;

iv) if existing pixels from **Type 1** and/or **Type 2** pixel positions have already been placed at position $\langle k, l \rangle$, and another pixel from a **Type 3** pixel position is warped there, discard the pixel from the **Type 3** pixel position;

v) if pixels from **Type 3** pixel positions have already been placed at position $\langle k, l \rangle$, and another pixel from a **Type 3** pixel position is warped there, add the pixel to the collision stack;

vi) if pixels from **Type 3** pixel positions have already been placed at position $\langle k, l \rangle$, and a pixel from a **Type 1** or **Type 2** pixel position are warped there, record the pixel from the **Type 1** or **Type 2** pixel position and discard the existing pixels. At which point, no further pixels from **Type 3** pixel positions are added to the collision stack.

[00216] A process that can be utilized to place warped pixels on a decimation grid in accordance with an embodiment of the invention is illustrated in FIG. 13. The process 1300 first determines [1302] whether the position $\langle k, l \rangle$ on the decimation grid to which the pixel is warped is occluded. When the pixel in position $\langle k, l \rangle$ is occluded in the reference LR imager, the pixel can be discarded [1308] in accordance with rule i) above. When the pixel is not occluded, a determination [1304] is made concerning whether there are any pixels from **Type 1** or **Type 2** pixel positions at position $\langle k, l \rangle$. In the event

that there are pixels from **Type 1** or **Type 2** pixel positions, then a determination (1306) is made concerning whether the pixel being warped to position $\langle k, l \rangle$ is from a **Type 1** or **Type 2** pixel position. When the pixel is from a **Type 3** pixel position, the pixel can be discarded (1308) in accordance with rule iv) above. When the pixel is from a **Type 1** or **Type 2** pixel position, then the pixel can be added (1310) to the collision stack in accordance with rule iii) above.

[00217] When there are no pixels from **Type 1** or **Type 2** pixel positions at position $\langle k, l \rangle$, then a determination (1312) is made concerning whether there are pixels from **Type 3** pixel positions at position $\langle k, l \rangle$. When there are no pixels from **Type 3** pixel positions, then the process knows that there are no pixels at position $\langle k, l \rangle$ and the pixel can be added (1310) to the location in accordance with rule ii) above. When there is at least one pixel from a **Type 3** pixel position at position $\langle k, l \rangle$, then a determination (1316) is made concerning whether the pixel being warped to position $\langle k, l \rangle$ is from a **Type 1** or **Type 2** pixel position. When the pixel is from a **Type 3** pixel position, the pixel is added (1310) to the collision stack in accordance with rule vi) above. When the pixel is from a **Type 1** or **Type 2** pixel position, then the pixels from **Type 3** pixel positions at position $\langle k, l \rangle$ are discarded and the pixel from the **Type 1** or **Type 2** pixel position is placed onto the decimation grid at position $\langle k, l \rangle$ in accordance with rule v) above. Although a specific process is illustrated in FIG. 13, any of a number of different processes incorporating either the rules listed above or a different set of rules can be utilized in the placement of pixels onto a decimation grid in accordance with embodiments of the invention.

[00218] Once collision stacks have been formed on a decimation grid, an estimated depth-based filtering can be performed to decide the sample that will be selected at each position $\langle k, l \rangle$ on the decimation grid where there is a collision stack. Each pixel in the collision stack at position $\langle k, l \rangle$ has associated parallax information ($target_parax$ and $target_paray$) and/or total shift information, as explained above. Relative-depth filtering can be performed using this parallax information.

[00219] The baseline for the target LR imager relative to the reference LR imager is known, and its vertical and horizontal components can be denoted by $btarget_x$ and $btarget_y$. In a number of embodiments, depth-based filtering of pixels in a collision stack, aims to select the pixel from the collision stack that would appear closest to the reference LR imager in the scene captured by the reference LR imager and, therefore, is the pixel which would occlude all the other pixels in the collision stack. The depth-based filtering is thus an ordering (or sorting) operation based on the estimated depth-of-field of each of the pixels in the stack.

[00220] In many embodiments of the invention, the following method of filtering for a collision stack including M pixels is utilized:

```

for s=1:M
    if target_parax~=0
        z(s)=|btarget_x/target_parax|;
    elseif target_paray~=0
        z(s)=|btarget_y/target_paray|;
    else
        z(s)=0;
    end
end

index_of_min=sort(z); % sort values in ascending order and obtain the index of
                      the % minimum value

```

[00221] Using the above pseudo-code, the index of the pixel in the collision stack that has the minimum estimated depth in the scene can be determined. This pixel is then selected as the value recorded at position (k,l) in the decimation grid for the target LR imager. Its parallax information is also recorded, since the gradient resulting at this position will have to be warped back into the frame of coordinates of the HR image estimate (i.e. the frame of coordinates corresponding to the reference LR imager).

[00222] Although specific processes for building pixel stacks and performing depth filtering are disclosed above, any of a variety of techniques that place a pixel from a pixel stack on a decimation grid of a target LR imager that are appropriate to a specific application can be utilized in accordance with embodiments of the invention.

DATA STRUCTURE

[00223] A data structure that can be utilized in the processing of pixels in the manner outlined above can be formed which, at each pixel position in the decimation grid can store the pixel values being shifted there, with the possibility that more than one pixel, or none, may land in the same position in this data structure. If multiple pixels are shifted to the same position in this data structure, these multiple pixels represent a "collision stack" in the warped image. The data structure can also include a binary flag for each warped pixel, which indicates whether the pixel has associated origination information. The data structure can also include filtered parallax information for each pixel, where the filtered parallax information can be obtained in the manner outlined above and is utilized in shifting the pixel at that location from HR image estimate.

EMPTY PIXEL POSITIONS ON DECIMATION GRID

[00224] After warping (populating) the (k,l) positions of the decimation grid for the target LR imager, in a manner similar to the manner outlined above, some pixel positions (k,l) may still remain empty. These pixels can be interpolated. In this way, a fully populated decimation grid can be formed, which is ready for the next stages of forward blurring and gradient computation. Alternatively, it is possible to leave these positions empty and avoid performing the operations that follow (i.e., blurring, gradient formation, inverse warping) across the vacant pixel positions, and setting the gradient to zero at these positions, for the target LR camera.

FORWARD BLURRING

[00225] Once the HR image estimate has been warped onto the decimation grid of a target LR imager, the warped HR image estimate can be blurred by a total blur including both lens blur and sensor blur. Processes for determining the total blur and for applying the total blur to the pixels on the decimation grid of a target LR imager can be performed in accordance with the blur processes described above.

