



All-in-one electrical atrial substrate indicators with deep anomaly detection

Luca Bindini ^{a,*}, Stefano Pagani ^b, Andrea Bernardini ^{c,d}, Benedetta Grossi ^e, Andrea Giomi ^c, Antonio Frontera ^f, Paolo Frasconi ^a

^a DINFO, Università di Firenze, Via di Santa Marta 3, Firenze, 50139, FI, Italy

^b MOX, Department of Mathematics, Politecnico di Milano, Piazza Leonardo da Vinci, 32, Milano, 20133, MI, Italy

^c Cardiology and Electrophysiology Unit, Santa Maria Nuova Hospital, Piazza di Santa Maria Nuova, 1, Firenze, 50122, FI, Italy

^d Department of Experimental and Clinical Medicine, Università di Firenze, Largo Brambilla, 3, Firenze, 50134, FI, Italy

^e Department of Biomedical Sciences, Humanitas University, Via Rita Levi Montalcini, 4, Milano, 20072, MI, Italy

^f Electrophysiology Unit, De Gasperis Cardio Center, Niguarda Hospital, Piazza dell'Ospedale Maggiore, 3, Milano, 20162, MI, Italy

ARTICLE INFO

Keywords:

Atrial fibrillation
Deep anomaly detection
Unsupervised learning
Artificial intelligence
Deep learning

ABSTRACT

Background and Objective: Electrophysiological studies based on high-density catheter mapping have become a cornerstone in ablation procedures for managing atrial fibrillation. Since their introduction, extensive efforts have been devoted to analyzing intracardiac atrial electrograms (EGMs) to identify potential arrhythmogenic regions, often employing machine learning techniques. The aim is to investigate the potential of deep anomaly detection algorithms as either complements or replacements for established electrophysiological indicators used to characterize EGMs associated with arrhythmic substrate.

Methods: We investigated three deep anomaly detection algorithms. Model outputs were used to generate consistent and robust scores for each signal in a completely unsupervised manner. We applied these techniques to 8 patients with a leave-one-out strategy.

Results: Our numerical experiments show that: (i) higher anomaly scores are correlated with higher EGM fractionation and duration and lower voltage, (ii) thresholding anomaly score percentiles and standard indicator values produce consistent classifications, and (iii) morphology analysis is more robust compared to a stratification provided by single standard indicator, without the need for determining arbitrary thresholds.

Conclusions: Our results demonstrate the effectiveness and robustness of deep anomaly detection algorithms in the characterization of anomalous cardiac EGMs. By providing an *all-in-one* method to assess pathological features, these NN models eliminate the limitations that arise from manually combining and visually comparing traditional indicators. Our electro-anatomical maps displaying anomaly scores could have significant implications for improving the accuracy and efficiency of ablation procedures aimed at managing cardiac arrhythmias, such as atrial fibrillation.

1. Introduction

Atrial fibrillation (AF) is the most common arrhythmia worldwide, with an estimated prevalence in adults between 2% and 4% [1] and an increasing incidence with age. Catheter ablation (CA) represents a cornerstone strategy for symptom improvement and rhythm control. CA aims to electrically isolate arrhythmogenic areas and is typically recommended for symptomatic patients as first option or when pharmacological therapy fails or has adverse effects. Since the abnormal electrical activity is frequently initiated within the pulmonary veins [2], the standard for first ablation procedure consists of pulmonary vein isolation (PVI). In patients with paroxysmal atrial fibrillation (PAF),

PVI resulted to be effective and superior to antiarrhythmic drugs, but in persistent AF (PsAF), success rates are lower [3–5].

In the attempt to improve success rate in PsAF, several researchers have supported the ablation of additional patient-specific drivers beyond PVI for PsAF, but results are mixed and there is currently no consensus on the most successful ablation strategy for these patients [6] and the identification of arrhythmogenic regions (which may guide CA) is still an open challenge.

The three main approaches being pursued to spot those regions are:

1. processing and interpretation of electro-anatomical (EA) data acquired during sinus rhythm, e.g., [7–11],

* Corresponding author.

E-mail address: luca.bindini@unifi.it (L. Bindini).

<https://doi.org/10.1016/j.bspc.2024.106737>

Received 14 April 2024; Received in revised form 18 July 2024; Accepted 3 August 2024

Available online 14 August 2024

1746-8094/© 2024 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

2. processing and interpretation of EA data acquired during fibrillation, e.g., [12–16],
3. non-invasive imaging-based numerical simulations providing in silico prediction of the induction and maintenance of re-entrant drivers, e.g., [17–19].

For ease of implementation, robustness and reproducibility in clinical routine the first option remains the preferred approach by electrophysiologists. However, the optimal processing and interpretation of electrograms (EGMs) signals is still controversial.

In Sinus Rhythm, various indicators of arrhythmic substrate have been suggested and studied, including signal amplitude (voltage), number of spikes (fractionation), and duration of the fractionated component [11,20,21]. Despite attempts to define more specific indicators, such as EGM amplitude-normalized area [22], there is still a lack of effective quantitative tools capable of characterizing EGM morphology in a synthetic manner.

In this paper, we aim to fill this gap through the application of deep learning techniques for the identification of anomalous signals.

We propose to explore unsupervised learning, particularly deep anomaly detection algorithms, as an *all-in-one* solution to the characterization of the arrhythmic substrate in SR, hypothesizing that anomaly detection can more effectively characterize the pathological features of electrograms than traditional signal biomarkers.

Anomaly detection is the problem of identifying the support (i.e., the typically low-dimensional region of the data space where most of the probability mass is concentrated) of some unknown distribution based on a training set of samples [23–26]. An anomaly detection algorithm output for each new data point an anomaly score that is high when the point has low probability of being in distribution according to the model. Unlike density estimation techniques, anomaly detection algorithms do not suffer from the curse of dimensionality problem and have been effectively applied in several industrial [27,28] and medical applications [29–33].

In this work, the use of anomaly detection is justified by the reasonable assumption that most of the collected EGMs are taken from healthy tissue and thus lie in the support of the distribution of EGM signals, while anomalous signals (i.e., with a high anomaly score), may be expected to be associated with unhealthy tissue. Unlike the supervised approach, which relies on hand-labeled training examples [34], unsupervised learning can discover classes of electrical anomalies without extensive human supervision. Compared to previous works here we aim, for the first time, to bring state-of-the-art deep anomaly detectors [35–37] as an innovative tool to automatically identify electrical anomalies. This approach has the potential to improve clinical practice and the success rate of CA procedures by providing a new class of ablation targets.

Our results indicate that there is a strong agreement between anomalous EGMs and previously known EGM target candidates confirms that anomaly detection is a plausible direction for further study, while discovered anomalous signals that do not belong to already known categories may provide useful insights and may reveal potentially useful targets to be clinically validated.

2. Background

Electroanatomical mapping allows electrophysiologists to visualize both a high-density reconstruction of the anatomical shape of the endocardium, and a detailed representation of its electrical activity in terms of thousands of spatially localized EGMs. In the following, we describe three common EGM indicators used in clinical practice, along with their respective implementations based on published algorithms from OpenEP library [38]. In Section 4 we will correlate each of them with anomaly detection scores.

2.1. Peak-to-peak voltage

This indicator measures the EGM amplitude and is usually binned to classify EGMs into categories: healthy, border zone, and low-voltage/scar. Typical ranges for bipolar EGMs are > 1.0 mV for healthy tissue, $[0.5, 1.0]$ for the border zone, < 0.5 mV for low-voltage, and < 0.1 mV for dense scar, but a gold standard to assess scar criteria is missing and the specificity of this biomarker is low [39,40]. Specifically, this indicator is influenced by a large number of confounding factors such as electrodes position with respect to the wavefront, electrodes spacing, size and contact and stimulation rate [41].

In patients with early-stage AF, the vast majority of voltage values are usually greater than the healthy cut-off, whereas, in some patients with a progression of the disease, reduced potentials are the most frequent. In addition, conduction anomalies, like pivot points or slow conduction corridors, can be found in high-voltage areas [42]. These characteristics hinder voltage's stratification of the different tissue areas.

2.2. Fractionation index

The previous indicator does not encode the information contained in the morphology of the signal. In this direction, signal fractionation has been employed to identify EGMs with complex morphologies. While a physiological EGM consists of one peak followed by a single deflection (see Fig. 1(a)), in a fractionated EGM there are multiple deflections (see Fig. 1(b)).

Fractionated EGMs have been the subject of numerous studies in the context of cardiac arrhythmias in the last decades, starting from the results of Konings et al. [12,43], complex fractionated atrial electrograms (CFAE), recorded during AF, have become a target of ablation [44]. However, CFAEs are highly dependent on the specific episode recorded by the catheter, and given the high functional characteristics of AF (close to chaotic behavior), clinical trials have shown that they do not constitute effective ablation targets [45–49]. During SR, fractionation typically results from functional phenomena in the form of wavefront collisions and conduction delay in proximity of slow conducting areas, causing dyssynchrony between the two electrodes [11,42,50]. While the former may mask local information, the latter may identify areas contributing to AF maintenance (this assumption is usually supported by correlating this indicator with low peak-to-peak voltage).

As detailed in a recent meta analysis [51], there is no exact universally recognized definition of fractionation and defining a numerical index requires choosing the values of several parameters. One simple formula for computing the fractionation index F of bipolar EGMs recorded in sinus rhythm can be derived as follows [38]:

$$M = \max_t s^2(t) \quad (1)$$

$$F = \left| \left\{ t \in \mathcal{P} : s^2(t) > \frac{M}{3} \right\} \right| \quad (2)$$

where \mathcal{P} is the set of peak positions in the signal. Based on this definition, we expect $P \leq 3$ in healthy tissue.

2.3. Duration index

EGM duration is a surrogate indicator of conduction properties: low amplitude-long duration signals are usually associated with slow conduction corridors [8,52], that facilitate the formation, depending on their severity, of stable or unstable localized re-entrant drivers [53]. This indicator is also investigated in the context of ventricular tachycardia [54,55] as a possible target of ablation. One formula for calculating the duration, following the OpenEP library implementation [38], can be derived as follows:

$$D = \left| \left\{ t : s^2(t) > \frac{M}{10} \right\} \right| \quad (3)$$

where M is defined in (1).




