Helix 2.0 image processing

White Paper

Authors

Justin Wanek, M.S., X-Ray Senior Image Quality Systems Engineer
Ping Xue, Ph.D., X-Ray Image Quality Principal Engineer
Germán Vera, M.S., X-Ray Image Quality Engineering Manager

Abstract

With the introduction of the FlashPad™ HD detector and an improved 100 micron pixel size, the Helix™ image processing chain was introduced to take advantage of the additional high frequency information in the image and provide images with exceptional detail. The Helix 2.0 image processing chain builds upon the success of Helix by including three new algorithms which provide local contrast enhancement in chest images, noise reduction with edge and detail preservation, and consistency in brightness and contrast for image display. Helix and Helix 2.0 represent a leap forward in image quality by consistently delivering images with appropriate detail, contrast, and latitude.

Introduction

The transformation of X-ray photons into an image that is suitable for visual interpretation consists of several essential steps, each which factor into the image quality of the final displayed image as shown in Figure 1. These elements are comprised of X-ray generation, acquisition, image processing and image display.

During the X-ray generation process, the exposure techniques determine the number and energy spectrum of X-ray photons that penetrate through the body and are incident on the detector. The X-ray photons that are absorbed by the detector are converted to a digital signal, which is linearly proportional to dose, and corrected for the detector characteristics in the acquisition stage. The X-ray spectrum as well as the characteristics of the detector such as the sensitivity, pixel size, electronic noise, and detective quantum efficiency all contribute to the signal, contrast, noise and resolution of the image output by the detector.

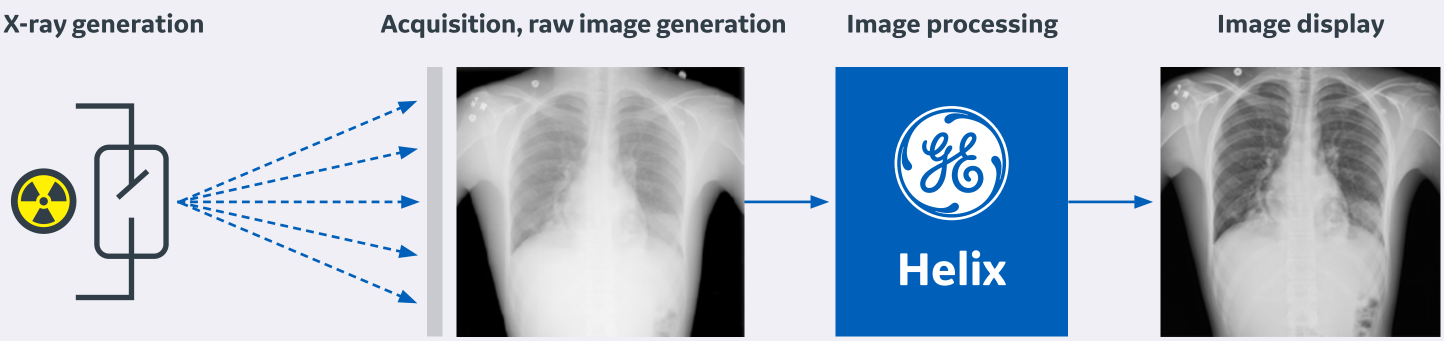

Figure 1: Schematic of the essential elements and processes required for X-ray imaging. Each step is a factor in the quality of the final displayed image.

Introduction *(continued)*

In the image processing stage, the raw image from the detector is enhanced to bring out clinically relevant details and produce an image that is suitable for viewing and interpretation. Enhancement of low contrast features and providing ideal latitude over different regions in the image is important for subjective assessment of image quality and aids in the visual identification of subtle features. Lastly, the image is typically rendered on a monitor, using a display lookup table that mimics the H-D response of film, or printed on film. The computation of the display lookup table is also critical to ensure the image is presented with the appropriate brightness and contrast on the display monitor.

Although each element plays an important role in the image quality of the displayed image, the design of intelligent image processing algorithms that can enhance subtle features consistently and deliver images with suitable contrast and brightness remains a challenging task. This is likely due to the large range of conditions (exposure techniques, field of view, anatomical positioning and thickness) that is encountered in general radiography, as well as the difficultly of designing generalized algorithms that can adapt to these variations. Image processing algorithms must adequately account for these differences to consistently produce images that are subjectively pleasing and allows the radiologist to visualize clinical details.

Helix is a collection of algorithms that work together as part of an image processing chain that aims to provide consistent enhancement and display. At its core are algorithms that provide different functionality such as anatomical segmentation, multi-resolution frequency enhancement, scatter reduction, equalization, noise reduction, and display presentation, but as a whole target to deliver images with appropriate detail and contrast consistently. Since the individual components of the Helix image processing chain can be elegantly controlled, it also provides a great amount of flexibility to obtain different looks to meet the wide range of preferences of its users.

Helix incorporates many improvements over its predecessor. With the introduction of the FlashPad HD detector and the 100 micron pixel size, the image processing chain was modified to provide additional enhancement of fine-detail in images. Edge handling and presentation of high contrast objects was also improved to allow better discrimination and visualization of edges. The determination of display brightness and contrast was also modified to provide robust consistency with changes in patient thickness and dose. Finally, automated software scatter reduction by the Autogrid algorithm provided equivalent contrast improvement to images acquired with a grid.[1]

Helix 2.0 is comprised of the same image processing elements of Helix, while incorporating three additional algorithms. An AI-based brightness/contrast algorithm harnesses the power of deep-learning to minimize inconsistencies in the final presentation compared to traditional methods that derive the display parameters from histogram methods. Helix 2.0 also provides additional capability to improve local contrast in chest images, and reduce noise using a filter that minimizes noise while preserving edges and details. The local contrast enhancement (LCE) algorithm increases local contrast for improved visualization of the lungs, heart, and spine regions. The detail preserving noise reduction (DPNR) algorithm is based on an edge preserving filter to selectively reduce noise in homogenous regions of the image, while maintaining subtle edge features. Ask a GE sales representative for information about which systems are available with Helix and Helix 2.0 image processing software.

