

Multi-task Deep Learning for Glaucoma Detection From Color Fundus Images

Lucas Pascal

EURECOM

Oscar J. Perdomo

Universidad del Rosario

Xavier Bost

Orkis, Aix-en-Provence

Benoit Huet

Median Technologies

Sebastian Otálora

University Hospital of Bern

Maria A. Zuluaga (✉ maria.zuluaga@eurecom.fr)

EURECOM

Research Article

Keywords:

Posted Date: February 24th, 2022

DOI: <https://doi.org/10.21203/rs.3.rs-1357713/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Multi-task deep learning for glaucoma detection from color fundus images

Lucas Pascal^{1,2}, Oscar J. Perdomo³, Xavier Bost², Benoit Huet⁴, Sebastian Otálora^{5,+}, and Maria A. Zuluaga^{1,*,+}

¹Data Science Department, EURECOM, Sophia Antipolis, 06410, France

²Orkis, Aix-en-Provence, 13290, France

³School of Medicine and Health Sciences, Universidad del Rosario, Bogotá, Colombia

⁴Median Technologies, Valbonne, 06560, France

⁵Support Center for Advanced Neuroimaging–University Institute of Diagnostic and Interventional Neuroradiology, Bern, 3010 Switzerland

*maria.zuluaga@eurecom.fr

+these authors contributed equally to this work

ABSTRACT

Glaucoma is an eye condition that, if not diagnosed in time, leads to loss of vision and blindness. While diagnosing campaigns are regularly launched, these require human experts to estimate in a limited time subtle changes in the shape of the optic disc from retinal fundus images. Automatic glaucoma detection methods are desirable to help with the screening of patients, reducing the diagnosis time by analyzing the increasing amount of images. Deep learning methods have been satisfactory to classify and segment diseases in retinal fundus images. To guarantee a good generalization, these methods require vast amounts of annotations, which can be highly problematic given that existing annotated glaucoma datasets contain at most few hundreds of samples. Furthermore, deep learning techniques are computationally greedy, which can be a problem in scenarios with limited resources. The objective of this work is to make optimal use of the few annotated data at disposal and obtain better generalization performance of glaucoma diagnosis deep learning models. We achieve this by designing and training a novel multi-task deep learning model that leverages the similarities of related eye-fundus tasks and measurements used in glaucoma diagnosis. The proposed multi-task deep learning model with independent optimizers (MTL-IO) simultaneously learns different segmentation and classification tasks, thus benefiting from their similarity. The framework uses a U-Net encoder-decoder convolutional network as backbone architecture. Our MTL-IO approach was trained in a retinal fundus glaucoma challenge dataset, including 1200 retinal fundus images (400 for training, 400 for validation, 400 for testing) from different cameras and medical centers. MTL-IO obtained a 92.61 ± 0.38 AUC performance compared to a 87.37 ± 1.51 obtained by the same backbone network trained to detect glaucoma, also outperforming two other state-of-the-art multi-task learning models. Our approach pairs with trained experts and uses ~ 3.5 times fewer parameters than training each task separately. The data and the code for reproducing our results are publicly available.

Introduction

Glaucoma is one of the leading causes of irreversible but preventable blindness in working-age populations¹, which relates to an abnormal fluid balance in the eye that causes an increase in internal ocular pressure. The increased pressure gradually damages the eye optic nerve. If not diagnosed, these induced damages may lead to permanent vision loss. In 2020, it affected approximately 11.2 million people^{2,3}.

While an early diagnosis is critical to prevent irreversible damages, patients affected by glaucoma usually do not present symptoms in the early stages of the disease. It is thus essential to develop inexpensive detection methods to massively and systematically control patients before the symptoms appear. One way to achieve this is by performing a visual examination of the posterior pole or retinal fundus image. Specialized cameras obtain the color fundus images in a short image acquisition time. The analysis of the fundus images is performed by ophthalmologists, where the most discriminant symptom for detecting glaucoma on fundus images is the presence of a "cupping," which is the retraction of the optic disc (OD) on the optic cup (OC). This cupping causes an increase in the vertical Cup-to-Disc ratio (vCDR), which is the height ratio between the OC and OD. However, establishing an accurate diagnosis from these images is particularly difficult and prone to error in the accurate estimation of vCDR.

Deep convolutional networks have shown to be beneficial in medical imaging and in tasks of disease classification in eye fundus⁴⁻⁷, learning relevant features and patterns directly from images. Over the last years, glaucoma detection using deep

learning models reached a remarkable performance at the pair with residents in ophthalmology^{3,8-11}, thus representing a viable alternative to support current visual assessment. However, automating glaucoma diagnosis suffers from lack of data. Existing annotated datasets therefore contain a few hundred samples, while deep learning models require extensive databases to guarantee a good generalization. Moreover, they include millions of trainable parameters, requiring significant computational resources for training and deployment^{12,13}. Therefore, it is essential to develop methods that can make the most from the limited resources: computational requirements and the available annotated images, thus operating in a low data size regime while guaranteeing a good generalization. However, automating glaucoma diagnosis with deep learning models suffers from a lack of data. Existing annotated datasets contain a few hundred samples, while deep learning models require extensive databases to guarantee a good generalization

