



OPEN

## DeepFake electrocardiograms using generative adversarial networks are the beginning of the end for privacy issues in medicine

Vajira Thambawita<sup>1,2,8</sup>✉, Jonas L. Isaksen<sup>3,8</sup>, Steven A. Hicks<sup>1,2</sup>, Jonas Ghouse<sup>3</sup>, Gustav Ahlberg<sup>3</sup>, Allan Linneberg<sup>3,4</sup>, Niels Grarup<sup>3</sup>, Christina Ellervik<sup>3</sup>, Morten Salling Olesen<sup>3</sup>, Torben Hansen<sup>3</sup>, Claus Graff<sup>5</sup>, Niels-Henrik Holstein-Rathlou<sup>6</sup>, Inga Strümke<sup>1</sup>, Hugo L. Hammer<sup>1,2</sup>, Mary M. Maleckar<sup>1,2</sup>, Pål Halvorsen<sup>1,2</sup>, Michael A. Riegler<sup>1,7,8</sup>✉ & Jørgen K. Kanters<sup>3,8</sup>✉

Recent global developments underscore the prominent role big data have in modern medical science. But privacy issues constitute a prevalent problem for collecting and sharing data between researchers. However, synthetic data generated to represent real data carrying similar information and distribution may alleviate the privacy issue. In this study, we present generative adversarial networks (GANs) capable of generating realistic synthetic DeepFake 10-s 12-lead electrocardiograms (ECGs). We have developed and compared two methods, named WaveGAN\* and Pulse2Pulse. We trained the GANs with 7,233 real normal ECGs to produce 121,977 DeepFake normal ECGs. By verifying the ECGs using a commercial ECG interpretation program (MUSE 12SL, GE Healthcare), we demonstrate that the Pulse2Pulse GAN was superior to the WaveGAN\* to produce realistic ECGs. ECG intervals and amplitudes were similar between the DeepFake and real ECGs. Although these synthetic ECGs mimic the dataset used for creation, the ECGs are not linked to any individuals and may thus be used freely. The synthetic dataset will be available as open access for researchers at OSF.io and the DeepFake generator available at the Python Package Index (PyPI) for generating synthetic ECGs. In conclusion, we were able to generate realistic synthetic ECGs using generative adversarial neural networks on normal ECGs from two population studies, thereby addressing the relevant privacy issues in medical datasets.

The use of artificial intelligence (AI) has increased in medicine over the past years. The goal of AI in medicine is to aid clinicians with decisions that are more accurate and to improve personalized medicine. The prominent prerequisite and foundation for AI is a large amount of high-quality clinical data.

With updates of the General Data Protection Regulation (GDPR) regulative in the EU, the free flow of data has been restricted to ensure patient consent and anonymity<sup>1</sup>. Even anonymized or de-identified data must not be shared between research groups in different countries, because combining few variables in an anonymized dataset, may allow for individual identification<sup>2</sup>. For example, knowing the zip code, birthday and sex is enough to identify 87% of US citizens<sup>3</sup>. The European GDPR rules are stricter than the US HIPAA rules for health data exchange<sup>4</sup>. EU demands that health data protection in a third country is essentially equivalent to that in the EU, which is not the case with the US HIPAA system<sup>5</sup>. All health data transfers require that informed consent is received from each patient, which makes most transatlantic collaboration impossible, if not planned in advance.

<sup>1</sup>SimulaMet, 0167 Oslo, Norway. <sup>2</sup>Oslo Metropolitan University, 0167 Oslo, Norway. <sup>3</sup>University of Copenhagen, 2200 Copenhagen N, Denmark. <sup>4</sup>Bispebjerg and Frederiksberg Hospital, 2400 Copenhagen NV, Denmark. <sup>5</sup>Aalborg University, Aalborg, Denmark. <sup>6</sup>Novo-Nordisk Foundation, Copenhagen, Denmark. <sup>7</sup>UiT The Arctic University of Norway, Tromsø, Norway. <sup>8</sup>These authors contributed equally: Vajira Thambawita, Jonas L. Isaksen, Michael A. Riegler and Jørgen K. Kanters. ✉email: vajira@simula.no; michael@simula.no; jkanters@sund.ku.dk

