

Optimization of Acquisition Geometry for Intra-operative Tomographic Imaging

Jakob Vogel¹, Tobias Reichl¹, José Gardiazabal¹, Nassir Navab¹, and Tobias Lasser^{1,2}

¹ Chair for Computer Aided Medical Procedures (CAMP), Technische Universität München, Germany

² Institute for Biomathematics and Biometry, HelmholtzZentrum München, Germany

Abstract. Acquisition geometries for tomographic reconstruction are usually densely sampled in order to keep the underlying linear system used in iterative reconstruction as well-posed as possible. While this objective is easily enforced in imaging systems with gantries, this issue is more critical for intra-operative setups using freehand-guided data sensing. This paper investigates an incremental method to monitor the numerical condition of the system based on the singular value decomposition of the system matrix, and presents an approach to find optimal detector positions via a randomized optimization scheme. The feasibility of this approach is demonstrated using simulations of an intra-operative functional imaging setup and actual robot-controlled phantom experiments.

1 Introduction

Three-dimensional imaging modalities such as X-ray CT, PET or SPECT have turned out to be indispensable tools for diagnosis in modern medicine. Their size requirements due to a big imaging gantry, however, prevents wide-spread use also in intra-operative therapy. In the operating room, the state of the art of imaging mostly encompasses the use of one- and two-dimensional modalities like gamma probes or ultrasound, possibly combined with previously acquired volumetric images.

A possible solution to provide intra-operative volumetric imaging are space-efficient acquisition setups using small detectors guided by humans (e.g. freehand ultrasound or freehand SPECT [8]) or robots (e.g. laparoscopic ultrasound with the Da Vinci robot [6] or C-arm X-ray imaging [2]), and using these measurements for tomographic reconstruction. In this case, it is essential to quickly generate useful acquisition trajectories such that a minimal number of acquisitions yields the best possible sampling of the region of interest, and thus an optimal reconstruction for the task at hand.

In this paper we will present an incremental optimization approach based on the singular value decomposition of the system matrix to compute optimized acquisition trajectories. The feasibility of this approach is demonstrated using simulations of an intra-operative functional imaging setup and actual robot-controlled phantom experiments.

1.1 Terminology

We consider a series expansion approach [4] for iterative reconstruction. In a volume of interest $\Omega \subset \mathbb{R}^3$ the unknown signal $f : \Omega \rightarrow \mathbb{R}$ is approximated using a set of n basis functions $b_i : \Omega \rightarrow \mathbb{R}$, such that $f \approx \sum_{i=1}^n x_i b_i$ with a corresponding coefficient vector $\mathbf{x} = (x_i) \in \mathbb{R}^n$. \mathbf{x} is informally referred to as the *reconstruction* of f . Popular choices for basis functions are for example voxels, Kaiser–Bessel functions, or wavelets.

The set of *measurements* from the detector is denoted $\mathbf{m} = (m_j) \in \mathbb{R}^m$. In order to relate these sensor readings with the image function f , we define physical *measurement models* \mathcal{M}_j that map f to the corresponding value $\mathcal{M}_j(f) = m_j$. In the case of X-ray CT, for example, the Radon transform is a suitable model.

Assuming the models \mathcal{M}_j are linear, the series expansion approach combines the discretization and the physical model:

$$m_j = \mathcal{M}_j(f) \approx \mathcal{M}_j\left(\sum_i x_i \mathbf{b}_i\right) = \sum_i x_i \underbrace{\mathcal{M}_j(\mathbf{b}_i)}_{=: a_{ji}}.$$

For a specific measurement model \mathcal{M}_j encapsulating a sensor pose, the a_{ji} form a unique row vector $\mathbf{a}_j^T = (a_{ji}) \in \mathbb{R}^n$ describing the coverage of each of the n basis functions when viewed from the respective sensor’s perspective, and we get $m_j = \langle \mathbf{a}_j^T, \mathbf{x} \rangle$ in case of a good reconstruction \mathbf{x} .

Using all m measurements m_j , this leads to a linear system $\mathbf{A}\mathbf{x} = \mathbf{m}$ with the *system matrix* $\mathbf{A} = (a_{ji}) \in \mathbb{R}^{m \times n}$. \mathbf{A} is generally not invertible, and to compute a reconstruction \mathbf{x} we solve a least-squares problem instead (typically with iterative methods):

$$\min_{\mathbf{x}} \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{m}\|^2.$$

1.2 Quality of Acquisition Geometry

In devices using imaging gantries, the acquisition protocol is designed such that the solution of the linear system is as well-posed as possible. For instance in X-ray CT, the X-ray source/detector pair rotates around the region of interest Ω , and measurements are being collected on an arc-shaped trajectory of at least 180° .