DECIMATION

[00226] Each position (k,l) in the warped and blurred HR image estimate is a pixel in an HR image simulated with respect to the frame of coordinates of a target LR imager. A simulated LR image with respect to the frame of coordinates of the target LR imager can be obtained by decimating the warped and blurred HR image estimate at regular decimating intervals corresponding to the SR factor r , starting at the origin of the decimation grid (see for example the decimation step 1012 in FIG. 10). While decimation can be performed as a separate process to obtain simulated LR images, as is discussed further below many SR processes in accordance with embodiments of the invention locally perform decimation as part of gradient computation.

GRADIENT COMPUTATION

[00227] Once the HR image estimate has been warped and blurred into the frame of coordinates of a target LR imager, a gradient computation similar to the gradient computation 1012 shown in FIG. 10 can be performed based upon the difference between the warped and blurred HR image estimate at the decimation positions and the input LR image pixels captured by the current target LR imager. In a number of embodiments, the gradient computation is performed for each pixel position (k,l) at the decimation positions on the warped and blurred HR image decimation grid corresponding to the target LR imager, and that are not occluded in the field of view of the reference LR imager. The aforementioned pixels on the decimation grid represent a

synthesized LR image in the field of view of the target LR imager. While processes in accordance with embodiments of the invention can decimate the warped and blurred HR image to obtain the synthesized LR image, the pixels from the warped and blurred HR image estimate in the high resolution neighborhood of the pixels in the synthesized LR image can be utilized in uncertainty processing. Therefore, processes in accordance with many embodiments of the invention do not decimate the warped and blurred HR image estimate until after the gradient is formed. In these processes, a comparison is performed between the synthesized LR image pixels and the input LR image pixels captured by the target LR imager. However, additional information from the warped and blurred HR image estimate is also utilized in the uncertainty processing.

[00228] Whether a pixel is occluded when shifted back into the field of view of the HR image estimate can be easily tested as each pixel has associated with it total shifting information. The gradient can be directly computed from a comparison of the pixels at the decimation positions corresponding to the target LR imager and the input LR image captured by the target LR imager. In many embodiments, the gradient is determined in a way that accommodates uncertainty.

ACCOMMODATING UNCERTAINTY IN GRADIENT CALCULATIONS

[00229] In several embodiments, uncertainty is accommodated by performing block comparisons for pixels (k,l) at the decimation positions corresponding to the target LR imager (i.e. using a local decimation of the warped and blurred HR image estimate). Many processes for calculating the gradient at a pixel (k,l) at a decimation position in accordance with embodiments of the invention involve assigning the pixel to one of two different cases depending upon whether the target LR imager is within the same color channel as the HR image estimate. Although a specific process is discussed below for accommodating uncertainty in the gradient calculation, any of a variety of processes for determining gradient can be utilized in a MAP estimation process in accordance with embodiments of invention.

CALCULATING GRADIENT WHEN SAME COLOR CHANNEL

[00230] **Case A** can be defined as the situation where the target LR imager is part of the same color channel as the HR image estimate, and the pixel (k,l) has origination information with respect to the target LR imager (i.e., is from a **Type 1** pixel position). When the pixel (k,l) has origination information from the target LR imager, its originating coordinates $(i[cam_count], j[cam_count])$ from the target LR imager cam_count are recorded. As noted above, uncertainty in the array camera processing pipeline can be accommodated by comparing the pixel position (k,l) with a block of pixels in the neighborhood of the pixel's originating coordinates in the LR image. As such, other positions around $(i[cam_count], j[cam_count])$ may be recorded as well (e.g., in a window centered at $(i[cam_count], j[cam_count])$), for the purpose of forming a gradient at position (k,l) . In several embodiments, the positions recorded for an LR imager centered at position $(i[cam_count], j[cam_count])$ are designated as s . Therefore, the set of positions that are used in the gradient calculation can be enumerated using the notation $(i[cam_count,s], j[cam_count,s])$. Expressing the positions in this way indicates the dependency of the LR position, with respect to which a gradient will be computed, on both the target LR imager that is selected, and the spatial position in that LR imager. Once the appropriate pixel positions are identified, the process for determining the gradient is the same in each case and is discussed below.

CALCULATING GRADIENT WHEN DIFFERENT COLOR CHANNEL

[00231] **Case B** can be defined as the situation where the pixel (k,l) at a decimation position does not fall within **Case A** above. The difference from **Case A** above comes from how the gradient is computed at these types of positions (k,l) . To determine the gradient at such positions two types of information can be determined:

- 1) the LR imagers to select for determining gradient at (k,l) , and
- 2) the coordinates to consider in the selected LR imagers.

[00232] To decide on the first question, a window is centered at position (k, l) in the warped and blurred HR image estimate, the neighboring pixel positions in that window are checked for certain characteristics, and their information is recorded according to several criteria. These include whether a neighboring pixel was warped from a **Type 1** or **Type 2** pixel position in the current HR estimate, as previously defined, and whether the source LR imager from which the pixel originates is in the same color channel as the reference LR imager. If a neighboring pixel position satisfies these two conditions, the source LR imager identifier, i.e., $camera_id$, is added to the list, and the source LR imager origination information corresponding to that neighboring pixel is recorded. The latter consists of the coordinates $(i[camera_id], j[camera_id])$ in the source LR imager from which the neighboring pixel originated. Around this originating pixel, as before, a neighborhood of pixel positions can be recorded in each selected source LR imager. The neighborhood of positions recorded in a source LR imager can be indexed by s , situated in a local neighborhood (e.g., a rectangular window) around position $(i[camera_id], j[camera_id])$ and inclusive of it, i.e., the pixel and its neighbors can be referred to as indexed by $(i[camera_id, s], j[camera_id, s])$. In contrast to **Case A** described above, several neighborhoods of pixels in several of the captured LR images can be utilized in gradient calculations of pixels in **Case B** in accordance with embodiments of the invention. In addition, the size of the neighborhood can be adapted based upon the reliability of the pixel at position (k, l) .

GRADIENT FORMATION

[00233] The process described above identifies blocks of pixels in the captured LR images that can be used in the computation of gradients for some, but possibly not all, of the decimation positions (k, l) in the warped and blurred HR image estimate. In many embodiments, the information that is utilized in calculating gradients includes lists of source LR imagers $[camera_id]$, and corresponding sets of coordinates $(i[camera_id, s], j[camera_id, s])$ from the source LR images, obtained as described above

in various cases. Gradient formation for pixel positions $\langle k, l \rangle$ when various options are taken in terms of the number of LR imagers considered and the number of $\langle i, j \rangle$ positions in those LR images is discussed below.

[00234] For forming gradients at some positions $\langle k, l \rangle$ on the decimation grid in the warped and blurred HR image (such as described by **Case A**), we can have a single LR imager selected to be visited (i.e., the originating camera), and multiple positions around the coordinates $\langle i, j \rangle$ in this LR imager. Let us denote by $p(k, l)$ the value of the pixel at $\langle k, l \rangle$. Alternatively, even for **Case A** we could use the same LR imager selection methodology for the purpose of forming gradients at position $\langle k, l \rangle$, as was described for **Case B**.