Fig. 1. Comparison between physiological and fractionated EGMs.

Table 1

Baseline characteristics of the study population.

Variables	Population (N = 8)
Age, y	65.5 \pm 7.7
Male, n (%)	5 (62.5%)
Paroxysmal AF, n (%)	3 (37.5%)
Persistent AF, n (%)	5 (62.5%)
LVEF, (%)	55.6 \pm 6.8
LA area (cm ²)	25.6 \pm 3.7
Dyslipidemia, n (%)	7 (87.5%)
Hypertension, n (%)	5 (62.5%)
Diabetes, n (%)	1 (12.5%)
Mild mitral regurgitation, n (%)	6 (75%)

Data are presented as: Mean \pm Standard deviation or n (percentage). AF, atrial fibrillation; LVEF, left ventricular ejection fraction; LA left atria.

3. Methods

3.1. Study design and enrollment

Eight patients scheduled for catheter ablation due to symptomatic paroxysmal or persistent atrial fibrillation (AF), were included in the study. The definitions of paroxysmal (PAF) or persistent (PsAF) and the clinical indications to proceed to ablation were given according to latest ESC Guidelines [1]. Exclusion criteria were a contraindication for catheter ablation or a previous ablation attempt, a known structural heart disease, and the presence of a thrombus in left auricula. This study complies with the Declaration of Helsinki and all participants provided written informed consent. The main baseline data patients' characteristics, AF type and cardiovascular risk factors, are shown in Table 1.

3.2. Electroanatomical mapping and EGMs acquisition

In order to allow EGMs acquisition, all patients enrolled underwent left atrium (LA) electro-anatomical mapping with the support of CARTO3 mapping system (Biosense Webster, Diamond Bar, CA). For each patient, a sinus rhythm electroanatomical map were created with PentaRay mapping catheter (Biosense Webster, Inc, Illinois, USA), a multipolar navigational diagnostic catheter provided with 20 acquisition electrodes. High-density maps were created with at least 3000 acquired EGMs homogeneously distributed over the entire LA surface. Electrodes contact with the endocardial surface was ensured with the CARTO3 Tissue Proximity Index (TPI™) feature, in order to avoid misleading bipolar voltage values. At the end of each procedure, sinus rhythm EGMs were exported from the mapping system, where they were filtered in a frequency range of 30 Hz–300 Hz.




Fig. 2. EGM region of interest.

3.3. Preprocessing

The bipolar EGM signals described in Section 3.2 were centered with respect to the timing corresponding to the minimum of the derivative (estimated with finite differences) of the corresponding unipolar signal as in the standard CARTO annotation

$$\operatorname{argmin}_t \frac{u(t+1) - u(t-1)}{2\Delta t}$$

where $u(t)$ is the corresponding unipolar signal sampled at time t and Δt is the length of the time discretization ($\Delta t = 1$ ms in our dataset). With the central point as a reference, we considered a window of 64 ms, which is adequate to illustrate the dynamics of the EGM (Fig. 2).

The signals were subsequently smoothed to reduce fluctuations due to signal acquisition system using a Gaussian convolutional kernel with $\sigma = 0.8$, and they were normalized between 0 and 1 in order to disregard the different scaling factor between different atrial zones and/or patients and thus consider only the morphology of the signal itself.

3.4. Deep anomaly detection models

Several approaches based on deep learning have been studied for anomaly detection [56]. All these methods learn a representation of the raw data and aim to identify the support of the data distribution so that anomalous (low probability) points can be subsequently identified.

We considered several deep neural network models with the objective of assigning a score to each signal (anomaly score). This score corresponds to the reconstruction error of the signal in autoencoder like architectures, even memory augmented, or to other metric as in the Deep SVDD architecture of Ruff et al. [37].

An autoencoder (AE) is a composite parameterized function $f(x) = d(e(x))$ where the encoder e maps data point, x , into a lower-dimensional representation $z = e(x)$, and the decoder d is trained to reconstruct the original signal as closely as possible. Using a low-dimensional z produces a bottleneck and the reconstruction error for points that are far from the support of the data distribution will be higher. Because of this property and their simplicity, autoencoders have been widely studied in the context of anomaly detection [26,56], using the reconstruction error as the anomaly score. An and Cho [35] showed that better anomaly detection rates can be obtained with variational autoencoders (VAE) [57], a model where z is stochastic, with prior distribution $p(z) = \mathcal{N}(0, I)$, likelihood $p(x|z)$ and posterior $q(z|x) = \mathcal{N}(\mu(x), \sigma(x))$.

Parameters are tuned to maximize the Kullback–Leibler divergence between the posterior and the prior of z , which can be seen as a regularizer and can be computed as $\sum_j (1 + \log \sigma_j(x)^2 - \mu_j(x)^2 - \sigma_j(x)^2)$.

In An and Cho [35] a VAE is used to compute an anomaly score consisting of the following “reconstruction probability”: first, sample L latents $z^{(l)}$ from the prior and obtain L generated points $x^{(l)}$ using the decoder; then feed these generated points to the encoder to obtain posterior parameters $\mu(x^{(l)})$, $\sigma(x^{(l)})$, and finally compute the anomaly score as

$$S = \frac{1}{L} \sum_{l=1}^L p(x^{(l)} | \mu(x^{(l)}), \sigma(x^{(l)})).$$

A memory-augmented autoencoder (MemAE) [36] reconstructs signals from a learned dictionary of basis vectors that operates as an external memory $M = [m_1, \dots, m_N]$ of size N . Encoding proceeds in two steps: first, an embedding vector, z , is computed with an encoder $f_e(x)$. Then, z is used to query the memory by computing a soft addressing vector α with elements:

$$\alpha_i = \frac{\exp(z^\top m_i)}{\sum_{j=1}^N \exp(z^\top m_j)}, \quad i = 1 \dots N,$$

being m_j the j th memory element. Following Gong et al. [36], sparse addressing is achieved by hard shrinkage:

$$\hat{\alpha}_i = \frac{\max(\alpha_i - \lambda, 0)\alpha_i}{|\alpha_i - \lambda| + \epsilon},$$

where $\lambda \in [1/N, 3/N]$ controls the amount of shrinkage. The embedding, \hat{z} , is computed as a weighted sum of the memory values:

$$\hat{z} = \frac{\hat{\alpha}}{\|\hat{\alpha}\|_1} M,$$

and it is passed to a decoder f_d to obtain a reconstruction $\hat{x} = f_d(\hat{z})$. The anomaly score is

$$S = \|x - \hat{x}\|^2.$$

Deep Support Vector Data Description (Deep SVDD) [37] aims to construct a discrimination function over the data space that separates normal from anomalous points. Like in the original SVDD method [58], Deep SVDD seeks the smallest hypersphere that encloses “most” training data so that anomalous points fall outside the hypersphere:

$$\begin{aligned} \min_{R, c, \xi_1, \dots, \xi_n} \quad & R^2 + \frac{1}{vn} \sum_{i=1}^n \xi_i \\ \text{s.t.} \quad & \|\phi(x_i) - c\|^2 \leq R^2 + \xi_i, \quad \xi_i \geq 0. \end{aligned}$$

However, ϕ is now a representation learned from data through a deep model instead of a fixed representation induced by a kernel. The anomaly score for a new data point x is computed as the distance from the hypersphere:

$$S = \|\phi(x; w) - c\|^2.$$

We used VAE, MemAE, and Deep SVDD because they represent distinct methods for performing deep anomaly detection and obtaining

an anomaly score. Specifically, VAE and MemAE rely on the reconstruction error to determine the anomaly score, effectively measuring how well the model can reconstruct the input signal. In contrast, Deep SVDD uses the distance in the latent space to assign an anomaly score, focusing on how far a data point deviates from the center of a learned feature space hypersphere. This diversity in approaches ensures comprehensive coverage of anomaly detection scenarios, leveraging different mechanisms to identify anomalies effectively.

3.5. Training details

We trained each architecture using Adam as optimizer with a learning rate of 0.005 for a maximum number of epochs equal to 100, stopping training with an early stopping criterion with patience of 3, monitoring the corresponding loss on a subset of the train-set (20%) as the validation set. Autoencoder-like architectures are fully 1D convnets, with convolutional layers (kernel sizes of 3), which reduce the size of the input and encode it in a latent space of size 16, the encoder has 6 convolutional layers and the decoder has 6 upsampling layers with convolutional layers after each upsampling layer.

We tried different architecture for encoders and decoders by varying latent space size (8, 16, 32), kernel size of convolutional filters (2, 3, 5) and number of convolutional filters (8, 16, 32). All combinations produce strongly correlated results (average weighted Kendall’s τ of different architectures scores > 0.8) suggesting that the choice of a particular combination of these hyperparameters does not significantly affect the model’s behavior.

We are interested in detecting anomalies in a way that can be generalized to new patients. For this reason, all models were trained and tested in a leave-one-patient-out fashion. Results reported in Section 4 are aggregates for the left-out data.

3.6. Ranking agreements

In the following, the agreements between rankings are computed using a *weighted* version of the Kendall’s statistic, denoted τ_w [59]. The motivation for a weighted coefficient is that we are not particularly interested in the ranking differences for normal EGMs (low scores) but we aim at focusing on the differences when the anomaly scores or the traditional indicators are high¹. An established approach to define a weighting function that puts more emphasis to high scores was presented in Vigna [60], and we briefly recall it here for clarity.

Given two vectors of values R and S assigned by two different scoring functions to the same set of EGMs, the weighted coefficient is defined as

$$\tau_w(R, S) = \frac{\langle R, S \rangle_w}{\|R\|_w \|S\|_w}, \quad (4)$$

where the weighted inner product is

$$\langle R, S \rangle_w = \sum_{i < j} w(i, j) \operatorname{sgn}(r_i - r_j) \operatorname{sgn}(s_i - s_j),$$

and its associated norm is $\|E\|_w = \sqrt{\langle R, R \rangle_w}$. Here the weights are defined as

$$w(i, j) = \frac{1}{\rho(r_i)} + \frac{1}{\rho(s_j)},$$

being ρ an integer-valued function returning the rank of its argument.

¹ Since voltage correlates negatively with normalcy, we changed signs before computing the correlation coefficients.

