Cross-Domain Detection of Pulmonary Hypertension in Human and Porcine Heart Sounds

Alex Gaudio[1,2], Noemi Giordano[3], Miguel Coimbra[2], Benedict Kjaergaard[4], Samuel Schmidt[5], Francesco Renna[2]

[1] Carnegie Mellon University, Pittsburgh, United States
[2] INESC TEC, Faculty of Sciences of the University of Porto, Porto, Portugal
[3] Politecnico di Torino, Torino, Italy
[4] Aalborg University Hospital, Aalborg, Denmark
[5] Aalborg University, Aalborg, Denmark

Aims: Detection of Pulmonary Hypertension (PH) via automated analysis of cardiac auscultation may offer a non-invasive, accurate, effective and efficient solution with low resource requirements. We detect PH in human and in porcine datasets and demonstrate generalization across the two datasets.

Methods: Extending our previous work, we train a deep convolutional network typically used for image analysis on a representation of aligned second heart sounds (S2). We train models on either porcine or human data and evaluate generalization to human data. The human dataset contains digital stethoscope recordings (PCG) of 42 patients (29 with PH and 13 without PH), with reference PAP measurements obtained via right heart catheterization (RHC). The porcine dataset contains PCG and seismocardiogram (SCG) recordings of ten pigs chemically induced to have PH, where PH was determined via RHC. Evaluation considers bootstrapped cross-validation of 12 independently trained models with randomized splits.

Results: Strong performance on human and porcine datasets is obtained. The area under the ROC curve (auROC) and area under the Precision-Recall curve (AP) on human data via 10-fold cross-validation are 0.928 and 0.972, respectively. On the porcine dataset, leave-one-out cross-validation gives 0.833 auROC and 0.851 AP. Moreover, we demonstrate transferability across domains, where training on the porcine dataset and evaluating on the human dataset gives 0.702 auROC and 0.848 AP.

Dataset	auROC	AP
Human (S2)	0.928	0.972
Porcine (S2)	0.833	0.851
Porcine (S2) $\rightarrow$ Human (S2)	0.702	0.848

Conclusion: Results show that it is possible to use porcine data for developing human AI models. Applying deep convolutional networks to aligned 1-d audio sequences gives strong performance generalizable across different recording media (PCG and SCG) and different animals (porcine and human).