I-ImaS

CTI:

Improved Adaptive Control by Anticipatory and Reinforcement-Learning options for the I-ImaS Controller Logic

Trieste, 10th – 11th January 2006

Current Progress Overview:

- Specifications for software modules
- Image pre-filtering and restoration
- Controller design stages: OOP, ORP, OCP
- OOP: modes of operation for control logic
- ORP: designing the desired response
- Quality and Dose approximation models
- OCP: RL-based anticipatory control model
- Taylor approximations for predicted gains
- Extensions: exposure and tissue tracking
- Design Specifications and Considerations
- Block diagram for the on-line processing

Current Specifications for the Software Modules:

- Image Analysis: boxes of 512x32 pixels, 12-bit or 14-bit words
- Image Resolution: 32µm SR (uniform), 5625x7500 pixels (max)
- Wedge Filters: 10 independent modules tiled vertically
- Timeframe: about 200ms total for “step-and-shoot” cycle, about 20ms for each wedge filter module (1 processor in h/w)
- Block Size: depends on system setup and total integration time

Image Processing includes:

1. Flat-field correction + Pre-processing filters
2. Textural features extraction (1st-order statistics)
3. Image quality evaluation phase
4. Adaptation step for quality improvement
5. Any cleanup/storage procedures for next cycle

Overview of complete image restoration process:

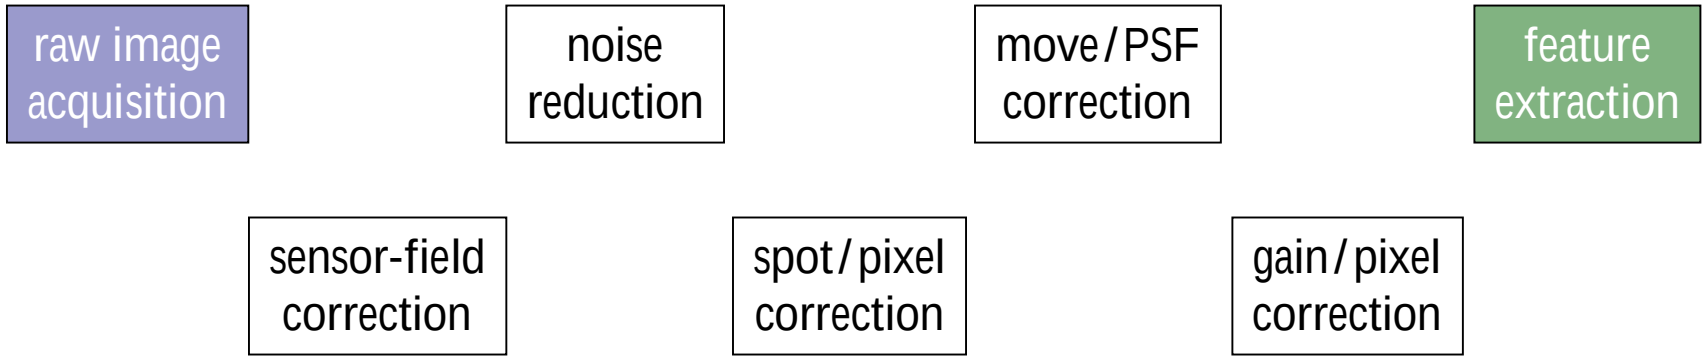

- Different stages embed different levels of complexity and processing time
- Exact ordering and sequence of the stages IS important for optimum results
- Ideal case: feature extraction comes after the last restoration stage
- Real case: embed as many stages as possible within the on-line loop

Design & Implementation Plan:

1. Split the complete sequence into pre- and post-processing modules
2. On-line processing loop includes pre-processing modules, plus feature extraction, plus control logic (x10 times, i.e. all wedge filters, within 200ms)
3. Remaining restoration modules are placed off-line as post-processing

Pre-Filtering Stage:

1. Flat-Field correction: adjust gain of pixels based on a pre-measured spatial template, compensate any constant “bad” pixels and camera geometry distortions.
2. Noise Reduction: apply pre-configured noise removal filters, designed as combined smoothing-sharpening Wiener (FIR) masks of small tap size, for statistical PSF correction.
3. Grayscale Mapping: Normalize and/or Modify histogram profile, according to the requirements of the next stage (textural features extraction).

Notes:

- If PSF is sharp enough, i.e. if relative target/sensor movement is negligible, the noise filter can be a simple 3x3 or 5x5 smoothing mask (96x96 or 160x160 μm averaging kernels)
- Histogram translations are necessary only if any of the final textural features uses modified grayscale, e.g. gamma-corrected.

Design of the desired system response:

1. **Optimal Operational Profile (OOP):** Defines what the system does at different "states" of operation. Currently there are two such states: (1) minimal exposure mode when no tissue, and (2) normal adaptation mode when tissue is detected.

2. **Optimal Response Profile (ORP):** Defines what the control should use as a guideline for optimal behavior. In this case, this should be a general template for image quality versus exposure, provided by one or more human experts.

3. **Optimal Control Profile (OCP):** Defines the adaptation procedure/model that will be used in order to achieve optimized results, according to the ORP. This is the analytical control model and adaptation algorithm.

Note: see deliverable D.9 for further details on operational profiles.

1. Optimal Operational Profile design (OOP):

- In theory, exposure should be minimal (“beam off”) when no tissue is detected.
- In practice, exposure should be as low as possible yet capable of discriminating extra-tissue from intra-tissue background.
- According to previous studies (D.8), Standard Deviation (STD) can be used as a “tissue detector” function and as a triggering mechanism for running in “idle” or “full” mode:

1. Scout scan produces a primal image of the underlying background, STD is calculated upon this primal image.
2. If STD<T (threshold) then there is no tissue under the specific wedge filter and the beam can be cutoff completely during the adaptation phase.
3. If STD>T (threshold) then the detailed imaging begins and the adaptation step is calculated using all the available textural features for optimizing the (adjusted) beam scan.

2. Optimal Response Profile design (ORP):

- The system should be able to improve the resulting quality of the image with respect to diagnostic needs and specifications.
- A general template for the image quality has to be defined externally by one or more human experts.
- Similar evaluations have already been conducted for the UCL mammographic images (WP3).