Helix image processing chain architecture

Collectively the algorithms that comprise the Helix image processing suite provide enhancement of the raw image from the detector to generate a diagnostic image with improved visibility of fine details. Internally, Helix consists of two parts that provide different functionality and outputs. The first part of the image processing chain is composed of algorithms that provide identification of the edges and region of the X-ray field, segmentation of the anatomical areas in the image, as well as basic image transformations to rotate or flip the image based on the acquisition protocol. There is no change to the pixel data other than any rotation and flip that may be applied. The output is the raw image itself which is saved on the system. In the second portion of the image processing chain, a collection of algorithms is used to modify the image pixel data and enhance clinical features for review. The algorithms provide different types of enhancements to the image such as multifrequency edge and contrast enhancement, tissue equalization, and brightness and contrast enhancement. In addition, there are algorithms for the derivation of a display lookup table that provides suitable contrast and brightness of the image for viewing. The output is the processed image and display lookup table for diagnostic review.

The advantage of this architecture is that the raw images saved on the system can be accessed at any time and reprocessed to generate additional processed images. This provides the user the flexibility to reprocess the same raw image with different looks or types of enhancement if desired.

Helix image processing

Identification of the X-ray field

As part of the Helix image processing chain, the X-ray field is determined by identifying edges of the collimator on the image. This identification is necessary to crop the X-ray field from the processed image for display, and aid in the segmentation of the anatomical regions in the image. An accurate identification of the X-ray field is significant since errors result in additional workflow for technologists to modify and adjust the field of view after image display on the system.

Two different techniques can be used to extract the collimator edges. In the first method, the positioner module as part of an integrated X-ray system provides the collimator vertices which is used to robustly estimate the collimator edges in the image. With the approximate location of the collimator edges, the true edges of the collimator can be identified in the raw image using a template-based correlation technique that is constrained to specific regions of the image. The X-ray field can then be obtained since it is enclosed by the collimator edges.

Because this method requires positioner feedback, it is used in cases when the system has a collimator with this feature such as in some GE's fixed X-ray systems when using wallstand and table receptors.

The second method determines the collimator edges based on image data alone and is completely automated. It is used in cases when no positioner feedback is available, such as in GE mobile X-ray systems and fixed X-ray systems when using digital cassette mode. This method extracts potential candidate edges from the image using an edge operator, then performs a validation process based on the characteristics of the lines to obtain 4 collimator edges.

With the introduction of Helix, significant modifications were made to improve the overall performance of the image-based collimation edge detection algorithm and provide consistent performance with variations in collimation and patient positioning. This is particularly true for pediatric cases, where there is significant variability in patient positioning and X-ray field size as shown in Figure 2.

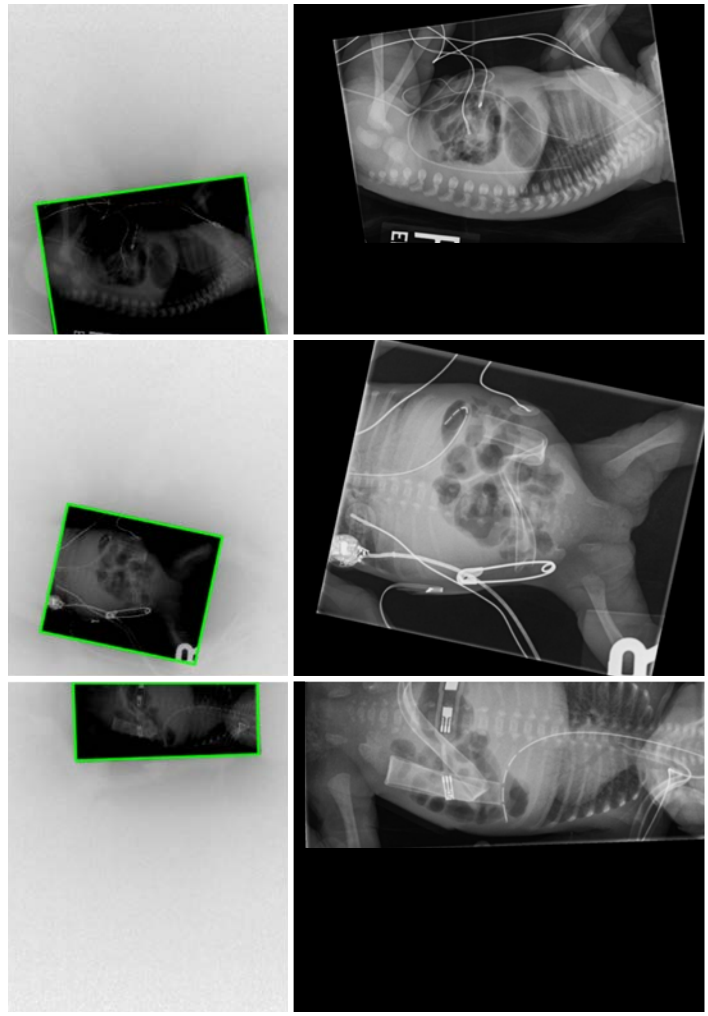

Figure 2: Examples of the image-based collimation edge detection algorithm on pediatric patients.

