Appendix: Content Summaries of Best Papers for the 'Sensors, Signals, and Imaging Informatics' Section of the 2020 IMIA Yearbook of Medical Informatics

Armanious K, Jiang C, Fischer M, Küstner T, Hepp T, Nikolaou K, Gatidis S, Yang B

MedGAN: Medical image translation using GANs

Comput Med Imaging Graph 2020;79:101684

Nearly 20 years ago, the European Cross-Language Evaluation Forum (CLEF) Campaign started to regard medical images as being analogous to "language". Concerning the diverse imaging modalities, there are a lot of image-based languages, and "translation" between languages is a frequent task in medical applications. This paper presents a novel approach to this problem, where generative adversarial networks (GANs) are used to generate synthetically images which are based on a template in a different language. The potential applications for this modeling approach include, but are not limited to, inter-modality translation, image de-noising, and motion-artifact correction. Using the MedGAN framework, all of these tasks can be done without task-specific training. A new generator architecture, CasNet, is introduced. CasNet concatenates several encoder-decoder pairs (similar to stacked U-nets) and captures high as well as low-frequency components of the desired target modality by combining the adversarial framework with non-adversarial losses. The authors analyzed individual loss functions and quantitatively showed the superiority of MedGAN over existing translation algorithms. The neural network was applied to three different tasks without any task-specific adaptation: positron emission tomography (PET) to computed tomography (CT) translation, PET image denoising, and magnetic resonance imaging (MRI) motion artifact correction. Interestingly, five experienced radiologists confirmed the equivalence of the synthetically generated images to original image recordings.

Chandra BS, Sastry CS, Jana S

Robust heartbeat detection from multi-modal data via cnn-based generalizable information fusion

IEEE Trans Biomed Eng 2019 Mar;66(3):710-7

This paper demonstrates an innovative way of signal fusion deploying convolutional neural networks (CNNs) for robust heartbeat detection. In the monitoring of cardiovascular diseases, especially in critical-care situations, accurate heartbeat detection is key, but often prone to errors when monitoring with one physiological signal e.g., the electrocardiogram (ECG) only. Multi-signal detectors exist, but do not systematically exploit inter-signal correlations. To fill this gap, the authors propose a CNN-based approach, which directly fuses information from multiple physiological signals for estimating heartbeat locations without the need for any intermediate detection. They employ ECG and blood pressure signals from the PhysioNet 2014 Challenge as well as the MIT-BIH arrhythmia database for network training. Their CNN learns a set of linear filters to extract features, temporarily fusing information from multiple signals. A fully connected network including a sigmoid output function maps these features for heartbeat prediction. Their trained networks perform with a score of 94% using blood pressure and ECG signals on the PhysioNet 2014 dataset, and have a performance score of 99.92% using two ECG channels on the MIT-BIH arrhythmia database. Hence these results compare well with previously reported database-specific results. In conclusion, their CNN-based approach is generalizable, robust, and efficient in detecting heartbeat locations from multiple signals. Furthermore, the authors suggest their technique as an accurate method to estimate heartbeats on sparse datasets.

de Vos BD, Berendsen FF, Viergever MA, Sokooti H, Staring M, Išgum I

A deep learning framework for unsupervised affine and deformable image registration

Med Image Anal 2019 Feb;52:128-43

Image registration is a frequent task in medical imaging and computer-aided diagnosis. This paper describes an entire framework to perform registration tasks using CNNs. The framework has innovative features such as the support of 1) model-based (affine) as well as elastic (deformable) image registration, 2) n-dimensional images, 3) unsupervised training, 4) multi-modal data, and 5) coarse-to-fine hierarchical architectures. They performed a comprehensive evaluation, demonstrating a considerable speedup with respect to the processing time of conventional methods yielding performance similar to state-of-the-art methods. The core contribution is the framework called Deep Learning Image Registration (DLIR). The DLIR training procedure is similar to a conventional iterative image registration framework, where fixed and moving images are compared, transformed, and resampled into the warped image, but adds a CNN that allows unsupervised training. In other words, it can recast conventional intensity-based image registration into a learning problem, and hence speed up the registration process essentially, once the learning is done. The framework is flexible, allowing for a variety of architectures by stacking multiple CNNs into larger compositions. Among others, the authors demonstrate and evaluate a four-dimensional (4D) task using the publicly available DIR-Lab 4D chest CT data.

Zhu T, Pimentel MAF, Clifford GD, Clifton DA

Unsupervised Bayesian inference to fuse biosignal sensory estimates for personalizing care

IEEE J Biomed Health Inform 2019 Jan;23(1):47-58

This paper addresses the difficulties and challenges of reliable, consistent, and real-time labeling of high data volumes arising from medical sensors used for diagnosis and patient-specific treatments. In this work, the authors present two fully Bayesian approaches to BCLA modeling (Bayesian continuous-valued label aggregator) using

Gibbs sampling to fuse continuous-valued labels of biosignal sensor data from independent or potentially correlated annotators in an unsupervised manner. They estimate the bias and precision of each annotator to infer the underlying ground truth. One of these models takes, for the first time, into account the correlation between annotators, allowing a grouping based on their decision-making process. The manuscript is notably detailed, giving a deep insight into the methods developed and their proof-of-concept. The authors performed a comprehensive validation based on two clinical datasets, comprising QT intervals of ECGs and Capnobase Respiratory Rate data, as well as a synthetic QT dataset. The validation demonstrated impressively the outperforming performance and robustness in dealing with missing values of both proposed models in comparison to existing approaches. This study thus provides a valuable advanced method for aggregating the labeling of several imperfect automated algorithms, generating highly reliable labels to better support and improve decisions in personalized care.