GRADIENT FORMATION CONSIDERING A SINGLE LR IMAGER

[00235] In one embodiment, for decimation positions $\langle k, l \rangle$ where a single LR imager is considered, the gradient is formed by taking the weighted sum of the difference between each pixel value q in the neighborhood of position $\langle i, j \rangle$ including the pixel at position $\langle i, j \rangle$ in the input LR image (i.e., for $camera_id = cam_count$), designated by $q(i(cam_count, s), j(cam_count, s))$, and the pixel value $p(k, l)$. The weighted formation of the gradient at position $\langle k, l \rangle$ in the warped HR image corresponding to the target LR imager denoted by cam_count , can be written as:

$$g(cam_count, k, l) = \sum_s weight(cam_count, s) * [q(i(cam_count, s), j(cam_count, s)) - p(k, l)] \quad [7]$$

[00236] The cam_count index in the gradient expression g indicates that the gradient was formed by warping into the coordinate frame of the target LR camera (i.e. cam_count). The determination of $weight(cam_count, s)$ is described below.

GRADIENT CALCULATION INVOLVING MULTIPLE LR IMAGERS

[00237] For decimation positions $\langle k, l \rangle$ where more than one input LR image is involved in the computation of the gradient, whether the pixel is part of the simulated LR image or not, e.g., falling under **Case B** respectively, there are lists of input LR images indexed by $camera_id$, and coordinates of neighborhoods of pixels $i(camera_id, s), j(camera_id, s)$ in each of the input LR images, as it was described above. A weighted gradient can be accumulated at position $\langle k, l \rangle$ in the simulated LR image corresponding to the input LR image cam_count , using the following expression:

$$g(cam_count, k, l) = \sum_{camera_id} \sum_s weight(camera_id, s) * [q(i(camera_id, s), j(camera_id, s)) - p(k, l)] \quad (8)$$

[00238] In Eq. (8), a weighted sum is taken of the difference between every pixel in the neighborhoods identified in each of the input LR images recorded in the list of LR imagers to be visited, and the value of the pixel at position $\langle k, l \rangle$. Again, the cam_count index in the gradient expression g indicates that the gradient was formed between the pixel values at the decimation positions from the HR image estimate, and the selected input LR images for that purpose (including the image captured by the target LR imager if appropriate, as described above).

[00239] Even though they were presented separately, Eq. (8) has as particular case Eq. (7) (obtained by making $camera_id = cam_count$), and only this latter expression will be further discussed.

DETERMINING WEIGHTS TO APPLY TO PIXEL COMPARISONS

[00240] To determine the values of $weight(cam_count, s)$ in Eq. (7) where $camera_id = cam_count$, or $weight(camera_id, s)$ in Eq. (8) the following approach is taken, which assumes that localized decimation is performed as part of the gradient

calculation process on the HR grid, instead of actually generating the decimated simulated LR images:

- 1) Take a block of size $L \times L$ centered at position $i(camera_id,s), j(camera_id,s)$ in the input LR image denoted by $camera_id$. Call this blk_LR .
- 2) Take a block of size $L*r \times L*r$ centered at position (k,l) in the blurred and warped HR image (r is the SR factor).
- 3) Decimate the block formed in point 2) above by r in each direction such that a block of size $L \times L$ is obtained with $p(k,l)$ at its center. Call this block blk_HR .
- 4) Compute the $weight(camera_id,s)$ as the bilateral weight (un-normalized) using the norm of the difference between blk_LR and blk_HR , i.e.,

$$weight(camera_id,s) = e^{-|blk_LR - blk_HR|/\sigma}, \quad [9]$$

where σ is a user-selectable parameter (typically as a percentage of the dynamic range of the imagers).

[00241] The sum of the weights for the component gradients that contribute to the gradient being formed at position (k,l) in the current warped HR frame is also recorded. This is:

$$wsum_grad(cam_count,k,l) = \sum_{camera_id} \sum_s weight(camera_id,s). \quad [10]$$

[00242] The $g(cam_count,k,l)$, and the sum of the weights $wsum_grad(cam_count,k,l)$ are decimated from the decimation grid using the decimation positions. The result is an LR gradient and LR sum of weights corresponding to the target LR imager. The accumulated gradient at any decimation position (k,l) (these positions correspond to the LR gradient image) is not normalized in the warped and blurred HR frame. Instead, the weighted gradient $g(cam_count,k,l)$, and the sum of the weights $wsum_grad(cam_count,k,l)$ computed as described above, are recorded at each

decimation position (k,l) visited in the warped and blurred HR image estimate, and these two sets of LR data are backward-simulated (i.e. applying a backward imaging transformation, which is the reverse of the forward imaging transformation) through upsampling, transpose blurring, and warping back to the frame of coordinates of the current HR estimate (corresponding to the reference LR imager). Only after the weighted gradient and the corresponding weights are in the frame of coordinates of the current HR estimate, are the total accumulated gradients, at the grid positions in the current HR estimate, normalized. Processes for backward simulating weighted gradients and weights and accumulating backward simulated weights in accordance with embodiments of the invention are discussed further below.

UPSAMPLING, BLURRING AND BACK WARPING THE CALCULATED GRADIENTS

[00243] For the decimation positions (k,l) in the warped and blurred HR image situated on the decimation grid, gradients and sum of weights can be computed as shown in Eqs. (7), (8) or (10). These values are situated on the decimation grid for the target LR imager, which is an HR grid corresponding to a high resolution image. The rest of the values in this HR grid are set to zero, thus effectively constructing an upsampled weighted gradient image, and the corresponding upsampled sum-of-weights matrix, corresponding to the warped HR image that is being processed. Both of these images are convolved with the transpose of the total blur H that was applied in the forward simulation previously described.

[00244] Following these operations, the resulting two matrices (upsampled and blurred weighted gradient, and upsampled and blurred sum of weights) are warped back into the system of coordinates of the current HR image estimate, by using the total shift information that was applied during the forward warping of the current HR image estimate into the frame of coordinates of the target LR imager to perform the backward or reverse warping (i.e. from the frame of coordinates of the current LR imager to the frame of coordinates of the HR image estimate).

[00245] It should be noted that the above operations do not have to be performed at the level of the entire HR grid, but rather a patch (local) implementation approach can be taken, with consideration given to the sizes of the local patches such that the operations described function correctly.

[00246] As noted above, in order to compute the likelihood gradient term the upsampled and blurred weighted gradient and upsampled and blurred sum of weights are warped back and accumulated at the corresponding position in the current HR image. The corresponding position in the current HR image for position (k,l) on the upsampled grid can be designated as (k',l') . This warping operation can use the total forward warp information generated during the warping of the current HR image estimate to the frame of coordinates of the target LR imager.

