HHS Public Access

Author manuscript

Funct Imaging Model Heart. Author manuscript; available in PMC 2019 December 03.

Published in final edited form as:

Funct Imaging Model Heart. 2019 June ; 11504: 37–45. doi:10.1007/978-3-030-21949-9_5.

GRÖMeR: A Pipeline for Geodesic Refinement of Mesh Registration

Jake A. Bergquist[1,2,3], **Wilson W. Good**[1,2,3], **Brian Zenger**[1,2,3], **Jess D. Tate**[1,2,3], **Rob S. MacLeod**[1,2,3]

[1]University of Utah, Salt Lake City, UT 84112, USA

[2]Scientific Computing and Imaging (SCI) Institute, Salt Lake City, UT 84112, USA

[3]Nora Eccles Harrison Cardiovascular Research and Training Institute (CVRTI), Salt Lake City, UT 84112, USA

Abstract

The electrical signals produced by the heart can be used to assess cardiac health and diagnose adverse pathologies. Experiments on large mammals provide essential sources of these signals through measurements of up to 1000 simultaneous, distributed locations throughout the heart and torso. To perform accurate spatial analysis of the resulting electrical recordings, researchers must register the locations of each electrode, typically by defining correspondence points from post-experiment, three-dimensional imaging, and directly measured surface electrodes. Often, due to the practical limitations of the experimental situation, only a subset of the electrode locations can be measured, from which the rest must be estimated. We have developed a pipeline, GRÖMeR, that can perform registration of cardiac surface electrode arrays given a limited correspondence point set. This pipeline accounts for global deformations and uses a modified iterative closest points algorithm followed by a geodesically constrained radial basis deformation to calculate a smooth, correspondence-driven registration. To assess the performance of this pipeline, we generated a series of target geometries and limited correspondence patterns based on experimental scenarios. We found that the best performing correspondence pattern required only 20, approximately uniformly distributed points over the epicardial surface of the heart. This study demonstrated the GRÖMeR pipeline to be an accurate and effective way to register cardiac sock electrode arrays from limited correspondence points.

Keywords

Registration; Surface meshes; Geodesic

1 Introduction

The detection and diagnosis of electrical abnormalities of the heart continues to be a topic of active research despite over a century of study [5]. The electrical signals produced by the heart are used extensively to assess cardiac health and diagnose adverse pathologies [11].

jbergquist@sci.utah.edu.

However, many clinical applications are limited by a lack of mechanistic understanding. To better understand these mechanisms, *i.e.,* how electrical signals relate to specific pathologies, researchers use *in situ* experiments with high-density electrode arrays to measure up to 1000 individual signals directly from the hearts and torsos of large mammals [2,3]. One example of these high-density arrays is the Utah High Density Epicardial Sock (UHDES), a flexible mesh that is stitched with 240–490 evenly spaced electrodes and stretched over the ventricles of the heart [1,3,8]. The signals are then analyzed in space and time to visualize electrical changes occurring during experimental manipulation of the heart [3,6]. To include space in this analysis, researchers must know the locations of each electrode on the epicardial surface. However, in practical experimental conditions, often only a subset of the electrode locations can be measured directly due to limitations of the *in situ* preparations, *e.g.,* physical access to the electrodes, heart motion, and limited range of motion of mechanical digitizers. This limited subset of electrodes must then be used to estimate the locations of the entire array using surface registration of a preexisting geometrical model or one derived from subsequent imaging. The resulting registration problem is challenging because of the limited subset of electrode locations recorded, the non-rigid nature of the sock placement, and variations in cardiac anatomy between the experiment and any subsequent imaging.

Many examples of similar surface-to-surface registration problems exist, in which one must align a template model to a sub-sampled and distorted specific shape. Various registration approaches have been employed to address such problems, such as rigid or affine iterative closest points (ICP), thin plate splines, model based, Procrustes, and combinations of these methods [3,4,7,9,10]. Each approach and the resulting software pipelines have limitations, often with small numbers of correspondence points and the ambiguities of non-rigid deformation [2,3]. We developed a registration pipeline, GRÖMeR, that performs global deformations of an epicardial sock electrode array, calculates a rigid registration informed by the correspondence points, and then non-rigidly deforms the points into their final locations according to the known positions of correspondence-point electrodes. The result is a smooth, non-rigid deformation of the sock dictated by the correspondence points that results in an accurately registered electrode array.