Overall Design Process:

1. A Quality Template has to be provided by the human experts as a guideline for the I-ImaS controller. This template is to be approximated by a combination of the textural features.
2. A Dose Template (model) has to be defined in order to assess the "cost" to the patient for any given exposure setup.
3. These two factors are to be combined together into one "optimality" value that the I-ImaS controller must maximize in every scanning cycle (adjust current mAs of scout scan).

2.1 Quality Template Approximation:

1. A set of images are produced by the I-ImaS prototype, in the full range of input settings (mAs: wedge filter positions, kVp: fixed)
2. Human experts rate the overall “quality” of the results according to the advancement or deterioration of the diagnostic value.
3. Evaluations from the human experts are averaged and one (combined) quality profile is generated for each image set.
4. Textural features are calculated upon the same image sets.
5. The desired quality profile is approximated by the weighted average of these textural features.
6. The weights of the weighted average are optimized using a LSE or MSE criterion over the desired output value (min.approx.error)
7. The (optimized) weight values are to be used in the final implementation of the control process.

Note: Any weight that is near-zero implies that the corresponding textural feature can be discarded as irrelevant to the quality.

2.2 Dose Template Approximation:

- The I-ImaS system should be able to enhance image quality together with dose reduction to the patient.
- Dose profile can be approximated adequately with analytical models for a given exposure (mAs: wedge filter position)
- Correlation between dose and mAs is more “linear” than between dose and kVp, thus any such approximation will be relatively accurate for the purposes of the I-ImaS controller.
- Beam geometry can be assumed as parallel or conical, according to the resulting deviations of dose estimation. In other words, if the difference is insignificant, simple non-conical dose models can be employed within the on-line processing cycle.
- For breast composition coefficients, i.e. “glandular” versus “fatty” percentage, a typical setup is usually adequate, since the dose estimation is to be used only as an additional optimization factor.

Note: Human experts could be asked to assert dose as a factor in overall quality, but they are not used to do it in practice (D.6).

2.3 Control Template Generation:

1. According to the Quality Template Generation, a weighted average (optimized) of textural features can be used to approximate the human expert's estimation on the diagnostic value of the current image.
2. According to the Dose Template Generation, the "cost" to the patient can be calculated analytically with respect to the current exposure configuration (mAs: wedge filters, kVp: fixed)
3. These two contradictory factors are combined with a proportional weight (Wq, Wd) into one single value, which is the target of the control/optimization process.

Note: Any weight that is near-zero implies that the corresponding textural feature can be discarded as irrelevant to the quality.

Quality Template:

$$C(\vec{f}) = C_0 + \sum_{i=1}^{|F|} w_i f_i \quad , \quad f_i : feature(i)$$

Dose Template:

$$D(v) = E(v) \cdot A_0 \cdot e^{-GT(\theta)} \quad A_0 \cdot e^{-GT_0/\cos\theta} \quad , \quad G \quad g \cdot \mu_{glandular} + (1-g) \cdot \mu_{fatty}$$

$$\theta = \arctan\left(\frac{\sqrt{(x-X_0)^2 + (y-Y_0)^2}}{dist\{source, sensor\}} \right)$$

$$E(v) = E_1 \cdot \log(kVp_0^2 \cdot mAs) + E_0$$
$$= 2E_1 \log(kVp_0) + E_1 \log(v) + E_0$$

Control Target Template:

$$Q_n = Q(v_n, \vec{f}_n) \quad W_C \cdot C(\vec{f}_n) + W_D \cdot D(v_n) \quad , \quad |W_C| + |W_D| = 1$$

A Note on Gain Correction:

- For Flat-Field correction, an analytical model similar to the one used for dose estimation could also be used for Gain:

$$G(v,\theta) = G_1 E(v) \cdot \left(1 - A_0 \cdot e^{-GT(\theta)}\right) + G_0$$

- This model assumes that the detected radiation is directly proportional to the remaining exposure, after tissue absorption.
- However, the use of collimators over the sensors can cause significantly lower detection rate, especially as the angle (θ) increases (i.e. away from the center of the sensor plane).
- Exact measurements have to be conducted with the I-ImaS prototype to establish the exact exponential law for the Gain estimation.

Effect of beam geometry to absorption and gain:

- Absorption: As angle (θ) increases, beam goes through increasingly thicker layers of tissue, thus more radiation is absorbed.
- Gain: Sensors placed far from the center (X_0,Y_0) receive significantly less photons, as a combined result of tissue profile, collimators and distance from target.

