



**HAL**  
open science

## A model is worth tens of thousands of examples

Thomas Dages, Laurent D. Cohen, Alfred M. Bruckstein

► **To cite this version:**

Thomas Dages, Laurent D. Cohen, Alfred M. Bruckstein. A model is worth tens of thousands of examples. 9th International Conference on Scale Space and Variational Methods in Computer Vision (SSVM 2023), May 2023, Santa Margherita di Pula, Italy. 10.1007/978-3-031-31975-4\_17. hal-04271603

**HAL Id: hal-04271603**

**<https://hal.science/hal-04271603>**

Submitted on 6 Nov 2023

**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.

# A model is worth tens of thousands of examples

Thomas Dages<sup>1</sup>, Laurent D. Cohen<sup>2</sup>, and Alfred M. Bruckstein<sup>1</sup>

<sup>1</sup> Department of Computer Science, Technion Israel Institute of Technology, Haifa, Israel

`{thomas.dages, freddy}@cs.technion.ac.il`

<sup>2</sup> Ceremade, University Paris Dauphine, PSL Research University, UMR CNRS 7534, 75016 Paris, France

`cohen@ceremade.dauphine.fr`

**Abstract.** Traditional signal processing methods relying on mathematical data generation models have been cast aside in favour of deep neural networks, which require vast amounts of data. Since the theoretical sample complexity is nearly impossible to evaluate, these amounts of examples are usually estimated with crude rules of thumb. However, these rules only suggest when the networks should work, but do not relate to the traditional methods. In particular, an interesting question is: how much data is required for neural networks to be on par or outperform, if possible, the traditional model-based methods? In this work, we empirically investigate this question in two simple examples, where the data is generated according to precisely defined mathematical models, and where well-understood optimal or state-of-the-art mathematical data-agnostic solutions are known. A first problem is deconvolving one-dimensional Gaussian signals and a second one is estimating a circle's radius and location in random grayscale images of disks. By training various networks, either naive custom designed or well-established ones, with various amounts of training data, we find that networks require tens of thousands of examples in comparison to the traditional methods, whether the networks are trained from scratch or even with transfer-learning or finetuning.

**Keywords:** Deep learning · Model-based methods · Sample complexity.

## 1 Introduction

Neural network-based machine learning has widely replaced the traditional methods for solving many signal and image processing tasks that relied on mathematical models for the data [14, 10]. In some cases, the assumed models provided ways to optimally address the tasks at hand and resulted in well-performing estimation and prediction methods with theoretical guarantees [21, 17, 7]. Nowadays, gathering raw data and applying gradient descent-like processes to neural network structures [15, 12, 19, 11] largely replaced modelling and mathematically developing provably optimal solutions.

It is commonly accepted that, if the networks are complex enough and when vast amounts of data are available, neural networks outperform traditionally

designed methods [16, 12] or even humans [9, 8, 4]. The required amount of data is called in statistical learning theory the sample complexity and is related to the VC-dimension of the problem [20], which is usually intractable for non trivial networks [2]. Instead, various rules of thumb have been used in the field to guess how many samples are needed: at least 10-50 times the number of parameters [1], at least 10 times per class in classification (and 50 times in regression) the data dimensionality [13] and at least 50-1000 times the output dimension [1].

However, these rules only suggest how much data is needed to get a “good” network, but they do not relate to the traditional data-generation model-based methods. A natural question hence arises: do the neural network-based solutions perform as well as, or even outperform, the processing methods based on traditional data-generation models when lots of data is available, and if so how much data is necessary? We address this question in two simple empirical examples, where the data is produced according to precisely defined models, and where well understood optimal or state-of-the-art mathematical solutions are available. The first is the deconvolution of Gaussian signals, optimally solved with the Wiener filter [21]. The second is the estimation of the radius and centre coordinates of a disk in an image, which can be elegantly solved using a Pointflow method [22]. This work aids engineers to decide when to use model-based classical methods or simply feed lots of data (if available) to deep neural networks.

Section 2 presents our comparison for the one-dimensional signal recovery, and Section 3 deals with estimation of disk characteristics in an image.

## 2 One-dimensional signal recovery

We suggest to first analyse a simple and well-understood problem in the one-dimensional case where the optimal solution is provingly known.

