iToBoS 2024 - Skin Lesion Detection with 3D-TBP

Overview

The goal of this competition is to develop state-of-the-art machine learning techniques for detecting multiple skin lesions in clinical images. These images resemble photographs taken with a standard camera or smartphone, often used for digital healthcare services. Your algorithm could be valuable in environments lacking specialized care and could serve as a critical initial step in facilitating early detection of abnormalities in suspicious lesions.

Description

Diagnosing skin cancer traditionally relies heavily on the expertise of dermatologists and the use of dermatoscopy. This non-invasive approach uses a dermoscope to enhance the view of sub-macroscopic structures in pigmented skin lesions, which vary widely across dermatological conditions. While dermoscopy has improved diagnostic precision, its accuracy is closely tied to the clinician's level of expertise. Further, the process of taking a dermoscopic image of every suspicious lesion is a very tedious process. This necessitates the need for computer-aided diagnosis (CAD) systems using conventional cameras, especially in environments where dermatological expertise is scarce. This would enable non-specialist practitioners, such as general physicians without dermatological training, to identify suspicious lesions with ease, thereby facilitating early detection of abnormalities and allowing for timely intervention and improved prognosis. Additionally, tracking changes in suspicious lesions over time becomes feasible, enabling researchers to study the progression of the disease and assess the effectiveness of various treatments.

Detecting lesions in different regions of the body is a critical initial step in gaining such valuable insights to make informed decisions regarding patient care and treatment. This competition, therefore, challenges you to develop state-of-the-art machine learning techniques for detecting multiple skin lesions in clinical images. These images are tiles extracted from anonymized 3D avatars generated by the Canfield VECTRA WB360 system that captures comprehensive images of each patient's entire skin surface. As such, this competition leverages 3D total body photography (TBP) to present a novel dataset comprising multiple scans of hundreds of patients across two different continents.

Your work will contribute towards advancing the timely diagnosis and treatment of skin cancer. Therefore, we urge all teams, regardless of their placement in the competition, to publish a manuscript on arXiv detailing their solution and to open-source their code.

Evaluation

This competition is evaluated on the mean average precision at different intersection over union (IoU) thresholds. The IoU of a set of predicted bounding boxes and ground truth bounding boxes is calculated as:

$$IoU(A,B) = \frac{A \cap B}{A \cup B}$$

The metric sweeps over a range of IoU thresholds, at each point calculating an average precision value. The threshold values range from 0.5 to 0.75 with a step size of 0.05. In other words, at a threshold of 0.5, a predicted object is considered a "hit" if its intersection over union with a ground truth object is greater than 0.5.

At each threshold value *t*, a precision value is calculated based on the number of true positives (TP), false negatives (FN), and false positives (FP) resulting from comparing the predicted object to all ground truth objects:

$$\frac{TP(t)}{TP(t) + FP(t) + FN(t)}$$

A true positive is counted when a single predicted object matches a ground truth object with an IoU above the threshold. A false positive indicates a predicted object had no associated ground truth object. A false negative indicates a ground truth object had no associated predicted object.

Important note: if there are no ground truth objects at all for a given image, ANY number of predictions (false positives) will result in the image receiving a score of zero, and being included in the mean average precision.

The average precision of a single image is calculated as the mean of the above precision values at each IoU threshold:

$$\frac{1}{|\text{thresholds}|} \sum_t \frac{TP(t))}{TP(t) + FP(t) + FN(t)}$$

Lastly, the score returned by the competition metric is the mean taken over the individual average precisions of each image in the test dataset.

iToBoS

Submission File

For each ID in the test set, you must predict a space delimited set of bounding boxes. For instance, in the example below,

`abc123,0 0 50 50 0`

indicates that image `abc123` has a bounding box at `x_min`=0, `y_min`=0, `x_max`=50, `y_max`=50 with `class_id`=0. Note that the `class_id` should always be 0 since there is only one object class.

The file should contain a header and have the following format. Each row in the file must contain the bounding boxes for **all** the lesions identified in an image:

```
image_id,prediction_string
abc123,0 0 50 50 0
efg456,
hij789,0 0 50 40 0 60 80 30 30 0
etc...
```

Timeline

- **January 27, 2025** - Entry deadline. You must accept the competition rules before this date in order to compete.
- **January 27, 2025** - Team Merger deadline. This is the last day participants may join or merge teams.
- **February 09, 2025** - Final submission deadline.

All deadlines are at 11:59 PM CET on the corresponding day unless otherwise noted. The competition organizers reserve the right to update the contest timeline if they deem it necessary.

Workshop

The iToBoS consortium will host an online-only workshop on **February 28, 2025**, featuring a recap of this competition and talks from leading researchers and industrial experts in the field. Winning teams and selected participants will also be invited to present their solutions. The workshop is open for all interested researchers, students, and enthusiasts to attend free of charge. The detailed schedule and format will be updated on the challenge page and on the iToBoS website shortly.