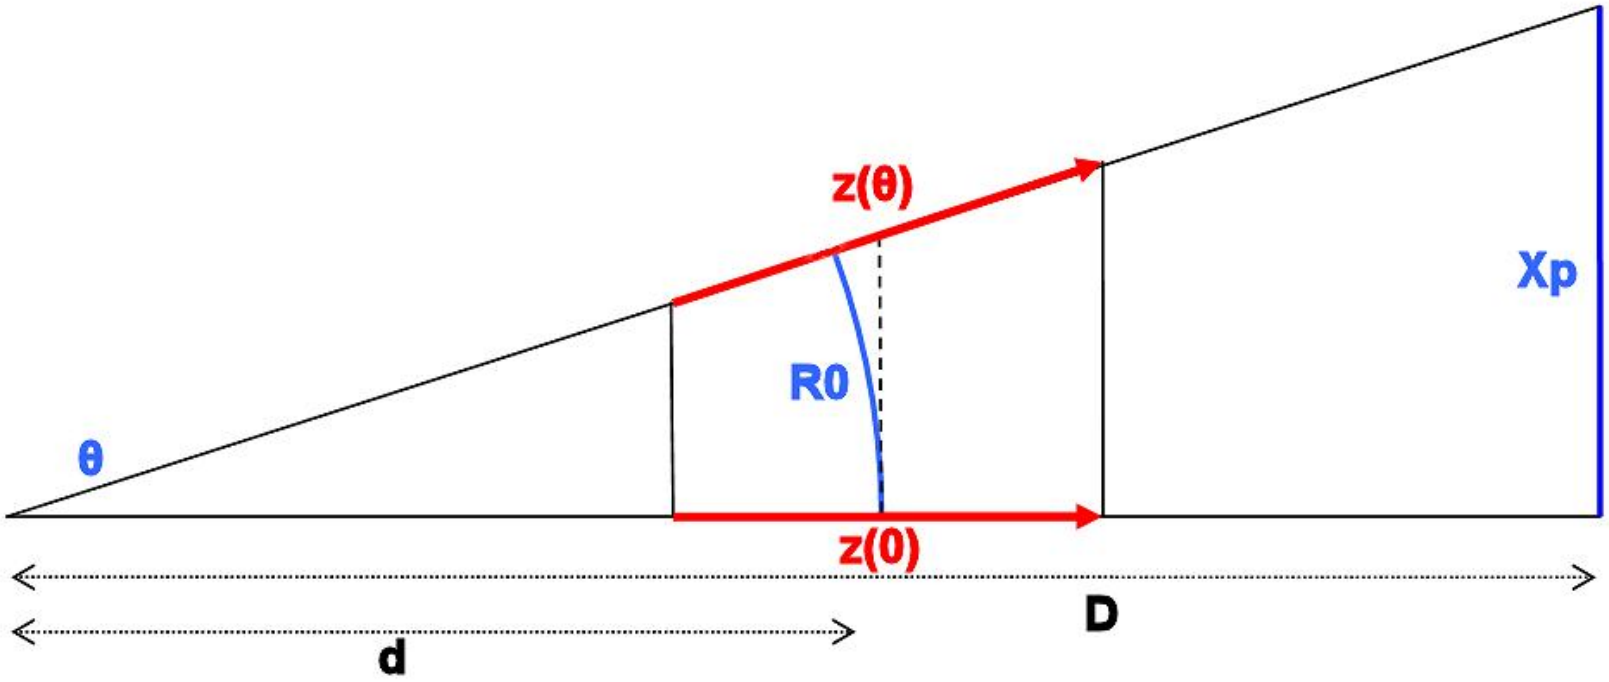

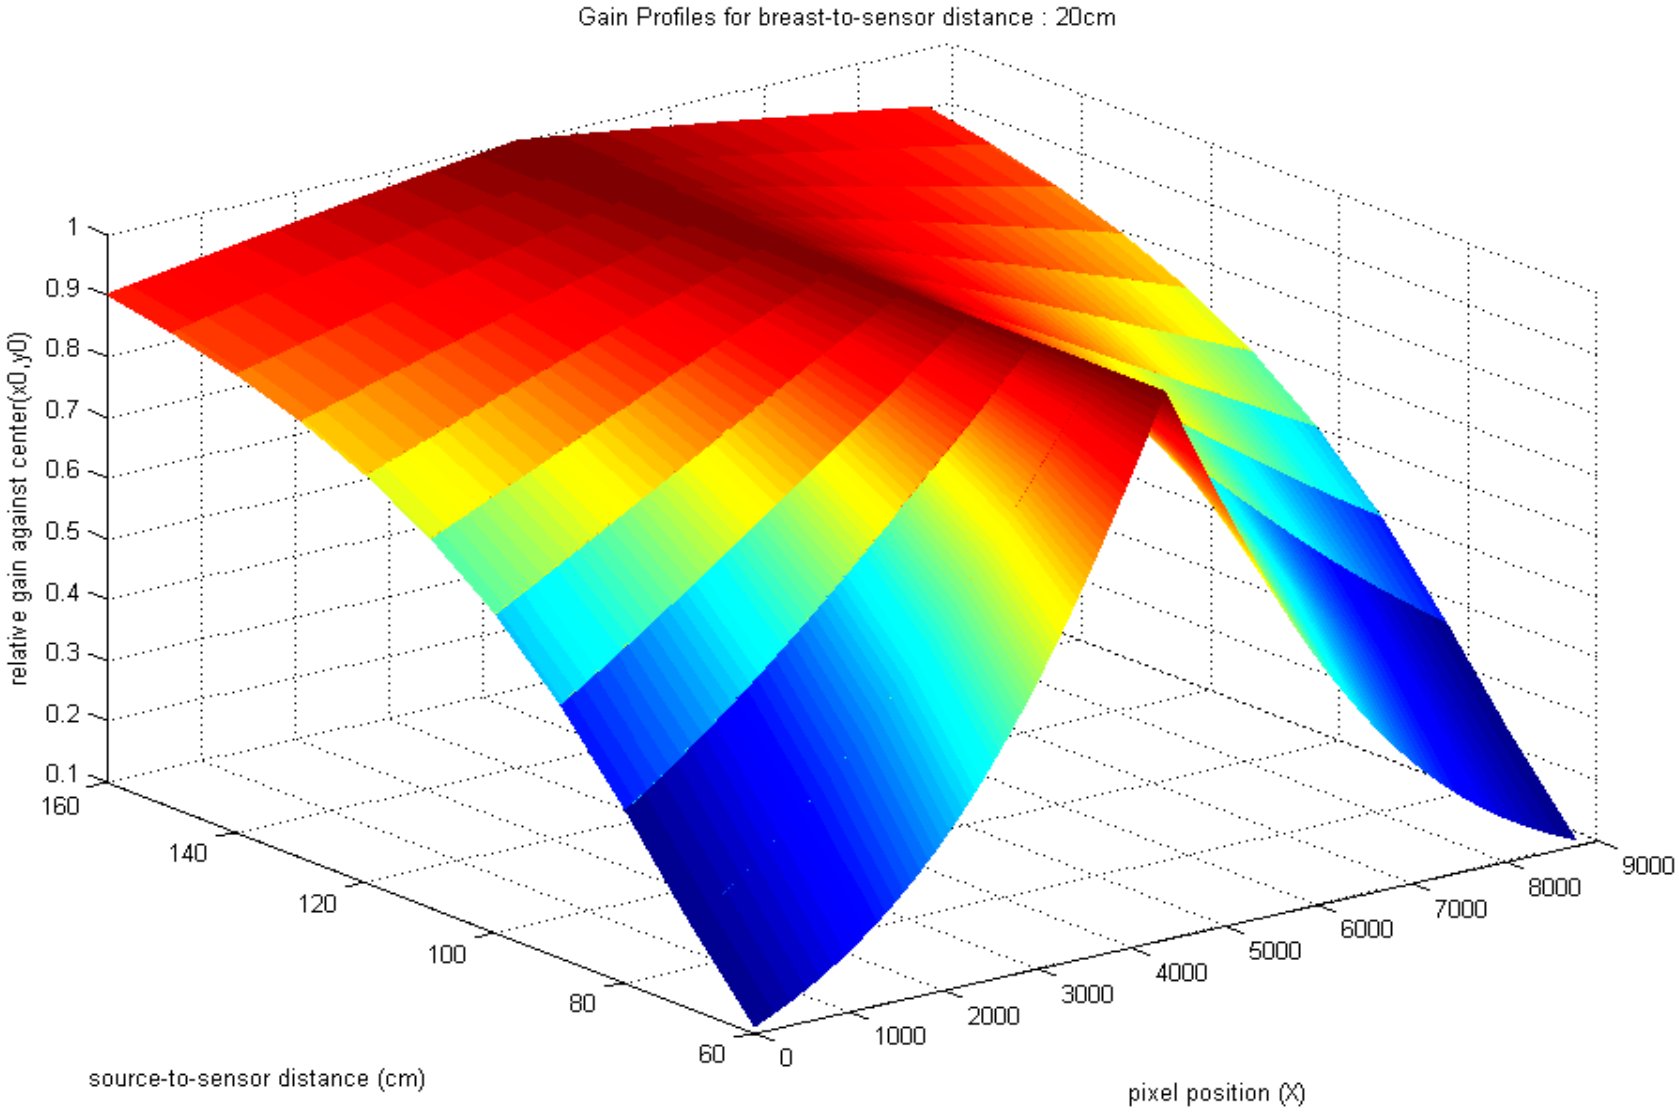

Note: Gain differences have been calculated for excessive breast composition model