The results of this study showed the GRÖMeR pipeline was able to register subject-specific and template UHDES arrays using a limited number of correspondence points. Among the correspondence point patterns considered, we found the most accurate registration across all target geometry scenarios with a relatively small number of measured correspondence points (20) distributed approximately uniformly over the heart. Similarly, we found that correspondence patterns that included both anterior and posterior coverage showed improved registration accuracy over posterior or anterior coverage alone.

2 Methods

Registration Pipeline

The GRÖMeR pipeline registers a template geometry to a target geometry using three inputs: (1) a template geometry, (2) correspondence points from the target geometry measured in the experiment coordinate system, and (3) a constraining surface derived from

post-experiment imaging. We assume that the target geometry exists on the surface of the constraining surface. The first step in GRÖMeR uses the positioning of the correspondence points to calculate a deformation along the long axis of the template geometry. This deformation is then applied to the template geometry to estimate the longitudinal stretch of the sock, which varies greatly across experiments. The longitudinally deformed template geometry is then rigidly registered to the constraining surface using a modified iterative closest points (ICP) algorithm [10]. This modified algorithm ensures that known correspondence points are included in registration calculations. The rigidly registered template is then snapped to the constraining surface by a nearest-node proximity. This snapped template geometry is deformed along the constraining surface as follows. For each correspondence pair, a Euclidean vector is constructed to move each correspondence point on the template geometry to the target geometry. This Euclidean vector is projected onto the constraining surface to calculate a geodesic path. An effect weighting (v) is then calculated using a radial basis function (Eq. 1) for each correspondence point and applied to each non-correspondence point on the template geometry, where d is the Euclidean distance between the correspondence point and a non-correspondence point, and r is the radial basis factor that determines the area of effect for a correspondence point. For this study, we selected $r = 50$ based on empirical assessment. Each non-correspondence point is then assigned a movement vector based on the sums of all correspondence point movement vectors and the applied radial basis effect. If the non-correspondence movement vector is normal to the local tangent space, that point remains static. Otherwise, the movement vector is projected onto the constraining surface to calculate a geodesic path for each non-correspondence point. Finally, the non-correspondence points are relaxed to maintain nearly average edge lengths across the registered geometry. The final result is a registered geometry, dictated by correspondence points and the constraining surface.

$$v = e^{-\left(\frac{d}{r}\right)^3} \tag{1}$$

Generation of Geometries

An essential goal of this study was to test the GRÖMeR pipeline under realistic conditions. To do so we evaluated it using a phantom 3D-printed heart, from which we generated target geometries to which we tried to register a template geometry. The template geometry was created by fitting the Utah High Density Epicardial Sock (UHDES) onto a plastic mold constructed from a heart harvested from a past experiment. All 247 electrode locations for this particular UHDES were measured using a mechanical digitizer (MicroScribe), and these electrode positions formed an open, cup-shaped surface, which was triangulated manually. For the constraining surface, a heart surface geometry was created according to Burton *et al.* [3] using a heart from an *in situ* preparation. Briefly, post-experiment the heart was excised and imaged with a seven Tesla MRI (Bruker BIOSPEC 70/30, Billerica, MA) using FISP (fast imaging with steady-state precession) and FLASH (fast low angle shot) imaging sequences. Utilizing images from both sequences, the heart volume was segmented using Seg3D open-source software (www.seg3d.org), and a heart surface mesh was extracted using the SCIRun problem solving environment (www.sci.utah.edu/cibc-software/scirun). To

simulate an experimental placement of the UHDES, we 3D printed the constraining surface as a phantom heart in PLA filament (Lulzbot Taz 3D printer) and then placed the UHDES on the 3D printed heart as in an *in situ* preparation. The locations of each electrode was again digitized (Fig. 1) and used as one target geometry. We then generated five additional target geometries by perturbing the phantom sock geometry as follows: uneven placement, in which posterior nodes were moved uniformly 20mm toward the apex (Fig. 1b); the Dilated or Contracted cases, in which all nodes (and the constraining surface) were dilated from or eroded toward, respectively, the center of the sock by 10 mm (Fig. 1c); and the Rotated cases, in which nodes were rotated about the long axis of the sock by 30° clockwise or counterclockwise (Fig. 1d). These target geometries were used as ground truths that we sought to recreate based on a limited sampling (Fig. 1).

