

White Paper

ALPHA Technology: Achieving Organ-based Reading in PET/CT and SPECT/CT for Improved Efficiency of Image Interpretation

Sean Zhou, PhD, Siemens Healthcare GmbH

Table of Contents

Introduction: The Quest for Reading Efficiency in PET/CT and SPECT/CT	1
Use Cases in Reading and Quantification	1
Challenge I: Multiple Time-point Study Alignment	1
Challenge II: Efficiency in PERCIST Quantification	2
ALPHA Technology	2
A Visual Recognition System	2
Learning from Examples	2
Reliability through Redundancy	3
1. Collection of a redundant set of local evidences	3
2. Enforcement of the spatial relationships among anatomical structures	4
3. Exploiting redundancy in scale space of images	4
ALPHA Answers	5
Answer for Challenge I: ALPHA Anatomical Registration	5
Answer for Challenge II: Auto-Detection of Reference ROIs	5
Batch Testing Results	5
Conclusion	7
About the Author	7
References	8

Introduction: The Quest for Reading Efficiency in PET/CT and SPECT/CT

Molecular imaging provides unparalleled capabilities for the detection, staging and monitoring of various diseases, including cancer and neurological disorders.

Today, as clinicians are pressed to interpret more and more studies each day, the quest for efficiency in reading as well as quantification is steadily growing.

Software-based automations can assist clinicians in reducing manual, laborious and non-essential work, thus allowing them to focus their time and attention on the most critical clinical questions at hand.

One example of “manual, laborious and non-essential work” is the preparation of multiple time-point studies for comparative reading.

In PET/CT, another example is the manual drawing of reference regions of interest (ROIs) in order to arrive at meaningful quantification when assessing disease progression or therapy response, e.g., by following the PERCIST guidelines.

In SPECT/CT, manual alignment of multiple time-points for assessment of therapy response during the chemo or radiation therapy is done for comparative reading.

Below is a more in-depth examination of today’s pain points and challenges facing clinicians along the directions represented by the

two examples above, and explain how ALPHA, a Siemens proprietary technology, provides answers for them.

Use Cases in Reading and Quantification

In order to track lesions over time and monitor tumor response to therapy, clinicians need to not only *visually* compare the same lesion(s) from two or more time-points, but also *quantitatively* measure the changes. In both scenarios, there are obstacles that slow them down.

Challenge I: Multiple Time-point Study Alignment

Conventional image registration algorithms have existed for many years now, which can help clinicians by automatically matching up studies from multiple time-points. However, these conventional algorithms often fail (see Figure 1) when imaging conditions change from one time-point to the other, for example, with different field-of-views, with patient posture changes (e.g., arms-up versus arms-down), or with different patient table setup (e.g., a thick versus a thin cushion). These variations can confuse a conventional image registration algorithm because their pixel-by-pixel comparison logic may find a “local minimum” by matching shoulder with pelvis, or by matching up the two tables instead of the patient anatomy.

When such failures occur, clinicians have to resort to manual alignment of the image volumes. This is a tedious task that slows them down and, when the reading load is high, may also wear them out mentally.

Figure 1. Conventional image registration methods often fail when two studies were scanned with drastically different fields-of-view (FOVs). Here, the top row shows a study around the neck while the bottom shows a study with extended torso coverage. “Optimal” pixel-based match resulted in matching the neck with the pelvic region. *Siemens-generated data.*




Figure 3. ALPHA in a nut shell.