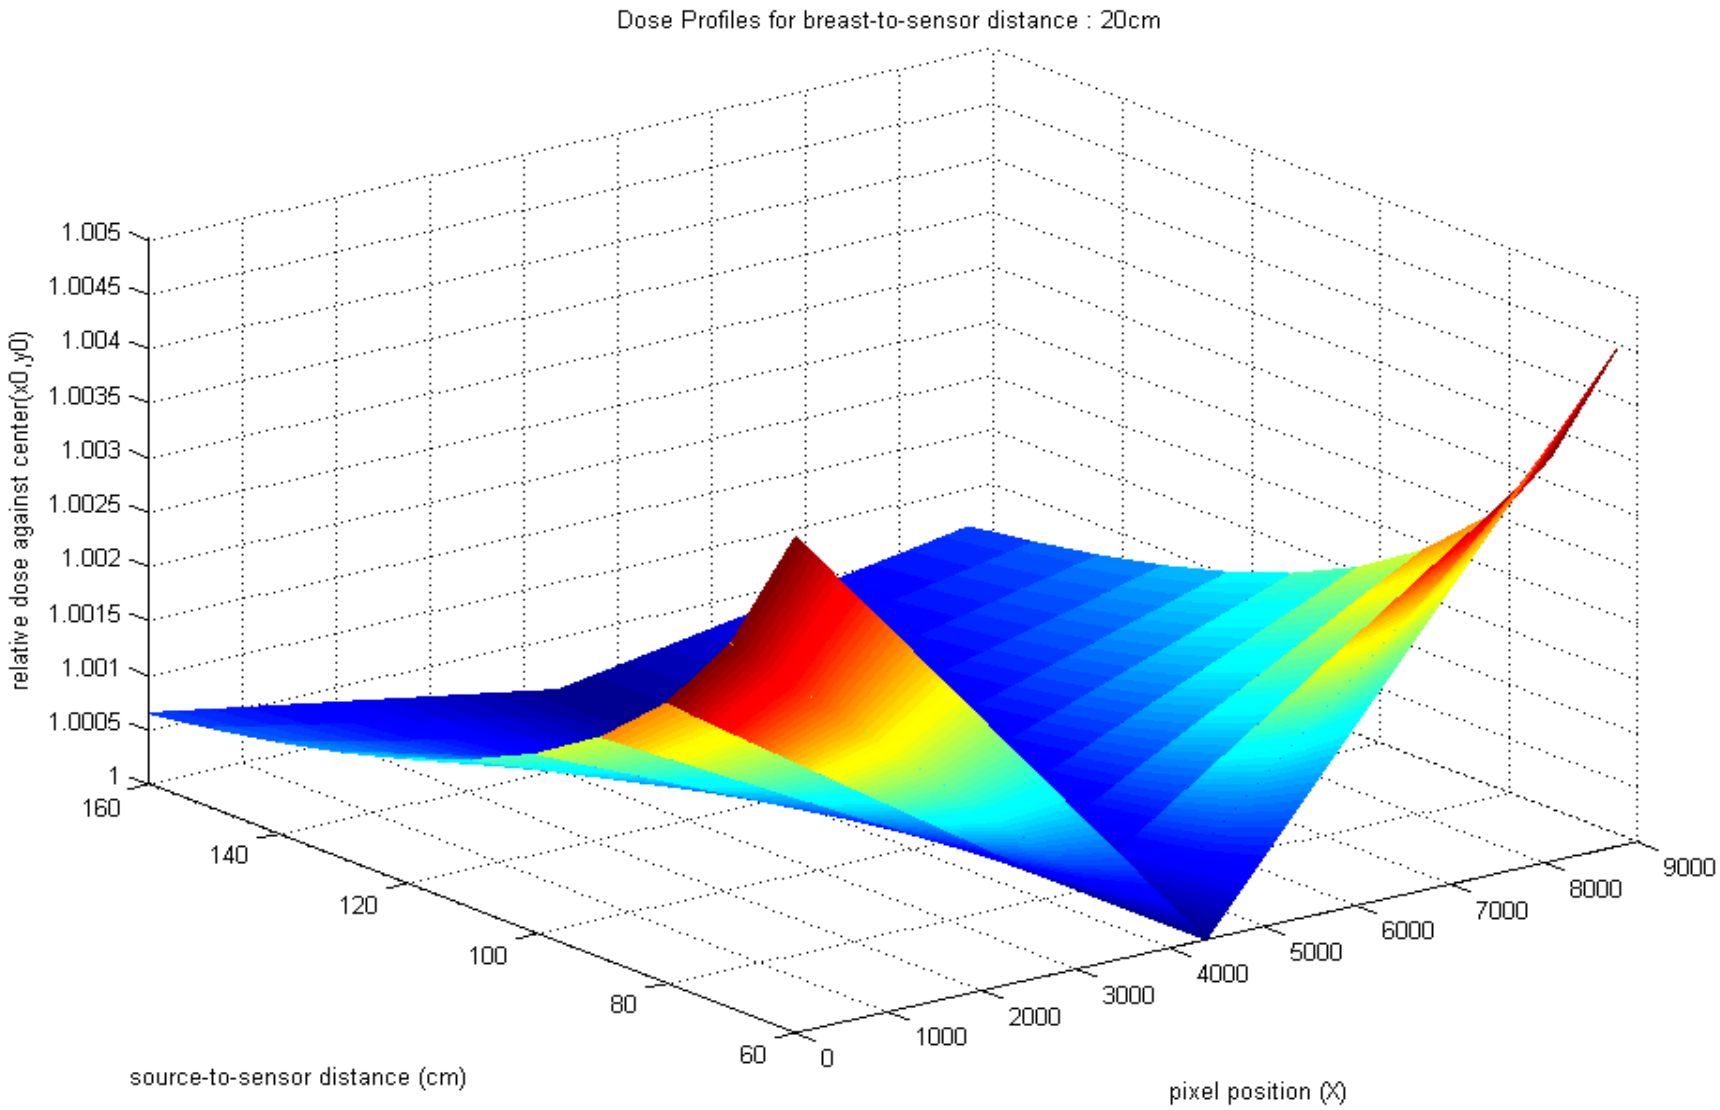

Note: Dose differences have been calculated for excessive breast composition model

3. Optimal Control Profile design (OCP):

- Having established analytical models for all input (mAs, features, dose) and output (combined quality/dose) factors, the OCP has to describe an optimized behavioral model for accomplishing the desired task.
- The OCP model has to be both adaptive and simple in order to implement a successful on-line algorithm for exposure control.
- Essentially, the final output from the control module will be a new setting (position) for each of the wedge filter.

Overall Design Process:

1. Describe a fully adaptive and generic model for the control loop.
2. Establish the control parameters that have to be fine-tuned.
3. Describe a consistent and complete calibration procedure.

Note: The following model is a generalization of the (linear) models introduced in the D.9 report and subsequent presentations.

Anticipatory Control:

- Control models for the I-ImaS system, introduced in D.9 report, employs “reactive” behavior – The system “tries” an exposure setting during the scout scan, assesses the results and then calculates a corrective response (feedback based on true error).
- Instead, if the system could maintain a “consistently improving” behavior via forward predictions and anticipated errors, an optimal stable state can be established much more efficiently.
- Anticipatory control requires “linking” of adaptation cycles, in order: (1) not to “forget” previous actions towards optimality, and (2) to create “traces” of adaptation trajectories and predict future behavior much more effectively.

⇒ Issue (1) refers to non-resetting scout scans, i.e. exploit the adaptation from cycle (n-1) as the seed for current cycle (n).

⇒ Issue (2) refers to the property of “memory” embedded in the control algorithm, i.e. combining results from previous cycles.

Generalization: Reinforcement Learning (RL) model

$$\rho_n(i) = (1-\gamma)\rho_{n-1}(i) + (1-\varepsilon)\left[\pi_n(i) - \alpha_{n-1}\right]$$

$$\rho_n(j) = (1-\gamma)\rho_{n-1}(j) + \left(\varepsilon / S\right)\left[\pi_n(j) - \alpha_{n-1}\right]$$

$$\alpha_n = \alpha_{n-1}\lambda + \pi_n(i)(1-\lambda)$$

- System preserves (S) possible actions that can be taken in each cycle. The action (i) is selected over the others (j) if its "propensity" (ρ) is the current maximum over the entire set (S).
- In each cycle, the real gain (π) is calculated for the previous action and it is compared against the predicted gain (α). The difference between these two is used to update the "propensities" (ρ) of all actions in the set (S).