Exploration of Correspondence Selection

A key parameter of any correspondence-based registration scheme is the choice and availability of correspondence points. We selected four patterns of correspondence points that mimicked realistic scenarios in our experiments: an *Anterior* group, *Posterior* group, evenly spaced *Vertical Strips*, and a *Uniformly Distributed* group over the entire target geometry. For each pattern, we selected either 8 or 20 correspondence points. The Anterior group simulated the limited access situation common during *in situ* experimental preparations. The posterior group simulated similarly limited access on the back of the heart. The vertical strips and distributed points were selected to explore the improvements possible through larger regions of correspondence points. Each correspondence point pattern was tested for each target geometry.

Analysis

The accuracy of the resulting registration was based on statistics of the Euclidean distances between each electrode of the registered and target geometries measured in millimeters (localization error). We also computed the percent of electrodes for which the localization error was below a threshold of 5 mm, termed the “percent of nodes below threshold” (PNBT). The mean, median, and standard deviation of the localization error and the PNBT were compared across the various target geometries and correspondence patterns.

3 Results

Registration Results

The accuracy achieved using the GRÖMeR pipeline depended on the locations and the number of correspondence points. The results summarized in Table 1 for sets of 8 correspondence points and Table 2 for sets of 20 correspondence points show the statistical results for each correspondence pattern and target geometry scenario.

The best performing correspondence pattern across many of the target geometries was the Uniformly Distributed correspondence pattern, closely followed by the Vertical Strips pattern. Within the target geometry perturbations, the dilated and uneven target geometries produced the highest localization error. The Distributed Correspondence pattern was the most robust to the target geometry variations according to all three metrics shown for both

numbers of correspondence points. The Distributed pattern with only 8 points was able to outperform 20 correspondence points in both Anterior and Posterior groupings. Figure 2 shows an example, comparing the Distributed-8 correspondence to the Anterior-20 for the uneven target geometry. Figure 2b shows the per node localization error mapped onto each node of the registered sock geometry, with higher error in red and lower error in blue. The maximum localization error was 8.56mm for the Distributed-8 correspondence set, and 20.15mm for the Anterior-20 set. Figure 2c shows the nodes above the 5 mm threshold in red and the nodes below the threshold in blue.

4 Discussion and Conclusions

In this study, we sought to develop a pipeline that could perform surface-to-surface registration based on a limited and often asymmetrically distributed set of correspondence points. In order to perform such a registration, we leveraged the availability of a high-resolution heart surface derived from MRI imaging over which the target sock geometry was stretched and over which our registration was constrained. We applied this novel registration approach to a cardiac sock electrode array, and tested it with a phantom heart model that we used to generate target geometries, which we then perturbed in realistic ways. We were able to achieve a high degree of fidelity using as few as 8 well-placed correspondence points. By selecting different correspondence point patterns, we were able to explore the robustness of the GRÖMeR pipeline under realistic sampling scenarios.

For all correspondence point patterns, contracting the target geometry and constraining mesh resulted in the most accurate registrations (highest PNBT and lowest localization errors), whereas dilating has the opposite effect. We consider that the localization error is a metric of the average electrode distance from the target geometry, whereas the PNBT is a metric of how well the registration performed overall compared to a predefined threshold. Therefore these results were expected, given that when the target geometry was dilated, the overall spacing of electrodes had to increase. Thus, error was magnified on this larger surface, and with the same threshold fewer nodes fell below the error threshold. The opposite is seen on the contracted target geometry. In these cases, the localization error provides a better interpretation for how the different correspondence patterns performed when compared to other target geometries in which the average electrode spacing is not eroded or dilated.

When examining the other target geometry perturbations, the uneven target geometry also showed high localization error and low PNBT. The local deformations were different on the anterior and the posterior sides and the non-uniform nature of the uneven perturbation likely contributed to the difficulty in registration. Correspondence patterns that did not cover locally deformed regions struggled to reconstruct these deformations. However, correspondence patterns that covered at least some part of these local deformation resulted in improved registrations. Tables 1 and 2 document these findings as does Fig. 2, which shows how a correspondence pattern that provides sampling closer to these local deformations resulted in better registration, as is expected.

