



HAL
open science

FULLY TRAINABLE GAUSSIAN DERIVATIVE CONVOLUTIONAL LAYER

Valentin Penaud–Polge, Santiago Velasco-Forero, Jesus Angulo

► **To cite this version:**

Valentin Penaud–Polge, Santiago Velasco-Forero, Jesus Angulo. FULLY TRAINABLE GAUSSIAN DERIVATIVE CONVOLUTIONAL LAYER. 29th IEEE International Conference on Image Processing (IEEE ICIP), Oct 2022, Bordeaux, France. hal-03722767

HAL Id: hal-03722767

<https://hal.science/hal-03722767>

Submitted on 13 Jul 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

FULLY TRAINABLE GAUSSIAN DERIVATIVE CONVOLUTIONAL LAYER

Valentin Penaud--Polge, Santiago Velasco-Forero, Jesus Angulo

Mines Paris, PSL University

Center for Mathematical Morphology

{valentin.penaud_polge, santiago.velasco, jesus.angulo} @minesparis.psl.eu

ABSTRACT

The Gaussian kernel and its derivatives have already been employed for Convolutional Neural Networks in several previous works. Most of these papers proposed to compute filters by linearly combining one or several bases of fixed or slightly trainable Gaussian kernels with or without their derivatives. In this article, we propose a high-level configurable layer based on anisotropic, oriented and shifted Gaussian derivative kernels which generalize notions encountered in previous related works while keeping their main advantage. The results show that the proposed layer has competitive performance compared to previous works and that it can be successfully included in common deep architectures such as VGG16 for image classification and U-net for image segmentation.

Index Terms— Convolutional Layer, Gaussian Derivatives, Local N-Jet

1. INTRODUCTION

The Gaussian kernel takes a major place in the image processing literature due to its proximity with the visual system and its scale-space representations. Several works have been published to determine the most adapted family of filters for scale-space representation by assuming constraints. All converged towards the Gaussian kernel and its derivatives to form the local N -jet [1–5]. Moreover, the Gaussian kernel coincides with the receptive field profiles of the human front-end visual system and from a mathematical viewpoint, the local N -jet is fundamental as it represents a local Taylor approximation of the studied image. This makes the local N -jet an efficient tool for image representation and description [6, 7].

For Convolutional Neural Networks, several approaches have been published considering either only the Gaussian kernel [8, 9] or the Gaussian kernel together with its derivatives [10–12]. The layers proposed in these references have the advantage to make the size of the filters independent of the

number of parameters. Tabernik et al. [8] proposes a convolutional layer whose filters are defined as mixtures of isotropic Gaussian kernels, also called mixture of displaced aggregation units (DAU). The proposed layer learns the mean of each Gaussian kernel and its weight in the mixture. The number of DAU in each filter is left as a hyperparameter and the scale of the DAU is either fixed [8] or learned by the network [13, 14]. In [9], it is proposed to compose free-form filters with Gaussian kernels to benefit from the diversity of variations of the free-form filters while adapting them to the shape of the structured filters. Nevertheless, as free-form filters are used, the number of parameters is similar to conventional convolutional layers. In [10], a layer using kernels defined to be linear combinations of a basis of isotropic Gaussian derivative kernels is proposed. More precisely, the basis of each layer is composed of local N -jets. Only the weights of the linear combinations are learned by the network while the order of the N -jets, their scales and their orientations are fixed and considered as hyperparameters. The proposed type of neural network is called Structured Receptive Field Neural Network (RFNN). The advantage of RFNN is shown when a small amount of training data is available as the number of parameters is reduced in comparison with standard CNN. RFNN has been reused in [12] to focus on the translation insensitivity of this network where the Gaussian derivative kernels incorporate Lipschitz continuity. Recently, a version of RFNN where the network learns the scales of the isotropic, non-oriented and non-shifted Gaussian derivative kernels has been published [15]. Finally, Lindeberg [11] presents a scale invariant neural network whose filters are similar to those used in [10]. To obtain scale in(equi)variance, Lindeberg reuses, in a modified manner, the principle of scale-channel networks [16].