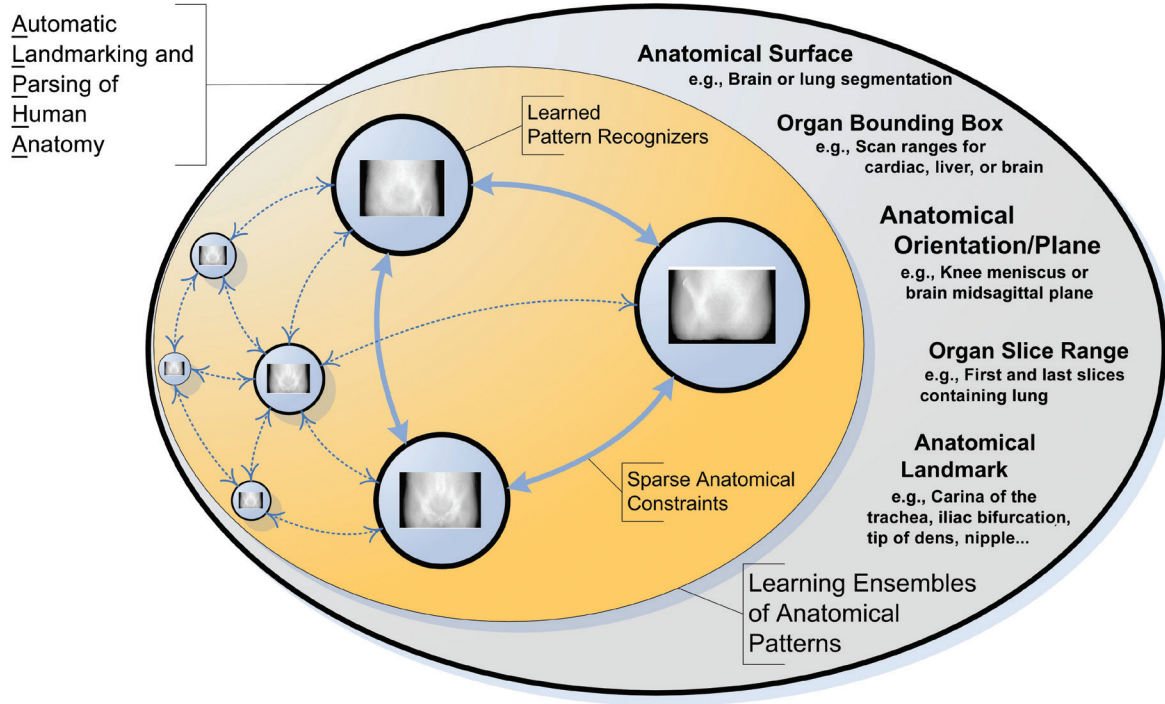



Figure 4. ALPHA learns from expert annotated examples.



Reliability through Redundancy

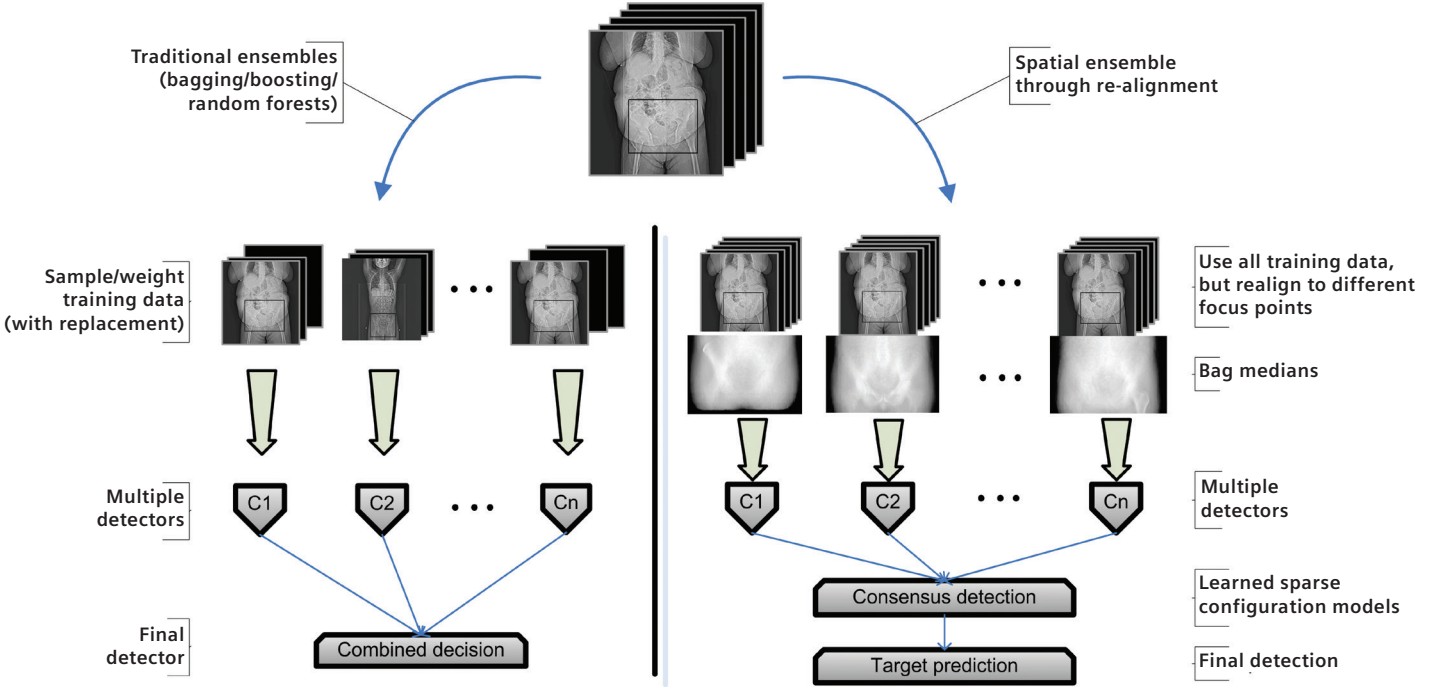
Designers for mechanical systems, such as an aircraft engine, use redundancy to detect possible failures and to boost system reliability. The way to implement redundancy is usually through duplication or triplication of critical functions or components. If a component fails, one or more redundant parts or alternative logic can identify, alert to or correct such a failure.