Multi-task learning (MTL)¹⁴ is a learning paradigm that aims to improve generalization by using the domain information contained in the training signals of related tasks as an inductive bias. In practice, this is done by training a shared model for all tasks. In deep MTL, the shared model consists in the parameters of a deep network¹⁵, hence, the resulting model is smaller than having separate networks for each task. Thanks to these features, MTL is a well-suited approach to automated glaucoma detection, where multiple tasks as OC and OD segmentation, and fovea localization are pre-requisite tasks for computer aided diagnosis (CAD) of retinal diseases¹⁶. The fovea localization task is related to the OD being located from the center of it by about 2-3 times the diameter of the OD¹⁷. Despite being related tasks, the use of MTL to simultaneously segment the OD and OC, locate the fovea and detect if the image is or not glaucomatous has not yet been explored, to the best of our knowledge. Instead, current state-of-the-art works treat each task separately through single task models (STL) or propose MTL approaches that do not exploit the full set of available tasks.

Among STL approaches, Cheng et al.¹⁸ proposed a super-pixel-based segmentation of the OD and OC for glaucoma screening, achieving a performance in terms of area under the curve (AUC) of 0.822. Fu et al.¹⁹ obtained 0.899 with a U-Net-based deep learning method and a transformation of the image to polar coordinates. Among the authors that have explored MTL techniques, Mojab et al. proposed a multi-task model for glaucoma detection composed of two modules for OD and OC segmentations and glaucoma prediction²⁰, obtaining 90.1 of F-score; the authors did not account for the dissimilarity between the distributions of the segmentation and prediction tasks. Chelaramani reported a novel MTL-based teacher ensemble method for knowledge distillation²¹. The proposed method requires a dataset with a variety of different eye pathologies, which may be difficult to obtain in practice.

This work aims to determine if the relation between tasks associated to glaucoma CAD, i.e. OD and OC segmentation, fovea location and glaucoma detection, can be exploited within an MTL framework to improve model generalization and accuracy for glaucoma detection in a low sample size, low computational resources regime. To this end, a deep MTL model is trained to leverage the similarities of the segmentation of the OD and OC tasks, together with localization of the fovea to detect the presence of glaucoma in retinal fundus images. The proposed MTL approach uses a U-Net encoder-decoder convolutional network as a backbone architecture and adapts it to handle the four tasks using independent optimizers (IO) that can simultaneously learn the segmentation and classification tasks. We denote it MTL-IO. We evaluate our method using the Retinal Fundus Glaucoma Challenge (REFUGE) dataset, including 1200 retinal fundus images (400 for training, 400 for validation, 400 for testing) from different cameras and medical centers, achieving better AUC performance than the same network trained for the single task of detecting glaucoma (92.91 ± 0.69 vs 90.09 ± 2.70). Our approach pairs with trained experts^{22,23} and uses approximately 3.5 times fewer parameters than training each task separately.

Results

This section presents the experimental results obtained on the REFUGE challenge dataset, comparing the proposed MTL-IO framework in different setups and against different baselines.

Multi-Task learning model with independent optimizers

Table 1 shows the classification and segmentation results of the proposed MTL-IO and the respective single task model (STL) for each of the tasks. Performance is measured in terms of the area under the curve (AUC) for the classification tasks, the Dice score (DSC) for the segmentation tasks, and the L2-distance (Fovea Error) for the localization task.

We compared our approach with two state-of-the-art multitask models: GradNorm²⁴ and PCGrad²⁵. GradNorm²⁴ adaptively balances the losses by gradient normalization, whereas PCGrad is based on estimating the right signs in the independent task gradients to avoid local minima. All models were trained three times to account for the variation in weights initialization. The standard deviation is reported for every performance measure. Model size, in terms of number of parameters (#P), and an iteration time (time), which represents the seconds required for a forward and backward pass in the framework, are also reported.

MTL-IO outperforms all other methods in glaucoma detection, OC segmentation and the fovea localization task. It ranks second in the OD segmentation task. In terms of model size, Mall MTL models use approximately $17.2e6$ parameters, making

them ~ 3.5 significantly lighter than the STL baseline, which uses 61.2e6 parameters. We estimate the parameters of STL as the sum of parameters of each single-task learner. In terms of computational time, the cumulative iteration time for STL is ~ 1.2 times slower than MTL-IO. MTL-IO's training iteration time is comparable to PCGrad, but much slower than GradNorm. This difference is explained by the use of the independent optimizers that incur in a computational overhead, which is compensated by the improved performance.

Figure 1 displays the ROC curves for the glaucoma detection task of STL and MTL-IO. It suggests that the multitask classifier benefits from the related tasks to achieve better performance than the single task of glaucoma detection (AUC=0.926 vs. 0.874) with all the operating points after 0.1 false positive rate.

We also analyzed the sensitivity of MTL-IO to different learning rates at training. Figures 2 and 3 respectively show each task best metric score and minimum loss values, for each of the explored learning rates on the validation set. MTL-IO obtains better results over the different learning rates on the glaucoma detection and fovea localization tasks. The STL model then performs better on the segmentation tasks. However, one can notice that it suffers a more important performance drop on the OC segmentation task when evaluating on the test set (see Table 1), suggesting less overfitting for MTL-IO.