None of these references investigated the possibility of a network learning all the parameters encountered in a mixture of anisotropic, shifted and oriented Gaussian derivatives. Moreover in all these references, some proposed layers have filters sharing a unique basis of Gaussian derivative kernels [10] for several coefficients of the linear combination, some proposed layers have filters sharing coefficients of the linear combinations with different scaled bases [11] and some layers uses one basis for each filter [8].

This work was granted access to the HPC resources of IDRIS under the allocation 2022-[AD011013367] made by GENCI. This work has been supported by Fondation Mathématique Jacques Hadamard(FMJH) under the PGMO-IRSDI 2019 program.




Fig. 1. (a): Schematic illustration of eq. (2). (b): Schematic illustration of eq. (4).

The aim of this paper is also to propose a set of layers whose filters can share a basis of Gaussian derivatives or have their own individual basis or “in-between” configurations. We will also propose two ways to linearly combine the Gaussian derivative kernels when dealing with multi-channels inputs.

2. TRAINABLE GAUSSIAN DERIVATIVE LAYERS

In this section we present the fully trainable Gaussian derivative convolutional layers. The layers are defined in a way that it is based on the local N -jet like RFNN and Lindeberg’s Gaussian derivative network to restrain the number of parameters. They also offer more complex filters, like it is the case in [9] but without the use of free-form filters. Finally they have adapted receptive fields as is the case in [8] and they generalize the notion of number of basis. First, we will recall the link between Gaussian derivatives and Hermite polynomials as it constitutes a convenient way of parameterising a deep learning convolution kernel into an anisotropic, shifted and oriented Gaussian derivative kernel.

Hermite Polynomials and Gaussian Derivatives. The one-dimensional centered Gaussian kernel and its p^{th} derivative are respectively given by $G(x, \sigma) = \frac{1}{(2\pi\sigma^2)^{\frac{1}{2}}} e^{-\frac{x^2}{2\sigma^2}}$ and $G_p(x, \sigma) = \frac{\partial^p}{\partial x^p} G(x, \sigma)$, where σ represents the standard deviation (or scale). Gaussian derivatives of any order can be expressed as a Hermite polynomial multiplied by a Gaussian kernel [17], and as Hermite polynomials are easily expressible, we will use them to formulate the anisotropic Gaussian derivative kernels without any use of derivation operations. The one-dimensional Hermite polynomial of order p is given by the following equation expressed in a form of a serie which makes it easier to use in practice $H_p(x) = \sum_{i=0}^{\lfloor p/2 \rfloor} \frac{(-1)^i p!}{i!(p-2i)!} (2x)^{p-2i}$. The relation between the one-dimensional Gaussian derivative of order p and the Hermite polynomial of order p is given by $G_p(x, \sigma) = \left(-\frac{1}{\sqrt{2}\sigma}\right)^p \sqrt{2\pi}\sigma H_p\left(\frac{x}{\sqrt{2}\sigma}\right) G\left(\frac{x}{\sqrt{2}\sigma}, \sigma\right)^2$. From

this, we can define the two-dimensional anisotropic Gaussian derivative kernel as follow: $G_{p,q}(x_1, x_2, \sigma_{x_1}, \sigma_{x_2}) = G_p(x_1, \sigma_{x_1}) G_q(x_2, \sigma_{x_2})$. For convenience, we will use the notation $\mathbf{x} = (x_1, x_2)$ and $\boldsymbol{\sigma} = (\sigma_{x_1}, \sigma_{x_2})$. The resulting filters can be oriented by an angle θ and shifted by $\boldsymbol{\mu}$ giving

$$G_{p,q}(\mathbf{u}, \boldsymbol{\mu}, \boldsymbol{\sigma}, \theta) = G_p(y_1, \sigma_{u_1}) G_q(y_2, \sigma_{u_2}), \quad (1)$$

with $(y_1, y_2) = (u_1 - \mu_1, u_2 - \mu_2)$ and

$$(u_1, u_2) = (x_1 \cos \theta + x_2 \sin \theta, -x_1 \sin \theta + x_2 \cos \theta)$$

. The bases used in the proposed layer are defined and denoted as $\{\phi^1, \dots, \phi^N\} = \{G_{p,q}(\mathbf{u}, \boldsymbol{\mu}, \boldsymbol{\sigma}, \theta) \mid p + q \leq K\}$ for a chosen $K \in \mathbb{N}$.