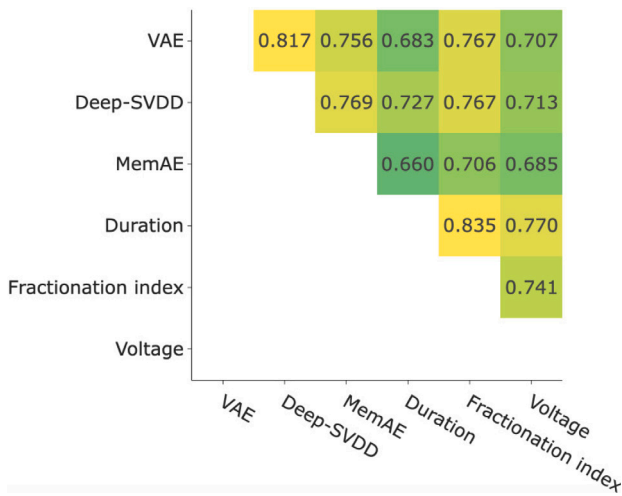


Fig. 3. Matrix of all τ_w statistics, values refer to all patients combined.

4. Results

In the following, we report comparisons among traditional indicators and anomaly detection scores, providing statistical analyses, electroanatomical visualizations, and showing, through examples, a detailed comparison of anomalies detected by anomaly scores and traditional indicators.

Both the fractionation and the duration indices were calculated following the OpenEP library [61] implementations as described in Sections 2.2 2.3. VAE and MemAE scores were calculated with our own implementations of the algorithms, while for Deep-SVDD we used the author's implementation available at <https://github.com/lukasruff/Deep-SVDD-PyTorch>.

4.1. Anomaly scores and indicators are all highly correlated

The weighted Kendall coefficients (4) among the three deep anomaly detection algorithms (see the upper-left triangle of Fig. 3) are all above 0.75. Although significance tests are not available for the weighted statistic τ_w , we computed p -values using the classic (unweighted) Kendall statistic τ , obtaining $p < 10^{-10}$ in all cases. This proves that, in spite of the differences in the architectures and in the algorithmic strategies used to identify outliers, all deep learning methods tend to be highly consistent. Agreements were also strong (with weighted coefficients above 0.74) for traditional indicators (lower-right triangle in Fig. 3, $p < 10^{-10}$ based on τ), and between traditional indicators and deep anomaly scores (top-right square in Fig. 3, $p < 10^{-10}$ based on τ).

The agreements between anomaly detectors and indicators were slightly higher when using deep-SVDD. This is reasonable since deep-SVDD learns a discriminating function based on *global* features of the EGMs, a trait which is certainly shared with the fractionation and the duration indices.

Algorithms based on reconstruction error, on the other hand, may just locally fail to reconstruct some short portions of EGMs and still produce a high score. To further investigate this aspect, we show in Fig. 4 a profile of the mean RMSE along temporal positions for signals having low Deep-SVDD anomaly scores (< 50 th percentile) and high reconstruction error (> 90 th percentile). Unsurprisingly, the reconstruction error is higher in the central region of the signal, where absolute voltage is typically also high.

4.2. Thresholding anomaly percentiles and indicator values produce consistent results

One difficulty when making decisions based on several indicators is the need to calibrate thresholds on each of them separately, and/or to decide a strategy to combine their values. For example, we might deem a certain EGM as non-physiological if *at least* one indicator takes on a value exceeding a certain threshold, but other options are available (e.g. if at least two indicators exceed their thresholds) and there is little guidance on how to select the proper criterion.

On the other hand, anomaly scores may be analyzed in terms of percentiles derived from the data. Based on the expected amount of non-physiological EGMs, we may simply put a threshold q on the percentile and highlight atrial locations whose score is above the q th percentile. This means choosing just one single threshold (instead of several) and there is no need to decide a combination strategy. To investigate if this procedure can detect outliers in a way that is coherent with indicators, we compute the average indicator values for EGMs whose anomaly score was above a certain percentile q :

$$\bar{v}(q) = \frac{\sum_i \mathbb{1}\{s_i > q\} v_i}{\sum_i \mathbb{1}\{s_i > q\}}$$

where $\mathbb{1}\{\cdot\}$ denotes the indicator function. We expect a monotonic trend for $\bar{v}(q)$ as q increases (increasing for the fractionation and the duration indices, decreasing for the voltage indicator). Results estimated in leave-one-patient-out mode are reported in Fig. 5 and perfectly match our expectation. To better characterize this relationship, we have shown in Fig. 6 three histograms with the distribution of different thresholds of fractionation index, duration, voltage and a combination of the three indicators as the anomaly score changes. As can be seen, the percentage of signals with high fractionation index increases as the anomaly percentile increases, reaching more than 50% of the signals. These histograms also confirm that voltage is the indicator with the least specificity compared to fractionation duration and the number of peaks, and the most complex to interpret. In the next section we investigate this aspect further.

4.3. Analysis of anomaly scores on low-voltage EGMs

Consistently with the common tenet in the literature that low-voltage EGMs are often associated with fibrotic tissue [7,62–64], we found that low-voltage EGMs largely produce high anomaly scores. We remind here that EGMs were normalized before feeding them into the anomaly detectors. This means that morphology of normalized signals alone contains sufficient information to deem anomaly.

In order to visually inspect morphologies, we plot in Fig. 7 the EGMs of all patients in four voltage ranges, colored by Deep-SVDD anomaly score percentile. The differentiation between low and high score signals is particularly evident when the voltage is below 0.1 mV or above 1mV. On the other hand, we observe a fairly large amount of EGMs with low anomaly in the borderline range 0.5–1mV, and to a lesser extent in the range 0.1–0.5 mV. It can be seen that the “consensus” morphology of these signals closely resembles the typical morphology of healthy tissue, confirming our starting assumption that most of the EGMs are physiological. In these cases, lower voltages might be simply due to other factors such as the direction of the propagation wave with respect to the bipolar electrodes, the angle of incidence of the catheter, or the contact strength with the atrial tissue [20]. These results show that anomaly detection approaches are more robust with respect to measurement errors, thanks to the ability to focus on the entire morphology and not on individual features.

In order to characterize differences with traditional indicators, all those signals with the anomaly score above 95th percentile colored by fractionation index can be seen in Fig. 8(a), and signals with fractionation index < 4 and high voltage are highlighted in Fig. 8(b).




Fig. 4. RMSE along temporal positions for VAE and MemAE.

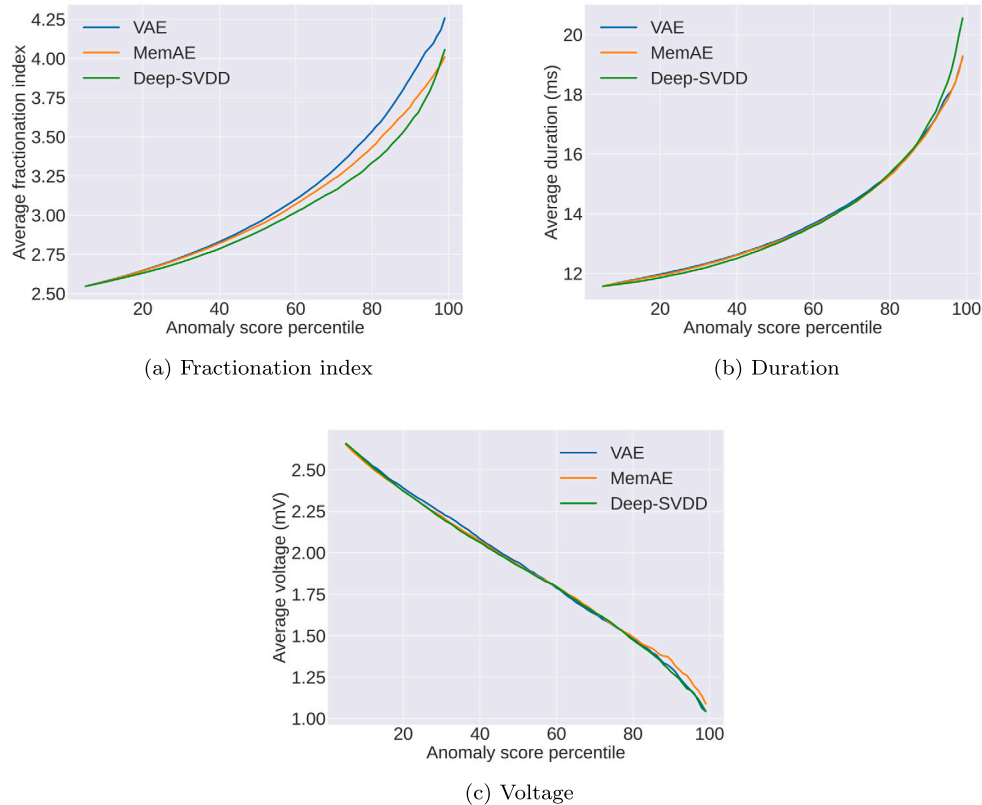


Fig. 5. Average values for traditional indicators change monotonically with the anomaly score.

The signature of this group of EGMs might be associated with some functional phenomena like pivot points.

Similarly, in Fig. 8(c) all signals with anomaly scores below the 50th percentile can be displayed, highlighting in Fig. 8(d) those with high fractionation index and low voltage. The signature of these EGMs is clearly biphasic with short duration, suggesting that the anomaly in low voltage probably arise from the positioning of the catheter.

4.4. Anatomical visualizations

In order to visualize anomalies in the anatomical context, we show in Fig. 9 and in Fig. 10 several maps of the left atrium of one of the patients in our dataset. Colors for the different models are associated with a normalized anomaly score: if s is the anomaly score from one of the models, color is proportional to $\frac{s}{q_{90}(S)}$ where $q_{\eta}(S)$ returns the η -th percentile of the whole set S of anomaly scores produced by that model over all EGMs for the patient being visualized.

In order to obtain a smooth reconstruction on the left atrium, we interpolated the points acquired during mapping onto a mesh of with a

finer granularity generated directly from CARTO3 system. Interpolation is performed with a weighted average of closest annotated points based on a Gaussian kernel with $\sigma = 1.5$.

