Computer Assisted Localization of a Heart Arrhythmia

Chris Vogl, Peng Zheng, Stephen P. Seslar, and Aleksandr Y. Aravkin

Abstract—We consider the problem of locating a point-source heart arrhythmia using data from a standard diagnostic procedure, where a reference catheter is placed in the heart, and arrival times from a second diagnostic catheter are recorded as the diagnostic catheter moves around within the heart.

We model this situation as a nonconvex feasibility problem, where given a set of arrival times, we look for a source location that is consistent with the available data. We develop a new optimization approach and fast algorithm to obtain online proposals for the next location to suggest to the operator as she collects data. We validate the procedure using a Monte Carlo simulation based on patients' electrophysiological data. The proposed procedure robustly and quickly locates the source of arrhythmias without any prior knowledge of heart anatomy.

Index Terms—Arrhythmia localization, feasibility problems, nonconvex optimization.

I. INTRODUCTION

Catheter ablation is the treatment of choice to diagnose and treat cardiac arrhythmias. Accurately determining the origin of a cardiac arrhythmia is of critical importance in catheter ablation procedures. In many instances, arrhythmias originate from a focal point source and the electrical signal spreads concentrically in all directions away from that point. The traditional method of localizing such arrhythmias involves a user-directed movement of a mapping catheter through a cardiac chamber — a somewhat haphazard "hunt and peck" method. The arrhythmia signal arrival time on the roving map catheter is compared against the signal arrival time on a stationary reference catheter until the *earliest relative timing* site is identified.

Based on inherent limitations in how humans recognize patterns, this process requires a significant area of the heart chamber to be mapped before the operator can start to hone in on the location of the arrhythmia source. The operator is essentially required to solve an optimization problem by hand

C. Vogl is with the Lawrence Livermore National Laboratory
P. Zheng and A. Aravkin are with the Department of Applied Mathematics at the University of Washington
S. Seslar is with Seattle Children's Pediatric Hospital and the University of Washington

to minimize the arrival time relative to that obtained by a stationary catheter.

We propose a computer-assisted mapping system that can effectively use all available information and point the operator to the 'next touch' location. This approach can substantially reduce the time and number of touch points needed to reliably determine the site of arrhythmia origin. We develop a method for arrhythmia localization and evaluate it using Monte Carlo simulations based on a deidentified electroanatomic map from a patients that underwent mapping and acutely successful ablation of a focal arrhythmia using the Rhythmia mapping system[1].

Related work. We found one prior approach to localizing arrhythmia using optimization [5]. The approach is similar in principle, and uses optimization to recommend the next point to sample by the operator. The key differentiating factor is that [5] assumes the existence of an anatomical map; the approach is predicated on being able to solve a linear regression problem that uses *every available nodal point* as a potential origin, and then pick the one that best fits the available data. In contrast, we make no assumptions about the existence of anatomical maps; our proposals are based only on observations taken by the operator, and we can help patients who have undergone no prior mapping studies. In addition, the algorithm of [5], for each recommendation, must solve a number of regression problems equal to the number of points in their mesh. In contrast, we only need to solve one feasibility problem to find the next point proposal.

The paper proceeds as follows. In Section II, we formulate arrhythmia localization as a feasibility problem and derive an associated optimization problem. In Section III we develop a fast algorithm, with practical considerations for dealing with noisy data and online application of arrhythmia localization. In Section IV we present results of the approach using patient arrival time data, and end with conclusions in Section V.

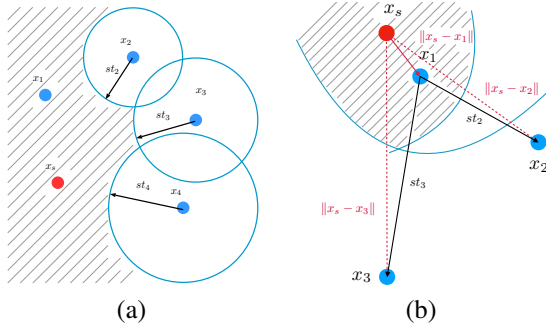[1]Boston Scientific, Marlborough, MA, USA.

Fig. 1. Simple timing relationships define a feasibility problem for locating the source location x_s of the arrhythmia. (a) The set Ω_1 described by (2). (b) The set Ω_2 described by (3).