[00247] Once the above process is completed, the backward-simulated weighted-gradient corresponding to the target LR imager is formed and accumulated at the proper positions (k',l') , along with its backward-simulated sum-of-weights matrix, in the current HR image grid.

ACCUMULATION OF GRADIENTS ACROSS ALL CAPTURED LR IMAGES

[00248] Using the same described process, forward-warping (1008), blurring (1010), gradient formation (1012), upsampling (1018), transpose blurring (1020), and backward-warping (1022) is repeated for all, or subsets, of the LR imagers, as shown in Fig. 10. The gradients and the corresponding weights accumulate accordingly at all positions (k',l') in the HR image estimate's system of coordinates. The accumulated gradients at each position (k',l') , resulting from gradient computation with respect to all captured LR images, can be filtered in various ways to form the final gradient. For example, the accumulated gradient at each position (k',l') can be normalized by the sum of the weights accumulated at that position, i.e., dividing the accumulated gradient value by the corresponding weight at that location (k',l') .

[00249] Normalizing the accumulated gradients with the weights accumulated at a position forms the final likelihood gradient term for the current iteration, for each position (k', l') in the current HR estimate (at iteration n). The likelihood gradient term can be combined with other gradient terms in the manner outlined below to determine a new estimate for the HR image (i.e., at iteration $n+1$).

THE INTRA-CHANNEL PRIOR GRADIENT TERM

[00250] An intra-channel prior term used in a MAP process in accordance with embodiments of the invention can be utilized to achieve any of a variety of image processing objectives. In a number of embodiments, the intra-prior term is used to stabilize the solution of the MAP estimation process. In many embodiments, the intra-channel prior term preserves edges in the estimated SR image, while smoothing noise in appropriate areas of the image (flat or smooth varying), and is adaptive to the local image content. In several embodiments, the intra-channel prior term seeks an invariance of the estimated SR image under an operator $(K(.))$. In a number of embodiments, the operator $(K(.))$ is, but is not limited to, a bilateral filter, or similar generalized non-local operator.

[00251] In many embodiments, $K(.)$ is a locally (spatially)-adaptive kernel, and Λ_1 is a function of the local SNR (estimated for each pixel position on the HR grid), or be the identity matrix. The SNR at a given position can be estimated using:

- a) the noise variance estimated at that position (as a function of the signal level and the analog gain of the sensor),
- b) the observed image variance at that position, and
- c) the blur model H^c .

[00252] Using the above formulation, the intra-prior term gradient is as follows:

$$\psi_{\Lambda_1}^{\text{intra}}(x^c) = 2\Lambda_1(I - U)(x^c - K(x^c)) = 2\Lambda_1[(x^c - K(x^c)) - U(x^c - K(x^c))] \quad (14)$$

where Λ_1 is a function of the local SNR (estimated for each pixel position on the high resolution grid), or can be the identity matrix. I is the identity matrix.

[00253] Given that the operator $K(\cdot)$ is local, Eq. (14) is implemented locally by filtering each pixel in the high resolution estimate x^c in a window W centered at that pixel position. Thus, to implement Eq. (14) and compute the intra-prior gradient at a position (k,l) (appropriately translated from 2D to vector coordinates, since x^c is a vector-form of the image) in x^c , in the local window W centered at position (k,l) , the filtering proceeds as follows. Assume $p(k,l)$ is the initial value of the pixels in x^c :

1) Filter the current estimate x^c using the filter operator $K(\cdot)$ (e.g., a bilateral filter). The resulting filtered pixels are denoted by $p1(k,l)$.

2) Form the block difference e containing pixels $e(k,l)=p(k,l)-p1(k,l)$, for all (k,l) in a local window W centered at (k,l) .

3) Convolve the block of difference signal $e(k,l)$ in window W , centered at (k,l) and obtained as in point 2) above, with the matrix operator $U_{W(k,l)}$ such that the gradient value at position (k,l) is:

$$g^{\text{intra}}(k,l) = 2\eta_{k,l}[(p(k,l) - p1(k,l)) - (U_{W(k,l)} * e)(k,l)], \quad [15]$$

where $*$ is the convolution operator, and $\eta_{k,l}$ is the value from the diagonal matrix Λ_1 in Eq. (14) corresponding to the 2D position (k,l) .

[00254] By way of example, a 5x5 window size W can be taken and the matrix operator $U_{W(k,l)}$ applied to the 2D local block e at position (k,l) has the form:

$$U_{W(k,l)} = \begin{bmatrix} \frac{w(k-2,l-2;k,l)}{Z(k-2,l-2)} & \frac{w(k-2,l-1;k,l)}{Z(k-2,l-1)} & \frac{w(k-2,l;k,l)}{Z(k-2,l)} & \frac{w(k-2,l+1;k,l)}{Z(k-2,l+1)} & \frac{w(k-2,l+2;k,l)}{Z(k-2,l+2)} \\ \frac{w(k-1,l-2;k,l)}{Z(k-1,l-2)} & \frac{w(k-1,l-1;k,l)}{Z(k-1,l-1)} & \frac{w(k-1,l;k,l)}{Z(k-1,l)} & \frac{w(k-1,l+1;k,l)}{Z(k-1,l+1)} & \frac{w(k-1,l+2;k,l)}{Z(k-1,l+2)} \\ \frac{w(k,l-2;k,l)}{Z(k,l-2)} & \frac{w(k,l-1;k,l)}{Z(k,l-1)} & \frac{w(k,l;k,l)}{Z(k,l)} & \frac{w(k,l+1;k,l)}{Z(k,l+1)} & \frac{w(k,l+2;k,l)}{Z(k,l+2)} \\ \frac{w(k+1,l-2;k,l)}{Z(k+1,l-2)} & \frac{w(k+1,l-1;k,l)}{Z(k+1,l-1)} & \frac{w(k+1,l;k,l)}{Z(k+1,l)} & \frac{w(k+1,l+1;k,l)}{Z(k+1,l+1)} & \frac{w(k+1,l+2;k,l)}{Z(k+1,l+2)} \\ \frac{w(k+2,l-2;k,l)}{Z(k+2,l-2)} & \frac{w(k+2,l-1;k,l)}{Z(k+2,l-1)} & \frac{w(k+2,l;k,l)}{Z(k+2,l)} & \frac{w(k+2,l+1;k,l)}{Z(k+2,l+1)} & \frac{w(k+2,l+2;k,l)}{Z(k+2,l+2)} \end{bmatrix} \quad [16]$$

where, $w(k,l;k',l')$ are the weights of the operator $K(\cdot)$ (e.g., bilateral filter, or its non-local block extension) computed at position (k,l) with respect to position (k',l') inside window W . In turn, the variables $Z(k,l)$ represent the sum of all the weights $w(k,l;k',l')$, where the indices (k',l') vary in the neighborhood W of (k,l) .