In case of freehand- or robot-guided data acquisition, however, the acquisition geometry is very sparsely and irregularly sampled, leading to very ill-posed problems. Fewer available measurements often result in under-determined systems, where a good sampling of the region of interest is mandatory. As communicating accurate instructions — particularly when orientation is involved — to a human operator is a complex and largely unsolved problem, we restrict ourselves to robot-controlled acquisitions in this work.

Ignoring application-specific approaches and focusing on a general setting, some of the approaches to characterize the ill-posedness of a system equation are:

1. *Column sums* [8]. Given m measurements, the linear system effectively contains m dot products between every row of the system matrix and the reconstruction \mathbf{x} . Vice versa, the i 'th column of the system matrix contains all the m contributions of the basis function coefficient x_i to the measurements. A low value of the sum over all of column values is consequently an indicator for under-sampling, and enforcing maximal column sums yields better reconstructions.
2. *Singular spectrum of A* [5]. Over- and under-determined linear systems are solved via least squares methods, and the solution is for example obtained from the normal equation $A^T A \mathbf{x} = A^T \mathbf{m}$. In order for this system to have full rank, and thus a well-defined solution, the eigenvalues of $A^T A$ need to be of sufficiently large magnitude. This spectrum is exactly the set of singular values of A , and optimizing them accordingly during the acquisition improves the numerical condition of the system.
3. *Null-space of A* [9]. If \mathbf{x} is a solution to $A \mathbf{x} = \mathbf{m}$, and if $\tilde{\mathbf{x}}$ is in the null-space of A , i.e. $\tilde{\mathbf{x}}$ is a solution to $A \tilde{\mathbf{x}} = \mathbf{0}$, then $\mathbf{x} + \tilde{\mathbf{x}}$ is also a solution of $A \mathbf{x} = \mathbf{m}$. To gain a unique solution for the inverse problem, the null-space has to be reduced to $\{\mathbf{0}\}$, which is typically not feasible. However, if all null-space vectors have a common sub-region of zero values, that region is uniquely determined by the acquisition geometry.

2 Methods

We generate a trajectory by iteratively selecting the best next sensor perspective out of a set of candidates. Such an approach requires two major components, generation of candidate perspectives, and quality estimation of each single candidate.

Using the notation from above, each sensor location is characterized by a unique row-vector $\mathbf{a}^T \in \mathbb{R}^n$. Describing the current state after k measurements by the system matrix $A \in \mathbb{R}^{k \times n}$ and the measurement vector $\mathbf{m} \in \mathbb{R}^k$, we are interested in finding an additional perspective, such that extending the linear system by the corresponding row \mathbf{a}_{k+1}^T and measurement m_{k+1} yields an equation

$$A' \mathbf{x} = \mathbf{m}' \quad \text{or} \quad \begin{pmatrix} A \\ \mathbf{a}_{k+1}^T \end{pmatrix} \mathbf{x} = \begin{pmatrix} \mathbf{m} \\ m_{k+1} \end{pmatrix}$$

with a better quality estimate $\eta(\mathbf{a}^T)$.

2.1 SVD-based Quality Estimation

In order to specify this aim accurately, we use a quality measure based on the singular value spectrum. Ideally, in comparison to A , the new system matrix A' will exhibit larger singular values. We enforce this objective by selecting new sensor poses that maximize the sum of singular values

$$\eta(\mathbf{a}^T) := \sum_{i=1}^{k+1} \sigma'_i$$

where $\sigma'_1 \geq \dots \geq \sigma'_{k+1} \geq 0$ are the singular values of the extended matrix A' , and \mathbf{a}^T is constrained to be generated by a legal sensor perspective.

2.2 Incremental Computation of the SVD

Computing the singular values of a matrix is a numerically expensive operation, in particular considering the size of a typical system matrix. Due to the incremental nature of the system, it is reasonable to use a known decomposition of A when computing the factors of the extended matrix A' during the step $k \rightarrow k+1$. We focus on underdetermined systems ($k < n$), and use the economy-sized definition $A = UDV_1^T$, where $U \in \text{SO}(k)$ is an orthogonal matrix, $D = \text{diag}(\sigma_1, \dots, \sigma_k) \in \mathbb{R}^{k \times k}$ a diagonal matrix holding the singular values $\sigma_1 \geq \dots \geq \sigma_k \geq 0$, and $V_1 \in \mathbb{R}^{n \times k}$ a matrix with orthonormal columns.