Detection of anatomical regions

A second important component of the Helix image processing chain is the segmentation of anatomical regions in the raw image. The algorithm identifies regions of unattenuated exposure (or raw radiation) in the image by examining both edge and signal information of the raw image, thereby generating a binary mask corresponding to the raw radiation. The mask is combined with the X-ray field of view to define the anatomical regions in the image.

It is important to identify pixels corresponding to the anatomy to determine ideal image processing parameters for algorithms downstream in the image chain. It is also critical for the determination of correct detector exposure and deviation index values as part of the DEI (detector exposure indicator) feature,[2] which provides end-users visual information on detector exposure in digital X-ray images.

To improve the overall accuracy of the identification of the anatomical regions in the image, Helix incorporates an additional step to refine the anatomy mask to remove miscellaneous objects such as lead markers within the X-ray field of view that are not connected to the anatomical region as shown in Figure 3. This step improves the overall consistency of image processing steps performed downstream and in the display of the processed image, particularly for images with a small field of view, by eliminating regions of the image that may have significantly different attenuation than that of the anatomy area.

Multi-frequency enhancement

Substantial changes were made to the Helix image processing chain to take advantage of the smaller pixel size of the FlashPad HD detector. Helix employs a multi-resolution decomposition technique to provide selective enhancement of features in an image based on their frequency content. In this approach, an input image is progressively filtered and downsampled by a factor of 2 to create a pyramid of images (levels) with decreasing resolution and size as shown in Figure 4. Since the images are progressively blurred, the frequency content contained in the images decreases from the higher to lower levels. At each level in the pyramid, a detail layer is computed by taking the difference of the image with the image generated at the level below after upsampling to match the size. In this manner, detail layers contain local pixel differences between two levels within a specific frequency range.

In the reconstruction phase, the detail layers are recombined through a summation process to generate an output image with selective enhancement of particular frequencies. The decomposition of the image into different frequency bands is advantageous for several reasons.

First, frequencies in the final image can be dictated by enhancing the detail layers differently. For example, if the top detail layers are enhanced more than the bottom layers, the final reconstructed image will have a larger emphasis of high frequencies relative to low frequencies. Second, the contrast within detail layers can be changed to enhance lower contrast features relative to larger ones, thereby increasing their visibility. This allows a significant amount of flexibility to generate images with the appropriate amount of edge detail and contrast.

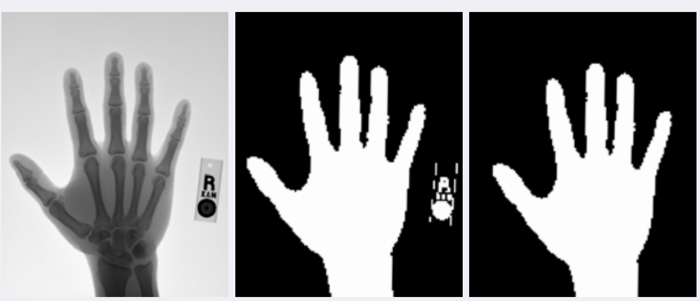

Figure 3: The signal and edge information of the raw image and the field of view mask are used to determine the anatomical region in the image. This region is further refined by removing other non-connected objects such as lead markers.

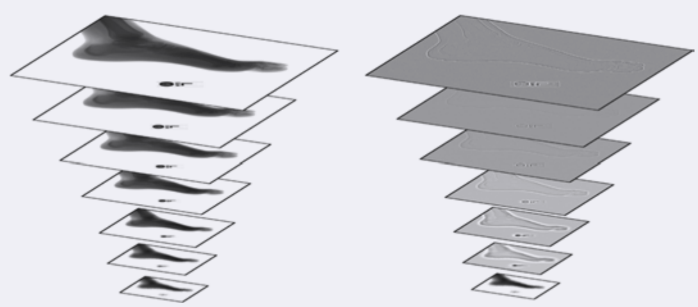

Figure 4: Multi-resolution decomposition technique progressively creates blurred images with decreasing size and resolution (left). Detail layers in the pyramid are generated by taking the difference between these blurred images at different levels and is comprised of a specific frequency band (right).

During the reconstruction phase, care must be taken in the enhancement of the "detail layers" to avoid overshoot/undershoot artifacts that can often occur at the border of high contrast objects. With Helix, the enhancements provided in each detail layer were elegantly controlled to minimize the presence of these types of artifacts and provide improved edge presentation as shown in Figure 5. This is particularly important in orthopedic applications, where overshoots and undershoots around metal implants can obscure the assessment of the placement and lamination of hardware in bone.

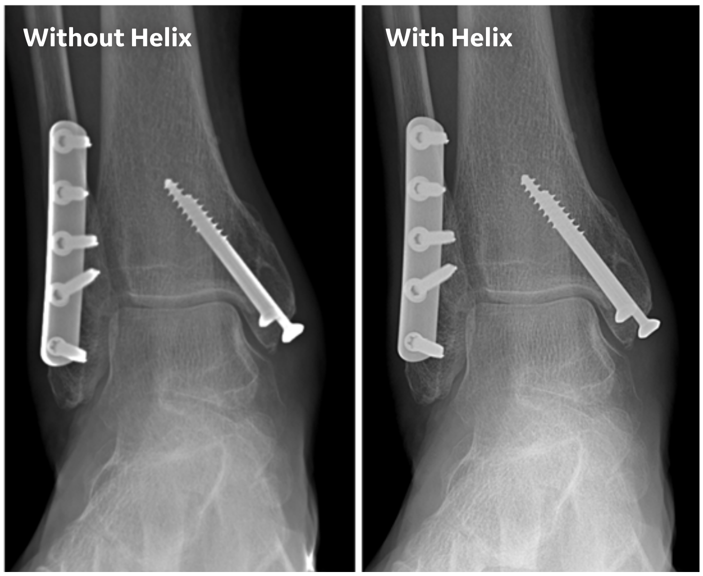

Figure 5: Simulated 200 micron images processed without Helix (left). 100 micron images from the FlashPad HD detector processed with Helix (right). Notice the metal overshoot artifact shown on the left and improved edge handling on the right.

