

Yolo4Apnea: Real-time Detection of Obstructive Sleep Apnea

Sondre Hamnvik^{1*}, Pierre Bernabé² and Sagar Sen²

¹Department of Informatics, University of Oslo, Norway

²Simula Research Laboratory, PB 135, 1325 Lysaker, Norway

sondrham@ifi.uio.no, {pierbernabe,sagar}@simula.no

Abstract

Obstructive sleep apnea is a serious sleep disorder that affects an estimated one billion adults worldwide. It causes breathing to repeatedly stop and start during sleep which over years increases the risk of hypertension, heart disease, stroke, Alzheimer's, and cancer. In this demo, we present **Yolo4Apnea** a deep learning system extending *You Only Look Once* (Yolo) system to detect sleep apnea events from abdominal breathing patterns in real-time enabling immediate awareness and action. Abdominal breathing is measured using a respiratory inductance plethysmography sensor worn around the stomach. The source code is available at <https://github.com/simula-vias/Yolo4Apnea>

1 Introduction

Obstructive sleep apnea (OSA) is a sleep-related breathing disorder that involves a decrease or complete halt in airflow despite an ongoing effort to breathe [Strollo Jr and Rogers, 1996]. It occurs when the muscles relax during sleep, causing soft tissue in the back of the throat to collapse and block the upper airway. This leads to partial reductions (hypopneas) and complete pauses (apneas) in breathing that last at least 10 seconds during sleep. Most pauses last between 10 and 30 seconds, but some may persist for one minute or longer. We can count from a few tens to several hundred sleep apneas per night. This can lead to abrupt reductions in blood oxygen saturation, with oxygen levels falling as much as 40 percent or more in severe cases. A recent literature-based study [Benjafield *et al.*, 2019], suggested that an estimated 936 million adults suffer from undiagnosed OSA worldwide. OSA increases risk for cardiac disease [Tadic *et al.*, 2020], stroke [McDermott and Brown, 2020], neurodegenerative diseases such as Alzheimer's [Lajoie *et al.*, 2020], and cancer [Brenner *et al.*, 2019].

Sleep laboratories perform *polysomnography* or a *sleep study* to diagnose sleep disorders. Polysomnography records brain waves, the oxygen level in your blood, heart rate and breathing, as well as eye and leg movements as illustrated in Figure 1. The overnight recordings are *manually analyzed*

by a sleep technician and annotated for sleep apnea events. The multi-sensor equipment and the skills needed for maintenance and sleep apnea annotation incur **high operating costs** for sleep laboratories making them economically challenging to operate. Moreover, this is a serious bottleneck in making sleep labs available to a billion people. Hence, there is dire need to automate detection of OSA events and assist a sleep technician achieve high accuracy and throughput.

Deep neural network (DNN) models such as *long short term memory* (LSTM) networks have been tried to automate prediction of OSA events from patterns of expansion and contraction seen in abdominal breathing signals [Van Steenkiste *et al.*, 2018]. The authors achieve a mean true positive rate of only 80% by observing fixed length $N (= 30)$ seconds of abdominal breathing. A systematic literature review [Mostafa *et al.*, 2019], investigates about 255 research papers that use deep learning to detect sleep apnea. Methods used to detect sleep apnea include multiple hidden layer neural networks, stacked auto-encoders, deep belief networks, 1D and 2D convolutional neural networks. The deep learning models have been used to train samples from different sensors such as ECG, oxygen saturation data, nasal pressure, and air flow. Table 1 in [Mostafa *et al.*, 2019] presents alternatives for reporting predictions and reporting results. Some models use 60 seconds intervals, while others use only 1 second as input length. Some models are binary detectors that detect apnea or non-apnea, while others predict AHI index. This non-standard variability in apnea detection makes it hard to compare results. Moreover, a study [Thorey *et al.*, 2019] observed that human experts themselves reach an accuracy of 75% in sleep apnea detection tasks. The widely differing accuracy of sleep apnea detection can be attributed to the *variable length of an apnea event* (as opposed to fixed length intervals) and *sleep apnea events spaced at variable distances*. There is a need for a *dynamic approach* to automatically detect variable length OSA events and perhaps go a step forward and do so in real-time such that immediate action can be taken.