Fully Trainable Gaussian Derivative Convolutional Layers. For all the configurations, the input and output tensors of the proposed layer as well as an intermediate tensor will respectively be denoted as I, O and J with $I \in \mathbb{R}^{l_{in} \times h_{in} \times f_{in}}$, $J, O \in \mathbb{R}^{l_{out} \times h_{out} \times f_{out}}$ and $l_{in}, h_{in}, l_{out}, h_{out}$ representing spatial dimensions and f_{in}, f_{out} representing the numbers of feature maps of the tensors. The relations between I and J and between J and O as well as the relations between their dimensions will be given in this subsection. Even though the contribution presented here can be seen as a set of Gaussian based layers, the primal aim is to propose a high level configurable layer which can be instantiated as any of the following configurations. The configurable levers are the following ones: i) the numbers of basis of Gaussian derivative kernels involved in the layer, ii) the number of linear combinations for each basis, iii) the type of linear combination and, iv) the orders of derivation of the Gaussian derivative kernels composing the bases. ¹ We will go through all the configurations, beginning from the closest one to RFNN layers. We begin with a layer having only one fully trainable basis and using a standard linear combination. Its only basis will be referred to as $\{\phi^n(\boldsymbol{\mu}^n, \boldsymbol{\sigma}^n, \theta^n) \mid n \in \llbracket 1, N \rrbracket\}$,

¹<https://github.com/Penaud-Polge/Fully-Trainable-Gaussian-Derivative-Layer>

# training samples	Layer Type					
	Free-Form	RFNN [10]	DAU-ConvNet [13]	Ours	Ours-few	Ours-separated
1000	78.8	79.4	79.9	80.5	79.7	78.9
5000	84.3	86.2	86.4	88.3	88.4	88.7
10000	86.4	88.4	89.1	90.5	90.1	89.3
30000	87.7	90.4	90.9	92.2	92.4	91.4
60000	89.3	91.2	92.1	93.1	92.8	91.4

Table 1. Test accuracy obtained on the Fashion-MNIST dataset with different sizes of training dataset by the RFNN model, the DAU-ConvNet model, a usual free-form model and several configurations of the fully trainable Gaussian derivative network. All the networks have equivalent architectures and training conditions. The best testing accuracy is highlighted in bold.

where N refers to the number of Gaussian derivative kernels composing the basis and where each ϕ^n is defined by (1). Together with the basis, the layer also owns a set of coefficients $\{\omega_{j,n}^i \mid n \in \llbracket 1, N \rrbracket, i \in \llbracket 1, f_{in} \rrbracket, j \in \llbracket 1, f_{out} \rrbracket\}$ used to compute the linear combination where the particular weight $\omega_{j,n}^i$ is used to compute the j^{th} feature map of J and is weighting the n^{th} kernel of the basis to convolve the i^{th} feature map of the input tensor I . In this case, the j^{th} feature map of J is computed by the following equation:

$$J_j = \sum_{i=1}^{f_{in}} \left(\sum_{n=1}^N \omega_{j,n}^i \phi^n(\boldsymbol{\mu}^n, \boldsymbol{\sigma}^n, \theta^n) \right) * I_i. \quad (2)$$

Fig.1 (a) illustrates eq.(2) by schematically showing the relation between I and J . If, now, the layer has several bases, for example $B \in \mathbb{Z}_+^*$ many bases, we propose that each basis should be used to compute $\frac{f_{out}}{B} \in \mathbb{Z}_+^*$ feature maps of J . In order to differentiate each basis, the b^{th} basis will be denoted as $\{\phi^{b,n}(\boldsymbol{\mu}^{b,n}, \boldsymbol{\sigma}^{b,n}, \theta^{b,n}) \mid n \in \llbracket 1, N \rrbracket\}$ and its corresponding weights are denoted by $\{\omega_{j,n}^{b,i} \mid n \in \llbracket 1, N \rrbracket, i \in \llbracket 1, f_{in} \rrbracket, j \in \llbracket 1, \frac{f_{out}}{B} \rrbracket\}$. Then, model from (2) becomes

$$J_{(b-1)\frac{f_{out}}{B}+j} = \sum_{i=1}^{f_{in}} \left(\sum_{n=1}^N \omega_{j,n}^{b,i} \phi^{b,n}(\boldsymbol{\mu}^{b,n}, \boldsymbol{\sigma}^{b,n}, \theta^{b,n}) \right) * I_i. \quad (3)$$