Checkpoint (epochs)	Fraction of DeepFake ECGs classified as Normal (%)	
	WaveGAN*	Pulse2Pulse
500	20.9	78.7
1000	69.5	81.2
1500	71.2	78.8
2000	<b>72.5</b>	79.7
2500	71.3	<b>81.6</b>
3000	65.3	81.5

**Table 1.** Quantitative difference between WaveGAN\* and Pulse2Pulse GAN in the initial training for determining the optimal network and optimal number of epochs. The best values are bolded for each GAN.

However, large-scale, publicly available open-access medical datasets are required for personalized medicine to improve data-heavy machine learning solutions in medicine.

Generating realistic synthetic data is an alternative solution to the privacy issue. Synthetic data should contain all the desired characteristics of a specific population, but without any sensitive content, making it impossible to identify individuals. Therefore, properly generated synthetic data is a solution to the privacy problem which enables data sharing between research groups.

An electrocardiogram (ECG) is a voltage time series that reflects the electric currents within the heart. An ECG is a widely used, easy applicable and inexpensive clinical screening procedure to detect cardiac diseases. With the use of multiple electrodes, 3D propagation of cardiac electric impulses is obtained and plotted as a standard 10-s 12-lead ECG.

In this paper, we showcase synthetic ECGs as an example of complex medical data. Synthetic ECGs have been a topic of interest and research for many years. McSharry et al.<sup>6</sup> and Sayadi et al.<sup>7</sup> proposed mathematical dynamical models to generate continuous ECG signals, but these models were restricted to only one lead and did not reflect the distribution found in the normal population, nor did they give any insight in the mechanisms behind any disease.

Generative adversarial networks (GAN) were introduced in 2014 by Goodfellow et al.<sup>8</sup> to generate synthetic data using multi-layer perceptrons. A GAN consists of two deep neural networks: a generator network, which creates signals (here ECGs) from random noise, and a discriminator network, which evaluates whether an ECG presented to it is real or fake. During training, a mix of real ECGs (from the underlying population) and generated DeepFake ECGs (from the generator) are presented to the discriminator, which assigns a score to the ECG (high score for real, low score for fake). As training proceeds, both the generator and the discriminator improve in performance until an equilibrium is reached<sup>9</sup>. Later, Radford et al.<sup>10</sup> developed a convolutional GAN to generate synthetic images, which is well suited for images.


Since ECGs are time series data, our initial approach was to use a WaveGAN<sup>11</sup> which is capable of generating sound signals. The classical WaveGAN is only able to output a single channel time series, so we modified the WaveGAN to generate 8 ECG channels (denoted WaveGAN\*) instead of audio signals. We then introduced a novel DeepFake ECG U-net generative model, called Pulse2Pulse, which was inspired by the WaveGAN<sup>11</sup>, and we compared our Pulse2Pulse GAN to the WaveGAN\*.

In this paper, we thus present two GANs with the ability to generate an unlimited number of 10-s 12-leads synthetic “DeepFake” ECGs as a solution to overcome the privacy issues related to real ECG data. These DeepFake ECGs can be openly distributed and freely downloaded as open access and used by other scientists to develop ECG algorithms.