Electroanatomical maps based on deep anomaly detection models seem to identify areas of electrical abnormalities with greater specificity, condensing the information contained in the state-of-the-art maps into a single map. In the reported case, voltage map characterize the whole LA as physiological, while duration map tends to overestimate the number of areas that are highlighted. Among the various anomaly detection methods, Deep-SVDD seems to provide the clearest characterization of the areas.

5. Discussion

The use of deep anomaly detection provides, in a completely unsupervised manner, a consistent and robust score that correlates with state-of-the-art biomarkers widely adopted in the medical literature, but without the need to manually combine them and decide on arbitrary thresholds. Our numerical results of Section 4.2 show that, on

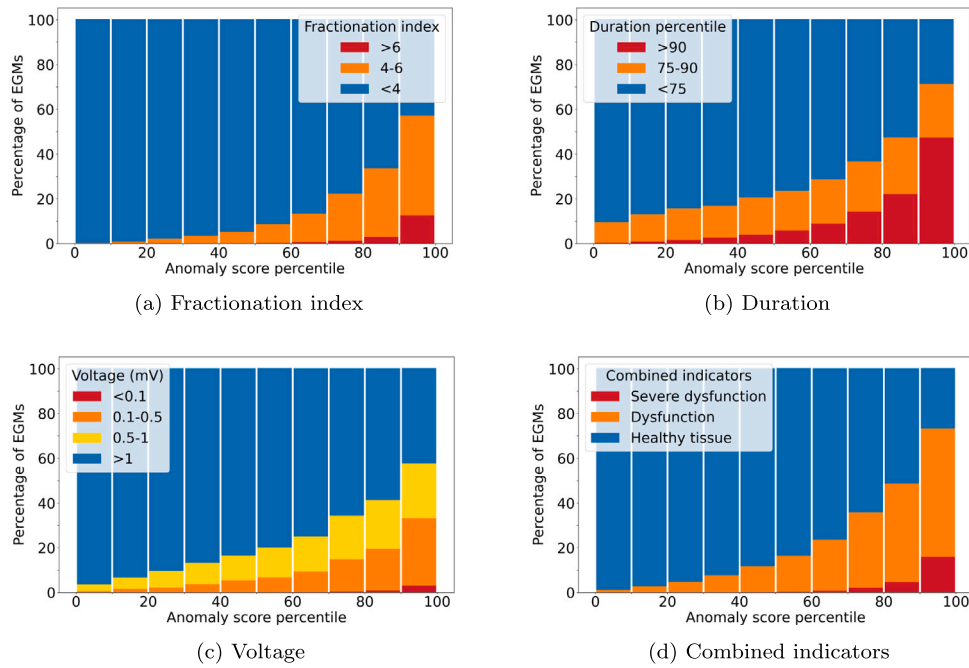


Fig. 6. In Fig. 6(a), the fractionation index was divided into three ranges, <4 , $4-6$, >6 . In Fig. 6(b), duration was divided according to its percentiles into three ranges <75 th percentile, 75 th– 90 th percentile and >90 th percentile. In Fig. 6(c), the voltage is divided into the four categories already described in Section 2.1. Fig. 6(d) put together all the indicators where a fractional index ≥ 4 , >90 th percentile of duration and a voltage <0.5 are considered as thresholds. Red (severe dysfunction) represents a signal that has all of them, orange (dysfunction) one or two and blue (healthy tissue) none of the three.

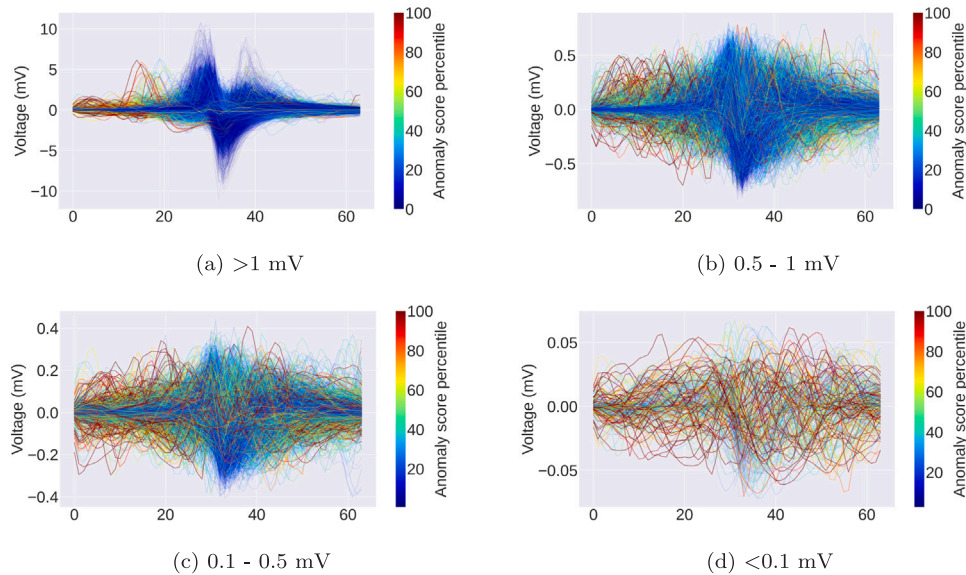


Fig. 7. EGMs of all patients colored by Deep-SVDD anomaly score percentile.

average, the higher the anomaly score is, the higher the fractionation (in terms of number of peaks and signal duration) and the lower the voltage.

The strong correlation (measured by weighted coefficients and un-weighted Kendall statistics) between the various deep models suggests good robustness on the ranking of the signals, as well as the strong correlation and monotonicity with the standard indicators.

A deeper analysis of EGM signatures reveals that voltage provides a poor subdivision of the signals, being extremely sensitive to measurement errors generated by a non-optimal positioning of the catheter. This generates false positives that actually correspond to low anomaly scores, as shown in Section 4.3. High anomaly score, on the other hand,

present pathological signatures, that can be associated to functional phenomena [65]. These results suggest that an anomaly score could be an *all-in-one* indicator replacing the standard ones. It, indeed, captures and fully exploits the morphology of the signal and is less sensitive to measurement errors.

As a step towards an effective clinical translation of this methodology, we construct electroanatomical maps of LA anomaly scores. These maps highlight areas characterized by electrical abnormalities through high anomaly scores. A visual inspection of these maps reveals that tissue stratification is more specific for deep anomaly detectors, and a cross-comparison of voltage, fractionation, and duration maps is required to obtain a similar results. This makes the reading of the deep anomaly electroanatomical map more immediate, which may enable to

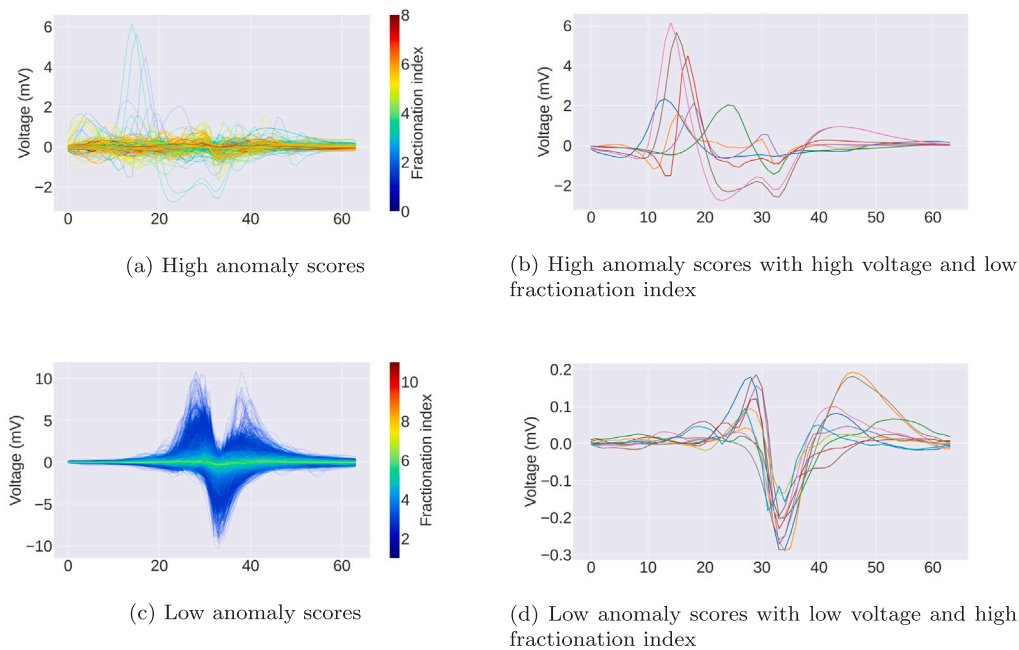


Fig. 8. Fig. 8(a) shows signals with high anomaly score (>95th percentile) colored by fractionation index. Fig. 8(b) shows only signals with fractionation index <4 and high voltage (> 1 mV). Fig. 8(c) shows signals with low anomaly score (<50th percentile) colored by fractionation index. Fig. 8(d) on the right shows only signals with fractionation index ≥ 4 and low voltage (<0.5 mV).

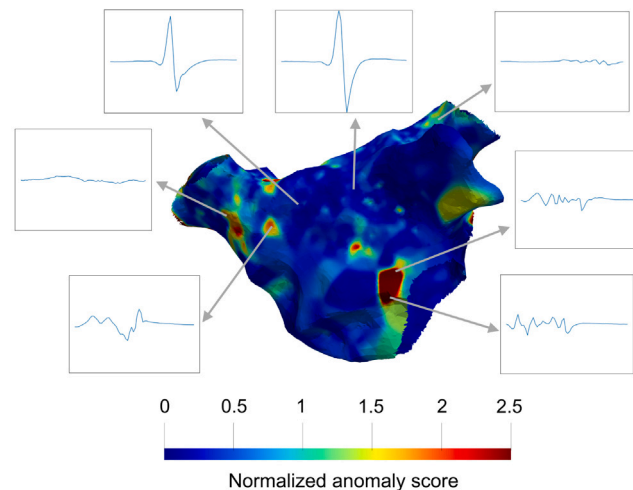


Fig. 9. Electroanatomical map showing the LA of a single patient colored according to the anomaly score of the Deep-SVDD method. Each signal is ± 2 mV range.

speed up the decision-making during ablation procedure and increase robustness with respect to signal acquisition problems.