II. FORMULATION

We are given a set of (ordered) relative arrival times

$$0 = t_1 \leq \ldots \leq t_m, \tag{1}$$

obtained by finding the differences between times recorded using a reference catheter and a diagnostic catheter at 3D locations $x_1, \ldots, x_m$ within the heart. **Simple Feasible Region.** We look for a source location x_s consistent with available observations. We assume that the local signal transmission speed s around the source x_s is known and given. Denoting the actual arrival times by $\hat{t}_i$, we get relations

$$\hat{t}_i \geq t_i, \quad i = 1, \ldots, m,$$

which translate to the constraints

$$\|x_s - x_i\|_2 \geq st_i, \quad i = 1, \ldots, m. \tag{2}$$

Finding a source that is consistent with the available observations is equivalent to finding x_s that satisfies (2). The feasible region from (2) is shown in panel (a) of Figure 1; the true source must lie outside of the union of the disks shown in the figure.
Coupled Feasible Region. A more powerful formulation incorporates first arrival information, similar to the fast marching method used to solve the Eikonal equation [4]. We look for x_s satisfying

$$\|x_s - x_i\| \geq \|x_s - x_1\| + st_i, \ i \geq 2. \tag{3}$$

The fast marching method and the inequalities (3) enforce a reverse triangle inequality: the time the signal takes to go from x_s to x_i is is larger than the time needed to go from x_s to x_1 and then x_1 to x_i; otherwise we would not have observed the given distribution of arrival times. The resulting region is shown in panel (b) of Figure 1.

III. ALGORITHM

Nonconvex feasibility problems can be solved using optimization techniques such as alternating optimization methods or Douglas-Rachford splitting [2].

These algorithms can take hundreds to thousands of iterations for simple problem instances [2, Table 1]. We propose a new relaxation that converges very rapidly, generating a feasible solution x_s within a few iterations in most instances.

In developing relaxations for (2) and (3), we use the ideas recently developed by [6]. We introduce auxiliary variables w_i to approximate each $x_s - x_i$, and minimize over both x_s and these auxiliary variables.

$$\min_{x_s, w} \quad f(x_s, w) := \frac{1}{2} \sum_{i=1}^{m} \|x_s - x_i - w_i\|^2 \tag{4}$$
$$\text{s.t.} \quad w \in \Omega,$$

with Ω a special set described below. The original sets defined by (2) and (4) are both subset of $\mathbb{R}^3$. The set describing (2) is a simple subset of $\mathbb{R}^{3m}$:

$$\Omega_1 := \bigoplus_{i=1}^{m} \{w_i : \|w_i\|_2 \geq st_i\}.$$

The set in (3) is also a subset of $\mathbb{R}^{3m}$, with more complex structure:

$$\Omega_2 := \{w_1 : w_1 \in \mathbb{R}^3\} \oplus$$
$$\bigoplus_{i=2}^{m} \{w_i : \|w_i\|_2 \geq \|w_1\|_2 + st_i\},$$

where $\oplus$ denote the direct sum. The original problem for calculating a projection onto a nonconvex set is difficult. By relaxing the formulation, the objective becomes more tractable, yielding a simple update rule using the structure of Ω. We also have a guarantee of optimality for the original feasibility problem based on the objective value of (4):

- Any x_s satisfying (2) gives a global minimizer of (4) with objective value 0, by $\overline{w}_i = x_s - x_i$.
- Any solution with zero objective value gives $\overline{x}_s$ feasible with respect to (2) or (3).

Problem (4) may have a nonzero optimal value, in which case the 'relaxed' solution x_s will not satisfy the original formulation. However, in practice we find a feasible point in each iteration. To solve (4), we minimize over x_s and w_i. Given $\{w_i\}$, we have a closed for solution for x_s:

$$x_s = \frac{1}{m} \sum_{i=1}^{m} (x_i + w_i). \tag{5}$$

To find $\{w_i\}$ given x_s, we have to solve

$$\min_{w \in \Omega} 0.5 \sum_{i=1}^{m} \|w_i - (x_s - x_i)\|_2^2.$$

When $\Omega = \Omega_1$, we have a closed form solution for the projection problem:

$$w_i = \frac{x_s - x_i}{\|x_s - x_i\|} \max(st_i, \|x_s - x_i\|_2). \tag{6}$$

When $\Omega = \Omega_2$, the projection is found by solving

$$\min_{\rho} \quad 0.5 \sum_{i=1}^{m} (\rho_i - r_i)^2 \text{ s.t. } \rho_i \geq \rho_1 + st_i, \quad i \geq 2,$$

which requires a specialized subroutine. The approach is summarized in Algorithm 1.