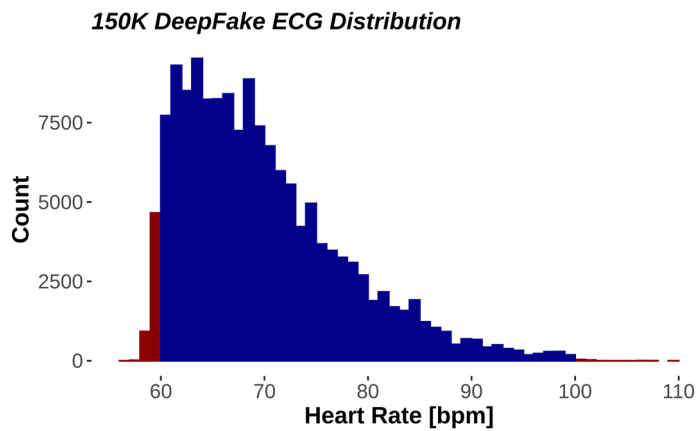
## Results

We used ECGs from two population studies (GESUS<sup>12</sup> and Inter99<sup>13</sup>). To avoid chimeras between normal and abnormal ECGs, we only trained the neural network with ECGs classified as normal by the MUSE 12SL (version 2.43). As shown in Table 1, both the WaveGAN\* and Pulse2Pulse improved during training expressed as the percentage of DeepFake ECGs classified by the commercial ECG interpretation program MUSE 12SL as normal ECGs. The Pulse2Pulse GAN trained faster than the WaveGAN\* and had a better performance (expressed as fraction of ECGs classified as normal by the MUSE) compared to the WaveGAN\* at their respective optimal number of training epochs (Table 1). Figure 1 shows a comparison of real and DeepFake ECGs, and the Supplementary Figure S1 shows twenty randomly chosen DeepFake ECGs. Figure 2 shows the distribution of heart rates in the DeepFakes. By clinical definition Normal ECGs heart rates are between 60 and 99 beats per minute. The MUSE 12SL<sup>14</sup> classified 129 DeepFakes (0.5%) as sinus tachycardia (fast heart rate  $\geq 100$ ) and 2863 (10.2%) as sinus bradycardia (slow heart rate  $< 60$ ). Figure 3 shows that the well-known established correlation between the QT interval and the RR interval<sup>15</sup> was preserved. All covariance structures can be seen in the Supplementary Figure S2.


The generated DeepFake ECGs can be downloaded at OSF.io (<https://osf.io/6hved/>) with the corresponding ground truth parameters for the QT, RR, PR and QRS intervals and the P, STJ, R, and T amplitudes (see Fig. 4 for ECG wave/interval naming terminology) delivered by the MUSE 12SL system. The DeepFake ECGs may be freely used for scientific use or commercial algorithm development if this paper is properly cited.




**Figure 1.** Comparison of examples of a real ECG (left lane) and a DeepFake ECG (right lane). See supplementary Figure S1 for 20 more randomly chosen pairs of real and DeepFake ECGs.



**Figure 2.** Distribution of heart rates in all 150,000 DeepFake electrocardiograms. Red fill denotes outside the normal heart rate range. Blue fill is within normal heart rate range (60–99).



**Figure 3.** Scatter plot of the QT/RR interval relationship where Real ECG are shown in blue and normal DeepFakes in red. DeepFake dots are nudged 1 ms to the left for visibility. Note that there are 121,977 normal DeepFakes and only 7233 Real ECG making the DeepFake distribution more pronounced. As seen by the correlation coefficient  $r^2$ , the real and the fake DeepFake ECGs are similarly distributed.



**Figure 4.** An ECG complex with the nomenclature of intervals (QT, QRS, P duration) and Amplitudes (STJ, R, T) and RR-interval (which can be converted to heart rate (HR) as  $HR = 60/RR$  interval).

		Real—normal (7233)				Pulse2Pulse—normal (121,977)				Pulse2Pulse—all (150,000)			
		Mean	Std	2.5%	97.5%	Mean	Std	2.5%	97.5%	Mean	Std	2.5%	97.5%
Heart rate	BPM	70	8	60	90	70	7	60	88	70	8	60	89
P duration	ms	105	12	82	130	117	17	86	152	118	17	84	152
QT interval	ms	395	21	352	436	395	20	354	436	395	22	352	436
QRS duration	ms	90	9	74	110	92	9	78	112	93	10	78	114
PR interval	ms	156	19	120	198	158	17	126	192	159	19	124	194
STJ amplitude (V5)	$\mu V$	2	27	-44	58	18	33	-44	87	16	36	-54	87
R amplitude (V5)	$\mu V$	1287	402	600	2163	1275	367	620	2026	1273	402	566	2094
T amplitude (V5)	$\mu V$	343	137	126	664	366	135	156	668	361	141	141	673

**Table 2.** Mean, standard deviation (std), 2.5%, and 97.5% percentile for standard ECG parameters in real and fake ECGs. *BPM* beats per minute.