Deep learning methods obtain results that correlate strongly with each other and with traditional indicators, as shown in Section 4.1, where Deep-SVDD obtains slightly higher correlations. Also, as explained in Section 4.1, the score produced by Deep-SVDD is based on *global* features of the signal unlike the other reconstruction methods that could have high scores even if they reconstruct poorly in one part of the signal.

All methods are extremely cost-effective, both in training and prediction, due to relatively small amount of parameters in the deep learning models (about 150k). Additionally, the average inference time is around 2 ms per batch (with a batch size of 512) on an NVIDIA RTX A6000 GPU and 10 ms per batch on an Intel Xeon Gold 6342 CPU. It is worth noting that variants with configurations of the autoencoder containing more filters or larger kernel sizes still achieve inference times in the order of tens of milliseconds.

Looking forward, areas of deep anomaly detection might correspond to specific pro-arrhythmic mechanisms which might improve the clinical understanding of atrial fibrillation. Moreover their analysis on large cohorts of patients may provide insights in describing atrial fibrillation progression and ablation strategies, adapting PVI lines to patient's characteristics. These aspects might improve patients' outcome as suggested by the promising results from the AEDUM study [66,67] based on EGM duration characterization and the ERASE-AF study [7] based on voltage characterization. In this direction, future computational and clinical studies will be addressed to investigate the potential link between deep anomaly electrical anomalies, electrical substrate and effective pro-arrhythmic behavior.

Moreover, automatic characterization of electrogram morphology, as recently shown in [68] for unipolar EGMs, may play a pivotal role in several electrophysiological studies, even those not finalized in constructing a complete electroanatomical maps, but only focusing on specific morphologies, like premature ventricular complexes ablation.

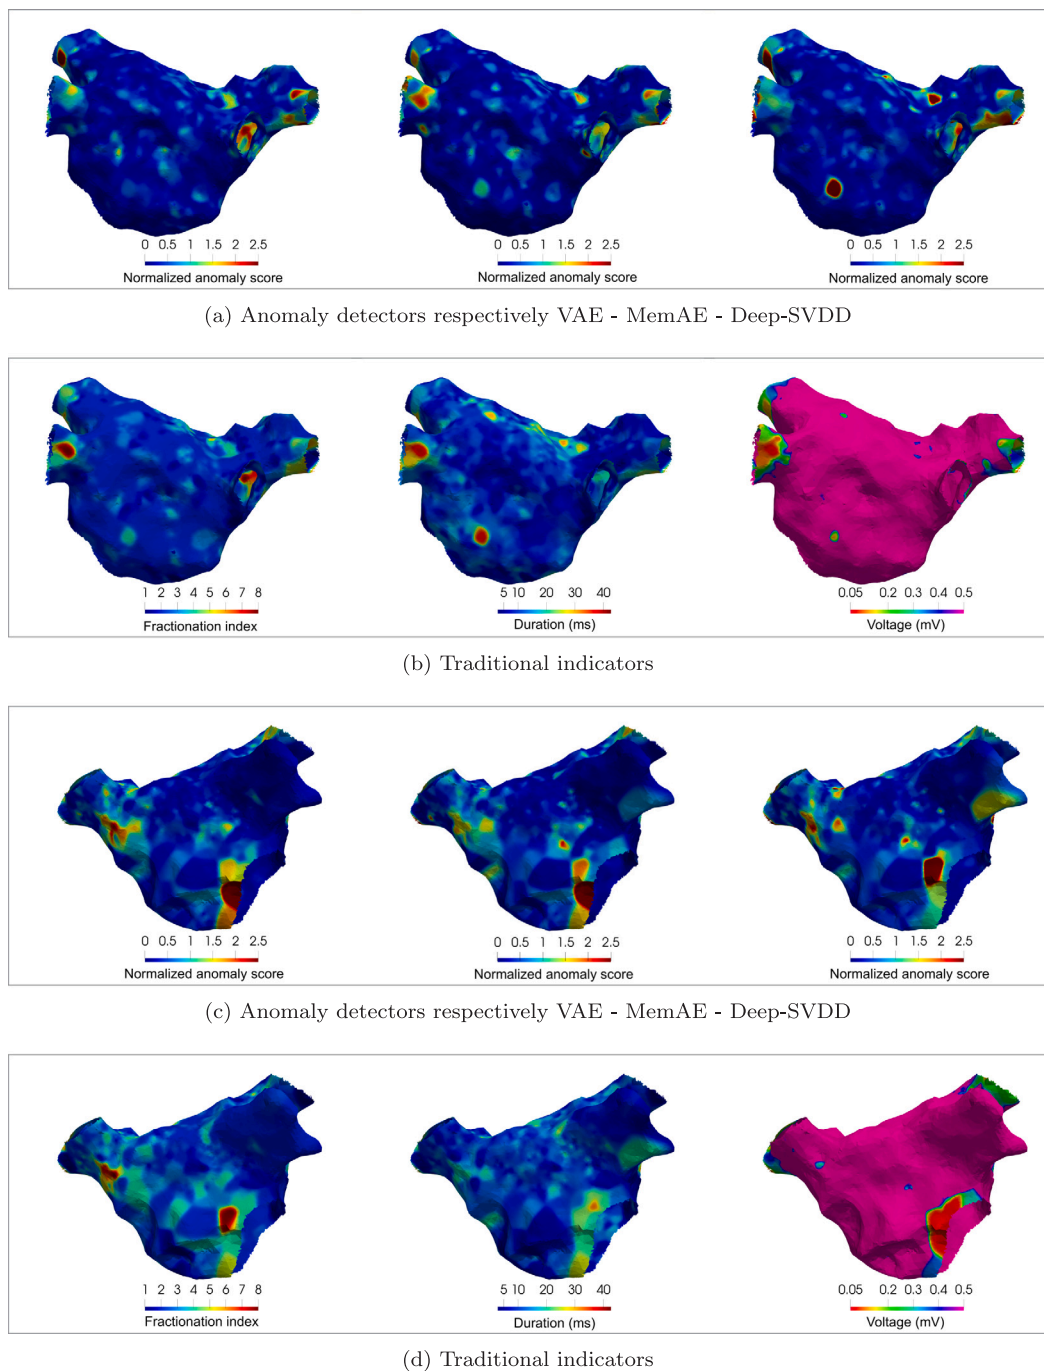


Fig. 10. Front and back view of the left atrium of a single patient colored according to the different indicators and anomaly scores of deep learning methods.

In conclusion, deep anomaly detection techniques provide novel and coherent indicators of electrical abnormalities in the EGM signature, outperforming state-of-the-art indicators.

CRediT authorship contribution statement

Luca Bindini: Writing – original draft, Writing – review & editing, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Conceptualization. **Stefano Pagani:** Writing – original draft, Writing – review & editing, Supervision, Methodology, Investigation, Formal analysis, Conceptualization. **Andrea Bernardini:** Writing – review & editing, Validation, Data curation, Resources, Conceptualization. **Benedetta Grossi:** Writing – review & editing, Visualization, Data curation. **Andrea Giomi:** Data curation, Resources.

Antonio Frontera: Validation, Supervision. **Paolo Frasconi:** Writing – original draft, Writing – review & editing, Validation, Supervision, Methodology, Investigation, Formal analysis, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Source code for the methods described in the paper will be released upon acceptance of the manuscript.

Acknowledgments

We thank Alessandro Paoletti Perini MD PhD for his valuable help in creating the dataset.

SP acknowledges the support by the MUR, Italian Ministry of University and Research, grant Dipartimento di Eccellenza 2023–2027. SP is a member of INdAM-GNCS and has been partially supported by the INdAM-GNCS Project 2023 “Mathematical models and numerical methods for computational medicine in cardiovascular diseases”.

Appendix A. Alternative tools

Activation maps are constructed by annotating local activation time (LAT), the relative (with respect to a given reference) timing at which the maximum negative EGMs variation is recorded. This map, depicts the temporal sequence of the times taken by the wavefront to cover the atrial tissue. It is specific to the repetitive patient’s rhythm (e.g., sinus or paced) during the acquisition. The LAT spatial distribution highlights abnormal electrical pathways like slow conduction corridors, whose severity and ramification characterizes the progression of atrial fibrillation [8]. Because they are measured at a specific rate and rhythm, activation maps provide only a partial view of the phenomenon, with knowledge gaps in areas with wavefront collisions and along directions different from the one of principal conduction. Multiple maps in different paced rhythms can bridge this gap while increasing the duration of the electrophysiological study [69]. Together with the activation map, substrate information derived from signal processing is added.

Appendix B. Enrollment details and ablative procedure

All patients underwent transthoracic echography at the admission in Hospital; a transesophageal or intracardiac echocardiography was performed in the procedural setting to exclude left auricular thrombosis and to guide the transeptal puncture. An electrical cardioversion was performed in patients who were not in sinus rhythm at the beginning of the CA procedure. All the ablation were performed under general anesthesia. Through right and left femoral veins, a deflectable decapolar catheter (Decanav, Biosense Webster, Diamond Bar, CA) was inserted in coronary sinus, a multipolar mapping catheter (Pentaray, Biosense Webster, Diamond Bar, CA), was used to map the left atria and a 3.5 mm open irrigated-tip catheter (Thermocool SF, Biosense Webster, Diamond Bar, CA) was chosen to perform radiofrequency ablation. To aid ablation catheter manipulation, a steerable long sheath (Agilis, Abbott, Chicago, IL) was used. After the transeptal puncture to reach the left atria, heparin (100 IU/kg) was administered to reach an activated clotting time of >300 s.