Algorithm 1 Source Finding Algorithm

1: **Input:** $\{x_i\}$, s
2: **Initialize:** $k = 0$, $x_s^0 = x_1$, $w_i^0 = x_s^0 - x_i$
3: **while** not converged **do**
4: $\quad x_s^{k+1} \leftarrow \frac{1}{m} \sum_{i=1}^{m} (x_i + w_i^k)$
5: $\quad w^{k+1} \leftarrow \text{proj}_\Omega(\{x_s^{k+1} - x_i\})$
6: $\quad k \leftarrow k + 1$
7: **Output:** x_s^k

Algorithm 1 terminates when the function value or step-size is less than a specified tolerance, or if we hit an iteration cap of 200. It is equivalent to proximal gradient descent on the value function for (4):

$$\widetilde{f}(w) = \min_{x_s} f(x, w).$$

See [6] for an analysis of such algorithms, including rates of convergence.

Robust modification. The data collection process by a diagnostic catheter is inherently noisy, and some trial points give anomalous timing data. These anomalies then give incorrect information about the feasibility region (2). To make the method more robust, we detect and remove potential outliers in the course of solving each optimization problem (4). The outliers naturally give constraints that are very hard to satisfy. We introduce a vector τ to indicate which constraints are easy to fit, and which are difficult. The few constraints that are the least consistent with the remaining data are likely outliers. This idea can be traced back to least trimmed squares [3]; see also [1] for a survey of modern applications.

To implement the approach, we remove a small number h of arrival times from consideration in each iteration. We sort $\{r_1 - st_1, \ldots, r_m - st_m\}$ from least to greatest, and for each index i in the smallest h residuals, we set $\tau_i = 0$, while all remaining τ_i are set to 1. When we update x_s, we modify (5) to $x_s = \frac{1}{m-h} \sum_{i=1}^{m} \tau_i(x_i + w_i)$. This strategy removes the influence of the potential outliers, and decreases the number of touches we need to find the source.

Online implementation. We start with several measurements obtained by a preliminary diagnostic catheter with 10 poles, see Figure 3. This sets up the first feasibility problem (2) or (3), which we solve by the reformulation (4) point to sample.
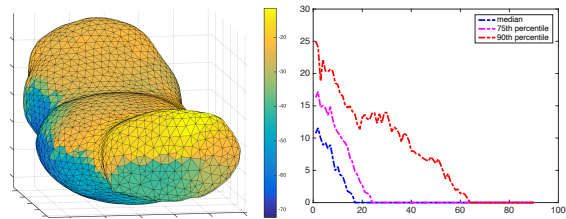

Fig. 2. Left: simple ventricle dataset with color-coded arrival times. Right: distance to source as a function of touches using simple feasible region (2). In this simple dataset, the proposed approach quickly finds the source.
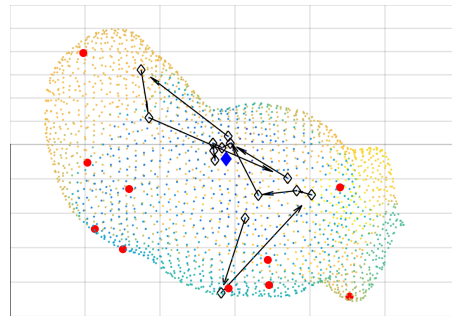

Fig. 3. Path to source for sample run using dataset in Figure 2. Initial placement of diagnostic catheter is shown by red dots; path to source is shown using diamond markers.

As we proceed, we consider the last 10 observations in forming each subsequent feasibility problem. If we accumulate a lot of data far away from the source, the feasibility problem becomes harder to solve; in particular the simple approximation of assuming a constant propagation speed s between the potential source x_s and all observations is not reasonable. Algorithm 1 typically finds a solution (i.e. the next potential x_s given current data) within 1-2 iterations for most problems.