### 2.1 Data model and optimal solution

The original data consists of real random vectors  $\varphi$  of size  $D$  that are centred, i.e.  $\mathbb{E}(\varphi) = 0$ , and with known autocorrelation  $R_\varphi = \mathbb{E}(\varphi\varphi^\top) \in \mathbb{R}^{D \times D}$ . However,  $\varphi$  is degraded by blur and noise producing the observed data  $\varphi^{data}$  as follows:

$$\varphi^{data} = H\varphi + n, \tag{1}$$

where  $H \in \mathbb{R}^{D \times D}$  is a known deterministic matrix and  $n$  is random additive noise independent from  $\varphi$  that is centred  $\mathbb{E}(n) = 0$  and with known autocorrelation matrix  $R_n \in \mathbb{R}^{D \times D}$ . It is well-known [21] that the best linear recovery of  $\varphi$  in the  $L^2$  sense, i.e. minimising the Expected Squared Error (ESE)  $ESE(\hat{\varphi}, \varphi) = \mathbb{E}(\|\hat{\varphi} - \varphi\|_2^2)$  with respect to the matrix  $M \in \mathbb{R}^{D \times D}$  such that  $\hat{\varphi} = M\varphi^{data}$ , is given by applying the Wiener filter  $W = R_\varphi H^\top (H R_\varphi H^\top + R_n)^{-1}$ , i.e.  $\hat{\varphi}^* = W\varphi^{data}$ . Moreover, if we further assume both  $\varphi \sim \mathcal{N}(0, R_\varphi)$  and  $n \sim \mathcal{N}(0, R_n)$  are Gaussian, then the Wiener filter  $W$  minimises the ESE over all possible recoveries including nonlinear ones. Furthermore, note that if  $R_\varphi$  is circulant,

i.e.  $\varphi$  is cyclostationary, and so is  $n$ , e.g. if  $n$  has independent entries implying  $R_n$  is diagonal, and if  $H$  is circulant, then  $W$  is also circulant and Wiener filtering is a pointwise multiplication in the Fourier domain given by applying the unitary Discrete Fourier Transform [DFT] with  $(k,l)$ -th entry  $[DFT]_{k,l} = \frac{1}{\sqrt{D}} e^{-i \frac{2\pi kl}{D}}$ .

In our tests, the dimensionality is  $D = 32$  and the problem is circulant. We use an interpretable symmetric positive-definite autocorrelation matrix  $R_\varphi$  parameterised by a large number  $\rho = 0.95$  to create high spatial correlation over a large support decaying with distance and  $H$  is a local smoothing convolution. The first lines of  $R_\varphi$  and  $H$  are  $(1 \ \rho \ \rho^2 \ \rho^3 \ \dots \ \rho^3 \ \rho^2 \ \rho)$  and  $(1 \ 1 \ 0 \ \dots \ 0 \ 1)$ . The noise is i.i.d.  $n \sim \mathcal{N}(0, \sigma_n^2 I)$  with  $\sigma_n = 0.1$ . We display example data in Figure 1, the designed  $H$ ,  $R_\varphi$ , and  $R_n$  along with their associated Wiener filter  $W$  in Figure 2.




Fig. 1. Two example signals  $\varphi^{data}$ , with their associated blur  $H\varphi$  and noise  $n$ .




Fig. 2. Chosen model matrices and associated optimal Wiener filter.

### 2.2 Neural models

We wish to evaluate the capabilities of neural networks by comparing them to humanly designed methods by classical experts using no training. Our criterion is the amount of random training samples  $N$  needed to reach or overtake human expertise. Working in the Gaussian case for the data model of equation (1), we create various random training datasets containing  $N$  data samples ranging in  $N \in \{10, 100, 1000, 10000, 100000\}$ . We train a variety of small Convolutional Neural Networks (CNNs) of various depths  $K \in \{0, 1, 2, 3\}$ . The depth of the network is measured as the number of successions of convolution-pointwise-nonlinearity layers. Each network ends with a final fully connected

layer  $A$  (with bias  $b_A$ ), i.e. a final unconstrained affine transformation. For simplicity, our CNNs will be single-channel only and without various architecture tricks, e.g. dropout, batch normalisation, or pooling. The network functions, denoted  $f_k$  for  $k \in \{1, \dots, K\}$  can thus be written as:

$$f_k(\varphi^{data}) = A\sigma \circ \tilde{C}_K \circ \sigma \circ \tilde{C}_{K-1} \circ \dots \circ \sigma \circ \tilde{C}_1(\varphi^{data}) + b_A, \quad (2)$$

where  $\tilde{C}_i(x) = C_i x + b_i$  is the  $i$ -th convolution layer comprising the circulant matrix  $C_i$  for the convolution and its additive unconstrained bias  $b_i$  and  $\sigma = \text{ReLU}$  the standard pointwise nonlinearity in neural networks. Note that a CNN with depth 0 degenerates to an unconstrained affine transformation in  $\mathbb{R}^D$  (no pointwise nonlinearity or convolution):  $f_0(\varphi^{data}) = A\varphi^{data} + b_A$ .

The networks are trained to minimise the Mean Squared Error (MSE)<sup>3</sup>, a proxy for the ESE, using the  $N$  generated samples. Denoting  $f_{k,N,\eta}$  the resulting networks (where  $\eta$  a hyperparameter of the optimisation algorithm), we have:

$$MSE_{train}(f_{k,N,\eta}) = \frac{1}{N} \sum_{i=1}^N \|f_{k,N,\eta}(\varphi_{train,i}^{data}) - \varphi_{train,i}\|_2^2, \quad (3)$$

where for a sample collection  $set$ ,  $\varphi_{set,i}$  and  $\varphi_{set,i}^{data}$  denote the  $i$ -th original and degraded samples. This quantity is to be compared with  $ESE(f_{k,N,\eta}(\varphi^{data}), \varphi)$ , which evaluates the performance on all possible data of a network trained on  $N$  instances only. Naturally, this quantity cannot be computed by hand and is approximated by another  $MSE$  calculation on a large test set using  $N_t$  test samples independently generated from the training ones:

$$MSE_{test}(f_{k,N,\eta}) = \frac{1}{N_t} \sum_{i=1}^{N_t} \|f_{k,N,\eta}(\varphi_{test,i}^{data}) - \varphi_{test,i}\|_2^2 \xrightarrow{N_t \rightarrow \infty} ESE(f_{k,N,\eta}(\varphi^{data}), \varphi). \quad (4)$$

In our tests,  $N_t = 100000$ . Note that implicitly in  $ESE(f_{k,N,\eta}(\varphi^{data}), \varphi)$  the network  $f_{k,N,\eta}$  is the given result of a minimisation algorithm. For randomised algorithms, it is thus to be understood as the expectation conditional to the learned network  $f_{k,N,\eta}$ :  $ESE(f_{k,N,\eta}(\varphi^{data}), \varphi) = \mathbb{E}(\|f_{k,N,\eta}(\varphi^{data}) - \varphi\|_2^2 \mid f_{k,N,\eta})$ .

We train our networks using Stochastic Gradient Descent with Nesterov momentum parameter equal to 0.9. We train the networks using various learning rates  $\eta \in \{0.0001, 0.0005, 0.001, 0.005, 0.01, 0.05, 0.1\}$  over 50 epochs, performing  $N_r = 50$  independent training trials per learning rate, and compute the final median performance per learning rate on a validation set generated independently of the train and test data comprising  $N_v = 100000$  validation samples:

$$MSE_{val}(f_{k,N,\eta}) = \frac{1}{N_v} \sum_{i=1}^{N_v} \|f_{k,N,\eta}(\varphi_{val,i}^{data}) - \varphi_{val,i}\|_2^2 \xrightarrow{N_v \rightarrow \infty} ESE(f_{k,N,\eta}(\varphi^{data}), \varphi). \quad (5)$$

<sup>3</sup> The loss function is actually scaled to  $\frac{1}{D} MSE_{train}$  as is commonly done in practice.