[00255] Although a 5x5 window size is discussed above, any of a variety of window sizes can be utilized in accordance with embodiments of the invention to obtain the matrix operator $U_{W(k,l)}$ used in the calculation of the intra-channel gradient term at pixel position (k,l) .

[00256] In other embodiments, different intra-channel priors can be utilized. In several embodiments, an alternative prior term can be used that can be formulated in similar terms to Eq. (1)-(2) above. As in the process described above, the output of the Pilot Fusion process is utilized as the initial HR estimate, $x^c(0)$. Alternatively, another initial estimate may be provided for the super-resolution processing (such as for example an interpolated high resolution version of the reference camera image). The iterative, HR-estimate updating process then proceeds as follows:

$$x_{ML}^c(n+1) = x^c(n) - \beta^n \nabla_{x^c(n)} J(x^c(n)), \quad [17]$$

where

$$\nabla_{x^c(n)}(J(x^c(n))) = VT(x^c(n)), \quad [18]$$

and

$$x^c(n+1) = F^{-1}\{T(F(x_{ML}^c(n+1)))\}, \quad [19]$$

where F is a sparsifying transform known in the art (e.g., a wavelet, or directional transform), F^{-1} is its inverse (it is understood that the transform can be designed such that an actual inversion operation need not be taken, e.g., an orthogonal transform for which its inverse is equal to its transpose), and T is a hard or soft thresholding operator known in the art, which is applied to the transform coefficients. The likelihood-term gradient in Eq. (18) is determined using processes in accordance with embodiments of the invention including processes described herein.

INTER-CHANNEL GRADIENT TERMS

[00257] The inter-channel prior term enforces a correlation of edge direction and strength across channels. Although an inter-channel prior term can be incorporated into a MAP process in accordance with many embodiments of the invention, in many embodiments it is computationally efficient to enforce edge direction and strength through the cross-channel of R and B with respect to a HR image estimate for the G channel obtained using the MAP process described above (excluding an inter-channel gradient term).

COMBINING THE COMPUTED GRADIENT TERMS

[00258] The gradient terms determined as described above (i.e., likelihood-term gradient, intra-prior gradient, and if available inter-prior gradient) are accumulated at

the end (possibly in a weighted fashion as dictated by the matrices Λ_1 and Λ_2 in Eq. (2)), to form the total gradient for the current iteration, as shown in Eq. (2). The total gradient is then scaled by the step size β^n for the iteration, and applied to the current HR image estimate, thereby obtaining a new HR image estimate in the current channel, as indicated in Eq. (1).

[00259] Using the approaches described for the determination of the gradients in Eq. (2), this is performed iteratively in Eq. (1), until a stopping criterion is reached (e.g., a norm on the variation of the estimate with iteration number falls below a certain threshold, or the maximum number of iterations is attained), as shown in Fig. 9.

COMBINING THE HR IMAGE ESTIMATES FROM EACH CHANNEL

[00260] Once HR image estimates have been obtained for each color channel, a final SR image can be generated by combining the HR images from each channel. In many embodiments, an estimation process similar to the process outlined above is used to generate a HR image for the Green channel. The Red and Blue channels can then be fused only, or can themselves be processed using a similar MAP process. Where the Red and Blue channels are super-resolved at lower resolution than the Green channel (possibly due to fewer LR imagers allocated to the R and B channels due to the lesser sensitivity of the human visual system to these channels compared to the Green channel), the Red and Blue channels can be upsampled (i.e., interpolated) to the full resolution of the Green channel and the three channels combined using any suitable technique for combining three color channels to produce a full color image. As part of the process, the Red and Blue channels can undergo conventional denoising and sharpening. In several embodiments, the Red and Blue channels are fused by enforcing a cross-channel correlation (e.g., in a color-difference sense and fused using the Red-Green and Blue-Green color difference). Or this can be done by fusing a combination of the R,G,B channel data such as fusing the $2 \times G-R-B$ and $G-B$ signals for example (pixel -

wise), instead of fusing the values of the R and B pixels on the HR grid. Thus, these signals are fused in a manner correlated with the HR image estimate for the Green channel generated in accordance with one of the processes outlined above. The fusing of the Red and Blue channels, or the Red-Green and Blue-Green color difference information, can be performed using a process similar to the Pilot Fusion processes outlined above or additionally using a selective MAP processing of the obtained fused R and B channels, after the color-difference fusion. The selectivity of the applied MAP processing can be driven by the HR occlusion maps that are created in the SR processing of the Green channel, such that the MAP processing is applied to the R and B fused data only in the occlusion zones on the HR grid onto which the color HR image is reconstructed. In another embodiment, the depth estimated for the pixels in the HR image is used to enhance the fusion of the R and B channels. The HR fusion of the R-G and B-G data and the reconstruction of R and B HR image information from that process, the enhanced fusion of the R and B channels using HR depth information, and the selective application of the MAP processing to the fused R and B HR images are discussed further below. It should be noted that the enhanced fusion of R and B and the optional MAP processing of the HR R and B images can be performed whether the R and B HR fusion is done channel-wise, or cross-channel (using color difference processing).

[00261] The grid resolution at which the R and B images are reconstructed (through fusion and possibly MAP processing) may be lower or the same as that of the resolution of the Green channel. In the former case once the R and B images have been fused at the given resolution, they are brought at the same grid resolution as the Green channel if the G,R,B HR color planes are output by the SR processing system. This final step of increasing the resolution from the fusion grid of the R and B channels to the grid resolution of the Green channel can be done by a conventional interpolation method (including cross-channel interpolation techniques). Alternatively, the R and B fusion can

be performed directly at the same grid resolution as that of the G channel, and using the already estimated (dense) HR G channel.

[00262] Enhanced fusion of the R and B channels can be performed similarly to the processing described above for the G channel, by using HR depth information in the process of fusion. In several embodiments, the enhanced fusion involves directing the interpolation of pixel values by assigning relative weights to the pixels based upon their HR depth information. The selective MAP processing for the restoration of the R and B channels can be carried out by using the HR occlusion map estimation generated by processes supporting dynamic focus functionality, in that for efficiency purposes the MAP processing can be carried out only in occlusion zones in the R and B channel. Although specific processes are disclosed above with respect to combination of different color channels, any of variety of processes can be utilized including processes that utilize alternative and/or additional channels of information to generate a high resolution image combining information from two or more channels in accordance with embodiments of the invention.

ALTERNATIVE CALCULATIONS OF GRADIENT TERMS

[00263] Although specific processes are described above for determining the likelihood, intra-channel, and/or inter-channel gradient terms utilized in improving an estimate of an HR image in accordance with embodiments of the invention, other techniques can also be used in the determination of gradient terms. By way of example, optimization processes based upon Eqs. (1) and (2) above that utilize alternative processes for determination of likelihood gradient in accordance with embodiments of the invention are discussed further below.