Updating the SVD after adding a single row or column to a matrix is a problem already investigated in the fields of Data Mining, Latent Semantic Analysis, and also Computer Vision. There are several approaches, depending on whether exact values are required or whether approximations suffice. Also, some applications use the dominant singular values only, and omit the smaller ones entirely.

In our case, we are interested in the full spectrum of the squared system matrix $A^T A$, and we use the exact method presented by Gu and Eisenstat [3], including the optimizations proposed by Chetverikov and Axt [1]. Using $A = UDV_1^T$, we can preliminarily decompose A' into

$$A' = \begin{pmatrix} A \\ \mathbf{a}^T \end{pmatrix} = \underbrace{\begin{pmatrix} U & \mathbf{0} \\ \mathbf{0}^T & 1 \end{pmatrix}}_{=: M} \underbrace{\begin{pmatrix} D & \mathbf{0} \\ \mathbf{z}^T & \zeta \end{pmatrix}}_{=: L} \underbrace{\begin{pmatrix} V_1^T \\ \mathbf{v}^T \end{pmatrix}}_{=: N^T} = \begin{pmatrix} UDV_1^T \\ \mathbf{z}^T V_1^T + \zeta \mathbf{v}^T \end{pmatrix},$$

where $\mathbf{z} = V_1^T \mathbf{a} \in \mathbb{R}^k$ is the projection of \mathbf{a} into the subspace defined by the rows of V_1 . The other unknowns, $\zeta \in \mathbb{R}$ and $\mathbf{v} \in \mathbb{R}^n$, can be solved from the equation $\mathbf{w} := \mathbf{a} - V_1 V_1^T \mathbf{a} = \zeta \mathbf{v}$ which is obtained from the last row of the decomposition of A' . If the additional vector \mathbf{a}^T is linearly independent of the rows of A , \mathbf{v} is orthogonal to all columns of V_1 , as required. Computing the SVD of the inner — relatively small — square matrix $L \in \mathbb{R}^{k \times k}$ using standard methods yields the decomposition $L = \tilde{U} \tilde{D} \tilde{V}^T$ with both, $\tilde{U}, \tilde{V} \in \text{SO}(k)$ orthogonal matrices. Using this and the preliminary decomposition $A' = MLN^T$, the economy-sized SVD of $A' = U' D' V_1'^T$ is given by $U' = M \tilde{U}$, $D' = \tilde{D}$, and $V' = N \tilde{V}$.

We have extended this algorithm to support the addition of linearly dependent rows. This situation appears to be uncommon in other settings where real-world measurements are used, but may appear when searching for additional sensor poses in a structured way. In this case, $\|\mathbf{w}\| \approx 0$ and $\zeta \approx 0$, as expected, but \mathbf{v} is usually no longer orthogonal to all columns of V_1 — leading to errors in the following step $k+1 \rightarrow k+2$. We detect this, and reinitialize \mathbf{v} by creating an orthogonal vector by means of applying the Gram-Schmidt orthogonalization procedure to a random initial vector.

Please note that these equations are valid for underdetermined systems only, as we expect such a setting in our application of functional imaging. Equivalent

rules can be developed for all other linear systems in a very similar way, but we omit them here for brevity. The reader is kindly referred to the two original publications [3, 1] for in-depth explanations and analyses.

2.3 Optimization and Trajectory Generation

Given a current system matrix A and a set of possible successor poses $C_{(1)}$, we can now select the best next pose out of that set using the techniques presented above. The remaining problem is that of generating that candidate set.

Given the singular value-based energy measure $\eta(\mathbf{a}^T)$, an optimal candidate row \mathbf{a}^T will be as orthogonal as possible to the existing rows, thus optimizing the coverage of the basis functions b_i . Such a position is typically rather far away from the current pose, and a global search will be required to find it.

Identifying an orthogonal row, however, is a complex operation, but random selection has been shown to be a good replacement with high average success rate [7]. We consequently select candidate positions arbitrarily within a space of possible poses. The latter is used to impose constraints caused by geometric limitations, maximal measurement distances, etc.

The entire optimization procedure will thus start at an initial location given by row-vector $\mathbf{a}_{(0)}^T$ and the corresponding system matrix $A_{(0)} = (\mathbf{a}_{(0)}^T) \in \mathbb{R}^{1 \times n}$ with known SVD. With $C_{(1)}$ denoting the randomly selected candidate set, the second sensor perspective is chosen as

$$\mathbf{a}_{(1)}^T = \arg \max_{\mathbf{a}^T \in C_{(1)}} \eta(\mathbf{a}^T),$$

yielding an extended system matrix $A_{(1)} \in \mathbb{R}^{2 \times n}$ with maximal sum over the singular value spectrum. That process is repeated until the required number of poses has been reached.