In a policy of reducing the number of trainable parameters, in the same spirit as separable convolution [18], we propose to separate the linear combination process. To describe this configuration, we will return to a layer having only one basis $\{\phi^n(\boldsymbol{\mu}^n, \boldsymbol{\sigma}^n, \theta^n) \mid n \in \llbracket 1, N \rrbracket\}$. We propose to split the set of coefficients $\{\omega_{j,n}^i \mid n \in \llbracket 1, N \rrbracket, i \in \llbracket 1, f_{in} \rrbracket, j \in \llbracket 1, f_{out} \rrbracket\}$ into two sets $\{\lambda_{j,n} \mid n \in \llbracket 1, N \rrbracket, j \in \llbracket 1, f_{out} \rrbracket\}$ and $\{\beta_n^i \mid n \in \llbracket 1, N \rrbracket, i \in \llbracket 1, f_{in} \rrbracket\}$ and eq. (2) becomes

$$J_j = \sum_{n=1}^N \lambda_{j,n} \left(\sum_{i=1}^{f_{in}} \beta_n^i \phi^n(\boldsymbol{\mu}^n, \boldsymbol{\sigma}^n, \theta^n) * I_i \right) \quad (4)$$

Fig.1 (b) illustrates this new configuration where an intermediate tensor is used before computing J . The split of the weights allows to decrease the number of weights from $f_{in} \times N \times f_{out}$ to $(f_{in} + f_{out}) \times N$. This configuration also offers

the possibility to consider several bases. In order to obtain O from J , we first apply a batch normalization [19] to J . Then, one can apply a nonlinear function F such as the rectified linear activation function or the sigmoid function which leads to $O = F\left(\gamma \frac{(J - \mu_{batch})}{\sqrt{\sigma_{batch}^2 + \epsilon}} + \beta\right)$, where γ and β are trainable parameters and where μ_{batch} and σ_{batch}^2 are either the mean and variance of the batch during training or the moving mean and the moving variance during inference. Finally, depending on the strides, a down-sampling operation can be applied, which gives: $l_{out} \leq l_{in}, h_{out} \leq h_{in}$.

3. EXPERIMENTS

In this section we will first compare the proposed fully trainable Gaussian derivative layer to the ones involved in RFNN [10] and in DAU-ConvNet [8]. Then, we will use the proposed layer in two common architectures to show its relevance: as a first layer of a VGG16 model [20] for cats and dogs classification, compared to a usual free-form convolutional layer and we will also use the proposed layer for a segmentation task with a fully trainable Gaussian derivative U-net and compare it with an equivalent classical U-net [21].

Comparison with related work. The aim of this experiment is twofold: First to compare several configurations of the proposed layer. More precisely, the different configurations will highlight the effect of using more or less trainable basis and the effect of the separation of the linear combination. Secondly, it will give a comparison between these different configurations and related work layers : the RFNN layer [10], the DAU-ConvNet layer [13] and the usual free-form layer with filters of size 3×3 . In order to demonstrate the positive effect of fully training the Gaussian derives, the values of the following parameters and quantities γ , μ_{batch} and $\sqrt{\sigma_{batch}^2 + \epsilon}$, used in the batch normalization stage, are respectively fixed to one, zero and one for this experiment. Only β is kept trainable, playing simply the role of a bias. For each type of layer, the same four-convolutional-layers architecture is used, ended by a global average pooling layer and a dense layer. The networks were trained using Fashion-MNIST [22] with different sizes of training dataset

Layer Type	Free-Form	Gaussian (Ours)
Validation Accuracy	97	98.1
Test Accuracy	96.4	97.2

Table 2. Accuracy obtained replacing and training the first layer of a pre-trained VGG16 on the Cats and Dogs dataset.