[00264] A process for determining a likelihood gradient term in an SR process in accordance with an embodiment of the invention is illustrated in FIG. 14A. The process 1400 starts by initializing (1402) the first estimate of the HR image using the output of Pilot Fusion processing, and obtaining (1404) any auxiliary input information such as

inter-LR imager geometric correction information as defined above, including scene-independent geometric correction information (such as geometric calibration) and scene dependent geometric correction information (such as parallax information), occlusion information, and imaging prior information (e.g., per-pixel SNR estimates in the HR frame). All these input data are discussed above with respect to other techniques for calculation of the likelihood gradient term. As before, not all of the inputs are necessary, although most embodiments utilize the current HR image estimate and the captured images at a minimum. If the parallax information is not available and/or parallax is only slight, the SR uncertainty processing may compensate for the lack of it. In addition, the absolute distortion and the relative distortion among cameras (i.e., the scene-independent geometric distortion) can be so small that only the scene dependent parallax inherent in the LR imagers manifest themselves and can be calibrated.

[00265] The HR image estimate is then blurred (1406) by the total blur H , which includes lens blur and the sensor blur. In the blurred HR image estimate, each pixel at position (k,l) is then visited for the purpose of forming a gradient at that position with respect to one or more of the input LR images. As before, it is assumed either that the LR images have been geometrically corrected, or that geometric correction information including geometric calibration data and parallax information (if available) are available. Thus, in the ideal case, the total shifting information is available for the pixels that originate from one of the LR imagers, as it was the case above. In case some of this information is not available, and depending on the severity of the parallax, the uncertainty search window size, which will be discussed below, may be adjusted (increased) accordingly.

[00266] The discussion of the process illustrated in FIG. 14A that follows assumes that geometric calibration data exists for the pixel at position (k,l) with respect to the other imagers in the array, and also that parallax information exists relating the LR

imagers in the array. Similar to the process for calculating the likelihood gradient term discussed above, there exist different cases in terms of how much information there is relating the position (k,l) in the blurred HR image estimate and the input LR images depending upon whether the pixel at position (k,l) originated from an LR imager in the same color channel as the HR image estimate. The gradient formation in each of these cases is described below involving selecting neighborhoods of pixels in the captured LR image(s) (1410), and accumulating weighted gradients (1412) at each pixel position (k,l) .

GRADIENT COMPUTATION

[00267] The computation of gradients in the process illustrated in FIG. 14A is conceptually illustrated in FIG. 14B. The computation of gradients depends upon whether a pixel has origination information from an LR imager from the same color channel as the blurred current HR image estimate.

GRADIENT COMPUTATION WHEN SAME COLOR CHANNEL

[00268] **Case A** can be defined as applying to pixel positions (k,l) in the blurred HR image estimate (1450) that originated from an LR imager (1452) within the same color channel as the HR image estimate. A neighborhood of pixels (1454) around the origination coordinates $(i[cam_count],j[cam_count])$ in the source LR imager cam_count can be recorded for the purpose of forming a gradient at position (k,l) . The neighborhood of recorded positions in a captured LR image can be indexed by s and the neighborhood of pixels indexed as $(i[cam_count,s],j[cam_count,s])$. In the illustrated embodiment, three **Case A** pixels (1456) are shown as being mapped onto three different LR images by way of illustration. The same process is used for determining the gradient at position (k,l) in both cases and is discussed below.

GRADIENT COMPUTATION WHEN NOT SAME COLOR CHANNEL

[00269] In the context of the process illustrated in FIG. 14, **Case B** can be defined as applying to pixel positions (k,l) in the blurred HR image estimate that do not have

origination information in a source LR imager from the same color channel as the current HR estimate (i.e., that of the reference LR imager), or do not have any origination information at all. The difference from **Case A** above comes from how the gradient is computed at these types of positions $\langle k, l \rangle$. To determine the gradient at such positions, two types of information are utilized:

- 1) the LR imagers to select for determining gradient at $\langle k, l \rangle$, and
- 2) the coordinates to consider in the selected LR imagers.

[00270] To decide on the first question, a window is centered at position $\langle k, l \rangle$ in the blurred HR image estimate (1450), the neighboring pixel positions in that window are checked according to specified criteria, and their information is recorded if several criteria are satisfied. These include whether a neighboring pixel falls under **Case A** described above. If the neighboring pixel position satisfies this condition, a neighborhood of pixels around the **Case A** pixel's originating position in the input LR image from which the **Case A** pixel originated is recorded (LR image from *camera_id*). The neighborhood of pixels can be referred to as $\langle i[\text{camera_id}, s], j[\text{camera_id}, s] \rangle$. The information used to compute gradients for pixels at positions $\langle k, l \rangle$ falling under **Case B** in the blurred HR image estimate typically includes a number of neighborhoods in the input LR images that are in the same color channel as the blurred HR image.

[00271] It should be noted that even in **Case A**, the selection of neighborhoods of pixels in multiple captured LR images can proceed in the same manner as outlined with respect to **Case B**.

GRADIENT FORMATION

[00272] A weighted gradient is accumulated at position $\langle k, l \rangle$ corresponding to the blurred HR image, assuming the current pixel value in the blurred HR estimate image is $p(k, l)$:

$$g(k,l) = \sum_{camera_id} \sum_s weight(camera_id,s) * [q(i(camera_id,s),j(camera_id,s)) - p(k,l)] \quad (20)$$

[00273] Eq. (20) shows how the weighted summation of gradients is taken over all LR imagers indexed by *camera_id*, and all the positions *(i,j)* in the identified neighborhoods in the corresponding captured LR images.

DETERMINING WEIGHTS TO BE APPLIED

[00274] To determine the values of *weight(camera_id,s)* in Eq. (20) the following approach is taken:

1) Take a block of size $L \times L$ centered at position *(i(camera_id,s),j(camera_id,s))* in the input LR image designated by *camera_id*. Call this *blk_LR*.

2) Take a block of size $L*r \times L*r$ centered at position *(k,l)* in the blurred current HR image (*r* is the SR factor).

3) Decimate the block formed in 2) above by *r* in each direction such that a block of size $L \times L$ is obtained with *p(k,l)* at its center. Call this block *blk_HR*.

4) Compute the *weight(camera_id,s)* as the bilateral weight (un-normalized) using the norm of the difference between *blk_LR* and *blk_HR*, i.e.,

$$weight(camera_id,s) = e^{-|blk_LR - blk_HR|/\sigma}, \quad (21)$$

where σ is a user-selectable parameter, typically assigned a value that is a function of the dynamic range of the input LR images.

[00275] The sum of the weights for the component gradients that contribute to the gradient being formed corresponding to position *(k,l)* in the blurred current HR image is also recorded as follows:

$$wsum_grad(k,l) = \sum_{camera_id} \sum_s weight(camera_id,s). \quad [22]$$

[00276] It should be noted that the above operations do not have to be performed at the level of the entire HR grid, but rather a patch (local) implementation approach could be taken, with consideration given to the sizes of the local patches such that the operations described function correctly.