Figure 4 shows an example of the segmentation of a Glaucomatous eye. The proposed MTL strategy provides a better segmentation in this challenging case, with a distinctive light dome in the middle of the eye, probably due to poor capture conditions. It is a glaucomatous case, although the vCDR does not suggest it.

Multi-Task learning model with independent optimizers and transfer learning

Transfer Learning is a widely adopted method to bias a model with prior knowledge on an input domain and lead it to better generalization on new data. In practice, Imagenet²⁶ pre-trained models have proven to be profitable on a large majority of vision tasks. In medical imaging, although the input domain is different from the Imagenet domain (natural images), the benefits are still noticeable²⁷, and particularly appreciated to compensate for the usual lack of training data. Its combination with Multi-Task Learning strategies studied here is thus relevant. As Imagenet only involves image classification, there exist no Imagenet pre-trained models for semantic segmentation. However, it is possible to use a pre-trained VGG-16²⁸ network for the encoding part of the U-Net in the pipeline, while the decoder is initialized from scratch.

When using the pre-trained VGG encoder the models had better performance, MTL reaching a 96.15 ± 0.14 AUC performance in comparison with 92.61 ± 0.38 of the MTL strategy with weights trained from scratch. These results are presented in Table 2.

Discussion

MTL-IO improves generalization by using a unique neural network to learn all tasks jointly. It outperforms the other MTL and STL baselines on three of the four proposed tasks, in particular on the glaucoma detection task by a large margin.

When comparing the performance of multi-task and single-task models, it is interesting that the other state of the art MTL methods GradNorm²⁴ and PCGrad²⁵ perform worse than the single task baselines on every task, highlighting a task interference issue. Instead, when using the proposed MTL-IO optimization scheme, the multi-task network can significantly reduce task interference and often improve performances compared to the single-task baselines.

In addition to the improved performance, MTL-IO has the advantage that it uses a unique convolutional network for all tasks. This means that it achieves a good performance while being more lightweight than single-task learners: STL is ~ 3.5 times larger in terms of parameters and ~ 1.2 times slower than MTL-IO. This is an important feature for real-world use, where resources are often constrained. Our experiments combining transfer learning suggested that the gains achieved by MTL-IO, both in terms of both generalization performance and computational efficiency by reusing the parameters of the OD, OC, and fovea localization branches into the glaucoma detection task, hold in smaller proportions. Although STL observes larger improvements, the MTL-IO remains the best performing at glaucoma detection, which is the main task. As such, it is possible to say that the two strategies, MTL and transfer learning, can be efficiently combined in real-world contexts to create better generalization performance on problems involving multiple tasks.

Despite the above-mentioned advantages, a disadvantage of MTL strategies relates to the extra effort that may be required from a user/expert to put them in place. While an STL strategy requires simple binary labels for training (i.e. presence or absence of glaucoma), MTL techniques also need pixel-wise annotations of the objects to segment and the location of the fovea. All of these annotation tasks are more time consuming and costly. In such setup, it is therefore necessary to assess what is the most critical criterion to optimize. If access to experts for image annotation is difficult, an STL classifier should be used. Instead, if lack of data and limited resources are an important constraint, MTL techniques should be favored.

Materials and Methods

Materials: REFUGE challenge dataset

In 2018 the *Retinal Fundus Glaucoma Challenge* (REFUGE) was launched as a satellite event at the 2018 MICCAI conference. For this event, 1200 retinal fundus images (400 for training, 400 for validation, 400 for testing) from different cameras and medical centers have been collected and annotated by human experts. Annotations were provided for four different tasks: glaucoma diagnosis, optic disc segmentation, optic cup segmentation, and fovea localization. For the diagnosis task, the ground truth is provided as binary labels, attesting to the presence of glaucoma. In the segmentation tasks, the regions defined by the OD (optic nerve head) and the OC (the white elliptic region located inside the optic disc) are provided as binary segmentations. In the fovea localization case, the ground truth is given as the fovea's (x, y) pixel location. All the methods developed and experiments were carried out in accordance with the relevant guidelines and regulations associated to this publicly available dataset.

Methods

In the following we describe the overall MTL deep learning architecture adopted, the loss functions used for each task and, finally, the independent optimizer (IO) strategy adopted.

Multitask deep learning architecture

We use a U-Net²⁹, an encoder-decoder convolutional network, with a VGG-16³⁰ structure and added skip connections between equivalent depths of encoder and decoder, which allow the decoder to recover fine-grained details through the multiple upscalings. This network is well known for solving efficiently biomedical segmentation tasks³¹. Our MTL approach uses this architecture for two segmentation tasks (OD and OC), one regression task (fovea coordinates) and one classification task (glaucoma diagnosis). The design of the MTL architecture is shown in Figure 5, and detailed in the following.

Optic disc and cup segmentation tasks. The OD and OC segmentation masks are through the convolutional layer after the shared decoder for each task. Similar to existing works⁹, the segmentations of OD and OC are refined through a post-processing step that keeps the main connected component in the prediction map to remove possible prediction noise around these elliptic regions.