With the introduction of the FlashPad HD detector, the spatial limiting resolution of GE's digital detectors were increased from 2.5 lp/mm to 5 lp/mm. Helix is able to enhance this high frequency information as shown in Figure 6. The consequence of this is better visualization of features such as bone trabecular detail and small hardware, which may aid in the assessment of subtle fractures or placement of hardware.

Helix provides end users with ability to select different levels of edge enhancement and control contrast enhancement within the multi-resolution algorithm. Users can select from 13 different edge levels and 30 tissue contrast enhancement levels per anatomy, view and patient size protocol. This allows a significant amount of flexibility to obtain different types of looks to accommodate the wide range of preferences of its users.

Tissue equalization

In addition to multi-resolution frequency enhancement, equalization of the image plays a critical role in allocating suitable gray scale range over different portions of the image and improving contrast and visibility in dense or thin regions of the anatomy. Helix uses an equalization process, referred to as Tissue Equalization (TE), in the image processing chain to modify the apparent thickness of thick and thin regions in the image. By changing the thickness, the contrast and visibility of these regions are increased, while maintaining suitable contrast in the primary area of interest as shown in Figure 7. To improve the contrast of thick regions, the relative thickness is reduced and brought closer to the normal thickness region or "region of interest." Similarly, to improve the visibility of thin regions, the thickness is increased relative to the "region of interest." This equalization process is applied to the low frequency component of the image to preserve high frequency details.

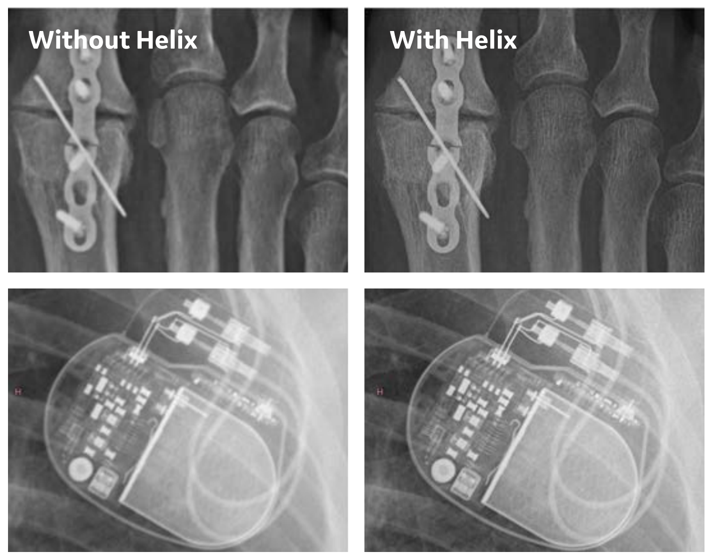

Figure 6: Simulated 200 micron images processed without Helix (left). 100 micron images from the FlashPad HD detector processed with Helix (right). Notice the additional bone trabecular and hardware detail in the images on the right.

The thick or thin regions of the image are defined by the under-penetrated and over-penetrated areas, which are specified as a percentage of the total anatomical area. The equalization amounts for the thin and thick regions are specified by the under-penetrated and over-penetrated strengths, respectively which vary from 0 to 100%. With Helix, users can select both over-penetrated (thin) and under-penetrated (thick) areas and over and under-penetrated equalization strengths. These parameters are configurable for each anatomy, view, and patient size to allow a wide range of customizable looks.

Figure 7: Example of images processed without (left) and with (right) Tissue Equalization (TE) (top). TE provides better visualization of the hardware in the thicker shoulder region as well as more contrast in the thinner soft-tissue regions in the Cervical-Spine images (bottom). The spine is better visualized in the lung and hip region with TE in the Lumbar-Spine image.

Helix image processing *(continued)*

Display

The last components of the image chain are performed to optimize the image for display and visual inspection. To mimic the film optical density response to exposure, the grayscale of the image is inverted and transformed using a gamma relationship. A suitable sigmoidal curve or lookup table for image display is then calculated from parameters derived by histogram analysis of the pixels within the anatomy region. Since this is performed on a per-image basis, each image is individually optimized for display. In addition, two supplementary lookup tables are calculated to generate a soft and hard look, in addition to the normal look provided by default. This flexibility allows a radiologist the option to switch between the normal, soft and hard looks while reviewing the image on PACs based on preference.

Helix was optimized to produce images with consistent brightness and contrast, even when there are large variations in exposure as shown in Figure 8. By determining characteristics of the histogram and accounting for changes in the histogram with dose, a consistent presentation can be achieved even with significant variations in the detector entrance exposure.