Using the Pulse2Pulse model from the optimal number of epochs (2500), we generated 150,000 DeepFake ECGs. To ensure that these ECGs were realistic, we uploaded the 150,000 ECGs to the GE MUSE system and analyzed them using the 12SL algorithm. We found that 81.3% of the 150,000 DeepFake ECGs were classified as “Normal ECG” (vs. 81.6% in the initial training). Table 2 compares real vs. DeepFake ECGs using eight ECG properties (heart rate, P duration, QT interval, QRS duration, PR interval, STJ amplitude, R amplitude, and T amplitude) extracted using MUSE 12SL. See Fig. 4 for ECG nomenclature). The real data included all ECGs from GESUS and Inter99 classified as “Normal ECG” which were used for training. DeepFake ECGs are presented both as all 150,000 generated ECGs and the subset classified as Normal ECG. Supplementary Table S4 summaries the most common reasons for classifying DeepFake ECGs as Non-Normal ECGs.

## Discussion

Although deep learning has previously been used for ECG analysis<sup>16,17</sup>, this study is the first study to generate realistic synthetic 10-s 12-lead DeepFake ECGs. We demonstrate that the characteristics of the real ECGs were preserved with the DeepFake ECGs.

In our study, nearly one fifth of the DeepFake ECGs were not recognized as Normal ECGs (Non-Normal) by the commercial MUSE 12SL ECG analyzer (no ECGs were rejected as being invalid). Many ECG parameters use hard boundaries in distinguishing between Normal and Non-Normal. For example, a normal heart rate is by definition located between 60 to 99 beats per minute. Since we trained our model only on Normal ECGs, the input distribution for the GAN was a truncated asymmetric distribution. Thus, the clinically defined boundaries are skewed compared to the normal distribution of heart rates. The left truncation (at low heart rates) will discard

more individuals than the right truncation (at high heart rates), and the final distribution of the real ECGs will be close to a truncated normal distribution with asymmetric truncations. The GAN will generally learn that heart rates outside 60–99 are not valid, but small deviations will occur as seen in Fig. 2 and Table 2. Since similar boundaries exist for many ECG parameters (for example PR interval > 220 ms or QRS Interval < 120 ms) sharp truncations occur with several ECG parameters. This could lead to the exclusion of some DeepFake ECGs, simply because the ECG intervals or amplitudes were marginally outside the normal range. Most ECG amplitudes and intervals were similar between real ECGs and DeepFake ECGs. It is noteworthy that the STJ amplitude and the P duration had the greatest deviation between real ECGs and DeepFake ECGs. This may be because both STJ and P amplitudes are small, and that the network may tend to focus on larger waves such as the R and T waves. Following this theory, the network would to some extent neglect the smaller waves and features thereby introducing a larger uncertainty. Future networks may improve the ECG generation using conditional GANs to give more attention to smaller signal features. The Pulse2Pulse model was able to preserve the covariance structure between different ECG features, as seen in the most important relationship the QT/RR relationship which is known to have prognostic importance<sup>18</sup>.

A challenging task is to define the optimal number of epochs for training. GANs tend to become unstable during the training process with the risk of the generator producing unrealistic output. To get an unbiased estimate on how well the trained GAN performs, we used the commercial MUSE 12SL system which automatically and reliably evaluates an ECG with a sensitivity of 99.9% and specificity of 100%<sup>19</sup>. Although the ECG discarded by the MUSE 12 SL may only have minimal abnormalities (like a heart rate of 59.9 bpm where 60 bpm is normal), the filtering of the DeepFake ECGs ensures that the best epoch is chosen without bias. It also ensures that the resulting ECGs are normal not only according to the discriminator, but also according to one of the most widely used ECG system in hospitals worldwide.

Personalized medicine depends on big data, which is frequently facilitated by international collaborations to ensure large datasets for both researchers and industry. However, privacy and general data protection regulation rules are major obstacles for sharing data between researchers from different institutions or countries or with the industry<sup>20</sup>.