Fovea localization. The fovea localization task is addressed as a segmentation task: from the ground truth coordinates of the fovea, a map is created, the center of such map represents the localization of the fovea. The map is a multivariate normal distribution centered in the coordinates (equal variances and null covariances). An example is shown in figure 6 (right). The network is trained to fit the maps with a task-respective convolutional layer on the shared decoder. The fovea coordinates are then predicted as the center of mass of the predicted saliency map. In this case, no refinement or postprocessing is performed since it may shift the center of mass.

Glaucoma detection task. The glaucoma detection task (classification) consists of two steps:

1. A prediction is obtained from a fully connected layer, branched after the U-Net encoder (FC classifier).
2. Similarly to some previous works⁹, a second prediction is obtained from a logistic regression classifier (Linear classifier), taking as input the vertical Cup-to-Disc Ratio (vCDR) obtained from the OD and OC segmentation tasks. The vCDR is computed as:

$$vCDR = \frac{OC_{height}}{OD_{height}}$$

with OC_{height} and OD_{height} the heights of the OC and OD, obtained from the segmentation branches.

The outputs before the binary outcome of each classifier are averaged. The final classification is obtained by using a threshold of 0.5 over this average.

Loss functions

Were we present the loss functions used for the optimization of the different tasks.

OD and OC segmentation. The OD and OC segmentation tasks both use a binary cross-entropy loss (BCE), averaged over every pixel i of the segmentation maps:

$$\mathcal{L}_{BCE}(p, y) = -\frac{1}{N_{pix}} \sum_{i=1}^{N_{pix}} y_i \log(p_i) + (1 - y_i) \log(1 - p_i)$$

with p, y and N_{pix} respectively the prediction, ground-truth and number of pixels.

Fovea localization. For the fovea localization task, the network is trained to fit the pre-processed saliency maps with a L1-loss, since the map values are not binarized:

$$\mathcal{L}_{L1}(p, y) = \sum_i |y_i - p_i|$$

Afterwards, the predicted fovea location is computed as the center of mass of the predicted saliency map.

Glaucoma classification. For the glaucoma classification task, a focal loss³² is used to better handle the unbalance between positive and negative samples (only 10% of positives):

$$\mathcal{L}_{Focal}(p, y) = (1 - p_t)^\gamma \log(p_t)$$

with

$$p_t = \begin{cases} p & \text{if } y = 1 \\ (1 - p) & \text{otherwise} \end{cases}$$

Concretely, this loss multiplies the usual binary cross-entropy term with a classification uncertainty term $(1 - p_t)$ to give more importance to uncertain classifications, i.e., these of low populated classes. We set the hyperparameter γ to 2 in our experiments.

MTL independent optimizer optimization strategy

In the following, we present the IO optimization strategy used in this work. It relies on the alternative optimization scheme, alternating independent gradient descent steps on the different task-specific objective functions, as proposed by Pascal et al.³³. We detail then main steps leading to this optimization scheme, and refer the interested reader to Pascal et al.³³ for more details.

The standard MTL optimization setup with an aggregated loss¹⁴ can be expressed as:

$$\mathcal{L}(w_t, \xi_t) = \sum_{k=1}^N c^{(k)} \cdot \mathcal{L}^{(k)}(w_t, \xi_t)$$

where $\mathcal{L}^{(k)}$ is the loss function associated to k^{th} out of N task, w_t the shared parameters, and ξ_t the data sample, at iteration t . $c^{(k)}$ are task-specific weighting coefficients, for which we assume uniform weighting, i.e. $c^{(k)} = 1$. If $g^{(k)}$ denotes the derivative of $\mathcal{L}^{(k)}$ with respect to the shared parameters w , the update rule for w at step $t + 1$ using stochastic gradient descent is:

$$w_{t+1} = w_t - \eta_t \cdot \sum_{k=1}^N g^{(k)}(w_t, \xi_t) \quad (1)$$

where η_t is the learning rate.

Recent works^{15,34,35} propose a variation to the update rule in equation 1, in which alternate independent update steps with respect to the different task-specific loss functions are executed, instead of aggregating all the terms at once. This strategy aims to minimize task interference and, hence improve generalization. The alternate update rule can be expressed as:

$$w_{t+1}^{(k)} = \begin{cases} w_t^{(N)} - \eta_t \cdot g^{(k)}(w_t^{(N)}, \xi_t), & k = 1 \\ w_t^{(k-1)} - \eta_t \cdot g^{(k)}(w_t^{(k-1)}, \xi_t), & \forall k > 1 \end{cases} \quad (2)$$

In this work, we adopt the approach from Pascal et al.³³. It uses a modified alternate update rule (eq. 2) that allows to use individual optimizers (IO) in the form of individual exponential moving averages for each task, to prevent state-of-the-art optimizers (e.g. Adam) from accumulating and mixing previous gradient descent directions of all the different tasks. The modified update rule can be expressed as:

$$w_{t+1}^{(k)} = \begin{cases} w_t^{(N)} - \eta_t \cdot \hat{m}^{(k)} \left(g^{(k)}(w_t^{(N)}, \xi_t) \right), & k = 1 \\ w_t^{(k-1)} - \eta_t \cdot \hat{m}^{(k)} \left(g^{(k)}(w_t^{(k-1)}, \xi_t) \right), & \forall k > 1 \end{cases} \quad (3)$$

where $\hat{m}^{(k)}$ is a task-specific exponential moving average mechanism. Here, the memory term introduced by $m^{(k)}$ only involves previous updates of task k . Such formulation is equivalent to using one independent optimizer per task, and is therefore denoted as MTL-IO. In this paper, we use MTL-IO to denote the complete pipeline.