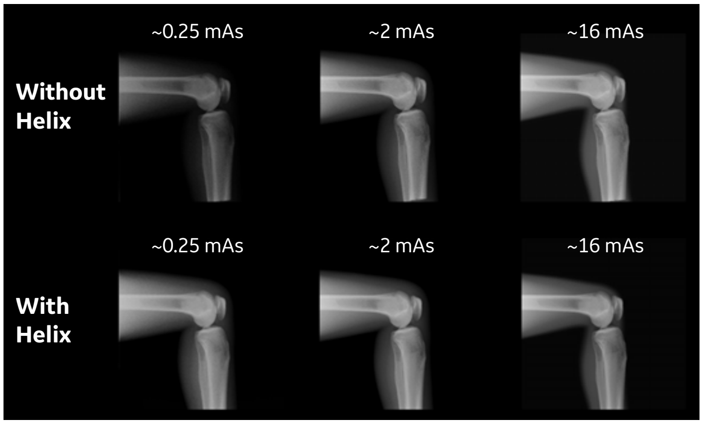

Figure 8: Helix provides consistent brightness and contrast despite large variations in detector entrance dose.

Improvements with Helix 2.0

Three additional algorithms were integrated into the Helix 2.0 image chain. These algorithms expand the capabilities of Helix by providing improved local contrast enhancement in chest images for lung and spine visualization, noise reduction with edge and detail preservation, and consistency in brightness and contrast for image display.

Local contrast enhancement

Chest X-ray imaging is the most frequently performed clinical X-ray exam and is used by clinicians to evaluate a multitude of conditions including breathing difficulties, chest pain, pneumonia, heart failure, and other medical conditions. Due to the frequency of chest X-rays and their use to diagnose many different clinical conditions, there is an increased importance in providing high quality images that offer ideal contrast in the lung, heart and spine regions to improve discrimination of subtle features and pathology. In addition, there can be a wide range in the preferences of radiologists, particularly in the look of chest X-rays. Any algorithm that enhances contrast must also provide the capability and flexibility to deliver a range of looks that meet the diverse preferences of clinicians.

The Local Contrast Enhancement (LCE) algorithm increases local contrast in chest X-ray images for better visualization of the lungs, heart, and spine regions as shown in Figure 9. The algorithm improves local contrast by estimating a low frequency component of the image and removing this signal to boost local contrast. The low frequency signal is estimated by the algorithm which has been optimized in conjunction with multiple configuration parameters that are specific to the patient size, view, and acquisition techniques. The result is an image with improved local contrast, particularly in the lung region, when compared to the original.

LCE is available for single energy frontal and lateral chest exams and the standard image of dual energy chest images. LCE can be customized to provide four levels of enhancement (none, low, medium, high). The strength level can be configured per view and per patient size, thereby offering the flexibility to tailor the strength of the enhancement to the preference of the radiologist. The algorithm supports large pediatric and all adult patient sizes, and is available for all receptors (digital cassette, table, and wallstand). The algorithm can also be used with and without use of a physical grid and autogrid processing.

The LCE feature was reviewed at five different hospitals over a four-month duration. The overall configuration and levels of enhancement provided by the LCE algorithm was optimized based on feedback from radiologists at each site.

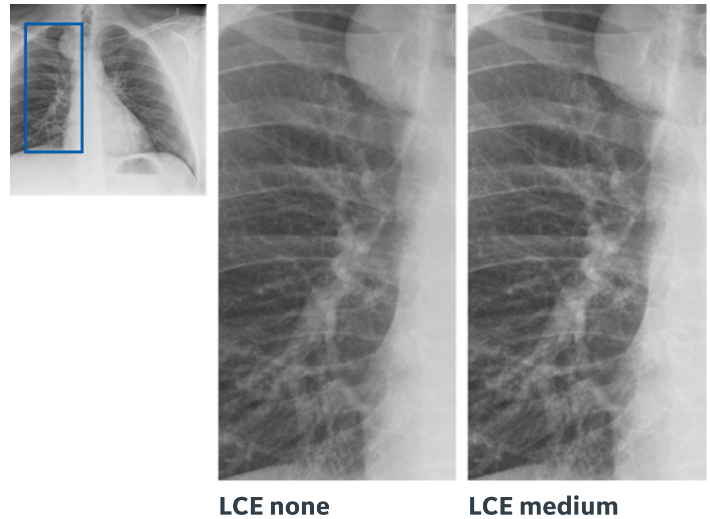

Figure 9: Example of the increased local contrast in lung obtained with LCE (medium strength) compared to no LCE application.

Detail preserving noise reduction

Noise in X-ray systems originates from several sources including quantum photon noise and detector electronic noise. For typical use-cases, the largest contributor of noise is quantum photon noise, which is random and appears as a granular or textured pattern in images. Depending on its magnitude, noise may reduce the visibility of certain features within the image particularly for low-contrast objects. The amount of photon noise in an image depends on the exposure level or number of X-ray photons incident on the detector. As exposure increases, noise in the image increases according to a square root relationship with the number of incident X-ray photons. However, if you consider the amount of noise relative to the signal level (or signal to noise ratio, SNR), the relative noise decreases (SNR increases) with a square-root relationship to the exposure.

Since the number of X-ray photons that reach the detector for a given exposure level depends on many factors including patient thickness, attenuation, and positioning, different regions of the same image can have markedly different levels of noise and SNR. In some regions of the image, SNR may be high and noise may not be visually apparent, while in other regions SNR may be low with visible noise that obscures details. In other applications such as pediatric imaging, limiting dose may be a primary concern, resulting in low SNR throughout the image. In addition, image post-processing algorithms that provide contrast or detail enhancement may change the noise characteristics of the image and often increase the visibility of noise. Care must be taken to account for these changes, while only filtering noise in locations of the image where it is needed.