Yolo4Apnea is real-time apnea detection system based on training *You Only Look Once* YOLO [Redmon *et al.*, 2016], a DNN architecture capable of real-time object detection. Yolo4Apnea can predict OSA events of variable size from 2D images of the time series of abdominal breathing achieving a frame rate upwards of 30 fps with varying levels of confidence. The confidence in percent helps a sleep technician

*Contact Author




Figure 1: Polysomnography to Detect Obstructive Sleep Apnea

certify automatically predicted OSA event annotations.

The rest of the paper is organized as follows. In Section 2, we briefly present the Sleep Heart Health Study used to train Yolo4Apnea. In Section 3, we present the architecture, evaluation, and target users of Yolo4Apnea.

2 Sleep Heart Health Study Dataset

The Sleep Heart Health Study (SHHS)¹ is a multi-center cohort study implemented by the National Heart Lung & Blood Institute to determine the cardiovascular and other consequences of sleep-disordered breathing. The study entailed performing *polysomnography* on patients (illustrated in Figure 1) to obtain EEG, EOG, EMG, thoracic (THOR), abdominal (ABDO) respiration, air flow, oximetry, ECG, HR, body position, and ambient light (ON/OFF) all measured while patients sleep overnight. In all, 6,441 men and women aged 40 years and older were enrolled between 1995 and 1998 to take part in SHHS Visit 1. A second polysomnogram (SHHS-2) was obtained in 3295 of the participants between 2001 and 2003. SHHS study provides annotated OSA events primarily on abdominal respiration data.

3 Architecture

Yolo4Apnea relies on YOLO, a real-time object detection system [Redmon *et al.*, 2016]. YOLO is a fully convolutional neural network (FCNN) that passes an image once through the FCNN and generates an output bounding box around one or more objects found in the image. YOLO is trained by optimizing sum squared error between predicted bounding box and a hand annotated bounding box around an object of interest. YOLO is specifically implemented using Darknet² which is an open source neural network framework written in C and CUDA. It is fast and can process up to 45 frames per second on a Titan X GPU³ allowing for real-time object detection in a video stream.

Yolo4Apnea leverages the ability to generate bounding boxes around OSA events in 2D images of abdominal breathing. It is based on training YOLO's FCNN architecture using 2D images of times series data containing abdominal breathing patterns as shown in Figure 2. The images are created by converting abdominal breathing signals from SHHS (available in the EDF format [Kemp and Olivan, 2003]) into 2D




Figure 2: Yolo4Apnea Architecture




Figure 3: (a) Thresholding in Yolo4Apnea (b) Different sleep apnea patterns

plots normalized between -1 and +1 on the Y-axis and de-ciseconds on the X-axis. In the full SHHS dataset, the median duration of an OSA event was about 22.8 seconds. Very rarely were OSA events longer than 90 seconds. Therefore, each training image containing an OSA event was scaled to accommodate a 90 second long breathing sequence. The *start* and *end* times of OSA events in the SHHS dataset (available in XML annotation files) is used to automatically create a bounding box to train YOLO as shown in Figure 2. Several thousand training images and bounding boxes were generated from the SHHS data to train YOLO. During training the difference between predicted start and end time (bounding box) and ground truth bounding boxes is used to improve the FCNN's performance.

The main reason for going from 1D time series to a 2D image was to directly enable the training of Yolo v3 that accepts 2D images without the need to create 1D version of Yolo's neural network architecture. Moreover, the transformation from 1D to 2D can be seen as a *kernel trick* to make sleep apnea features over variable lengths of time more discernible to the human eye and consequently to Yolo.

¹<https://sleepdata.org/datasets/shhs>

²<https://pjreddie.com/darknet/>

³<https://www.nvidia.com/en-us/geforce/products/10series/titan-x-pascal/>




Figure 4: Receiver Operating Characteristic Curve for Yolo4Apnea

3.1 Evaluation

Yolo4Apnea predicts an apnea with different confidence levels. In Figure 3 (a), we illustrate how thresholding of confidence level generates multiple bounding boxes. By increasing the threshold to a minimum of 80% confidence, we are able to accurately predict OSA events. Yolo4Apnea is able to automatically predict atypical OSA events as seen in Figure 3 (b). These atypical patterns may arise due to laboured breathing.