In conclusion, by constructing synthetic signals from real patients which retain the same clinical information as was present in the real dataset, we have paved a new way to overcome privacy and ethical<sup>21</sup> concerns for data sharing. The synthetic data generated by our Pulse2Pulse GAN are not linked to any specific patients but to the entire population, and therefore the ECGs prove useful for data scientists and the industry in developing novel algorithms for ECG analysis. The approach is not limited to ECGs but could be generalized to all medical multi-channel data, e.g., electroencephalography and electromyography. Therefore, the DeepFake ECGs generated from the Pulse2Pulse model can be used as a replacement to overcome the privacy constraints in real medical datasets.

## Methods


GAN models were first introduced by Goodfellow et al.<sup>8</sup>. In a GAN, two deep neural networks termed the generator (G) and the discriminator (D) are combined to achieve the generation task. The main goal of the generator is to produce a data sample input [ECG(z)] from random noise (z) to present to the discriminator. The discriminator is tasked with differentiating between real and fake data, thus forcing the generator to improve performance. The generator and discriminator are trained together in a competition (minmax game). When a steady state is reached, the training halts and the generator will generate realistic synthetic ECGs.

**Data preparation.** We used two combined datasets: the Danish General Suburban Population Study<sup>12</sup> (GESUS) and the Inter99 study<sup>13</sup> (CT00289237, ClinicalTrials.gov). GESUS consists of 8939 free-living subjects, and Inter99 consists of 6667 free-living subjects with an available digital ECG. To avoid generation of hybrid ECGs with mixed ECG abnormalities not occurring in real persons (e.g., to both be in sinus rhythm and atrial fibrillation at the same time which is impossible), we excluded ECGs who were not classified as normal (n = 8348) leaving 7233 Normal ECGs for training.

A 10-s 12-lead ECG consists only of 8 independent channels since 4 of the channels are simply trigonometric rotations of the two first channels. Therefore, the input ECG signal is 5000 × 8 data points (corresponding to 10 s with 500 samples per sec × 8 channels). We calculated the missing four channels with trigonometric functions to create the classic 12-channels ECG from 8-channels ECG.

**WaveGAN\*.** The input to WaveGAN\* is a 1D 100 × 1 random noise vector sampled from the uniform distribution (mean = 0, std = 1) which passes through six deconvolution blocks to generate the desired output of 5000 × 8 samples (Fig. 5a). The deconvolution blocks were built from a series of four layers: an up-sampling layer, a constant padding layer, a 1D-convolution layer, and a ReLU activation function consecutively. This implementation is deeper than the original architecture, which uses five deconvolution blocks used to generate synthetic music samples. Table S1 has comprehensive details of our WaveGAN\* generator network.

**Pulse2Pulse.** The implementation of the Pulse2Pulse architecture (Fig. 5) is inspired by the U-Net architecture<sup>22</sup>, which was used for image segmentation. However, our Pulse2Pulse implementation is different from the original U-net implementation because the Pulse2Pulse implementation uses 1D convolutional layers for ECG signal generation as opposed to the 2D convolutional layers used for the original image segmentation task. The Pulse2Pulse network takes an 8 × 5000 noise vector, i.e., the same dimension as the output ECG. The noise is passed through six down-sampling blocks followed by six up-sampling blocks as illustrated in Fig. 5b. Each down-sampling block consists of a 1D-convolution layer followed by a Leaky ReLU activation. The up-sampling block is similar to the deconvolution block used in WaveGAN\*. In down-sampling, we have used Leaky ReLU



**Figure 5.** Model architectures of the generators and the discriminator used to generate synthetic ECGs. WaveGAN\* uses a 1D noise vector with 100 points. Pulse2Pulse uses a 2D noise vector with size of  $8 \times 5000$  as input, same as the output ECG size.

instead of the ReLU layer used in the up-sampling to match the down-sampling operations to the discriminator. In addition to the up-sampling and down-sampling, the major modification is a bypass option, which concatenates the down-sampling block features with the up-sampling block features (represented by the black arrows in Fig. 5b). To facilitate for this concatenation, we doubled the input size of up-sampling blocks compared to WaveGAN\* up-sampling blocks. More details about the Pulse2Pulse generator network are shown in the Supplementary Table S2.