The validation set is used to choose the best learning rate for each amount of training data  $\eta^*(N)$  by taking:

$$\eta^*(N) = \underset{\eta}{\operatorname{argmin}} \operatorname{MEDIAN}_r(\operatorname{MSE}_{\text{val}}(f_{k,N,\eta})), \quad (6)$$

where  $\operatorname{MEDIAN}_r$  takes the median over the  $r \leq N_r$  best independent runs on the validation set per  $\eta$ . Given that a significant amount of runs do not converge or get trapped early in a poor local minimum depending on the random initialisation, choosing  $r \ll N_r$  ensures that only the networks finding a good local minimum are considered. The final performance of CNNs  $\operatorname{SCORE}_{k,r}(N)$  for each amount of data  $N$  is then the median of the test performance over those selected  $r$  trials<sup>4</sup> of the final test score at the chosen learning rate  $\eta^*(N)$ :

$$\operatorname{SCORE}_{k,r}(N) = \operatorname{MEDIAN}_r(\operatorname{MSE}_{\text{test}}(f_{k,N,\eta^*(N)})). \quad (7)$$

We display the evolution of the networks’ performance on the amount of training data  $N$  in Figure 3 for each depth  $k$ , with detailed scores in Table 1, along with the performance of the Wiener filter. Regardless of  $N$ , the Wiener filter outperforms the neural models as expected by the theory, but their performance converges to the Wiener’s one when a lot of data is available, with similar performance when at least 10000 training samples are available. We can thus consider this study as providing a criterion that a model would be preferable if data is limited to fewer than 10000 samples to train on.



**Fig. 3.** Median test scores for CNNs with depth  $k \in \{0, 1, 2, 3\}$  on  $r = 10$  selected runs ( $k = 0$  is just a linear layer). Vertical bars represent the standard deviation of the MSE of these runs. The right figure is a zoomed-in plot of the left one for large  $N$ .

### 3 Two-dimensional geometric estimation

We next analyse a more complicated yet well-understood problem based on Euclidean geometry. The goal is to estimate basic geometric properties on simple data: the radius and centre location of a random disk in an image. It was shown in [6] that this seemingly trivial task is more complex than expected for neural models even when focusing on radius estimation of centred disks.

<sup>4</sup> Selected on the validation set.

$N$	0	10	100	1000	10000	100000
Wiener	<b>1.743</b>	—	—	—	—	—
Linear ( $k = 0$ )	—	6.204	2.295	1.811	1.751	1.748
CNN ( $k = 1$ )	—	12.386	3.662	1.924	1.799	1.762
CNN ( $k = 2$ )	—	15.614	3.789	2.051	1.842	1.767
CNN ( $k = 3$ )	—	20.395	4.911	2.167	1.869	1.771

**Table 1.** Median MSE scores  $SCORE_{k,r}$  of the CNNs on  $r = 10$  selected runs, compared to the theoretically optimal Wiener filter.

### 3.1 Data model


The original data now consists of  $D \times D$  random two-dimensional grayscale images of disks. Images are centred at  $(0, 0)$ , and for a pixel  $x \in [-\frac{D-1}{2}, \frac{D-1}{2}]^2$ :

$$\varphi(x) = \begin{cases} b & \text{if } \|x - c\|_2 > r \\ f & \text{if } \|x - c\|_2 \leq r, \end{cases} \quad (8)$$

where  $r$  is the circle’s radius,  $c = (c_x, c_y)$  its centre, and  $f$  (resp.  $b$ ) is the foreground (resp. background) intensity. These parameters are independently<sup>5</sup> and uniformly chosen at random:  $r \sim \mathcal{U}([\frac{\varepsilon_r}{2}, \frac{D-1}{4}], (1 - \frac{\varepsilon_r}{2})\frac{D-1}{4}]$  with  $\varepsilon_r = 0.4$ ,  $c \sim \mathcal{U}([(D-1)\frac{\varepsilon_c}{2} - \frac{D-1}{2}, (D-1)(1 - \frac{\varepsilon_c}{2}) - \frac{D-1}{2}]^2)$  with  $\varepsilon_c = 0.5$ ,  $b \sim \mathcal{U}([0, 1])$ , and  $f \mid b \sim \mathcal{U}([0, 1] \setminus [b - \delta, b + \delta])$  with  $\delta = \frac{50}{255}$  the minimum contrast<sup>6</sup>. However,  $\varphi$  is degraded with blur and noise giving the observed data  $\varphi^{data}$  as follows:

$$\varphi^{data} = g_{\sigma_b} * \varphi + n, \quad (9)$$

where  $g_{\sigma_b}(x) = \frac{1}{2\pi} \exp(-\frac{\|x\|_2^2}{2\sigma_b^2})$  is a Gaussian convolution kernel, and  $n$  is i.i.d. white noise  $n \sim \mathcal{N}(0, \sigma_n I_{D^2})$ . We plot example data in Figure 4. The task is to estimate the three geometric numbers  $(r, c) = (r, c_x, c_y)$  from  $\varphi^{data}$ .