in order to study their behavior when a small amount of training samples is available. Each network was trained with a batch size of 32 and a cosine learning rate decay defined as $lr_k = lr_0 \left(\alpha + (1 - \alpha) \left(1 + \cos \left(\frac{2\pi \min(k, K_{max})}{K_{max}} \right) \right) \right)$. For RFNN, $\sigma = 1.5$ has been used for the first layer and $\sigma = 1$ for the others as they were the values used in [10] for the MNIST dataset using a similar architecture. Concerning the DAU-ConvNet, σ was let being learned by the network for each unit and the number of units for each filter took the value 10 as it corresponds to the number of filters composing the bases of the fully trainable Gaussian derivative network and RFNN. The networks representing three different configurations of the proposed layer are the following one: two networks, referred to as ‘Ours’ and ‘Ours-few’, using layers with non-separated linear combinations, one having half the number of basis (‘Ours-few’) than the other (‘Ours’) for each layer. And one network whose layers use separated linear combinations, with the same number of basis as ‘Ours’, is referred to as ‘Ours-separated’. Three trials have been realized and the corresponding accuracy values have been averaged for the two smallest sizes of dataset in order to counter the effect of the random choice of images among the full training dataset. Results are presented in Tab.1 and show that the proposed layers obtain higher accuracy values compared to the other types of layers. The performance of ‘Ours-separated’ highlights the benefits of the separated linear combination, as it allows comparable performances compared to the other configurations while decreasing the number of parameters. ‘Ours-separated’ uses a third less parameters than ‘Ours’.

VGG16. The purpose of this experiment is to study the behavior of the proposed layer when it is used as a first layer of a deep architecture. We considered a VGG16 pretrained on Imagenet [23], and we perform transfer learning for Kaggle Cats and Dogs dataset. The obtained network achieves an accuracy of 92.8% on the test dataset. We then replaced the first layer of the network by a fully trainable Gaussian derivative layer, froze the rest of the network and trained it. Results presented in Tab.2 show that the proposed layer is adapted to be used as the first layer of a convolutional neural network as it obtains better performances compared to a classical free-form layer. The next experiment will investigate the legitimacy of using an architecture using exclusively the proposed layer.

Gaussian U-net. This experiment aims to test the fully trainable Gaussian derivative layer for a segmentation of nuclei

Network Type	U-net	Gaussian U-net
Test Accuracy	94.7	94.4
# Trainable parameters	97 202	25 298

Table 3. Accuracy obtained on the histopathology test dataset and # of trainable parameters of Gaussian and classical U-net.

in histopathology images. The used dataset has been generously provided by the authors of [24]. It is composed of 50 images of size 512×512 with their ground truths. For each of the 50 images, 16 images of size 128×128 have been extracted, forming a dataset of 800 images with their associated ground truths (600 for training, 200 for testing). We compare the performance of a classical U-net (with filters of 5×5) by replacing the convolutional layers by our proposed Gaussian layer, called Gaussian U-net. Results in Tab.3 shown that the fully trainable Gaussian derivative layer can compete with classical free-form layers while using nearly 25% of parameters. Fig.2 shows an example of predictions for both U-nets, highlighting they are quite similar despite the reduction in the number of parameters.

4. CONCLUSION

We presented a configurable fully trainable Gaussian derivative layer generalizing previous works while keeping their advantages. It allows for more diversified shapes but still with the possibility to reduce the number of parameters. The conducted experiments showed the behavior of the proposed layer, its competitiveness with related previous work and that it offers enough diversity to deal with challenging tasks in computer vision.




Fig. 2. First row: Input images. Second row: Ground truths. Third row: Classical U-net predictions. Fourth row: Gaussian U-net predictions.