Implementation details

All methods were implemented in Pytorch 1.2, and ran on NVIDIA Titan XP graphic cards.

Data and Materials Availability Statement

The data used to run train our models and run experiments is available, upon registration from the REFUGE Challenge (<https://refuge.grand-challenge.org/Home2020/>). All code to reproduce the results of this article is available in a GitHub repository (https://github.com/robustml-eurecom/glaucoma_mtl). The code can be anonymously downloaded from the following link: https://github.com/robustml-eurecom/glaucoma_mtl/archive/refs/heads/main.zip.

References

1. Weinreb, R. N., Aung, T. & Medeiros, F. A. The pathophysiology and treatment of glaucoma: a review. *Jama* **311**, 1901–1911 (2014).
2. Tham, Y.-C. *et al.* Global prevalence of glaucoma and projections of glaucoma burden through 2040: a systematic review and meta-analysis. *Ophthalmology* **121**, 2081–2090 (2014).
3. Shibata, N. *et al.* Development of a deep residual learning algorithm to screen for glaucoma from fundus photography. *Sci. reports* **8**, 1–9 (2018).
4. Litjens, G. *et al.* A survey on deep learning in medical image analysis. *Med. image analysis* **42**, 60–88 (2017).
5. Christopher, M. *et al.* Performance of deep learning architectures and transfer learning for detecting glaucomatous optic neuropathy in fundus photographs. *Sci. reports* **8**, 1–13 (2018).
6. Shibata, N. *et al.* Development of a deep residual learning algorithm to screen for glaucoma from fundus photography. *Sci. reports* **8**, 1–9 (2018).
7. Graziani, M. *et al.* Improved interpretability for computer-aided severity assessment of retinopathy of prematurity. In *Medical Imaging 2019: Computer-Aided Diagnosis*, vol. 10950, 109501R (International Society for Optics and Photonics, 2019).
8. Chen, X. *et al.* Automatic feature learning for glaucoma detection based on deep learning. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 669–677 (Springer, 2015).
9. Orlando, J. I. *et al.* Refuge challenge: A unified framework for evaluating automated methods for glaucoma assessment from fundus photographs. *Med. image analysis* **59**, 101570 (2020).
10. Hemelings, R. *et al.* Deep learning on fundus images detects glaucoma beyond the optic disc. *Sci. Reports* **11**, 1–12 (2021).
11. Gheisari, S. *et al.* A combined convolutional and recurrent neural network for enhanced glaucoma detection. *Sci. reports* **11**, 1–11 (2021).
12. Justus, D., Brennan, J., Bonner, S. & McGough, A. S. Predicting the computational cost of deep learning models. In *2018 IEEE international conference on big data (Big Data)*, 3873–3882 (IEEE, 2018).
13. Strubell, E., Ganesh, A. & McCallum, A. Energy and policy considerations for modern deep learning research. In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, 13693–13696 (2020).
14. Caruana, R. Multitask learning. *Mach. learning* **28**, 41–75 (1997).
15. Pascal, L., Michiardi, P., Bost, X., Huet, B. & Zuluaga, M. A. Maximum roaming multi-task learning. In *35th AAAI Conference on Artificial Intelligence*, vol. 35, 9331–9341 (2021).
16. Xie, R. *et al.* End-to-end fovea localisation in colour fundus images with a hierarchical deep regression network. *IEEE Transactions on Med. Imaging* **40**, 116–128 (2021).
17. Welfer, D., Scharcanski, J. & Marinho, D. R. Fovea center detection based on the retina anatomy and mathematical morphology. *Comput. methods programs biomedicine* **104**, 397–409 (2011).
18. Cheng, J. *et al.* Superpixel classification based optic disc and optic cup segmentation for glaucoma screening. *IEEE transactions on medical imaging* **32**, 1019–1032 (2013).
19. Fu, H. *et al.* Joint optic disc and cup segmentation based on multi-label deep network and polar transformation. *IEEE transactions on medical imaging* **37**, 1597–1605 (2018).