**Fig. 4.** Four examples of clean  $\varphi$  and degraded  $\varphi^{data}$  disk images.

### 3.2 Expert engineer’s solution

Unlike in the Wiener case, the optimal estimator minimising the ESE is not so trivial to find. Instead, we choose a method called Pointflow designed by an expert engineer that perfectly tackles the problem at hand.


<sup>5</sup> Except  $f$  and  $b$  which are slightly correlated to ensure a minimal contrast  $|f - b| > \delta$ .

<sup>6</sup> In our tests, we take  $D = 201$  implying that  $r \sim \mathcal{U}([10, 40])$  and  $c \sim \mathcal{U}([-50, 50]^2)$ .

Pointflow [22, 3] is an elegant subpixel level contour integrator and edge detector in images requiring no learning whatsoever. It consists in defining potential vector fields  $V$  along which random points  $P$  flow:  $\frac{dP}{dt}(t) = V(P(t))$ , such that end trajectories lie on edges of the image  $I$ . The vanilla Pointflow [22] uses two fields  $V_+$  and  $V_-$  from the edge attraction  $V_a$  and rotating  $V_r$  fields based on the image gradients as follows:

$$V_a = \nabla \|\nabla I_b\|_2, \quad V_r = \nabla I_b^\perp, \quad V_\pm = \frac{1}{2}(V_a \pm V_r), \quad (10)$$

where  $I_b = g_{\sigma_{Pf}} * I$  is a blurred version of  $I$  with a Gaussian kernel  $g_{\sigma_{Pf}}$ . Various stopping conditions and uses of  $V_\pm$  exist to detect edges in natural images, however on our data containing a single circular edge per image, we need only consider the basic ones. Indeed, the possible cases for trajectories are: it loops ( $C_l$ ), it leaves the image domain ( $C_o$ ), or it is stuck in an area with small magnitude  $\|V\|_2$  ( $C_s$ ). Flowing initially from  $V_+$ , if we loop ( $C_l$ ), then the point has reached the circle and it suffices to reflow along  $V_+$  to extract just the circle's contour. If we end up outside the image domain ( $C_o$ ), which is rare in our data, we reflow from the exit point along  $V_-$ . If we are in a low flow magnitude area ( $C_s$ ), which is not rare, then we discard the trajectory. In total,  $N_{Pf} = 200$  points are randomly uniformly sampled in the image domain and used for Pointflow, and a fraction of them end up flowing on the disk's edge with subpixel precision, as the other ones lead to discarded trajectories. For more details on our implementation of Pointflow see Appendix A. For some illustrations of Pointflow results on our data see Figure 5.



**Fig. 5.** Pointflow in practice. (5a): clean data. (5b): degraded data. (5c) and (5d): pointflow fields sampled every five pixels. (5e): initial flows of points without reflowing with groundtruth boundary in green. (5f): all final trajectories that have reflowed in a closed loop ( $C_l$ ) with groundtruth boundary in green and all regressed circles using least squares on each looped trajectory in blue (used for estimating the centre).

To estimate the disk's radius and centre from pointflow contours ( $C_l$ ), we can simply compute the average length of the reflowed closed contours and divide it by  $2\pi$ . To compute the centre's coordinates, we could compute for each closed trajectory the average of its points, and then average over these estimations. However, this method empirically did not best perform on validation data, so we refined it by applying least-squares regression on the equation of a circle to estimate from it its location per trajectory and then average the estimations. Note that the least-squares regression did not provide a better estimation of the radius so we keep the crude length integration strategy for it.



### 3.3 Neural models

As in the one-dimensional case, the expert’s method is to be compared with a convolutional neural network. Although the learning problem seems trivial, it is actually harder than expected for networks, as has been shown in [6] even when the circles are centred. Empirically, we were not able to train correctly a small custom model similar to those previously used having just three layers and even many channels per layers and no further deep learning tricks. To overcome this limitation, we use famous networks in the deep learning literature: Alexnet [12], VGG [19] and Resnet<sup>7</sup> [11]. To adapt the model to our task, we change the final fully connected layer to have 3 outputs only.