**Discriminator.** The same discriminator was used by WaveGAN\* and Pulse2Pulse to discriminate between real and fake ECGs (Fig. 5c). We used seven convolution layers (the original WaveGAN<sup>11</sup> has five layers), and each convolution layer is followed by a Leaky ReLU activation and the phase shuffle layer introduced in the original WaveGAN paper<sup>11</sup>. The discriminator takes an ECG as input ( $5000 \text{ samples} \times 8 \text{ channels}$ ) and outputs a score how close the ECG are to be determined fake or real. Complete details about our discriminator network are given in the Supplementary Table S3.

**Training.** The models were trained on a Ubuntu workstation with two Xeon processors and a GeForce NVIDIA RTX 2080ti running the Pytorch deep learning framework<sup>23</sup>. We ran all our experiments (generators + discriminator) using the Adam<sup>24</sup> optimizer with a learning rate of 0.0001,  $\beta_1$ -value of 0.5, and  $\beta_2$ -value of 0.9. As loss function, we used gradient clipping WGAN-GP<sup>25</sup>, to ensure faster and better convergence. Similar to the audio generation paper of WaveGAN<sup>11</sup>, we updated (backpropagated) the discriminator five times per update of the generator. We used a batch size of 32, which is half of the original batch size of 64 used in the original WaveGAN paper, because we used larger networks than the WaveGAN networks. We kept the training process until 3000 epochs ( $\sim 10$  days computing time) because we experienced unstable training curves for both WaveGAN\* and Pulse2Pulse afterwards.

**DeepFake ECGs.** For evaluation of our two GAN models, we initially generated 10,000 ECGs from every 500 epochs until 3000 epochs from each GAN model. The DeepFake ECGs were transferred to the MUSE system and evaluated by the MUSE 12SL algorithm v. 2.43<sup>14</sup>, and we used the fraction of DeepFake ECGs described as Normal as the metric (because we only used Normal Real ECGs for the training). Using the best epoch for the best GAN, we generated 150,000 DeepFake ECGs. These DeepFakes were also evaluated by the MUSE 12SL.

### Data availability

The Normal DeepFake ECGs are available at OSF (<https://osf.io/6hved/>) with corresponding MUSE 12SL ground truth values freely downloadable and usable for ECG algorithm development. The DeepFake generative model is available at <https://pypi.org/project/deepfake-ecg/> to generate only synthetic ECGs.

### Code availability

The complete source code of all networks discussed in paper are available at GitHub (<https://github.com/vlbthambawita/deepfake-ecg>).