20. Mojab, N., Noroozi, V., Philip, S. Y. & Hallak, J. A. Deep multi-task learning for interpretable glaucoma detection. In *2019 IEEE 20th International Conference on Information Reuse and Integration for Data Science (IRI)*, 167–174 (IEEE, 2019).
21. Chelaramani, S., Gupta, M., Agarwal, V., Gupta, P. & Habash, R. Multi-task knowledge distillation for eye disease prediction. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 3983–3993 (2021).
22. Azuara-Blanco, A., Burr, J., Thomas, R., MacLennan, G. & McPherson, S. The accuracy of accredited glaucoma optometrists in the diagnosis and treatment recommendation for glaucoma. *Br. J. Ophthalmol.* **91**, 1639–1643 (2007).
23. Li, Z. *et al.* Efficacy of a deep learning system for detecting glaucomatous optic neuropathy based on color fundus photographs. *Ophthalmology* **125**, 1199–1206 (2018).
24. Chen, Z., Badrinarayanan, V., Lee, C.-Y. & Rabinovich, A. GradNorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In *International Conference on Machine Learning*, 794–803 (PMLR, 2018).
25. Yu, T. *et al.* Gradient surgery for multi-task learning. *arXiv preprint arXiv:2001.06782* (2020).
26. Deng, J. *et al.* Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255 (Ieee, 2009).
27. Tajbakhsh, N. *et al.* Embracing imperfect datasets: A review of deep learning solutions for medical image segmentation. *Med. Image Analysis* **63**, 101693 (2020).
28. Jaderberg, M., Simonyan, K., Zisserman, A. *et al.* Spatial transformer networks. *Adv. neural information processing systems* **28**, 2017–2025 (2015).
29. 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).
30. Simonyan, K. & Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
31. Falk, T. *et al.* U-net – deep learning for cell counting, detection, and morphometry. *Nat. Methods* **16**, 67–70 (2019).
32. Lin, T.-Y., Goyal, P., Girshick, R., He, K. & Dollár, P. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, 2980–2988 (2017).
33. Pascal, L., Michiardi, P., Bost, X., Huet, B. & Zuluaga, M. A. Optimization strategies in multi-task learning: Averaged or independent losses? *arXiv preprint arXiv:2109.11678* (2021).
34. Maninis, K.-K., Radosavovic, I. & Kokkinos, I. Attentive Single-Tasking of Multiple Tasks. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 1851–1860 (2019).
35. Bragman, F. J., Tanno, R., Ourselin, S., Alexander, D. C. & Cardoso, J. Stochastic filter groups for multi-task cnns: Learning specialist and generalist convolution kernels. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1385–1394 (2019).

Acknowledgements

Sebastian Otálora thanks the Colombian science ministry Minciencias for partially funding his Ph.D. studies through the call "756 - Doctorados en el exterior". The French government supports Maria A. Zuluaga through the 3IA Côte d'Azur Investments in the Future project managed by the National Research Agency (ANR) (ANR-19-P3IA-0002).

Author contributions statement

L.P., M.A.Z. and S.O. conceived the methodology and experiments, L.P. conducted the experiments, L.P., O.J.P., S.O. and M.A.Z. analyzed the results, L.P., O.J.P and S.O. wrote the original draft, S.O. and M.A.Z. supervised the work. All authors reviewed, edited, and approved the manuscript.

Competing interests statement

The authors declare not to have any competing interest.

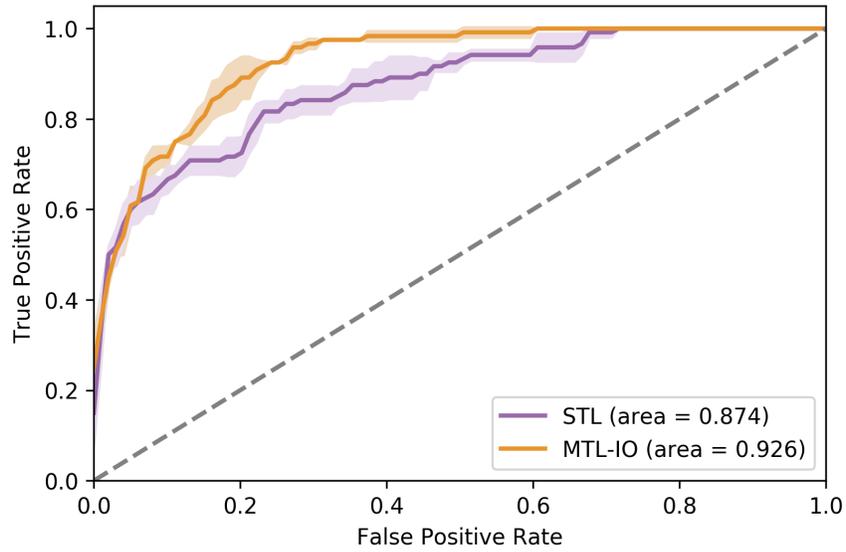


Figure 1. Receiver operating characteristic (ROC) curve for the glaucoma detection task for the single task learning model (STL) and our multi-task learning (MTL-IO) approach.