References

[1] G. Hindricks, T. Potpara, N. Dagres, E. Arbelo, J.J. Bax, C. Blomström-Lundqvist, G. Boriani, M. Castella, G.A. Dan, P.E. Dilaveris, L. Fauchier, G. Filippatos, J.M. Kalman, M. La Meir, D.A. Lane, J.P. Lebeau, M. Lettino, G.Y.H. Lip, F.J. Pinto, G.N. Thomas, M. Valgimigli, I.C. Van Gelder, B.P. Van Putte, C.L. Watkins, ESC Scientific Document Group, P. Kirchhof, M. Kühne, V. Aboyans, A. Ahlsson, P. Balsam, J. Bauersachs, S. Benussi, A. Brandes, F. Braunschweig, A.J. Camm, D. Capodanno, B. Casadei, D. Conen, H.J.G.M. Crijns, V. Delgado, D. Dobrev, H. Drexler, L. Eckardt, D. Fitzsimons, T. Folliguet, C.P. Gale, B. Gorenek, K.G. Haessler, H. Heidbuchel, B. Jung, H.A. Katus, D. Kotecha, U. Landmesser, C. Leclercq, B.S. Lewis, J. Mascherbauer, J.L. Merino, B. Merkely, L. Mont, C. Mueller, K.V. Nagy, J. Oldgren, N. Pavlović, R.F.E. Pedretti, S.E. Petersen, J.P. Piccini, B.A. Popescu, H. Pürerfellner, D.J. Richter, M. Roffi, A. Rubboli, D. Scherr, R.B. Schnabel, I.A. Simpson, E. Shlyakhto, M.F. Sinner, J. Steffel, M. Sousa-Uva, P. Suwalaki, M. Svetlosak, R.M. Touyz, N. Dagres, E. Arbelo, J.J. Bax, C. Blomström-Lundqvist, G. Boriani, M. Castella, G.A. Dan, P.E. Dilaveris, L. Fauchier, G. Filippatos, J.M. Kalman, M. La Meir, D.A. Lane, J.P. Lebeau, M. Lettino, G.Y.H. Lip, F.J. Pinto, G. Neil Thomas, M. Valgimigli, I.C. Van Gelder, C.L. Watkins, T. Delassi, H.S. Sisakian, D. Scherr, A. Chasnoits, M.D. Pauw, E. Smajić, T. Shalhanov, P. Avraamides, J. Kautzner, C. Gerdes, A.A. Alalaz, P. Kampus, P. Raatikainen, S. Boveda, G. Papiashvili, L. Eckardt, V. Vassilikos,

Z. Csanádi, D.O. Arnar, J. Galvin, A. Barsheshet, P. Caldarola, A. Rakisheva, I. Bytyçi, A. Kerimkulova, O. Kalejs, M. Njeim, A. Puodziukynas, L. Groben, M.A. Sammut, A. Grosu, A. Boskovic, A. Moustaghfir, N.D. Groot, L. Poposka, O.G. Anfinsen, P.P. Mitkowski, D.M. Cavaco, C. Siliste, E.N. Mikhaylov, L. Bertelli, D. Kojic, R. Hatala, Z. Fras, F. Arribas, T. Juhlin, C. Sticherling, L. Abid, I. Atar, O. Sychov, M.G.D. Bates, N.U. Zakirov, 2020 ESC Guidelines for the diagnosis and management of atrial fibrillation developed in collaboration with the European association for Cardio-Thoracic Surgery (EACTS), *Eur. Heart J.* 42 (5) (2021) 373–498, <http://dx.doi.org/10.1093/eurheartj/ehaa612>, URL: <https://academic.oup.com/eurheartj/article/42/5/373/5899003>.

[2] M. Haïssaguerre, P. Jais, D.C. Shah, A. Takahashi, M. Hocini, G. Quiniou, S. Garrigue, A. Le Mouroux, P. Le Métayer, J. Clémenty, Spontaneous initiation of atrial fibrillation by ectopic beats originating in the pulmonary veins, *N. Engl. J. Med.* 339 (10) (1998) 659–666, <http://dx.doi.org/10.1056/NEJM199809033391003>.

[3] J.P. Piccini, R.D. Lopes, M.H. Kong, V. Hasselblad, K. Jackson, S.M. Al-Khatib, Pulmonary vein isolation for the maintenance of sinus rhythm in patients with atrial fibrillation: A meta-analysis of randomized, controlled trials, *Circ. Arrhythm. Electrophysiol.* 2 (6) (2009) 626–633, <http://dx.doi.org/10.1161/CIRCEP.109.856633>, URL: <https://www.ahajournals.org/doi/10.1161/CIRCEP.109.856633>.

[4] F.D. Ramirez, D.H. Birnie, G.M. Nair, A. Szczotka, C.J. Redpath, M.M. Sadek, P.B. Nery, Efficacy and safety of driver-guided catheter ablation for atrial fibrillation: A systematic review and meta-analysis, *J. Cardiovasc. Electrophysiol.* 28 (12) (2017) 1371–1378, <http://dx.doi.org/10.1111/jce.13313>, URL: <https://onlinelibrary.wiley.com/doi/10.1111/jce.13313>.

[5] A.C. Perino, G.C. Leef, A. Cluckey, F.N. Yunus, M. Askari, P.A. Heidenreich, S.M. Narayan, P.J. Wang, M.P. Turakhia, Secular trends in success rate of catheter ablation for atrial fibrillation: The SMASH-AF cohort, *Am. Heart J.* 208 (2019) 110–119, <http://dx.doi.org/10.1016/j.ahj.2018.10.006>, URL: <https://linkinghub.elsevier.com/retrieve/pii/S0002870318303028>.

[6] R. Parameswaran, A.M. Al-Kaisey, J.M. Kalman, Catheter ablation for atrial fibrillation: current indications and evolving technologies, *Nat. Rev. Cardiol.* 18 (3) (2021) 210–225, <http://dx.doi.org/10.1038/s41569-020-00451-x>, URL: <https://www.nature.com/articles/s41569-020-00451-x>.

[7] Y. Huo, T. Gaspar, R. Schönbauer, M. Wójcik, L. Fiedler, F.X. Roithinger, M. Martinek, H. Pürerfellner, B. Kirstein, U. Richter, et al., Low-voltage myocardium-guided ablation trial of persistent atrial fibrillation, *NEJM Evid.* 1 (11) (2022) EVIDo2200141.

[8] A. Frontera, S. Pagani, L.R. Limite, A. Peirone, F. Fioravanti, B. Enache, J. Cuellar Silva, K. Vlachos, C. Meyer, G. Montesano, A. Manzoni, L. Dedé, A. Quarteroni, D.G. Lațcu, P. Rossi, P. Della Bella, Slow conduction corridors and pivot sites characterize the electrical remodeling in atrial fibrillation, *JACC Clin. Electrophysiol.* 8 (5) (2022) 561–577, <http://dx.doi.org/10.1016/j.jacep.2022.01.019>, URL: <https://linkinghub.elsevier.com/retrieve/pii/S2405500X2200127X>.

[9] M.S. van Schie, N.M. de Groot, Clinical relevance of sinus rhythm mapping to quantify electrophathology related to atrial fibrillation, *Arrhythm. Electrophysiol. Rev.* 11 (2022).

[10] C. Dallet, C. Roney, R. Martin, T. Kitamura, S. Puyo, J. Duchâteau, C. Dumas-Pomier, G. Ravon, L. Bear, N. Derval, et al., Cardiac propagation pattern mapping with vector field for helping tachyarrhythmias diagnosis with clinical tridimensional electro-anatomical mapping tools, *IEEE Trans. Biomed. Eng.* 66 (2) (2018) 373–382.

[11] A.S. Jadidi, E. Duncan, S. Miyazaki, N. Lellouche, A.J. Shah, A. Forclaz, I. Nault, M. Wright, L. Rivard, X. Liu, et al., Functional nature of electrogram fractionation demonstrated by left atrial high-density mapping, *Circ. Arrhythm. Electrophysiol.* 5 (1) (2012) 32–42.

[12] K. Konings, C. Kirchhof, J. Smeets, H. Wellens, O.C. Penn, M.A. Allesie, High-density mapping of electrically induced atrial fibrillation in humans., *Circulation* 89 (4) (1994) 1665–1680.

[13] F. Atienza, J. Almendral, J. Jalife, S. Zlochiver, R. Ploutz-Snyder, E.G. Torrecilla, A. Arenal, J. Kalifa, F. Fernández-Avilés, O. Berenfeld, Real-time dominant frequency mapping and ablation of dominant frequency sites in atrial fibrillation with left-to-right frequency gradients predicts long-term maintenance of sinus rhythm, *Heart Rhythm* 6 (1) (2009) 33–40.

[14] S.M. Narayan, D.E. Krummen, K. Shivkumar, P. Clopton, W.J. Rappel, J.M. Miller, Treatment of atrial fibrillation by the ablation of localized sources: CONFIRM (conventional ablation for atrial fibrillation with or without focal impulse and rotor modulation) trial, *J. Am. Coll. Cardiol.* 60 (7) (2012) 628–636.

[15] S.M. Narayan, T. Baykaner, P. Clopton, A. Schricker, G.G. Lalani, D.E. Krummen, K. Shivkumar, J.M. Miller, Ablation of rotor and focal sources reduces late recurrence of atrial fibrillation compared with trigger ablation alone: extended follow-up of the CONFIRM trial (conventional ablation for atrial fibrillation with or without focal impulse and rotor modulation), *J. Am. Coll. Cardiol.* 63 (17) (2014) 1761–1768.