Such a random path will obviously contain large hops, and to minimize the acquisition path length and time to scan it is essential to post-process it. We use a two-stage approach to sort the poses into a useful sequence, first partitioning all positions into local clusters, and then reordering each of them individually using an approximative Traveling-Salesman-solver based on the Minimum Spanning Tree heuristic. If steps of large size remain, we insert ‘safe positions’, and recombine the partitions, yielding a smooth acquisition path that can for example be traced by a robot in reasonable time.

3 Experiments and Results

3.1 Experiments

In order to test the optimization procedure, we created a Matlab script computing trajectories. Path planning has been performed on a laptop computer with an Intel Core i7 CPU with 4 processing cores at 2.3 GHz and 8 GB of memory.

The experiments are based on an intra-operative functional imaging setup similar to freehand SPECT [8], using a tracked gamma detector to generate localized reconstruction of a radioactivity distribution. Assuming a region of interest surrounded by a bounding box of $(10 \text{ cm})^2 \times 5 \text{ cm}$, we used two alternative parameterizations of two degrees of freedom, generating locations on a hemisphere around the base plane’s center, then retracting the sensor along that direction to reach the surface of the bounding box. The first parameterization used spherical coordinates, the other directly created unit directions in Cartesian space. Of the six sides of the cuboid, we only measure on the top plane, as well as on two orthogonal side planes, thus imitating the spatial constraints of an intra-operative situation where patient body and operating bed place constraints on accessibility.

Using a rather coarse discretization of that region of $10 \times 10 \times 5$ voxels, we were able to generate a path of 300 positions in about 140 seconds. The coarse discretization is reasonable as the small system matrix shows similar behavior as observed when working with its full-size equivalent. Furthermore, when using a robot to traverse the trajectory, additional data can be recorded during the movements to record considerably more than just 300 measurements.

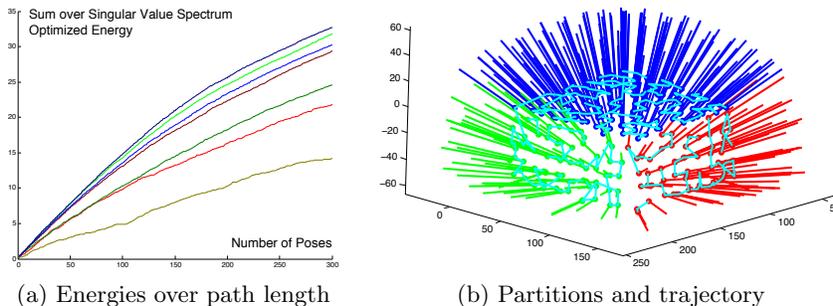


Fig. 1: Evaluation. (1a) Energy curves denote, from bottom up, a human measurement, random sampling with 1 candidate and spherical coordinates (SC), with 1 candidate and Cartesian coordinates (CC), 5 candidates SC, 5 candidates CC, 12 candidates SC, 12 candidates CC. (1b) Sampling of the (invisible) region of interest; acquisition locations on three bounding box planes (red/green/blue = left/front/top), and robot trajectory (cyan).

3.2 Results

Singular-Values-Spectrum over Path Length In a first experiment, we compared the energy $\eta(\mathbf{a}^T)$ at different evolution stages for different sampling strategies. The result is shown in figure (1a). We used candidate sets of size 1, 5, and 12 and the two mentioned parameterizations. In general, more samples yielded better results, and sampling in Cartesian space turned out to be slightly superior —

this parameterization does not show any clustering, as can be observed around the poles when using spherical coordinates.

We also evaluated the evolution of the energy $\eta(\mathbf{a}^T)$ for a trajectory recorded while a human performed a standard protocol. In terms of our energy, this path is the worst, potentially due to the fact that the probe is primarily translated, while the orientation vector remains almost static for a considerable amount of time. Also, the number of measurements per side is not proportional to the surface area.

Simulated Measurements Next, we generated ground truth volumes and simulated measurements, and compared the reconstruction results for several trajectories, generated by humans and the optimization procedure. The reconstructions were considerably underdetermined (300 measurements for 500,000 unknowns), solved by MLEM [4] (20 iterations), and show SPECT-typical displacement errors. Nevertheless, the reconstructions based on trajectories created by the proposed method show better separation between the hot spots, and considerably less ‘activity bleeding’. Examples are shown in figures (2a) and (2b).

Real Robot-Guided Measurements Finally, we fed our trajectories to a robot arm guiding the gamma detector. An example image showing the partitioned sets of probe positions, and the robot trajectory (omitting the intermediate safe positions) is shown in figure (1b). Results are given in figures (2c), (2d).