Of the correspondence point patterns, the Posterior performed poorly in all cases. A possible explanation is that the posterior surface of the heart and therefore the target geometry is,

generally, smooth and flat. The anterior surface by contrast, particularly along the junction between the left and right ventricles, has a more complex shape, which is captured by the sock geometry. Thus, correspondence patterns that had at least some sampling in the anterior performed better than those that did not, because without anterior sampling it was difficult to accurately reconstruct the anterior geometry. Correspondence patterns with sampling over both the anterior and posterior were superior, because these patterns could capture local deformations across the majority of the sock geometry.

We also compared the performance of correspondence points spread evenly over the entire sock geometry rather than unevenly constrained to one side. An evenly distributed, or four strips of correspondence points across the target geometry, resulted in improved registration results compared to the Anterior and Posterior patterns. This result is particularly apparent in Fig. 2, in which the Distributed-8 correspondence pattern resulted in a superior registration accuracy than did the Anterior-20 correspondence pattern. The increase in PNBT of 44.77% and decrease in mean localization error of 4.16mm demonstrates the benefit of even a modest number of correspondence points distributed across the target geometry and highlights that simply increasing the number of correspondence point measurements on the anterior does not prove superior to a more even coverage. The Vertical strips pattern behaved similarly to the distributed, resulting in superior registration as compared to the anterior and posterior patterns, likely due to the increased correspondence coverage of the Distributed and Vertical strips patterns. In some cases, the Vertical strips 20 pattern resulted in a slightly higher PNBT than even the Distributed 20 pattern. However, at the 8 correspondence points level, these two patterns perform similarly. The results of this analysis highlight the motivation to collect more distributed correspondence points to get accurate registration of the complete sock electrode array. In the future we plan to address these technical limitations of our experimental preparation to achieve a correspondence distributions with errors that approache those of the ideal distributed pattern.

This study has demonstrated the novel GRÖMeR pipeline to be an accurate and effective way to register surface electrode arrays from limited correspondence points. Using the GRÖMeR pipeline, we were also able to suggest ways to improve the registration process by evenly distributing and perhaps increasing the number of recorded correspondence points. In the future, we will use this tool to register our cardiac sock electrodes and test ideal correspondence points to register the UHDES array.

References

1. Aras K, Burton B, Swenson D, MacLeod R: Sensitivity of epicardial electrical markers to acute ischemia detection. J. Electrocardiol 47(6), 836–841 (2014). 10.1016/j.jelectrocard.2014.08.014 [PubMed: 25242529]
2. Bear LR, et al.: Forward problem of electrocardiography: is it solved? Circ. Arrhythmia Electrophysiol 8(3), 677–684 (2015). 10.1161/CIRCEP.114.001573
3. Burton BM, Aras KK, Good WW, Tate JD, Zenger B, MacLeod S: Image-based modeling of acute myocardial ischemia using experimentally derived ischemic zone source representations. J. Electrocardiol 51(4), 725–733 (2018). 10.1016/j.jelectrocard.2018.05.005 [PubMed: 29997022]
4. Chen Y, Zhao J, Deng Q, Duan F: 3D craniofacial registration using thin-plate spline transform and cylindrical surface projection. PLoS ONE 12(10), 1–19 (2017). 10.1371/journal.pone.0185567
5. Einthoven W: Le telecardiogramme. Arch. Int. de Physiol 4, 132–164 (1906)

6. Rodenhauser A, et al.: PFEIFER: preprocessing framework for electrograms intermittently fiducialized from experimental recordings. J. Open Source Softw 3, 472 (2018). 10.21105/joss.00472 [PubMed: 30259008]
7. Sablatnig R, Kampel M: Model-based registration of front- and backviews of rotationally symmetric objects. Comput. Vis. Image Underst 87(1–3), 90–103 (2002). 10.1006/cviu.2002.0985
8. Shome S, Lux R, Punske B, MacLeod R: Ischemic preconditioning protects against arrhythmogenesis through maintenance of both active as well as passive electrical properties in ischemic canine hearts. J. Electrocardiol 40(4), S5–S6 (2007). 10.1016/j.jelectrocard.2007.06.012
9. Tate J, et al.: Reducing error in ECG forward simulations withimproved source sampling. Frontiers Physiol. 9, 1304 (2018). 10.3389/fphys.2018.01304
10. Zhang Z: Iterative point matching for registration of free-form curves and surfaces. Int. J. Comput. Vis 13(2), 119–152 (1994). 10.1007/BF01427149
11. Zipes D, Jalife J, Stevenson W (eds.): Cardiac Electrophysiology, From Cell to Bedside, 7th edn Elsevier, Philadelphia (2014)