Received: 22 June 2021; Accepted: 26 October 2021

Published online: 09 November 2021

### References

- Voigt, P. & von dem Bussche, A. *The EU General Data Protection Regulation (GDPR)* (Springer, 2017). <https://doi.org/10.1007/978-3-319-57959-7>.
- de Montjoye, Y.-A., Radaelli, L., Singh, V. K. & Pentland, A. S. Identity and privacy. Unique in the shopping mall: On the reidentifiability of credit card metadata. *Science* **347**, 536–539 (2015).
- El Emam, K., Jonker, E., Arbutckle, L. & Malin, B. A systematic review of re-identification attacks on health data. *PLoS ONE* **6**, e28071 (2011).
- Bradford, L., Aboy, M. & Liddell, K. International transfers of health data between the EU and USA: A sector-specific approach for the USA to ensure an 'adequate' level of protection. *J. Law Biosci.* **7**, 055 (2020).
- Hallinan, D. *et al.* International transfers of personal data for health research following Schrems II: A problem in need of a solution. *Eur. J. Hum. Genet. EJHG* **29**, 1502–1509 (2021).
- McSharry, P. E., Clifford, G. D., Tarassenko, L. & Smith, L. A. A dynamical model for generating synthetic electrocardiogram signals. *IEEE Trans. Biomed. Eng.* **50**, 289–294 (2003).
- Sayadi, O., Shamsollahi, M. B. & Clifford, G. D. Synthetic ECG generation and Bayesian filtering using a Gaussian wave-based dynamical model. *Physiol. Meas.* **31**, 1309–1329 (2010).
- Goodfellow, I. *et al.* Generative adversarial nets. *Adv. Neural Inf. Process. Syst.* **27**, 9 (2020).
- Osborne, M. J. & Rubinstein, A. *A Course in Game Theory* (MIT Press, 1994).
- Radford, A., Metz, L. & Chintala, S. *Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks*. ArXiv151106434 Cs (2016).
- Donahue, C., McAuley, J. & Puckette, M. Adversarial audio synthesis. in *International Conference on Learning Representations* (2019).
- Juhl, C. R., Miller, I. M., Jemec, G. B., Kanters, J. K. & Ellervik, C. Hidradenitis suppurativa and electrocardiographic changes: A cross-sectional population study. *Br. J. Dermatol.* **178**, 222–228 (2018).
- Ghouse, J. *et al.* Rare genetic variants previously associated with congenital forms of long QT syndrome have little or no effect on the QT interval. *Eur. Heart J.* **36**, 2523–2529 (2015).
- GE Healthcare. *Marquette™ 12SL™ ECG Analysis Program Physician's Guide 2056246–002 Revision C.* (2015).
- Jensen, B. T. *et al.* Beat-to-beat QT dynamics in healthy subjects. *Ann. Noninvas. Electrocardiol.* **9**, 3–11 (2004).
- Attia, Z. I. *et al.* An artificial intelligence-enabled ECG algorithm for the identification of patients with atrial fibrillation during sinus rhythm: A retrospective analysis of outcome prediction. *Lancet* [https://doi.org/10.1016/S0140-6736\(19\)31721-0](https://doi.org/10.1016/S0140-6736(19)31721-0) (2019).
- Hicks, S. A. *et al.* Explaining deep neural networks for knowledge discovery in electrocardiogram analysis. *Sci. Rep.* **11**, 10949 (2021).
- Jensen, B. T. *et al.* QT dynamics in risk stratification after myocardial infarction. *Heart Rhythm* **2**, 357–364 (2005).
- Froelicher, V. F., Marcus, R. & Heidenrich, P. Prognostic value of computer electrocardiography in veteran outpatients. *Fed. Pract.* **21**, 11–20 (2004).
- El Emam, K., Rodgers, S. & Malin, B. Anonymising and sharing individual patient data. *BMJ* **350**, 1139 (2015).
- Ienca, M. *et al.* Considerations for ethics review of big data health research: A scoping review. *PLoS ONE* **13**, e0204937 (2018).
- Ronneberger, O., Fischer, P. & Brox, T. U-net: Convolutional networks for biomedical image segmentation. in *International Conference on Medical image computing and computer-assisted intervention* 234–241 (Springer, 2015).
- Paszke, A. *et al.* *PyTorch: An Imperative Style, High-Performance Deep Learning Library*. ArXiv191201703 Cs Stat (2019).
- Kingma, D. P. & Ba, J. *Adam: A Method for Stochastic Optimization*. ArXiv14126980 Cs (2017).
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V. & Courville, A. C. Improved training of Wasserstein GANs. *Adv. Neural Inf. Process. Syst.* **30**, 5767–5777 (2017).

### Author contributions

V. T., J. L. I., S. A. H., M. A. R., and J. K. K. conceived the experiment(s), V. T., J. L. I., and J. K. K. conducted the experiment(s), V. T., J. L. I., S. A. H., M. A. R., and J. K. K. analyzed the results. All authors reviewed the manuscript.

### Competing interests

The authors declare no competing interests.

### Additional information

**Supplementary Information** The online version contains supplementary material available at <https://doi.org/10.1038/s41598-021-01295-2>.

**Correspondence** and requests for materials should be addressed to V.T., M.A.R. or J.K.K.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2021