- [16] R. Dubois, A.J. Shah, M. Hocini, A. Denis, N. Derval, H. Cochet, F. Sacher, L. Bear, J. Duchateau, P. Jais, et al., Non-invasive cardiac mapping in clinical practice: Application to the ablation of cardiac arrhythmias, *J. Electrocardiol.* 48 (6) (2015) 966–974.
- [17] P.M. Boyle, T. Zghaib, S. Zahid, R.L. Ali, D. Deng, W.H. Franceschi, J.B. Hakim, M.J. Murphy, A. Prakosa, S.L. Zimmerman, et al., Computationally guided personalized targeted ablation of persistent atrial fibrillation, *Nat. Biomed. Eng.* 3 (11) (2019) 870–879.
- [18] S. Zahid, H. Cochet, P.M. Boyle, E.L. Schwarz, K.N. Whyte, E.J. Vigmond, R. Dubois, M. Hocini, M. Haissaguerre, P. Jais, et al., Patient-derived models link re-entrant driver localization in atrial fibrillation to fibrosis spatial pattern, *Cardiovasc. Res.* 110 (3) (2016) 443–454.
- [19] J. Zhao, B.J. Hansen, Y. Wang, T.A. Csepe, L.V. Sul, A. Tang, Y. Yuan, N. Li, A. Bratasz, K.A. Powell, et al., Three-dimensional integrated functional, structural, and computational mapping to define the structural “fingerprints” of heart-specific atrial fibrillation drivers in human heart ex vivo, *J. Am. Heart Assoc.* 6 (8) (2017) e005922.
- [20] E. Anter, M.E. Josephson, Bipolar voltage amplitude: what does it really mean? *Heart Rhythm* 13 (1) (2016) 326–327.
- [21] A. Frontera, L.R. Limite, S. Pagani, A. Hadjis, M. Cireddu, S. Sala, G. Tsitsinakis, G. Paglino, G. Peretto, F. Lipartiti, et al., Characterization of cardiac electrogram signals in atrial arrhythmias, *Min. Cardiol. Angiol.* 69 (1) (2021) 70–80.
- [22] C. Mendonca Costa, G.C. Anderson, V.M. Meijborg, C. O’Shea, M.J. Shattock, P. Kirchhof, R. Coronel, S. Niederer, D. Pavlovic, T. Dhanjal, et al., The amplitude-normalized area of a bipolar electrogram as a measure of local conduction delay in the heart, *Front. Physiol.* 11 (2020) 465.
- [23] B. Schölkopf, J.C. Platt, J. Shawe-Taylor, A.J. Smola, R.C. Williamson, Estimating the support of a high-dimensional distribution, *Neural Comput.* 13 (7) (2001) 1443–1471, <http://dx.doi.org/10.1162/089976601750264965>.
- [24] M. Markou, S. Singh, Novelty detection: A review—Part 2: Neural network based approaches, *Signal Process.* 83 (12) (2003) 2499–2521, <http://dx.doi.org/10.1016/j.sigpro.2003.07.019>, URL: <http://linkinghub.elsevier.com/retrieve/pii/S0165168403002032>.
- [25] V. Chandola, A. Banerjee, V. Kumar, Anomaly detection: A survey, *ACM Comput. Surv.* 41 (3) (2009) 15:1–15:58, <http://dx.doi.org/10.1145/1541880.1541882>.
- [26] R. Chalapathy, S. Chawla, Deep learning for anomaly detection: A survey, 2019, <http://dx.doi.org/10.48550/arXiv.1901.03407>, URL: <http://arxiv.org/abs/1901.03407>.
- [27] J. Liu, G. Xie, J. Wang, S. Li, C. Wang, F. Zheng, Y. Jin, Deep industrial image anomaly detection: A survey, *Mach. Intell. Res.* 21 (1) (2024) 104–135, <http://dx.doi.org/10.1007/s11633-023-1459-z>.
- [28] X. Zhou, W. Liang, S. Shimizu, J. Ma, Q. Jin, Siamese neural network based few-shot learning for anomaly detection in industrial cyber-physical systems, *IEEE Trans. Ind. Inform.* 17 (8) (2021) 5790–5798, <http://dx.doi.org/10.1109/TII.2020.3047675>.
- [29] A. Kasencas, P. Sanchez, P. Schrempf, C. Wang, W. Clackett, S.S. Mikhael, J.P. Voisey, K. Goatman, A. Weir, N. Pugeault, S.A. Tsaftaris, A.O. O’Neil, The role of noise in denoising models for anomaly detection in medical images, *Med. Image Anal.* 90 (2023) 102963, <http://dx.doi.org/10.1016/j.media.2023.102963>, URL: <https://www.sciencedirect.com/science/article/pii/S1361841523002232>.
- [30] J. Wolleb, F. Bieder, R. Sandkühler, P.C. Cattin, Diffusion models for medical anomaly detection, in: L. Wang, Q. Dou, P.T. Fletcher, S. Speidel, S. Li (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*, in: *Lecture Notes in Computer Science*, Springer Nature Switzerland, Cham, 2022, pp. 35–45, http://dx.doi.org/10.1007/978-3-031-16452-1_4.
- [31] W.H.L. Pinaya, M.S. Grahame, R. Gray, P.F. da Costa, P.D. Tudosiu, P. Wright, Y.H. Mah, A.D. MacKinnon, J.T. Teo, R. Jager, D. Werring, G. Rees, P. Nachev, S. Ourselin, M.J. Cardoso, Fast unsupervised brain anomaly detection and segmentation with diffusion models, in: L. Wang, Q. Dou, P.T. Fletcher, S. Speidel, S. Li (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*, in: *Lecture Notes in Computer Science*, Springer Nature Switzerland, Cham, 2022, pp. 705–714, http://dx.doi.org/10.1007/978-3-031-16452-1_67.
- [32] M. Gu, Y. Zhang, Y. Wen, G. Ai, H. Zhang, P. Wang, G. Wang, A lightweight convolutional neural network hardware implementation for wearable heart rate anomaly detection, *Comput. Biol. Med.* 155 (2023) 106623, <http://dx.doi.org/10.1016/j.combiomed.2023.106623>.
- [33] Y. Tian, G. Maicas, L.Z.C.T. Pu, R. Singh, J.W. Verjans, G. Carneiro, Few-shot anomaly detection for polyp frames from colonoscopy, in: A.L. Martel, P. Abolmaesumi, D. Stoyanov, D. Mateus, M.A. Zuluaga, S.K. Zhou, D. Racoceanu, L. Joskowicz (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*, in: *Lecture Notes in Computer Science*, Springer International Publishing, Cham, 2020, pp. 274–284, http://dx.doi.org/10.1007/978-3-030-59725-2_27.
- [34] S. Liao, D. Ragot, S. Nayyar, A. Suszko, Z. Zhang, B. Wang, V.S. Chauhan, Deep learning classification of unipolar electrograms in human atrial fibrillation: Application in focal source mapping, *Front. Physiol.* 12 (2021) URL: <https://www.frontiersin.org/articles/10.3389/fphys.2021.704122>.
- [35] J. An, S. Cho, Variational Autoencoder Based Anomaly Detection Using Reconstruction Probability, Technical Report, SNU Data Mining Center, 2015, URL: <https://www.semanticscholar.org/paper/Variational-Autoencoder-based-Anomaly-Detection-An-Cho/061146b1d7938d7a8dae70e3531a00fceb3c78e8>.
- [36] D. Gong, L. Liu, V. Le, B. Saha, M.R. Mansour, S. Venkatesh, A. Van Den Hengel, Memorizing normality to detect anomaly: Memory-augmented deep autoencoder for unsupervised anomaly detection, in: 2019 IEEE/CVF International Conference on Computer Vision (ICCV), IEEE, Seoul, Korea (South), 2019, pp. 1705–1714, <http://dx.doi.org/10.1109/ICCV.2019.00179>, URL: <https://ieeexplore.ieee.org/document/9010977/>.
- [37] L. Ruff, R. Vandermeulen, N. Goernitz, L. Deecke, S.A. Siddiqui, A. Binder, E. Müller, M. Kloft, Deep one-class classification, in: *Proceedings of the 35th International Conference on Machine Learning*, 2018, pp. 4393–4402, URL: <https://proceedings.mlr.press/v80/ruff18a.html>.
- [38] S.E. Williams, C.H. Roney, A. Connolly, I. Sim, J. Whitaker, D. O’Hare, I. Kotadia, L. O’Neill, C. Corrado, M. Bishop, S.A. Niederer, M. Wright, M. O’Neill, N.W. Linton, OpenEP: A cross-platform electroanatomic mapping data format and analysis platform for electrophysiology research, *Front. Physiol.* 12 (2021) 646023, <http://dx.doi.org/10.3389/fphys.2021.646023>.
- [39] I. Sim, M. Bishop, M. O’Neill, S.E. Williams, Left atrial voltage mapping: defining and targeting the atrial fibrillation substrate, *J. Interv. Cardiac Electrophysiol.* 56 (2019) 213–227.
- [40] G. La Rosa, J.G. Quintanilla, R. Salgado, J.J. González-Ferrer, V. Cañadas-Godoy, J. Pérez-Villacastín, J. Jalife, N. Pérez-Castellano, D. Filgueiras-Rama, Anatomical targets and expected outcomes of catheter-based ablation of atrial fibrillation in 2020, *Pacing Clin. Electrophysiol.* 44 (2) (2021) 341–359, <http://dx.doi.org/10.1111/pace.14140>, URL: <https://onlinelibrary.wiley.com/doi/10.1111/pace.14140>.
- [41] G.R. Wong, C.J. Nalliah, G. Lee, A. Voskoboinik, S. Prabhu, R. Parameswaran, H. Sugumar, R.D. Anderson, A. McLellan, L.H. Ling, et al., Dynamic atrial substrate during high-density mapping of paroxysmal and persistent AF: implications for substrate ablation, *JACC Clin. Electrophysiol.* 5 (11) (2019) 1265–1277.
- [42] A. Frontera, L.R. Limite, S. Pagani, M. Cireddu, K. Vlachos, C. Martin, M. Takigawa, T. Kitamura, F. Bourier, G. Cheniti, T. Pambrun, F. Sacher, N. Derval, M. Hocini, A. Quarteroni, P. Della Bella, M. Haissaguerre, P. Jais, Electrogram fractionation during sinus rhythm occurs in normal voltage atrial tissue in patients with atrial fibrillation, *Pacing Clin. Electrophysiol.* 45 (2) (2022) 219–228, <http://dx.doi.org/10.1111/pace.14425>, URL: <https://onlinelibrary.wiley.com/doi/10.1111/pace.14425>.
- [43] K. Konings, J. Smeets, O.C. Penn, H. Wellens, M.A. Allesie, Configuration of unipolar atrial electrograms during electrically induced atrial fibrillation in humans, *Circulation* 95 (5) (1997) 1231–1241, URL: <https://www.ahajournals.org/doi/10.1161/01.CIR.95.5.1231>.
- [44] K. Nademanee, J. McKenzie, E. Kosar, M. Schwab, B. Sunsaneewitayakul, T. Vasavakul, C. Khunnawat, T. Ngarmukos, A new approach for catheter ablation of atrial fibrillation: Mapping of the electrophysiologic substrate, *J. Am. Coll. Cardiol.* 43 (11) (2004) 2044–2053, <http://dx.doi.org/10.1016/j.jacc.2003.12.054>, URL: <https://linkinghub.elsevier.com/retrieve/pii/S0735109704005716>.
- [45] D.H. Lau, B. Maesen, S. Zeemering, P. Kuklik, A. van Hünnik, T.A.R. Lankveld, E. Bidar, S. Verheule, J. Nijs, J. Maessen, H. Crijns, P. Sanders, U. Schotten, Indices of bipolar complex fractionated atrial electrograms correlate poorly with each other and atrial fibrillation substrate complexity, *Heart Rhythm* 12 (7) (2015) 1415–1423, <http://dx.doi.org/10.1016/j.hrthm.2015.03.017>.
- [46] S. Conti, A. Verma, Ablation of complex fractionated electrograms may not be useful for catheter ablation of persistent atrial fibrillation, *Heart Rhythm* 13 (10) (2016) 2101–2103, <http://dx.doi.org/10.1016/j.hrthm.2016.06.037>.
- [47] J. Vogler, S. Willems, A. Sultan, D. Schreiber, J. Lüker, H. Servatius, B. Schäffer, J. Moser, B.A. Hoffmann, D. Steven, Pulmonary vein isolation versus defragmentation: the CHASE-AF clinical trial, *J. Am. Coll. Cardiol.* 66 (24) (2015) 2743–2752.
- [48] E. Buch, M. Share, R. Tung, P. Benharash, P. Sharma, J. Koneru, R. Mandapati, K.A. Ellenbogen, K. Shivkumar, Long-term clinical outcomes of focal impulse and rotor modulation for treatment of atrial fibrillation: A multicenter experience, *Heart Rhythm* 13 (3) (2016) 636–641.
- [49] Y. Waranugraha, A. Rizal, D. Setiawan, L.J. Aziz, Additional complex fractionated atrial electrogram ablation does not improve the outcomes of non-paroxysmal atrial fibrillation: A systematic review and meta-analysis of randomized controlled trials, *Indian Heart J.* 73 (1) (2021) 63–73, <http://dx.doi.org/10.1016/j.ihj.2020.11.004>, URL: <https://www.sciencedirect.com/science/article/pii/S0019483220302704>.
- [50] E. Vigmond, A. Pashaei, S. Amraoui, H. Cochet, M. Haissaguerre, Percolation as a mechanism to explain atrial fractionated electrograms and reentry in a fibrosis model based on imaging data, *Heart Rhythm* 13 (7) (2016) 1536–1543.
- [51] L.J.M.E. van der Does, N.M.S. de Groot, Inhomogeneity and complexity in defining fractionated electrograms, *Heart Rhythm* 14 (4) (2017) 616–624, <http://dx.doi.org/10.1016/j.hrthm.2017.01.021>, URL: <https://www.sciencedirect.com/science/article/pii/S1547527117300681>.