Detail preserving noise reduction *(continued)*

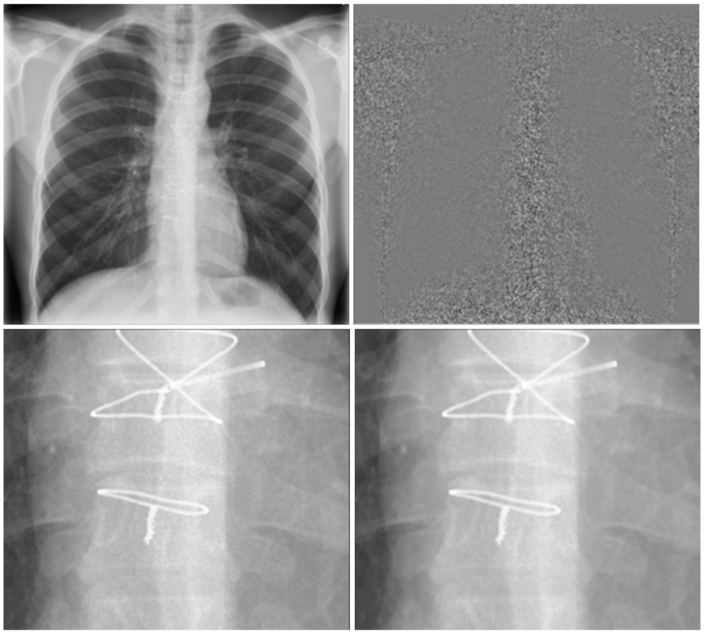

Figure 10: (top) Example of a chest image with high DPNR level and the noise that is removed. The amount of noise reduction spatially varies based on the SNR model. (bottom) Zoomed in portion of the spine displaying the noise removal with DPNR.

Helix employs a detail preserving noise reduction algorithm (DPNR) to adaptively remove noise in an image based on a SNR model as shown in Figure 10. Using the SNR model, locations of the image that have higher levels of noise relative to the signal are increasingly filtered (low SNR), while locations with higher SNR are filtered less or not at all. Since the SNR model is determined on a per-image basis, the noise reduction filter adapts to the noise characteristics of a particular image.

In addition to adaptively filtering locations of the image that have lower SNR, DPNR uses an edge-preserving filter to reduce noise while maintaining edge details as shown in Figure 10. Other filtering techniques that use general smoothing methods can effectively reduce noise, but suffer from loss of fine detail because edges are blurred. DPNR uses a method to minimize intensity variations between the input and filtered image, while penalizing any smoothing that would occur across significant edges. In this manner, DPNR maintains the same local average signal level in the output image compared to the input image, while efficiently reducing noise and preserving subtle edge details.

DPNR is available for single energy and standard dual energy images. It can be customized to provide four different levels of noise reduction (none, low, medium, high), which can be configured per anatomy, view and patient size. The algorithm supports all patient sizes and can be used with all receptors.

AI brightness/contrast

One of the primary challenges of any image processing suite for general radiography is producing consistent display of output images in terms of brightness and contrast. Inconsistency in the presentation of an image may arise from multiple sources including differences in patient body habitus, positioning, field of view, exposure, and algorithm handling as shown in Figure 11. These inconsistencies may require additional workflow for technologists or radiologists if adjustments to image brightness or contrast are needed on the system or on a PACS workstation.

One approach to derive a display curve is to use histogram-based methods. Typically, pixels corresponding to relevant regions in the image are identified and used to derive a histogram. Parameters characterizing the width, center, and potentially other properties of the histogram are computed to determine an appropriate display curve expressed as a lookup table. This lookup table commonly has a sigmoid shape that resembles the H-D response of film. Although histogram methods are straightforward, they can be sensitive to “outlier” pixels from metal or other highly attenuating objects, or errors in the exclusion of raw radiation regions. The lookup table will also be sensitive to differences in the pixel data itself which may change due to variations in patient positioning and FOV. As a result, derivation of the display lookup table and the corresponding contrast and brightness of the displayed image may be inconsistent when comparing images acquired in different patients, or even within the same patient with different positioning.

Helix 2.0 employs a novel AI deep learning-based algorithm to improve image presentation consistency. The AI Based Brightness Contrast Algorithm (AI BC) uses a convolutional neural network to classify an X-ray image based on its brightness and contrast to achieve consistent presentation.[4] The AI model was trained using more than 30,000 unique images from different anatomy/views. The AI algorithm classifies an image into one of multiple pre-trained WW and WC labels depending on the overall brightness and contrast of the image. Based on the output classification, the WW and WC of the image are then adjusted to achieve the default contrast and brightness for a particular anatomy and view, as shown in Figure 12. If a custom configuration is used, the brightness and contrast settings will be further adjusted based on settings of the custom look. This flexibility allows any target brightness and contrast setting to be obtained, but with the AI BC algorithm, these settings are achieved more consistently. The final WC and WW is then used to define a lookup table for display of the image on the system and PACs workstation.

Figure 11: An example of inconsistent image presentation due to the presence of a highly attenuating metal implants in 2 patients.

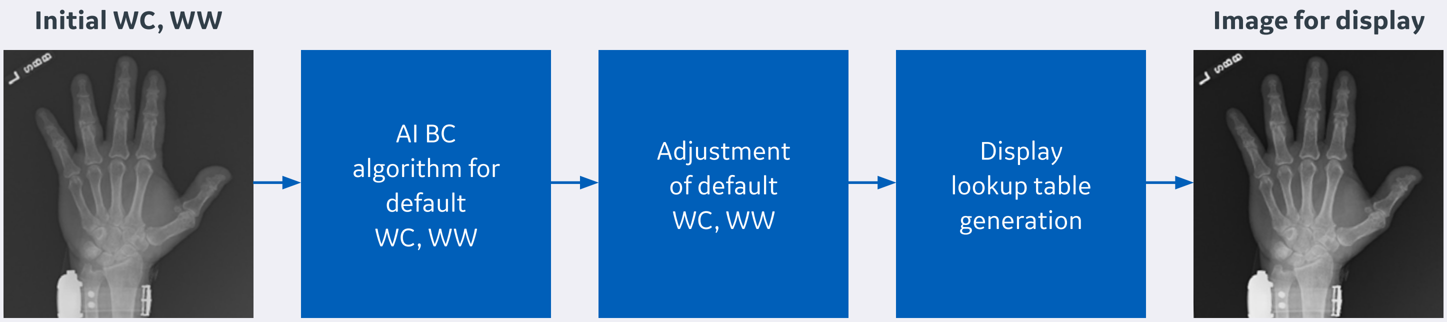

Figure 12: A flow diagram depicting an abridged model of the AI B/C algorithm.