IV. RESULTS

We run a simulation using real patient data. Timing and mesh data for three patients (two ventricular chambers and one atrium) are available from a diagnostic study, see Figures 2-5. The 'ground truth' of the arrhythmia source is inferred by the earliest arrival time observed during the entire data collection process. We start with 10 observations made by a diagnostic catheter, and use the online version of the algorithm to locate the source. Monte Carlo sampling is used to randomly initialize the initial 10 readings; we track the distance of estimates to source as a function of touches, and report the results across the simulations by using median, 75th, and 90th quantiles of the distance to source (mm) as a function of touches. A sample run of the algorithm is shown in Figure 3. The algorithm proposes the
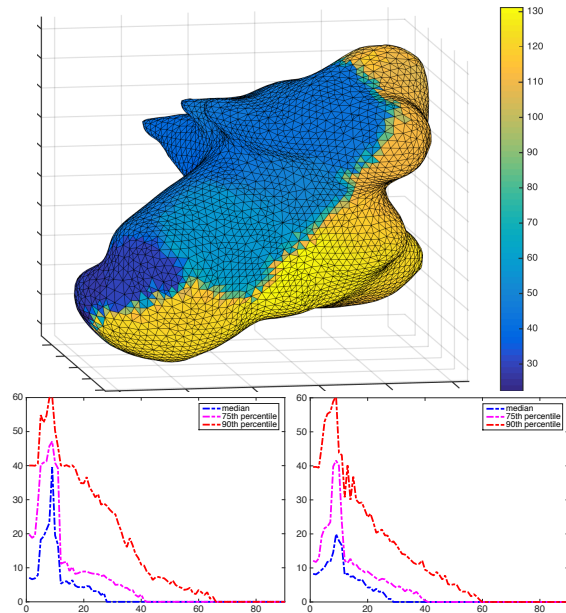
Fig. 4. Top panel: ventricle dataset with color-coded arrival times. Bottom left: distance to source as a function of touches using simple feasible region (2). Bottom right: distance to source as a function of touches using coupled feasible region (3). Both approaches perform similarly on this example.
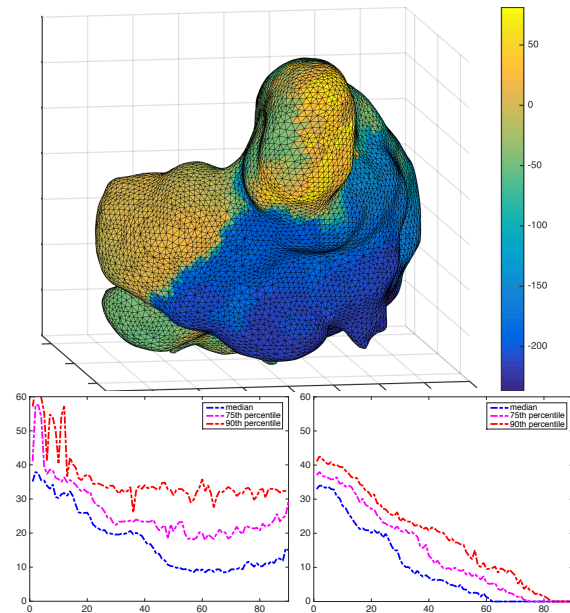

Fig. 5. Top panel: atrial dataset with color-coded arrival times. Bottom left: distance to source as a function of touches using simple feasible region (2). Bottom right: distance to source as a function of touches using coupled feasible region (3). Accounting for coupling in the signal propagation gives a significant improvement on this challenging dataset.

next point to sample after each touch. We then sample the closest point with data to each x_s proposed by the algorithm; in practice the operator can simply attempt to move the catheter to a proposed point to obtain the next measurement. The procedure takes 12 touches to locate the true source for this run. For each touch, we need 1 or 2 iterations of Algorithm 1 to solve the feasibility problem. The objective values are numerically 0, which means each x_s we find satisfies (2).

Simulation results are plotted for each of the three datasets in Figures 2-5. For the first two datasets, using the simple feasible region (2) works as well as using the coupled encoding (3). For the more complex atrial dataset in Figure 5, adding more geometrical constraints pays off, and the results of (3) are significantly better. In all cases, Algorithm 1 easily solves the nonconvex feasibility problems required to tell the operator where to sample next.

V. CONCLUSION

We formulated the arrhythmia localization problem as a sequence of nonconvex feasibility problems, and developed an efficient algorithm to solve these problems. The approach can accommodate different physical models of signal propagation in the heart; we compared two different models in this paper. The resulting approach opens a path to a new computationally-guided clinical approach to localize

point-source arrhythmias, making online suggestions based on minimal information: we assume no prior knowledge about the heart's anatomy. The next steps are to develop and test in the clinical setting with existing arrhythmia mapping systems.

VI. ACKNOWLEDGEMENTS

REFERENCES

[1] A. Aravkin and D. Davis. A smart stochastic algorithm for nonconvex optimization with applications to robust machine learning. *arXiv preprint arXiv:1610.01101*, 2016.

[2] G. Li and T. K. Pong. Douglas–Rachford splitting for nonconvex optimization with application to nonconvex feasibility problems. *Math. Prog.*, 159(1-2):371–401, 2016.

[3] P. J. Rousseeuw and C. Croux. Alternatives to the median absolute deviation. *Journal of the American Statistical association*, 88(424):1273–1283, 1993.

[4] J. A. Sethian. A fast marching level set method for monotonically advancing fronts. *Proceedings of the National Academy of Sciences*, 93(4):1591–1595, 1996.

[5] T. Weber, H. A. Katus, S. Sager, and E. P. Scholz. Novel algorithm for accelerated electroanatomic mapping and prediction of earliest activation of focal cardiac arrhythmias using mathematical optimization. *Heart rhythm*, 14(6):875–882, 2017.

[6] P. Zheng and A. Aravkin. Fast methods for nonsmooth nonconvex minimization. *arXiv preprint arXiv:1802.02654*, 2018.