- The predicted gain (α) is a "reference point" that is also updated in each cycle, in order to keep track of the actual efficiency of the system.

RL-based Control Model:

- The full RL-model contains a total of 5 control parameters:
  - ρ_0 : initial value for all propensities
  - α_0 : initial value for any reference points
  - λ : "persistency" parameter for reference points
  - ε : "experimentation" parameter for actions
  - γ : "forgetfulness" parameter for actions

- Example: For $\lambda=0$ the system uses only the previous gain (i.e. scout scan) calculation as a reference point for the current cycle.
- Example: For $\gamma=1$ and $\varepsilon=0$ the system employs a strict linear "reactive" control model, the one proposed in report D.9.
- Any other value introduces the notion of inter-cycle "memory" for the system, i.e. exploit previous actions and results to "predict" a better response (feedback) for the current adaptation step.
- Assert: "scout" mAs at cycle (n+1) is the "adapted" mAs at (n).

I-ImaS Implementation of RL-based Control:

- The system includes at least 3 distinct actions for each of the K wedge filters, each having M wedge steps:
  - wedgeF(K,i-1) : decrease mAs by -20% (-1/M)
  - wedgeF(K,i) : keep the current mAs
  - wedgeF(K,i-1) : increase mAs by +20% (+1/M)

- The controller has to choose one of these options for each (K) separately, selecting the propensity (ρ) that exhibits the current maximum for each wedge filter.
- If $\lambda > 0$, the controller takes into account roughly the $1/\lambda$ previous actions (cycles) when estimating the current efficiency.
- If $\gamma < 1$, the controller becomes increasingly "stiff" on deviating from the action taken during the previous adaptation cycles.
- If $\varepsilon > 0$, the controller becomes more resilient to current gain fluctuations, i.e. exhibit more "spread" localization of gain errors.

I-ImaS Implementation of RL-based Control:

Parameter Summary:

- Propensities (ρ): “weights” for wedge filter moving actions
- Current gain (π): combined image quality/dose factor (current)
- Reference point (α): trace of previous adaptation results

- If the image quality template is detailed enough, then the reference point (α) can be fixed at Q_{best}, in order to force the controller to always compare the current performance with a global optimum.
- The RL-based model can be roughly associated with a typical ARMA{1/(1-γ), 1/λ} model, as an iterative realization of a linear filtering system.
- The update of the propensities ρ(j) for the non-selected actions requires an estimation of the corresponding gain values π(j), as only the π(i) can be calculated analytically from the results of the scout scan.