A qualitative assessment of the AI BC algorithm can be performed by visually comparing a variety of images processed with the algorithm as shown in Figure 13. There are several images that display inconsistent brightness and contrast relative to other images for that particular anatomy/view without AI BC. Correction of brightness and contrast and overall improvement in the consistency of the images can be observed for the images with AI BC applied.

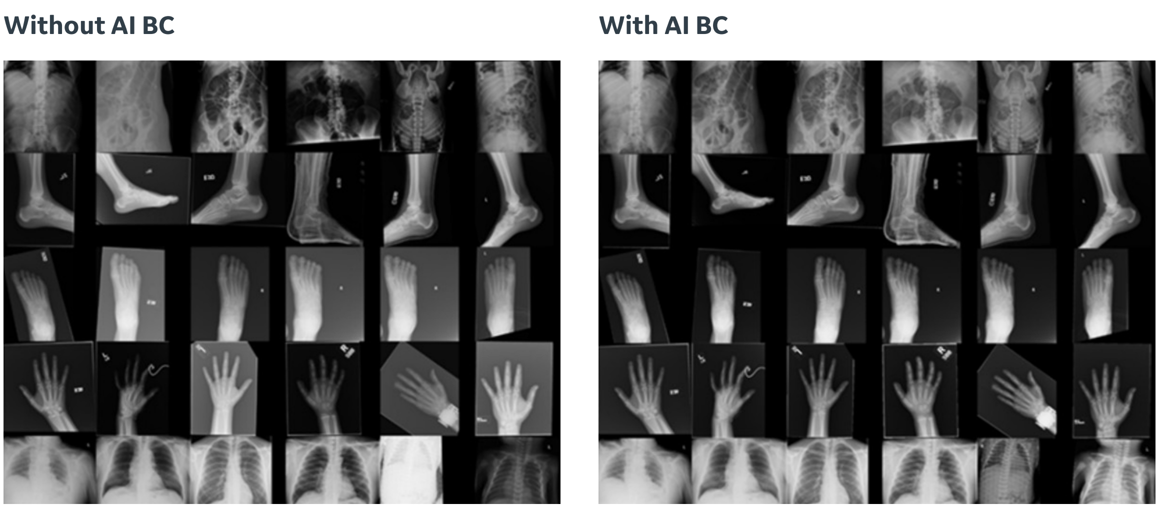

Figure 13: A montage of images comparing the brightness and contrast of images generated using a histogram method (without AI BC) and with application of the AI BC algorithm (with AI BC). Note the improvement in consistency for each anatomy/view using the AI BC algorithm.

Helix image processing *(continued)*

AI brightness/contrast *(continued)*

The performance of the AI BC algorithm can also be evaluated using quantitative metrics. Small ROI regions are selected in common locations such as bone, soft tissue, and raw radiation (air) in multiple images of the same anatomy and view from different patients. For each ROI, a mean signal level is computed and the standard deviation of the mean values among different patients is determined as a measure of consistency as shown in Figure 14. In this example of 12 wrist images, the standard deviation of the bone, soft tissue, and raw radiation was reduced on average by 53% with the AI BC algorithm compared to images with display derived using a histogram method, indicating substantial improvement in the consistency of the displayed brightness levels. In addition, variation in the contrast between the bone and tissue ROIs was reduced by 39%, indicating improved contrast consistency using AI BC. The ability to provide consistent image presentation particularly for challenging cases may reduce workflow required by both technicians and radiologist to adjust images.

In summary, the AI BC algorithm represents a significant advancement because of its capability to produce images with consistent brightness and contrast, especially in challenging cases. The AI BC algorithm is available for the most frequently imaged anatomy/view combinations and for all patient sizes.