5. REFERENCES

- [1] Luc MJ Florack, Bart M ter Haar Romeny, Jan J Koenderink, and Max A Viergever, “Scale and the differential structure of images,” *Image and vision computing*, vol. 10, no. 6, pp. 376–388, 1992.
- [2] Luc Florack, Bart ter Haar Romeny, Max Viergever, and Jan Koenderink, “The Gaussian scale-space paradigm and the multiscale local jet,” *International Journal of Computer Vision*, vol. 18, no. 1, pp. 61–75, 1996.
- [3] Tony Lindeberg, “Generalized Gaussian scale-space axiomatics comprising linear scale-space, affine scale-space and spatio-temporal scale-space,” *Journal of Mathematical Imaging and Vision*, vol. 40, no. 1, pp. 36–81, 2011.
- [4] Jan J Koenderink, “The structure of images,” *Biological cybernetics*, vol. 50, no. 5, pp. 363–370, 1984.
- [5] Jan J Koenderink and Andrea J van Doorn, “Representation of local geometry in the visual system,” *Biological cybernetics*, vol. 55, no. 6, pp. 367–375, 1987.
- [6] Anders Boesen Lindbo Larsen, Sune Darkner, Anders Lindbjerg Dahl, and Kim Steenstrup Pedersen, “Jet-based local image descriptors,” in *ECCV*. Springer, 2012, pp. 638–650.
- [7] Antoine Manzanera, “Local jet feature space framework for image processing and representation,” in *2011 Seventh International Conference on Signal Image Technology & Internet-Based Systems*. IEEE, 2011, pp. 261–268.
- [8] Domen Tabernik, Matej Kristan, and Aleš Leonardis, “Spatially-adaptive filter units for deep neural networks,” in *CVPR*, 2018, pp. 9388–9396.
- [9] Evan Shelhamer, Dequan Wang, and Trevor Darrell, “Blurring the line between structure and learning to optimize and adapt receptive fields,” *arXiv preprint arXiv:1904.11487*, 2019.
- [10] Jorn-Henrik Jacobsen, Jan Van Gemert, Zhongyu Lou, and Arnold WM Smeulders, “Structured receptive fields in CNNs,” in *CVPR*. IEEE, 2016, pp. 2610–2619.
- [11] Tony Lindeberg, “Scale-covariant and scale-invariant Gaussian derivative networks,” *Journal of Mathematical Imaging and Vision*, pp. 1–20, 2021.
- [12] Ganesh Sundaramoorthi and Timothy E Wang, “Translation insensitive cnns,” *arXiv preprint arXiv:1911.11238*, 2019.
- [13] Domen Tabernik, Matej Kristan, Jeremy L Wyatt, and Aleš Leonardis, “Towards deep compositional networks,” in *23rd ICPR*. IEEE, 2016, pp. 3470–3475.
- [14] Domen Tabernik, Matej Kristan, and Aleš Leonardis, “Spatially-adaptive filter units for compact and efficient deep neural networks,” *International Journal of Computer Vision*, vol. 128, no. 8, pp. 2049–2067, 2020.
- [15] Silvia L Pinteá, Nergis Tömen, Stanley F Goes, Marco Loog, and Jan C van Gemert, “Resolution learning in deep convolutional networks using scale-space theory,” *IEEE Transactions on Image Processing*, vol. 30, pp. 8342–8353, 2021.
- [16] Ylva Jansson and Tony Lindeberg, “Exploring the ability of CNNs to generalise to previously unseen scales over wide scale ranges,” in *25th ICPR*. IEEE, 2021, pp. 1181–1188.
- [17] Bo Yang, Jan Flusser, and Tomáš Suk, “Steerability of hermite kernel,” *International Journal of Pattern Recognition and Artificial Intelligence*, vol. 27, no. 04, pp. 1354006, 2013.
- [18] Daniel Haase and Manuel Amthor, “Rethinking depth-wise separable convolutions: How intra-kernel correlations lead to improved mobilenets,” in *CVPR*, 2020, pp. 14600–14609.
- [19] Sergey Ioffe and Christian Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” in *International conference on machine learning*. PMLR, 2015, pp. 448–456.
- [20] Karen Simonyan and Andrew Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014.
- [21] Olaf Ronneberger, Philipp Fischer, and Thomas Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.
- [22] Han Xiao, Kashif Rasul, and Roland Vollgraf, “Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms,” *arXiv preprint arXiv:1708.07747*, 2017.
- [23] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in *CVPR*. IEEE, 2009, pp. 248–255.
- [24] Peter Naylor, Marick Laé, Fabien Reyat, and Thomas Walter, “Nuclei segmentation in histopathology images using deep neural networks,” in *14th ISBI*. IEEE, 2017, pp. 933–936.