We compute the accuracy of Yolo4Apnea by classifying a test set of 10 patients randomly sampled from the SHHS dataset. The predictions are converted to a binary list of predictions spanning the whole recording at a 10 Hz sample rate and compared to a binary representation of the sleep technicians annotation file with the same sample rate. Figure 4 shows the receiver operating characteristic curve of our classifier, with Yolo4Apnea achieving an area under ROC (AUROC) of 0.71. In future, we will use a larger part of the SHHS dataset for training Yolo4Apnea. This will help increase our precision in classifying OSA and the corresponding AUROC. Moreover, there is a large variation in sleep breathing patterns in patients which informs us about a need to train the model extensively for robustness.

3.2 Target Users

General practitioners (GP): Patients with apnea often complain of fatigue and loss in productivity. A simple method for apnea detection can assist GPs screen for apnea quickly. The GP can then recommend a patient to a more advanced analysis at a sleep lab.

Patients at home: The real-time detection of OSA events can trigger actions that stimulate breathing while sleeping. For instance, the automatic diffusion of essential oils [Otaghi *et al.*, 2017] in the air can stimulate breathing or automatically moving the head using a smart pillow⁴ can relieve restricted breathing.

References

[Benjafield *et al.*, 2019] Adam V Benjafield, Najib T Ayas, Peter R Eastwood, Raphael Heinzer, Mary SM Ip, Mary J Morrell, Carlos M Nunez, Sanjay R Patel, Thomas Penzel, Jean-Louis Pépin, et al. Estimation of the

global prevalence and burden of obstructive sleep apnoea: a literature-based analysis. *The Lancet Respiratory Medicine*, 7(8):687–698, 2019.

[Brenner *et al.*, 2019] Ronen Brenner, Shaye Kivity, Marina Peker, Daniel Reinhorn, Lital Keinan-Boker, Barbara Silverman, Irena Liphstiz, Tamara Kolitz, Carmit Levy, Dekel Shlomi, et al. Increased risk for cancer in young patients with severe obstructive sleep apnea. *Respiration*, 97(1):15–23, 2019.

[Kemp and Olivan, 2003] Bob Kemp and Jesus Olivan. European data format ‘plus’(edf+), an edf alike standard format for the exchange of physiological data. *Clinical Neurophysiology*, 114(9):1755–1761, 2003.

[Lajoie *et al.*, 2020] Annie C Lajoie, Anne-Louise Lafontaine, R John Kimoff, and Marta Kaminska. Obstructive sleep apnea in neurodegenerative disorders: Current evidence in support of benefit from sleep apnea treatment. *Journal of Clinical Medicine*, 9(2):297, 2020.

[McDermott and Brown, 2020] Mollie McDermott and Devin L Brown. Sleep apnea and stroke. *Current Opinion in Neurology*, 33(1):4–9, 2020.

[Mostafa *et al.*, 2019] Sheikh Shanawaz Mostafa, Fábio Mendonça, Antonio G Ravelo-García, and Fernando Morgado-Dias. A systematic review of detecting sleep apnea using deep learning. *Sensors*, 19(22):4934, 2019.

[Otaghi *et al.*, 2017] Masoume Otaghi, Samiramiss Qavam, Siros Norozi, Milad Borji, and Mohsen Moradi. Investigating the effect of lavender essential oil on sleep quality in patients candidates for angiography. *Biomedical and Pharmacology Journal*, 10(1):473–478, 2017.

[Redmon *et al.*, 2016] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016.

[Strollo Jr and Rogers, 1996] Patrick J Strollo Jr and Robert M Rogers. Obstructive sleep apnea. *New England Journal of Medicine*, 334(2):99–104, 1996.

[Tadic *et al.*, 2020] Marijana Tadic, Cesare Cuspidi, Guido Grassi, and Giuseppe Mancina. Obstructive sleep apnea and cardiac mechanics: how strain could help us? *Heart Failure Reviews*, pages 1–9, 2020.

[Thorey *et al.*, 2019] Valentin Thorey, Albert Bou Hernandez, Pierrick J Arnal, and Emmanuel H Durling. Ai vs humans for the diagnosis of sleep apnea. In *2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 1596–1600. IEEE, 2019.

[Van Steenkiste *et al.*, 2018] Tom Van Steenkiste, Willemijn Groenendaal, Dirk Deschrijver, and Tom Dhaene. Automated sleep apnea detection in raw respiratory signals using long short-term memory neural networks. *IEEE journal of biomedical and health informatics*, 23(6):2354–2364, 2018.

⁴<https://www.smartnora.com/>