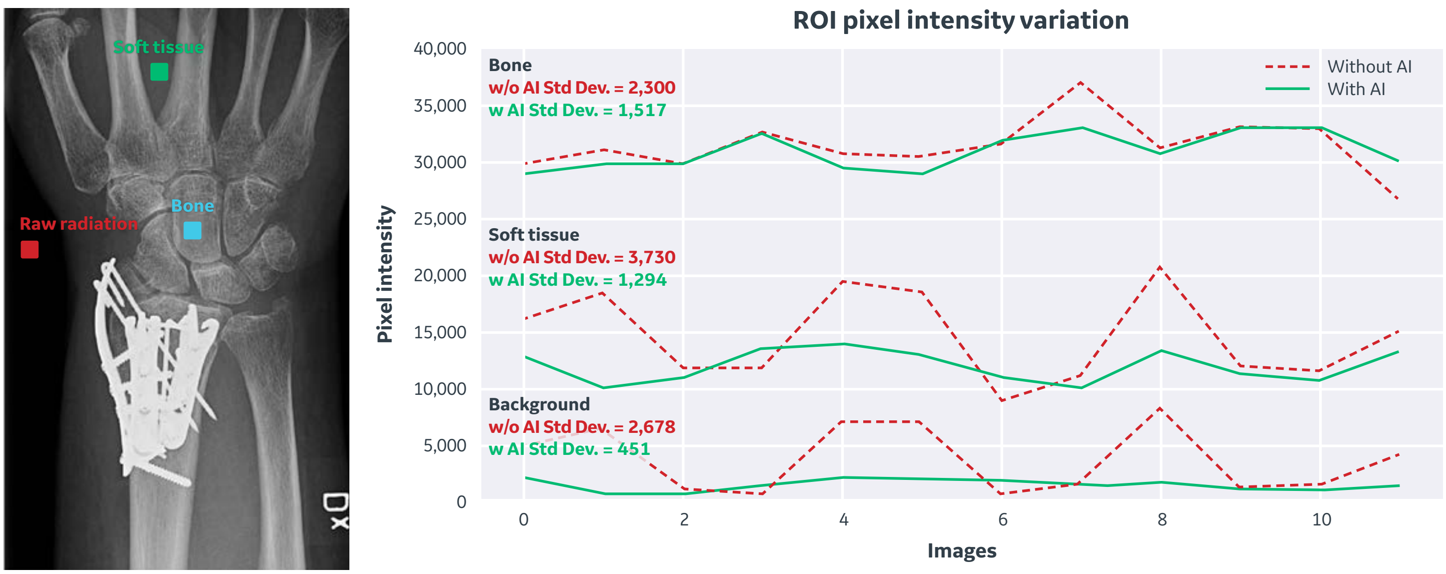

Figure 14: Example of 3 ROIs (Bone, Soft Tissue, Raw Radiation) placed in common locations for 12 wrist PA images acquired in different patients (left). Measurements of the mean signal levels in the 3 ROIs (right) in different patients. Ideally the curves would be a flat line with minimal variation for images of patients with similar bone density and characteristics. The standard deviation (std dev) of the mean signal level among the 12 patients decreases with AI BC, indicating less variation in the display brightness and contrast.

Conclusions

The Helix image processing chain is comprised of individual elements that provide different functionality but work together to deliver optimized contrast and detail enhancement with consistent image display. With the introduction of the FlashPad HD detector and the improved 100 micron pixel size, algorithms were modified to take advantage of the additional high frequency information in the image and provide images with exceptional detail. Improvements were also integrated into several algorithms to improve the overall consistency in the output image quality. Although Helix is designed to work out of the box, many of the algorithms can be individually controlled and configured. Therefore, Helix has a significant amount of flexibility to attain looks that meet the wide range of preferences of its users. Please refer to the operator manual to obtain additional information on how to configure custom looks.

Three new algorithms were introduced as part of Helix 2.0 image processing chain. These algorithms provide additional capability by offering improved local contrast enhancement in chest images, noise reduction with edge and detail preservation, and improved consistency of brightness and contrast for image display. Overall, GE's Helix 2.0 builds upon the improvements of Helix and provides a flexible image processing suite that offers images with exquisite detail, balance and consistency for its customers.

Helix 2.0 image processing examples

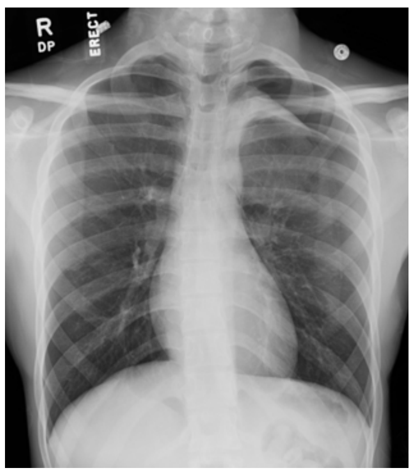

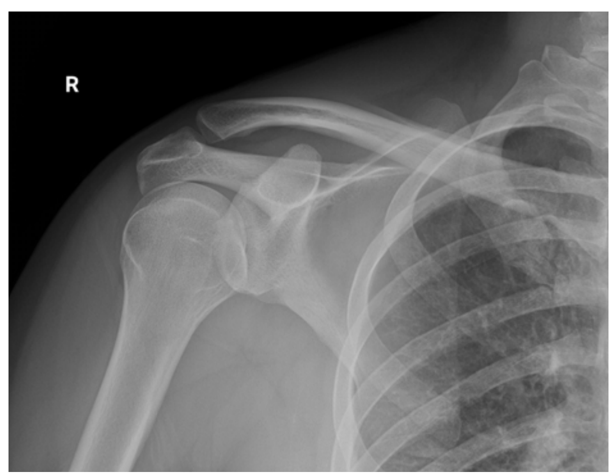

References

1. Mickevicius, K., Xue, P., & Vera Gonzalez, G.V. (2019). Contrast Enhancements and Workflow Improvements Provided by AutoGrid. [White paper]. GE Healthcare.
2. Jeon, S. (2015). Detector Exposure Indicator In GE X-ray Systems. [White paper]. GE Healthcare.
3. Guillard, M., Xue, P., & Gonzalez, G.V. (2017) High resolution for improved visualization on new mobile X-ray systems Optima XR240amx. [White paper]. GE Healthcare.
4. Aboobacker, N. A. M., Gonzalez, G. V., Zhang, F., Wanek, J., Xue, P., Rao, G., & Ye, D. H. (2021, February). Improving presentation consistency of radiographic images using deep learning. In Medical Imaging 2021: Physics of Medical Imaging (Vol. 11595, p. 115952G). International Society for Optics and Photonics.

June 2022
JB20664XX