NORMALIZATION OF ACCUMULATED GRADIENTS

[00277] Referring back to FIG. 14A, the accumulated weighted gradient at position (k,l) in the current HR image estimate (i.e., $g(k,l)$) is normalized [1416] by the accumulated sum of the weights $wsum_grad(k,l)$, by dividing $g(k,l)$ by $wsum_grad(k,l)$. The normalized gradient at position (k,l) is stored [1418] and the procedure advances [1420] across all positions on the blurred HR image estimate until [1422] every position on the HR grid has been considered.

[00278] Once the gradient image corresponding to the current blurred estimate has been formed as above, the resulting gradient image is blurred [1424] by the transpose of the total blur which was used in the forward simulation path (i.e., sensor plus lens blur) to provide the final likelihood gradient term for the current iteration, for each position (k,l) in the HR image estimate. The likelihood gradient term can be used in the manner outlined above, in combination with intra-channel and/or inter-channel gradient terms where available, to determine a new estimate for the HR image. Alternatively, any of a variety of other techniques can be utilized for determining the likelihood gradient term appropriate to specific applications in accordance with embodiments of the invention.

DEALING WITH PROBLEM PIXELS

[00279] During image processing, a complex consideration often exists concerning how to diminish the impact of “problem” pixels on an output image. Problem pixels are those pixels that are not appropriate members of the data set under consideration. Problem pixels can arise for a variety of reasons including “defective” pixels in an imager. The term “defective” pixel refers to a pixel in an imager that does not produce useful output information and includes so called “hot”, “cold”, “warm”, and “cool” pixels, which refer to specific types of pixel failures. The set of pixels that can be treated as problem pixels also includes phantom pixels in neighborhood operations when that neighborhood extends beyond the edge of an image or the phantom pixels that are exposed when a part of an image is shifted physically leaving no information in its original position (this class is referred to as “occlusion set” pixels). Typically a problem pixel requires special case processing, which slows throughput due to case checking and alternate processing branches. The special cases branches that are often involved in special case processing complicate, and increase the burden associated with image processing. Special case processing can also produce artifacts in the output image.

[00280] SR processes in accordance with embodiments of the invention weight the contributions that the LR image samples from the LR images make to the synthesized output image. In many embodiments, assigning a weighting of zero to LR image samples that are determined to be “outliers” means that they are ignored during the remainder of the SR processing. In several embodiments, the mechanism of assigning a weighting of zero to an LR image sample is utilized to deal with all classes of problem pixels in a unified manner. Whenever such pixels are encountered during upstream processing of the LR image samples, problem LR image samples are flagged by assigning them a unique value in place of their original value, if any, or setting an associated flag. For example, LR image samples can be flagged by maps of defective

pixels, in noise processing algorithms where they are flagged as outliers, at boundaries of images or of segmented portions of images, and in the occlusion sets of translated image elements. Subsequent SR processing can disregard all flagged problem LR image samples in the generation of the output SR image. Treating problematic LR image samples in this way not only provides for uniform treatment of such LR image samples, but also reduces many of the artifacts that can occur when processing such LR image samples with special case processing branches, and reduces the computational load on the processing pipeline.

[00281] As can be readily appreciated, any type of pixel determined to be problematic including but not limited to the problem pixel types referenced above can be identified at one or more processing stages that occur prior to the SR process in an image processing pipeline in accordance with an embodiment of the invention. In addition, any of a variety of flagging techniques appropriate to the requirements of the SR process can be utilized to communicate to the SR process that the problem LR image samples should be disregarded in the formation of the output SR images.

ARRAY CAMERAS WITH MULTIPLE FIELDS OF VIEW

[00282] The array cameras described above assume a fixed focal length for each of the imagers in the imager array. The sub-pixel phase shift between the cameras encodes much of the information from the scene in the LR images captured by the cameras, making the LR images amenable to resolution recovery through multi-channel signal processing methods such as the SR processes described above. An alternative technique for constructing an array camera in accordance with embodiments of the invention is to create an array, which includes imagers having different fields of view. Changing the field of view of an imager results in a change in magnification of the image captured by the underlying sensor. In an array camera, this difference in magnification between imagers in the array allows variation in the degree of SR processing required to recover a desired resolution, which can be beneficial both

from the point of view of needing less computation and from the point of view of resolution recovery quality. By varying the magnification in a controlled way across the lens array, the computation required to recover a desired resolution through SR processing can be reduced.

[00283] An imager's field of view can be determined by its optics. The focal length of the imager defines the 'size' of the image captured on the sensor. A lens with a larger focal length typically captures an image which is proportionately larger than that captured by a lens with the smaller focal length. For example, the image of an object produced by a lens with a focal length of 100 mm will be exactly double the size of an image formed by a lens of focal length 50 mm at the same distance from the object. The relationship between focal length and the field of view captured means that the focal length of an imager and the size of the images captured by the imager are directly proportional. The curvature of an imager's lenses also impacts its field of view. The field of view of an imager can be increased without changing the focal length of the lens by using more curved and aspheric surfaces.

[00284] Embodiments of 5 x 5 and 4 x 4 imager arrays including imagers with different fields of view that are suitable for use in array cameras in accordance with embodiments of the invention are illustrated in FIGS. 15A and 15B respectively. The 5 x 5 imager array 1500 illustrated in FIG. 15A includes 25 imagers, which are assigned so that 17 imagers sense luma (or green light (G)), and 8 imagers sense chroma (red light (R) or blue light (B)). The 17 imagers sensing luma are sub-partitioned into three categories, based on their respective fields of view (fov) – category 1 (G1) with $\frac{1}{4}$ the maximum field of view (fov_{max}), category 2 (G2) with $\frac{1}{2}$ the maximum field of view, and category 3 with the maximum field of view. The following table compares the resolution of the array camera shown in FIG. 15A, where each imager is a VGA camera, at each focal length with the resolution of a conventional 8 Megapixel fixed focal length camera

employing a Bayer filter and using digital zoom to provide an image of an equivalent field of view.