For each architecture, we either train the networks from scratch (SC), or initialise the weights, except those of the final fully connected layers, to those publicly available obtained by classification on Imagenet [18], as is commonly done in the field. The pretrained weights can be either frozen for transfer-learning (TL) or retrained as well for finetuning (FT). Although the task and data are fundamentally different from ours, it is generally believed that the wide variety of natural images encourages the famous networks to learn features that generalise quite well to most reasonable tasks.

Once again, the *MSE* loss is used for training<sup>8</sup>. As it is significantly more expensive to train such networks compared to the tiny ones in the Wiener case, we only perform  $N_r = 1$  run per learning rate configuration, ranging in  $\eta \in \{0.000001, 0.00005, 0.00001, 0.0005, 0.0001, 0.005, 0.001, 0.05, 0.01, 0.5, 0.1\}$ , with a batch size of 10 for 50 epochs. As previously, the optimal learning rate is chosen on the performance on validation data. Both the test and validation data use  $N_v = N_t = 100000$  independently randomly sampled images, whereas the training sets have  $N \in \{10, 100, 1000, 10000, 100000\}$  ones.

### 3.4 Results

We present the results in Figure 6 and Table 2. Although the networks are simultaneously trained for both the radius and centre location estimation, we also present the MSE on the estimation of each geometric concepts separately.

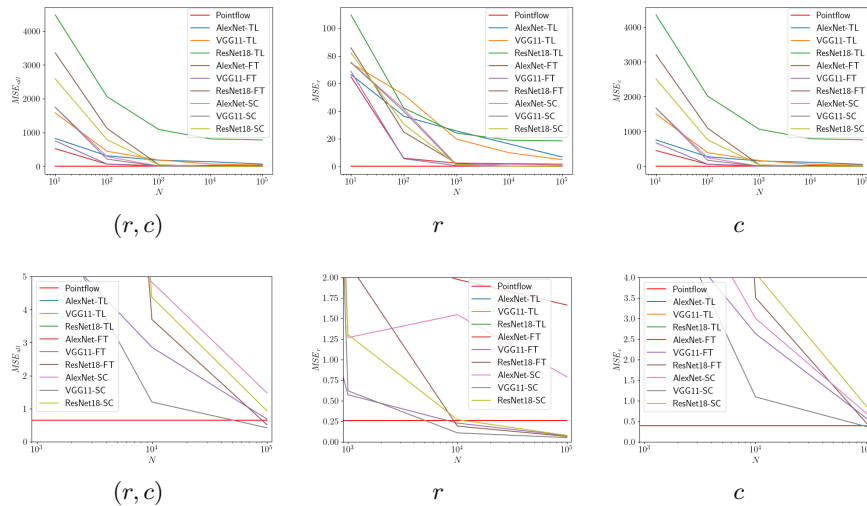
First, the transfer-learning network Resnet-TL is not able to correctly estimate the radius or centre’s location, meaning that its learned features on classification of Imagenet are not able to handle our simple data: they are not so general after all. Likewise, the prelearned features of Alexnet-TL and VGG-TL do not generalise well to this toy problem, requiring significantly more data than the maximum available to compare with the simple data agnostic pointflow. However, when the networks are entirely trained, either finetuned or from scratch, they are either flatly beaten by pointflow when using small amounts of training data or on par or slightly outperform it when  $N \geq 10000$ . The only networks

<sup>7</sup> We use the simplest ones VGG11 and ResNet18, as larger ones are here unnecessary.

<sup>8</sup> To help the networks converge, the radius and centre coordinates are scaled to  $[-1, 1]$  using  $r_s = \frac{8}{D-1}(r - \frac{D-1}{8})$  and  $c_s = \frac{2}{D-1}c$ . In all plots and numbers provided in this paper, the results are rescaled to the original scale:  $r$  and  $c$  and not  $r_s$  and  $c_s$ .

beating pointflow overall are VGG-SC and Resnet-FT when  $N = 100000$ , but more (VGG-FT, VGG-SC, Resnet-FT, and Resnet-SC) significantly outperform pointflow on the radius estimation when  $N \geq 10000$ . The difference between finetuning and training from scratch seems to only appear when small amounts of training data are available, and then finetuning is better. However, in these cases, both approaches pale in comparison to the reference data agnostic method.