The ALPHA engine is designed with multiple layers of redundancy as well², to achieve robustness against diseases or imaging artifacts, and reliability against algorithmic failure. ALPHA exploits redundancies in several ways:

1. **Collection of a redundant set of local evidences.** For example, to determine a 3D registration matrix between two PET/CT or SPECT/CT studies (or an oriented 3D MR scan range for knee meniscus), four points would suffice. ALPHA detects and uses ~20 landmark points. The way that these landmark points are learned is analogous to the traditional ensemble learning methods, such as bagging⁶, which have been shown, theoretically, to improve learning performance. The difference is that ALPHA uses a spatial re-alignment scheme (Figure 5) that makes use of *all* the training data, \mathcal{L} , instead of a subset of the training data (i.e., “bags” or bootstrap sample sets, $\mathcal{L}^{(B)}$) for each sub-task. More specifically, the classic bagging predictor for a given input \mathbf{x} and a predictor $\psi(\mathbf{x}, \mathcal{L})$ is formulated as

$$\psi_B(\mathbf{x}) = av \psi(\mathbf{x}, \mathcal{L}^{(B)}),$$

Figure 5. Spatial ensemble learning to exploit redundancies and dependencies for improved robustness.



where \mathbf{av} denotes averaging or voting among the predictors. In ALPHA, the bootstrapping is done by spatial re-alignment of the full training set to different landmarks on the target. Denoting the re-alignment process as $\mathcal{A}_i \odot \mathcal{L}$, with \mathcal{A}_i representing the i^{th} alignment parameters and \odot the alignment operation, the formulation is

$$\psi_A(\mathbf{x}) = \text{vote } \psi(\mathcal{A}_i \odot \mathcal{L}),$$

where vote is a specific voting strategy (described below).

2. Enforcement of the spatial relationships among anatomical structures. In the aforementioned voting function, ALPHA embeds anatomical constraints among the landmarks using a sparse spatial configuration model. The vote received by landmark \mathbf{p}_i is denoted by $\chi(\mathbf{p}_i | \mathbf{P}_g)$, where \mathbf{P}_g is a voting group. The vote is defined as \mathbf{p}_i 's likelihood of being accepted or predicted by \mathbf{P}_g based on the conditional distribution estimated using the annotated training set. Assuming Gaussianity with mean $\mathbf{m}_{i|g}$ and covariance $\Sigma_{i|g}$, the vote is

$$\chi(\mathbf{p}_i | \mathbf{P}_g) = \frac{1}{(2\pi)^{D/2} |\kappa \Sigma_{i|g}|^{1/2}} e^{-(\mathbf{p}_i - \mathbf{m}_{i|g})^T \kappa \Sigma_{i|g}^{-1} (\mathbf{p}_i - \mathbf{m}_{i|g})}$$

where D is the dimensionality of the image and κ is a scaling parameter used to adapt to under-represented non-Gaussian variations (e.g., due to large anatomical articulation). ALPHA applies anatomical constraints within each of a large number of groups, where each group contains only a very small number of landmarks (either two or three), instead of a global model containing all landmarks. The benefit of this sparse and highly redundant formulation is that even when diseases or artifacts alter/occlude much of the target anatomy, it can still be detected or inferred using limited evidences.

3. Exploiting redundancy in scale space of images. ALPHA detectors are built in multiple scales, but with reduced dependencies among them. In other words, while the traditional strategy for exploiting scale space redundancy is to implement a "coarse-to-fine" search to speed up the detection algorithm, ALPHA exploits such redundancy to improve robustness.




Figure 6. ALPHA registration can align multiple time-point studies even when they have very different FOVs. *Siemens-generated data.*

ALPHA Answers

Because ALPHA can automatically recognize anatomical structures in medical images, it can potentially automate many workflow steps to save clinicians time and efforts. Below, we discuss how ALPHA-based features answer the two problems discussed earlier.

Answer for Challenge I: ALPHA Anatomical Registration

ALPHA's answer for the registration challenge is a rather straightforward one: In the current implementation, ALPHA detects up to 28 landmarks in each study that are covering the whole body, filters them based on an anatomical consistency check and uses the overlapping landmarks to align studies from different time-points.

Because ALPHA registration is based on recognition of anatomical structure just like the way a human observer would do it, instead of low-level pixel matching, it is robust to all the variations mentioned above (see Figure 6).

Furthermore, as ALPHA's capability is learned from examples and trained to recognize landmarks in MR images as well, it can even align studies from different modalities (see Figure 7).