Category	fov as % of fov_{max}	# of VGA Imagers	# of Luma Pixels in Legacy Cam within fov	Resolution Recovery Strength per Linear Dim.
1	$\frac{1}{4}$	1	240,000	0.78
2	$\frac{3}{4}$	8	2,160,000	7.03
3	1	8	4,000,000	13.02

Table 1: Three Categories of Multi Field of View VGA Cameras for a 5 x 5 array configuration

[00285] The above table illustrates a comparison between a 3200 x 2400 8 Megapixel legacy camera spanning a field of view 'fov'. The field of view is spanned by all 4000 pixels along the diagonal of the 3200 x 2400 frame. Considering the first row of Table 1, the number of pixels along this diagonal that span just a quarter of the full field of view (i.e. $fov/4$) would be approximately 1000 pixels ($4000/4$). Given that the frame has a 4/3 aspect ratio, this would imply that the sub-rectangle that spans a diagonal of 1000 pixels center around the center of the frame is approximately 800 x 600 pixels. This rectangle has 480,000 pixels of which only half have green (luma) pixels (assuming legacy camera employs a Bayer filter). Thus, the fourth column in first row of Table 1 shows the number of luma pixels in the legacy camera within the fov of ($fov/4$) is 240,000 pixels. The final column is the ratio of the number of pixels in the legacy camera capturing this field-of-view ($fov/4$) to the number of pixels in the array camera, where each LR imager has 640 x 480 pixels, capturing the same field of view ($fov/4$). The ratio provided in the final column of table 1 lists the strength of resolution recovery required through SR processing for the array camera to achieve the same resolution as the image obtained from a conventional camera using digital zoom. A value less than 1

indicates that the resolution of the individual VGA camera is more than that of the legacy camera in that region. As can be seen from table 1, the array camera has greater resolution in the central foveal region than the 8 Megapixel conventional camera. The Category 2 and Category 3 comparisons reveal that a strength of resolution recovery greater than 1 is needed for the array camera to match the resolution of equivalent images captured with an 8 Megapixel conventional camera using digital zoom. The number of imagers in Category 3 falls a little short of what is required. However, the performance of the array camera is probably acceptable since the resolution of any camera degrades with increasing field of view and the marginal decrease posted by the Category 3 imagers at fields of view greater than 75% may be tolerable in the overall context. When performing SR processing, the SR processing is performed using processes similar to the processes outlined above, where the LR images are drawn from the LR imagers within the category of interest (i.e. the LR imagers are all from imagers having the same FOV). Processing a reduced number of LR images enables a reduction in the processing overhead of performing SR processing.

[00286] Turning now to the 4 x 4 imager array illustrated in FIG. 15B. The 4 x 4 imager array 1510 illustrated in FIG. 15B includes 16 imagers, of which 8 are dedicated to imaging the luma spectral band of the scene. In many embodiments, a requirement of filter placement symmetry is imposed. Therefore, only two categories of field of view are included in the imager array. Category 2 (G2) imagers have 2/3 the maximum field of view, and category 3 imagers (G3) have the maximum field of view. The following table compares the resolution of the array camera shown in FIG. 15B, where each imager is a SVGA camera, at each focal length with the resolution of a conventional 8 Megapixel fixed focal length camera employing a Bayer filter and using digital zoom to provide an image of an equivalent field of view.

Category	fov as % of fov_{max}	# of VGA Imagers	# of Luma Pixels in Legacy Cam within fov	Resolution Recovery Strength per Linear Dim.
2	0.71	4	1,920,000	4.00
3	1	4	4,000,000	8.33

Table 2: Three Categories of Multi Field of View VGA Cameras for a 4 x 4 array configuration

[00287] Category 2 cameras recover resolution adequately in the central 66% of the field of view, while category 3 imagers fall short of the required resolution recovery target, because of a lack of a sufficient number of imagers. However, the legacy camera also degrades in resolution at field heights above 70% and therefore the perceived effect of the reduction in resolution of the array camera is substantially reduced.

[00288] In many embodiments, small field of view differences (of the order of 10% or so) are utilized in order to assist the SR processing. An impediment to resolution through SR processing across multiple images in an imager array can be the presence of null spaces in the field of view covered by the different LR images. A null space is a convergence in the sampling diversity of the object space across the different LR imagers in the array such that the sampling diversity is reduced to 0. By varying the magnification of the different lens stacks of each LR imager in the imager array (for example, by 10% or more) the distribution of the null spaces is more randomized and reduces the likelihood that all of the LR imagers will converge on their sampling distribution simultaneously at some point in the object space. The SR processing can accommodate variations in the FOV of each of the LR imagers by filtering the pixels to account for the different magnification when placing the pixels on a high resolution grid using calibration information and appropriate forward imaging models for each of the LR imagers.

[00289] Although specific imager array configurations are illustrated in FIGS. 15A and 15B alternative imager array configurations can be used to construct multi field array cameras that use SR processes to combine the LR images from imagers within the array having the same field of view to produce synthesized images in accordance with embodiments of the invention.

EXTRACTING STEREO PAIRS

[00290] The eye is able to get a sense of depth of an object by focusing on an object from different points of view (the two eye positions). The distance between the eyes is referred to as the baseline. The angle made by the ray from the eye (focusing on the far object) with the base line in combination with the baseline is enough to judge the distance to the object. This is referred to as triangulation.

[00291] Array cameras in accordance with embodiments of the invention provide multiple view of the same scene, all offset from one another by small differences in viewing angels. A stereo pair of high resolution images of a scene can be obtained using an array camera by performing an SR processes similar to any of the SR processes described above using two different LR imagers (or virtual LR imagers) as the reference imager in the initial Pilot Fusion processes. The distance between the two reference imagers forms the baseline of the resulting HR stereo pair.

[00292] A process for generating a HR stereo pair in accordance with an embodiment of the invention is illustrated in FIG. 16. The process 1600 includes capturing (1602) LR images using an array camera in accordance with an embodiment of the invention. A pair of reference LR imagers are selected. Each LR imager can be an LR imager in the array camera or can be a virtual LR imager. An SR process in accordance with an embodiment of the invention is then used to generate (1604) an SR image in the frame of coordinates of each of the LR imagers. The pair of SR images is then output (1606) as a HR stereo pair. Although a specific process is illustrated in FIG. 16 for generating stereo pairs, other processes that utilize the LR images captured by one or more array

cameras can be utilized to generate HR stereo pairs in a manner appropriate to a specific application in accordance with an embodiment of the invention.

[00293] An advantage of utilizing an imager array to generate stereo pairs is the ability to choose multiple combinations of LR imagers when generating the stereo pair. This can be particularly useful when generating multiple stereo pairs each with its own vantage viewpoint within the captured light field. Instances in which multiple stereo pairs can be useful include but are not limited to generating images for driving auto-stereoscopic lenticular displays. Auto-stereoscopic lenticular displays can generate 3D images, without the requirement that a viewer wear special glasses, using multiple stereo pairs. The generation of multiple stereo pairs using an imager array can also be useful in a variety of other applications in accordance with embodiments of the invention.

[00294] Although the present invention has been described in certain specific embodiments, many additional modifications and variations would be apparent to those skilled in the art. It is therefore to be understood that the present invention may be practiced otherwise than specifically described, including various changes in the size, shape and materials, without departing from the scope and spirit of the present invention. Thus, embodiments of the present invention should be considered in all respects as illustrative and not restrictive.

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 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 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 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.