Approximation of adjacent gains π(j):

Analytical models have already been established for calculating the dose and the combined “quality” factor (Q), thus the Taylor series approximation may be employed just for the textural features (f):

1. The scout scan at (v) provides the new feature values for cycle (n) and the gain ($\pi_n(i)$) for the current exposure setting can be calculated.
2. The effects of a scout scan at (v-h) and (v+h) are approximated for the resulting textural features, using a Taylor series.
3. Since all textural features and dose calculations are now available for the current (scout) step, the propensities of all possible actions (mAs : decrease/same/increase) can be updated.
4. Using the new (updated) values of all the propensities, the best action for the adaptation step of cycle (n) is chosen to be the one that has the current maximum (ρ).

Taylor series approximation:

Problem Overview:

- The scout scan can give accurate calculations only for the current exposure setting (i) of the scout beam.
- The RL-based model requires an estimation for all other gains, i.e. the corresponding feature values if the scout scan used any other (j) setting.
- Estimation all remaining gains (j) can be realized by a simple approximation via Taylor series:

$$f_n(i) = f_n(v) = \{scout\ scan\ values\}$$

$$\hat{f}_n(j) = \hat{f}_n(v \pm h) = f_n(v) + h^1 \frac{\partial^{(1)} f_n(v)}{1! \cdot \partial v} + \ldots + h^Z \frac{\partial^{(Z)} f_n(v)}{Z! \cdot \partial v^Z}$$

Features approximation for RL-based control:

I-ImaS Overview:

- For a Z-order Taylor series approximation, keep the Z+1 previous values for each calculated textural feature (default: Z=1 or Z=2).
- Calculate the current scout scan values of the textural features.
- Use the "history" of each feature (f) from previous cycles, along with the values from the current scout scan (n):

$$f_n(i) = f_n(v) = \{scout\ scan\ values\}$$

$$\hat{f}_n(j) = \hat{f}_n(v \pm h) \quad f_n(v) + h\frac{\Delta f_n(v)}{\Delta v} + h^2\frac{\Delta^{(2)} f_n(v)}{2\Delta v^2}$$

- For reduced complexity, the Taylor approximation for each of the textural features (f) is a 1-order or 2-order series.

Calculating Taylor approximations on-line:

$$\frac{\Delta f_n(v)}{\Delta v} = \frac{f(v_n) - f(v_{n-1})}{v_n - v_{n-1}}$$

$$\frac{\Delta^{(2)} f_n(v)}{\Delta v^2} = \frac{f(v_{n+1}) - 2f(v_n) + f(v_{n-1})}{(v_{n+1} - v_n)(v_n - v_{n-1})}$$

- Calculation of derivatives requires different values for scout scan exposure (v), in order to create a “trace” of the quality surface.
- If scout scan mAs remains the same two or more cycles, derivatives should be calculated between the current and the most recent non-equal scout scan.
- Alternatively, derivatives can be fixed to their previous value and updated only when possible, i.e. at new mAs settings.

Note: For (S) possible actions and (Z)-order Taylor approximation, the controller keeps track of its performance over a (Z+1)xS grid.

Extension: mAs + Tissue tracking (2-D model)

- In practice, the image quality surface changes over two dimensions: the mAs AND the scanning position, i.e. tissue.
- In order to “guide” the adaptation step more accurately, a 2-D model is required for the Taylor approximation phase.
- The 2-D approximation requires tracking the controller’s performance over a “grid” that models the actual image quality surface, i.e. the different mAs scout scan settings AND the underlying tissue.
- As each tissue sample is scanned only with one mAs setting, all the other reference points on the “trace” grid are approximated.
- 2-D Taylor series requires the calculation of many more (partial) derivatives, using only a limited number of reference points.
- Using tissue “trace” together with mAs “trace”, results from the current scout scan can give a very accurate estimation of the expected textural feature values for mAs settings other than that of the scout scan, hence the results of any “adapted” beam.

Extension: mAs + Tissue tracking (2-D model)

$$f_n(i) = f(x,v) = \{\textit{scout scan values at column } x\}$$

$$\hat{f}_n(j) = \hat{f}(x, v \pm h) \quad f(x,v) \neq h\frac{\Delta_v f(x,v)}{\Delta v} + h^2\frac{\Delta_v^2 f(x,v)}{2\Delta v^2}$$

$$\hat{f}_{n+1}(i) = \hat{f}(x+1, v) \quad f(x,v) + 1\frac{\Delta_x f(x,v)}{(1)} + 1^2\frac{\Delta_x^2 f(x,v)}{2(1)^2}$$

$$\hat{f}_{n+1}(j) = \hat{f}(x+1, v \pm h) = f(x,v) + 1\frac{\Delta_x f(x,v)}{(1)} + h\frac{\Delta_v f(x,v)}{\Delta v} +$$

$$+ \frac{1}{2}\left(1^2\frac{\Delta_{xx}^2 f(x,v)}{(1)^2} + 1h\frac{\Delta_{xv}^2 f(x,v)}{(1)\Delta v} + h^2\frac{\Delta_{vv}^2 f(x,v)}{\Delta v^2}\right)$$

Note: As in the 1-D case, many derivatives may require subjective updates, using only non-equal reference points from the past.

Overview of RL-based Control Procedure:

1. Keep the (Z+1) most recent vectors of the textural feature values, along with the corresponding exposure setting and calculated gain value (optimality/efficiency).
2. Use the adaptation setting of the previous cycle (n-1) to initialize the scout scan setting for the current cycle (n).
3. Execute the scout scan and acquire the resulting textural features vector C(f), the estimated dose D(v) and the combined quality value Q(v,f).
4. Use the Taylor series approximation for estimating the corresponding results as if the scout scan was executed at any setting (j : v±h) other than the one actually used (i : v).
5. Use the RL rules to update the propensities (ρ) of each action possible in (S) and any reference point (α).
6. Use the updated propensities (ρ) to select the action that maximizes the likelihood of an “optimal adjustment”, i.e. a new setting for mAs (wedge filter position) that improves the current status.

Summary of Adaptive Controller Design:

- **Optimal Operational Profile (OOP):** Use different modes of operation, namely for areas were tissue is or is not detected via specific textural features (e.g. SD).

- **Optimal Response Profile (ORP):** Use specific quality templates, asserted by human experts, for quantifying the notion of desired/optimal response by the controller.