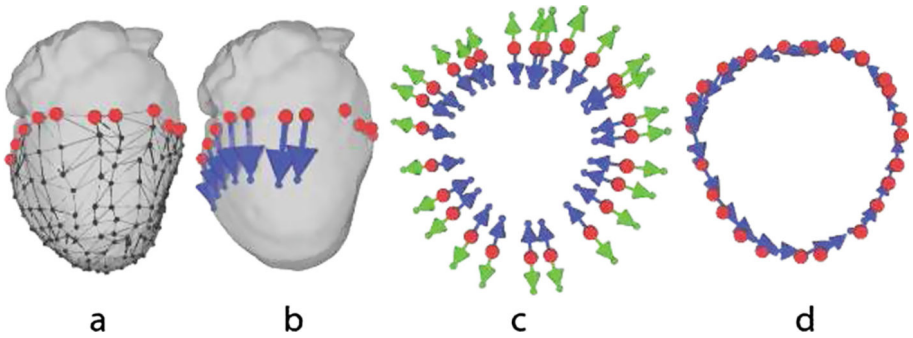

Fig.1.

Generation of target geometries. Digitized target geometry (black) on the constraining heart surface (gray) (a). Deformations used to generate the other target geometries are demonstrated on the top ring of nodes (red). Each node is an electrode. Uneven sock placement, in which the anterior electrodes are set and posterior electrodes are moved toward the apex by 20 mm (b). View from the base of the heart: electrodes and the heart mesh are either eroded (blue) or dilated (green) by 10 mm (c). Nodes are rotated by 30 degrees clockwise or counterclockwise. Counterclockwise shown in (d). (Color figure online)

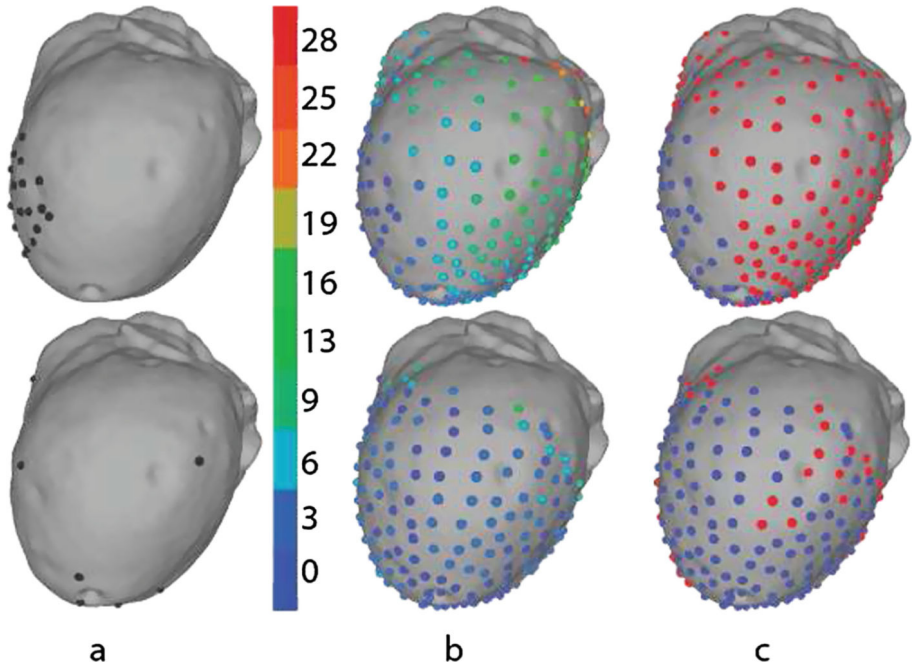

Fig.2.

Registration comparison between the Anterior 20 points (top) and the Distributed-8 points (bottom). The anterior surface of the constraining heart surface was oriented to the left and the posterior to the right. The correspondence points used are shown in black (a). Localization error (in millimeters) was mapped onto the registered electrodes (b) with high error in red and low error in blue. Electrodes with localization error below 5 mm were colored blue and with localization error at or above the same threshold were colored red (c). (Color figure online)

Table 1.