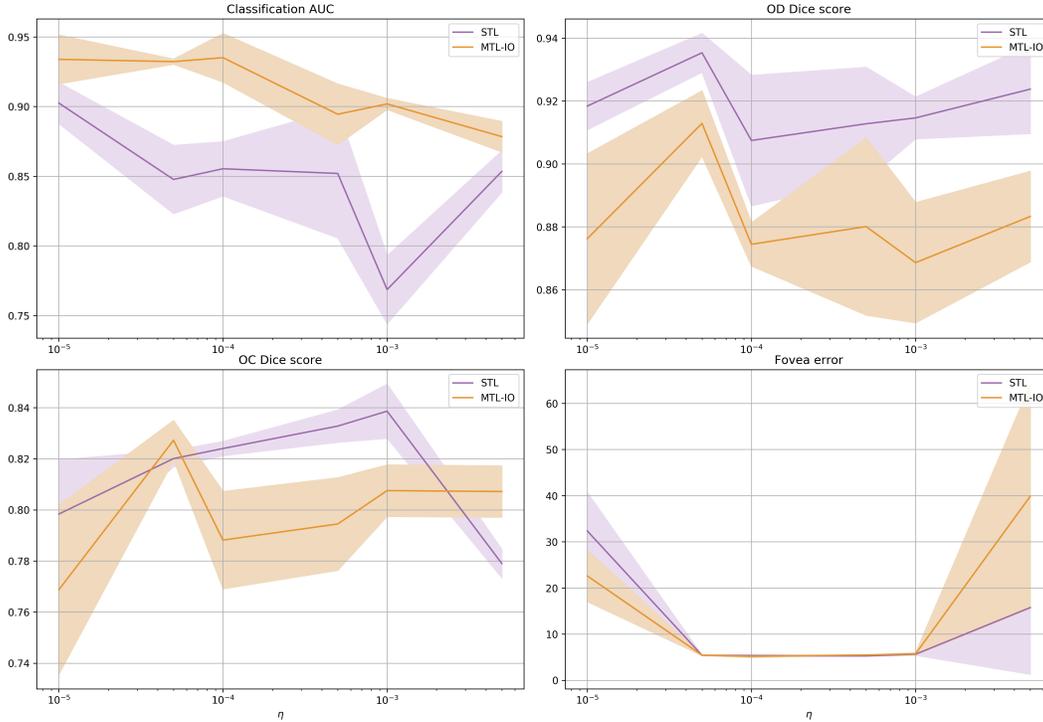


Figure 2. Performance vs. learning rate. Per task performance as a function of the learning rate (η) for the single task learning model (STL) and our multi-task learning (MTL-IO) approach. Standard deviation in shaded colour.

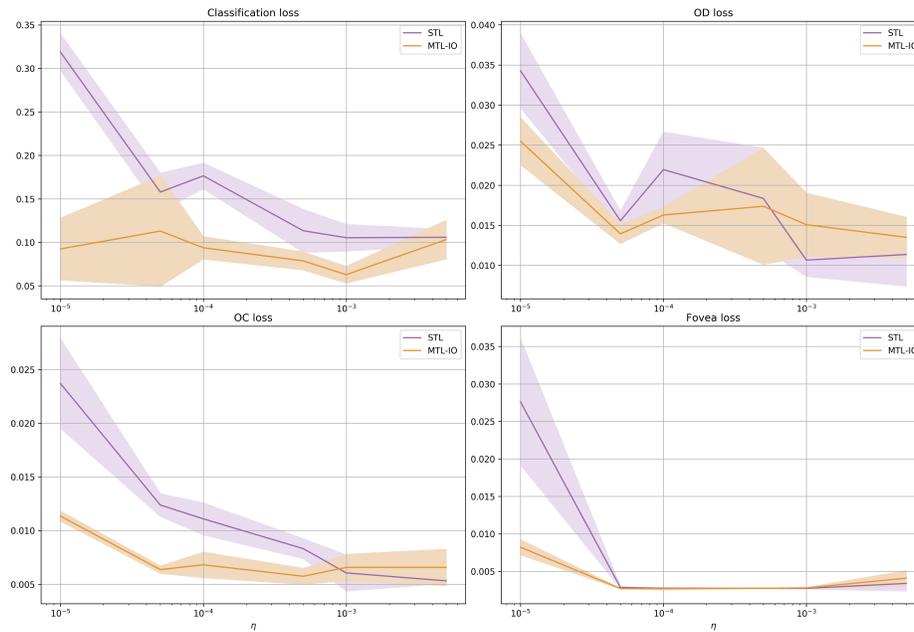


Figure 3. Loss vs. learning rate. Final loss values as a function of the learning rate (η) for the single task learning model (STL) and our multi-task learning (MTL-IO) approach. MTL-IO tends to have lower loss values, suggesting the benefit of learning also from other tasks. Standard deviation in shaded colour.

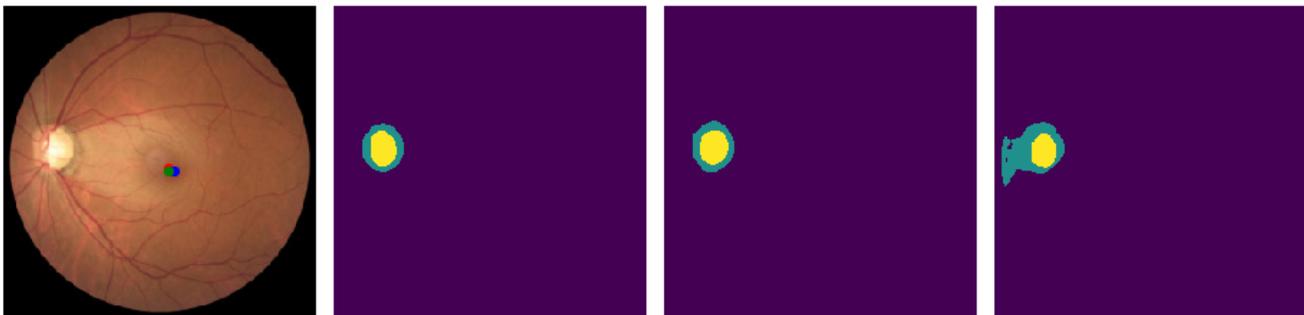


Figure 4. Glaucomatous image from the test set (first image), where the three dots represent the fovea location's ground truth (red), the MTL-IO prediction (blue) and the STL prediction (green), followed by OD (green) and OC (yellow) ground truth (second), MTL-IO (third) and STL (fourth) segmentation masks.