- **Optimal Control Profile (OCP):** Implement a flexible RL-based adaptation model that exploits previous control actions (history) and the current scout scan, in order to produce a good estimation of the system's performance for every possible "adaptation" action (mAs : decrease/same/increase).

Summary of Adaptive Controller Design:

- Relation between "propensities" and scout scans: Since the textural features are calculated only for the scout scans, the "gain" from the last adaptation action (n-1) can be evaluated only after the next scout scan at cycle (n). Thus, the RL propensities are updated before choosing the maximum, in order to complete the previous RL-model cycle (n-1) before the current action (n).

- Using more than 3 actions: If the timeframe permits, the Taylor series approximation can be extended to ±2h points, i.e. predicting the results on quality for a two-step wedge filter adjustment (mAs : decr.x2/decr.x1/same/incr.x1/incr.x2).

- Calibrating the RL controller: Since the RL model is much more flexible than the strict linear, calibration will include only the determination of optimal values for its five parameters.

Design Specifications and Considerations:

- Sampling Box Size (SBS): Depends on total image size, sensor tiles orientation, number of wedge filters and integration time. For (X)x(Y) = 180x240mm = 5625x7500 pixels, Y-axis scanning direction, X-axis segmentation for wedge filter zones, 200 msec step-and-shoot cycle time and 3 sec total integration time, the sampling box size is:

  SBS = (5625/10 x 7500/(3000/200)) = 563x500 pixels

- Relation to Textural Features: Preliminary results are based on 50x50 SBS, but quality of 1st-order statistics improves as SBS increases. Hence, true SBS can be much larger, as long as calculation time for kernels is within the available timeframe.

- Relation to Tissue Coverage: For 32µm Spatial Resolution, a SBS of 563x500 pixels covers a tissue area of 18x16 mm, i.e. only 10x15 tiles of "adapted" exposure for complete image.

Design Specifications and Considerations:

- Approximations for Textural Features: For large sampling box sizes, 2nd-order Taylor series approximations become more important. However, they may require very “distant” reference points for calculating 2nd-order derivatives.

- Timeframe limitations: Strict time limits may require only partial sampling within the SBS, as well as 1st-order only Taylor series approximations, since the step-and-shoot cycle period must be divided equally between the 10 wedge filters (in case of a single DSP unit in the IDAQ module).

- Tissue and Exposure Tracking: Although large SBS increases the “smoothness” of 1st-order textural features, OCP tracking performance decreases due to low resolution of quality surface.

Stability and Robustness considerations:

- Since the adaptation procedure is iterative, stability is an essential issue to be considered here.
- In order to avoid excessive increase of the propensity of any one action, the “forgetfulness” and “experimentation” parameters must be set to values such that: $(1-\gamma)<<(1-\varepsilon) \Rightarrow \gamma>>\varepsilon$
- In order to enhance system’s robustness over any possible tissue profile, the locality of the control and the resiliency to random fluctuations (e.g. due to noise) must be carefully reflected to the exact value of the “persistency” parameter (λ) for updating the reference points (σ).
- In general RL-based models, there is also the option for “World Resetting”, i.e. re-initialization of the whole system when large deviation is detected between expected (σ) and real gains (π).
- In practice, employing only a short-term memory for the system ($\gamma>>\varepsilon$, $\lambda \rightarrow 0$) should be adequate for making the system both stable and robust.