- [52] A. Frontera, R. Mahajan, C. Dallet, K. Vlachos, T. Kitamura, M. Takigawa, G. Cheniti, C. Martin, J. Duchateau, A. Lam, et al., Characterizing localized reentry with high-resolution mapping: evidence for multiple slow conducting isthmuses within the circuit, *Heart Rhythm* 16 (5) (2019) 679–685.
- [53] S. Pagani, L. Dede', A. Frontera, M. Salvador, L. Limite, A. Manzoni, F. Lipartiti, G. Tsitsinakis, A. Hadjis, P. Della Bella, A. Quarteroni, A computational study of the electrophysiological substrate in patients suffering from atrial fibrillation, *Front. Physiol.* 12 (2021) <http://dx.doi.org/10.3389/fphys.2021.673612>, URL: <https://www.frontiersin.org/articles/10.3389/fphys.2021.673612>.
- [54] P. Rossi, F.M. Cauti, M. Niscola, M. Calore, V. Fanti, M. Polselli, A. Di Pastena, L. Iaia, S. Bianchi, A novel ventricular map of electrograms duration as a method to identify areas of slow conduction for ventricular tachycardia ablation: The VEDUM pilot study, *Heart Rhythm* 18 (8) (2021) 1253–1260.
- [55] P. Rossi, F.M. Cauti, M. Niscola, M. Magnocavallo, M. Polselli, S. Capone, D.G. Della Rocca, J. Rodriguez-Garrido, G. Piccirillo, I. Anguera, et al., Ventricular electrograms duration map to detect ventricular arrhythmia substrate: the VEDUM project study, *Circ. Arrhythm. Electrophysiol.* 16 (8) (2023) 447–455.
- [56] G. Pang, C. Shen, L. Cao, A.V.D. Hengel, Deep learning for anomaly detection: A review, *ACM Comput. Surv.* 54 (2) (2021) 38:1–38:38, <http://dx.doi.org/10.1145/3439950>, 00650.
- [57] D.P. Kingma, M. Welling, Auto-encoding variational Bayes, in: Y. Bengio, Y. LeCun (Eds.), 2nd International Conference on Learning Representations, 2014, URL: <http://arxiv.org/abs/1312.6114>.
- [58] D. Tax, R. Duin, Support vector data description, *Mach. Learn.* 54 (1) (2004) 45–66.
- [59] G.S. Shieh, A weighted Kendall's tau statistic, *Statist. Probab. Lett.* 39 (1) (1998) 17–24, [http://dx.doi.org/10.1016/S0167-7152\(98\)00006-6](http://dx.doi.org/10.1016/S0167-7152(98)00006-6), URL: <https://www.sciencedirect.com/science/article/pii/S0167715298000066>.
- [60] S. Vigna, A weighted correlation index for rankings with ties, in: Proceedings of the 24th International Conference on World Wide Web, WWW '15, 2015, pp. 1166–1176, <http://dx.doi.org/10.1145/2736277.2741088>.
- [61] S.E. Williams, C.H. Roney, A. Connolly, I. Sim, J. Whitaker, D. O'Hare, I. Kotadia, L. O'Neill, C. Corrado, M. Bishop, S.A. Niederer, M. Wright, M. O'Neill, N.W. Linton, OpenEP: A cross-platform electroanatomic mapping data format and analysis platform for electrophysiology research, *Front. Physiol.* 12 (2021) 646023, <http://dx.doi.org/10.3389/fphys.2021.646023>, URL: <https://www.frontiersin.org/article/10.3389/fphys.2021.646023>.
- [62] S. Rolf, S. Kircher, A. Arya, C. Eitel, P. Sommer, S. Richter, T. Gaspar, A. Bollmann, D. Altmann, C. Piedra, G. Hindricks, C. Piorkowski, Tailored atrial substrate modification based on low-voltage areas in catheter ablation of atrial fibrillation, *Circ. Arrhythm. Electrophysiol.* 7 (5) (2014) 825–833, <http://dx.doi.org/10.1161/CIRCEP.113.001251>, URL: <https://www.ahajournals.org/doi/10.1161/CIRCEP.113.001251>.
- [63] Z. Ling, J. McManigle, V. Zipunnikov, F. Pashakhanloo, I.M. Khurram, S.L. Zimmerman, B. Philips, J.E. Marine, D.D. Spragg, H. Ashikaga, H. Calkins, S. Nazarian, The association of left atrial low-voltage regions on electroanatomic mapping with Low Attenuation Regions on cardiac computed tomography perfusion imaging in patients with atrial fibrillation, *Heart Rhythm* 12 (5) (2015) 857–864, <http://dx.doi.org/10.1016/j.hrthm.2015.01.015>, URL: <https://www.sciencedirect.com/science/article/pii/S1547527115000442>.
- [64] A.S. Jadidi, H. Lehrmann, C. Keyl, J. Sorrel, V. Markstein, J. Minners, C.I. Park, A. Denis, P. Jaïs, M. Hocini, C. Potocnik, J. Allgeier, W. Hochholzer, C. Herrera-Siklody, S. Kim, Y.E. Omri, F.J. Neumann, R. Weber, M. Haïssaguerre, T. Arentz, Ablation of persistent atrial fibrillation targeting low-voltage areas with selective activation characteristics, *Circ. Arrhythm. Electrophysiol.* 9 (3) (2016) e002962, <http://dx.doi.org/10.1161/CIRCEP.115.002962>, URL: <https://www.ahajournals.org/doi/10.1161/CIRCEP.115.002962>.
- [65] A. Frontera, M. Takigawa, R. Martin, N. Thompson, G. Cheniti, G. Massoulié, J. Duchateau, J.Y. Wielandts, E. Teijeira, T. Kitamura, M. Wolf, N. Al-Jefairi, K. Vlachos, S. Yamashita, S. Amraoui, A. Denis, M. Hocini, H. Cochet, F. Sacher, P. Jaïs, M. Haïssaguerre, N. Derval, Electrogram signature of specific activation patterns: Analysis of atrial tachycardias at high-density endocardial mapping, *Heart Rhythm* 15 (1) (2018) 28–37, <http://dx.doi.org/10.1016/j.hrthm.2017.08.001>, URL: <https://linkinghub.elsevier.com/retrieve/pii/S1547527117309621>.
- [66] P. Rossi, F.M. Cauti, M. Polselli, M. Magnocavallo, M. Niscola, V. Fanti, L.R. Limite, A. Evangelista, A. Bellisario, R. De Paolis, et al., Ablation of persistent atrial fibrillation based on atrial electrogram duration map: methodology and clinical outcomes from the AEDUM pilot study, *J. Interv. Cardiac Electrophysiol.* (2024) 1–12.
- [67] P. Rossi, M. Magnocavallo, F.M. Cauti, M. Polselli, M. Niscola, D.G. Della Rocca, A. Del Greco, L. Iaia, R. Quaglione, P. Gianfranco, et al., Functional substrate analysis in patients with persistent atrial fibrillation, *J. Interv. Cardiac Electrophysiol.* (2024) 1–11.
- [68] Z. Ye, M.S. van Schie, L. Pool, A. Heida, P. Knops, Y.J.H.J. Taverne, B.J.J.M. Brundel, N.M.S. de Groot, Characterization of unipolar electrogram morphology: a novel tool for quantifying conduction inhomogeneity, *EP Eur.* 25 (11) (2023) eua4324, <http://dx.doi.org/10.1093/europace/eua4324>.
- [69] C.H. Roney, J. Whitaker, I. Sim, L. O'Neill, R.K. Mukherjee, O. Razeghi, E.J. Vigmond, M. Wright, M.D. O'Neill, S.E. Williams, et al., A technique for measuring anisotropy in atrial conduction to estimate conduction velocity and atrial fibre direction, *Comput. Biol. Med.* 104 (2019) 278–290.