Answer for Challenge II: Auto-Detection of Reference ROIs

ALPHA support in automatic placement of reference ROIs in the liver and in the descending aorta (see Figure 8) is not as simple as it seems: ALPHA detects multiple landmarks in and around the liver and the descending aorta, and uses all of them to infer and confirm the final placement of the two ROIs. As a result, the reference ROI detectors are very reliable despite changes in image contrast, and highly robust to abnormalities such as calcifications in the aorta or lesions in the liver (see Figure 9).

Batch Testing Results

A test was conducted on 400 randomly selected PET/CT studies. An ALPHA algorithm was invoked to detect 40 anatomical landmarks, including soft tissue landmarks, bone landmarks and vasculatures. The overall detection *sensitivity* (the likelihood a landmark present in the image will be detected) was 97%. The overall *accuracy* (given a landmark is detected, how likely it is to be accurately placed), was 99%.*




Figure 7. ALPHA can automatically align PET/CT and SPECT/CT with MR images based on shared anatomical landmarks because it uses the CT technology. *Siemens-generated data.*

*Data on file. Clinical experience may vary based on patients and image acquisition factors.

Figure 8. ALPHA automatically places reference ROIs in the liver and descending aorta to assist the clinician in implementing PERCIST-based quantification workflows. *Siemens-generated data.*





Figure 9. Examples of ALPHA-based PERCIST reference ROI detection: descending aorta (top row) and liver (bottom row). *Siemens-generated data.*



Conclusion

ALPHA technology treats a medical image not merely as a collection of pixels, but rather as a configuration of various anatomical structures. Based on its capability to learn human anatomy from different modalities, ALPHA can provide meaningful help for a clinician to improve his or her workflow. However, as with all computer algorithms, ALPHA results should not be taken as the ground truth, and a careful review and verification should be performed before drawing clinical conclusions.

About the Author

Dr. Sean Zhou studied Electrical Engineering (B. Sci. '93), Enterprise Management (B. Econ. '93) and Economics (PhD Candidacy) at Tsinghua University in China until 1995, and continued his graduate study in the United States in the areas of Computer Vision and Pattern Recognition, receiving a Master's degree from University of Cincinnati and a PhD from University of Illinois at Urbana Champaign.

Since 2002, he has been working for Siemens, first at Siemens Corporate Research in Princeton, New Jersey; then at Siemens Healthcare in Malvern, Pennsylvania, USA.

Dr. Sean Zhou is currently a principal key expert and the head of the US Innovation team for Siemens Medical Imaging IT Solutions.

References

- 1 **P. Ghosh and M. Kelly**, "Expanding the power of PET with PERCIST", White Paper, Siemens Healthcare, 2010
- 2 **X. S. Zhou, Y. Zhan, Z. Peng, M. Dewan, B. Jian, A. Krishnan, M. Harder, R. Schwarz, L. Lauer, H. Meyer, S. Grosskopf, U. Feuerlein, and H. Ditt**, "Reliability and redundancy: reducing error cost in medical imaging", in: B Krishnapuram, S. Yu, and R. Rao (Eds.): Cost-Sensitive Machine Learning, Chapman and Hall/CRC, 2011
- 3 **X. S. Zhou, Y. Zhan, V. C. Raykar, G. H. Valadez, L. Bogoni, Z. Peng, M. Dewan, M. Wolf, Y. Shinagawa, S. Park, M. Salganicoff, L. Raghupathi, and P. Devarakota**, "Mining anatomical, physiological and pathological information from medical images." SIGKDD Explorations 14(1):25-34, 2012
- 4 **Y. Zhan, M. Dewan, M. Harder, A. Krishnan, X. S. Zhou**, "Robust automatic knee MR slice positioning through redundant and hierarchical anatomy detection," IEEE Trans. Med. Imaging 30(12): 2087-2100, 2011
- 5 **Y. Tao, Z. Peng, A. Krishnan, X. S. Zhou**, "Robust learning-based parsing and annotation of medical radiographs." IEEE Trans. Med. Imaging 30(2): 338-350, 2011
- 6 **L. Breiman**, "Bagging predictors," Machine Learning, vol. 24, pp. 123–140, 1996

Trademarks and service marks used in this material are property of Siemens Healthcare GmbH.

All photographs © Siemens Healthcare GmbH.
All rights reserved.

Note: Original images always lose a certain amount of detail when reproduced.

Siemens Healthcare Headquarters

Siemens Healthcare GmbH
Henkestr. 127
91052 Erlangen
Germany
Telephone: +49 9131 84-0
siemens.com/healthcare

Global Business Line

Siemens Healthcare GmbH
Molecular Imaging
2501 N. Barrington Road
Hoffman Estates, IL 60192
USA
Telephone: +1 847 304 7700
siemens.com/mi