4 Discussion and Conclusion

We have presented a method to generate an optimized trajectory for tomographic reconstruction in intra-operative settings. This optimality is defined based on the singular value spectrum, and the corresponding measure is computed using fast incremental updates of the SVD. Since this only depends on the system matrix, this approach is applicable for any imaging modality. Even though our experiments used a regular bounding box, more complicated geometries given as a polygon mesh, for instance obtained by laser-scanning the patient, could also be used to generate trajectories.

Several improvements can be made. Runtime can be improved by an efficient implementation as well as by exploiting the ‘broken-arrowhead’ structure of the inner matrix while computing its SVD, thus enabling the computation of longer trajectories. Furthermore, it would be interesting to investigate other energy measures based on the kernel or column sums.

The most interesting point will be to convert this approach to a real-time path planning application considering the kinematics of a robot. While the energy measure is sufficiently fast, a naïve local search for candidates is prone to get stuck at local maxima. A strategic planner will need to consider now what other positions to visit later, while still guaranteeing both full coverage and smooth motion.

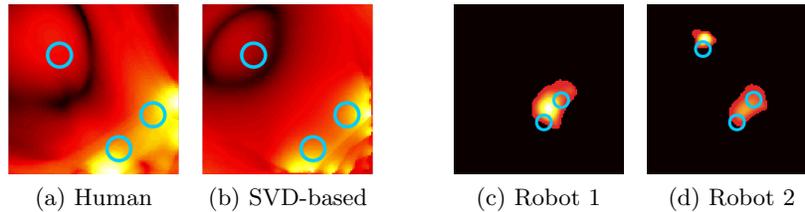


Fig. 2: Reconstruction results, looking top-down, with blue circles denoting ground truth locations of activity seeds. (2a)–(2b) show the logarithms of the per-column standard deviations of the intensities in a *simulated* setting. The path generated by our method yields a better separation and improved circumscription of the hot regions. (2c)–(2d) show the results of two *real* acquisitions performed by a robot following a trajectory created by our method.

Acknowledgements This work was partially funded by DFG SFB 824, DFG Cluster of Excellence MAP and EU FP7 EndoTOFPET-US.

References

1. Chetverikov, D., Axt, A.: Approximation-free Running SVD and its Application to Motion Detection. *Pattern Recognition Letters* 31(9), 891–897 (2010)
2. Ganguly, A., Fiesemann, A., Marks, M., Rosenberg, J., Boese, J., Deuerling-Zheng, Y., Straka, M., Zaharchuck, G., Bammer, R., Fahrig, R.: Cerebral CT Perfusion Using an Interventional C-Arm Imaging System: Cerebral Blood Flow Measurements. *American Journal of Neuroradiology* 32, 1525–1531 (2011)
3. Gu, M., Eisenstat, S.C.: A Stable and Fast Algorithm for Updating the Singular Value Decomposition. Research Report RR-966, Yale University, New Haven (1994)
4. Herman, G.T.: *Fundamentals of Computerized Tomography*. Springer, London, second edn. (2010)
5. Lasser, T., Ntziachristos, V.: Optimization of 360° Projection Fluorescence Molecular Tomography. *Medical Image Analysis* 11(4), 389–399 (2007)
6. Schneider, C., Guerrero, J., Nguan, C., Rohling, R., Salcudean, S.: Intra-operative “Pick-Up” Ultrasound for Robot Assisted Surgery with Vessel Extraction and Registration: A Feasibility Study. In: Taylor, R., Yang, G.Z. (eds.) *Information Processing in Computer-Assisted Interventions*. LNCS, vol. 6689, pp. 122–132. Springer, Berlin (2011)
7. Strohmer, T., Vershynin, R.: A Randomized Solver for Linear Systems with Exponential Convergence. In: Díaz, J., Jansen, K., Rolim, J.D.P., Zwick, U. (eds.) *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, LNCS, vol. 4110, pp. 499–507. Springer, Berlin (2006)
8. Wendler, T., Herrmann, K., Schnelzer, A., Lasser, T., Traub, J., Kutter, O., Ehlerding, A., Scheidhauer, K., Schuster, T., Kiechle, M., Schwaiger, M., Navab, N., Ziegler, S., Buck, A.: First demonstration of 3-D lymphatic mapping in breast cancer using freehand SPECT. *European Journal of Nuclear Medicine and Molecular Imaging* 37(8), 1452–1461 (2010)
9. Zeng, G.L., Gullberg, G.T.: Null-Space Function Estimation for the Three-Dimensional Interior Problem. In: *Proc. of Fully 3D*. pp. 241–245. Potsdam (2011)