Numerical results from registration scenarios given 8 correspondence points. For each registration scenarios, the mean and median localization error (mm) are shown as well as the percent of nodes below the threshold localization error of 5 mm. All errors are listed as ± one standard deviation.

Target		Anterior	Posterior	Distributed	Vertical strips
Basal	Mean	6.52 ± 3.3	13.87 ± 9.09	4.13 ± 2.08	4.45 ± 2.18
	Median	6.55 ± 3.3	12.92 ± 9.09	4.03 ± 2.08	4.49 ± 2.18
	PNBT	33.6%	21.05%	66.8%	59.51%
Dilated	Mean	10.92 ± 5.3	19.02 ± 12.26	6.36 ± 2.97	7.15 ± 3.16
	Median	10.73 ± 5.3	17.49 ± 12.26	6.51 ± 2.97	7.15 ± 3.16
	PNBT	13.77%	14.98%	32.79%	23.48%
Contracted	Mean	3.46 ± 1.94	7.28 ± 4.26	2.8 ± 1.52	3.14 ± 1.56
	Median	3.25 ± 1.94	7.07 ± 4.26	2.7 ± 1.52	3.22 ± 1.56
	PNBT	76.92%	36.44%	91.5%	88.26%
Uneven	Mean	9.05 ± 5.28	15.62 ± 9.18	3.87 ± 2.16	5.32 ± 3.03
	Median	8.34 ± 5.28	15.58 ± 9.18	3.78 ± 2.16	5.1 ± 3.03
	PNBT	21.86%	16.19%	74.49%	47.77%
Rotated +30°	Mean	6.19 ± 3.24	10.49 ± 6.32	3.51 ± 1.95	4.02 ± 1.88
	Median	6.01 ± 3.24	9.87 ± 6.32	3.45 ± 1.95	4.14 ± 1.88
	PNBT	34.82%	24.29%	79.76%	67.21%
Rotated −30°	Mean	6.99 ± 4.74	12.71 ± 8.3	4.18 ± 2.24	4.49 ± 2.16
	Median	6.2 ± 4.74	11.41 ± 8.3	4.09 ± 2.24	4.5 ± 2.16
	PNBT	44.13%	21.46%	63.56%	58.3%

Table 2.

Numerical results from registration scenarios given 20 correspondence points. Otherwise identical layout as Table 1

Target		Anterior	Posterior	Distributed	Vertical strips
Basal	Mean	6.12 ± 3.42	13.53 ± 9.95	3.25 ± 1.93	3.26 ± 1.93
	Median	6.06 ± 3.42	11.96 ± 9.95	3.1 ± 1.93	3.19 ± 1.93
	PNBT	38.87%	24.29%	80.57%	82.59%
Dilated	Mean	9.95 ± 5.66	18.77 ± 13.9	4.79 ± 2.71	4.8 ± 2.87
	Median	10.06 ± 5.66	16.75 ± 13.9	4.9 ± 2.71	4.86 ± 2.87
	PNBT	19.84%	18.62%	50.61%	53.44%
Contracted	Mean	3.17 ± 1.94	7.82 ± 5.38	2.33 ± 1.43	2.51 ± 1.52
	Median	3.13 ± 1.94	7.18 ± 5.38	2.2 ± 1.43	2.52 ± 1.52
	PNBT	81.38%	34.41%	95.55%	93.52%
Uneven	Mean	8.03 ± 5.27	16.37 ± 12.24	3.39 ± 2.43	3.26 ± 2.34
	Median	7.12 ± 5.27	14.18 ± 12.24	3.14 ± 2.43	2.94 ± 2.34
	PNBT	26.72%	21.05%	77.33%	80.97%
Rotated +30°	Mean	6.89 ± 4.77	10.37 ± 7.04	2.94 ± 1.79	3.15 ± 1.85
	Median	6.17 ± 4.77	9.22 ± 7.04	2.84 ± 1.79	3.02 ± 1.85
	PNBT	36.84%	27.94%	85.02%	85.02%
Rotated —30°	Mean	6.42 ± 4.49	13.29 ± 9.91	3.42 ± 1.97	3.15 ± 1.85
	Median	6.14 ± 4.49	11.07 ± 9.91	3.41 ± 1.97	3.03 ± 1.85
	PNBT	42.91%	24.29%	78.95%	82.59%