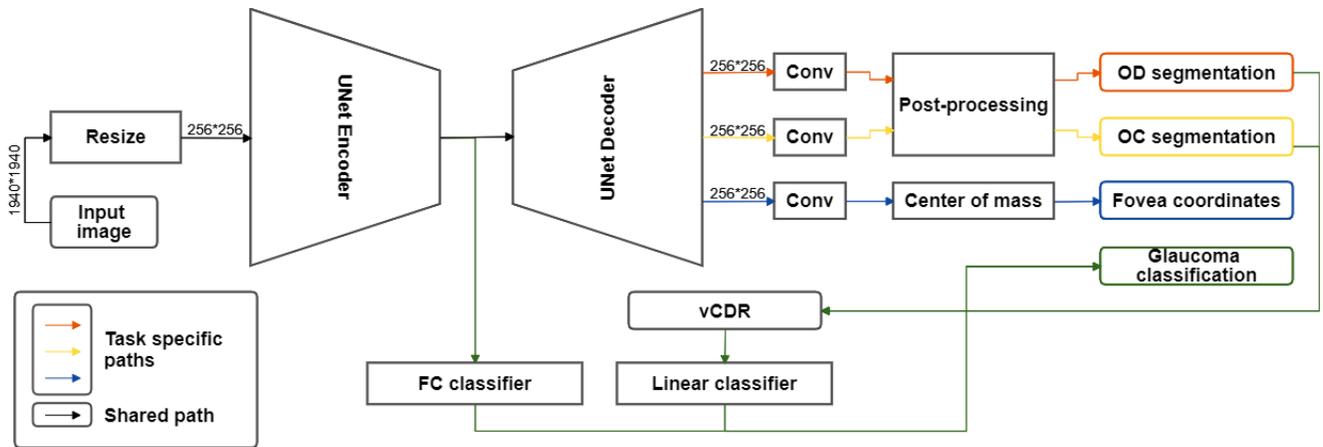


Figure 5. Multi-task learning framework for glaucoma detection, OD and OC segmentation, and fovea localization. The framework uses a U-net as its backbone architecture.

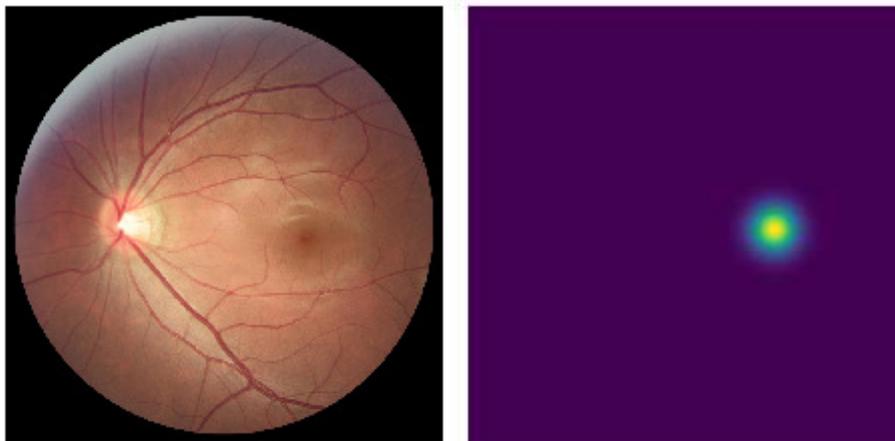


Figure 6. Example of a retinal fundus image (left) and the correspondent saliency map centered on the fovea coordinates (right).

Models			Tasks			
			Glaucoma	OD	OC	Fovea
Model	#P \times 1e6	time	AUC (\uparrow)	DSC (\uparrow)		Fovea Error (\downarrow)
STL	61.2	0.686	87.37 \pm 1.51	93.87 \pm 0.74	80.62 \pm 1.25	5.96 \pm 0.17
MTL-IO (ours)	17.2	0.557	92.61 \pm 0.38	91.61 \pm 0.64	81.21 \pm 0.99	5.71 \pm 0.24
GradNorm	17.2	0.260	86.31 \pm 1.35	80.32 \pm 1.10	68.93 \pm 4.78	28.48 \pm 0.65
PCGrad	17.2	0.556	85.51 \pm 1.71	85.23 \pm 1.52	71.08 \pm 7.40	34.45 \pm 0.33

Table 1. Results for the test set in the REFUGE dataset over the four tasks.

Model	Tasks			
	Glaucoma	OD	OC	Fovea
	AUC (\uparrow)	DSC (\uparrow)		Fovea Error (\downarrow)
STL	94.30 \pm 1.68	95.29 \pm 0.01	85.86 \pm 0.21	5.42 \pm 0.06
MTL-IO	96.15 \pm 0.14	94.24 \pm 0.38	83.95 \pm 0.90	5.22 \pm 0.18

Table 2. Results for the test set in the REFUGE dataset over the four different tasks. The results in this table correspond to the models trained with transfer learning. The same U-net architecture with an imagenet-pretrained VGG-encoder is used in all cases.