Block diagram for the on-line processing

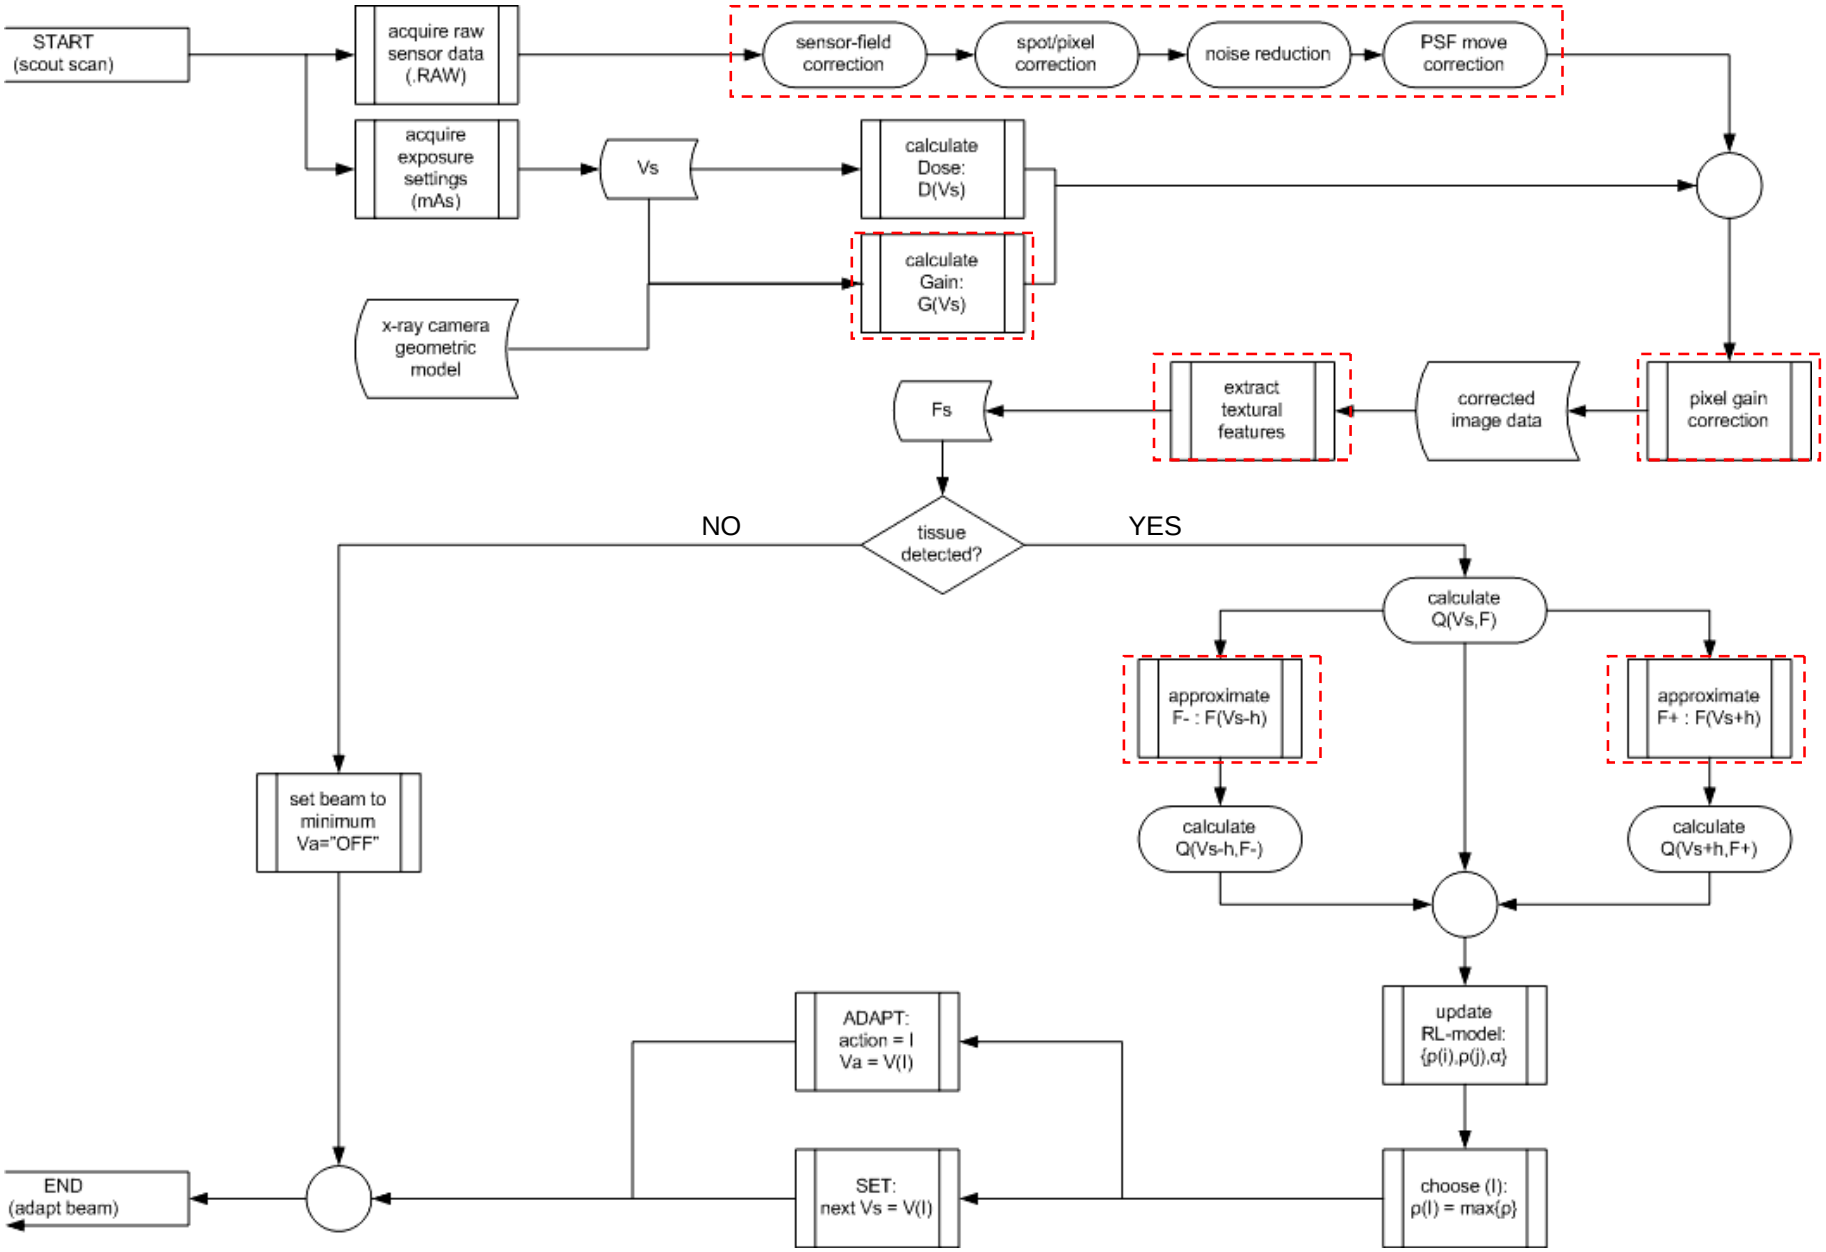

Future Work (CTI):

Next major workpackage involvements in:

1. WP8 – “X-ray camera design and manufacture” (Tasks: 8.2, 8.3), started in Dec/05.

2. WP9 – “System display system and camera control” (Tasks: 9.2, 9.3), starting in Jan/06.

3. Additional work: consultation on implementing the image processing, on-line control and efficient code for filtering modules.

4. As soon as a prototype of the complete acquisition system is ready, test runs are necessary for the verification, calibration & optimization of all the modules involved in imaging and control.

Suggestive References:

[23] *I-ImaS, Workpackage 3 – Deliverable D.8*, “Translating information signatures to a sequence of well-defined processing functions”, Feb.2005

[24] *I-ImaS, Workpackage 3*, “Update on current progress and report for deliverable D.8”, CTI presentation for 3rd I-ImaS meeting, London, 12-13 Oct 2004

[25] I-ImaS, Workpackage 3 – *Deliverable D.9*, “Different approaches to providing intelligence to the sensor/imaging system”, Mar.2005

[26] I-ImaS Workpackage 3, “Update on current progress and deliverable report D.8”, CTI presentation for 4th I-ImaS meeting, Oslo, 14-15 Feb 2005

[27] “D.9 – Considerations from Trieste”, WP3 communication, May-Aug 2005.

[28] I-ImaS, Workpackage 6 – *Deliverable D.17*, “Report on trade-offs for possible sensor/ASIC architectures for the chosen application”, Apr.2005

[29] I-ImaS, Workpackage 5 – *Deliverable D.14*, “Top-level system designs”, Mar.2005

[32] Tom M. Mitchell, “Machine Learning”, McGraw-Hill, 1997.

[33] Murray R. Spiegel, John Liu, “Mathematical Handbook of Formulas and Tables”, 2nd/Ed., McGraw-Hill, 1999.