From this experiment, we conclude that a realistic neural network is worth at least tens of thousands of examples on a fairly simple task (toy problem vs real world challenge) compared to an expert engineer. We thus provide a criterion that a data model is preferable if fewer than 10000 training samples are available.



**Fig. 6.** Test scores of the data-agnostic Pointflow and of the best transfer-learned or finetuned networks. Left: MSE computed on both the radius and the centre’s location estimation. Middle: same but only on the radius estimation. Right: same but only on the centre location’s estimation. We zoom-in in the bottom set of figures.

## 4 Conclusion

We analysed the amount of data required by neural networks, either shallow custom ones or deep famous ones, trained from scratch, finetuned, or transfer-learned, to compete with optimal or state-of-the-art traditional data-agnostic methods based on mathematical data generation models. To do so, we mathematically generated data, and fed various amounts of samples to the networks for training. We found that tens of thousands of data examples are needed for the networks to be on par or beat the traditional methods, if they are able to. For mathematical accuracy, we did not investigate real-world problems, which are

$N$	Pointflow	Alexnet			VGG			Resnet		
		TL	FT	SC	TL	FT	SC	TL	FT	SC
$(r, c)$	0	<b>0.66</b>								
	10	826	524	1748	1581	754	1748	5256	3358	2590
	100	314	66	291	448	69	215	2064	1124	793
	1000	183	23	9.0	189	6.5	7.6	1092	48	45
	10000	140	31	4.8	71	2.9	1.2	814	3.7	4.4
	100000	67	17	1.5	40	0.68	<b>0.42</b>	826	<b>0.51</b>	0.93
$r$	0	<b>0.26</b>								
	10	66	65	75	75	69	75	110	86	82
	100	36	5.9	42	52	5.6	40	42	25	31
	1000	26	2.4	1.3	20	0.57	0.62	24	2.2	1.3
	10000	16	2.0	1.5	9.9	<b>0.23</b>	<b>0.11</b>	19	<b>0.19</b>	0.27
	100000	6.8	1.7	0.69	4.9	<b>0.075</b>	<b>0.051</b>	19	<b>0.066</b>	<b>0.078</b>
$c$	0	<b>0.40</b>								
	10	759	454	1673	1506	675	1673	5142	3206	2507
	100	277	60	249	393	63	175	2020	1099	762
	1000	157	21	7.7	169	5.9	6.9	1067	46	43
	10000	118	29	3.0	61	2.6	1.1	795	3.5	4.1
	100000	57	6.4	0.79	36	0.57	<b>0.37</b>	807	0.45	0.86

**Table 2.** MSE scores of the networks compared to Pointflow on test data, computed on both  $r$  and  $c$  (top), just  $r$  (middle), or just  $c$  (bottom). Networks were trained on joint prediction of  $r$  and  $c$ . For the separate  $r$  (resp.  $c$ ) scores, the selected networks were those providing the best  $r$  (resp.  $c$ ) error on validation data.

commonly harder with less accurate or non-existent mathematical data generation models, but more data should be needed in those complex tasks. We have empirically derived a simple criterion, enabling researchers working on tasks where data is not easily available, to choose whether to use model-based traditional methods, by using either preexisting or newly created data generation models, or simply feed data to deep neural networks.

## A Pointflow implementation details

The pointflow dynamics are implemented by discretising time and approximating the time derivative with a forward finite difference scheme, although it could be improved with a Runge-Kutta 4 implementation [5]. Given the small magnitudes of the fields, we found that a large time step  $dt = 50$  works well. We define three thresholds,  $\tau_l = 0.9$  for  $C_l$ ,  $\tau_s = 10^{-6}$  for  $C_s$ , and  $\tau_{len} = 0.001$ . We consider having looped  $C_l$  if a point reaches a previous point within squared Euclidean distance  $\tau_l$  while having on the trajectory between the looping points at least one point with squared distance to them of at least  $\tau_l$ . A trajectory is stuck if it reaches a point where the current flow  $V$  has small magnitude  $\|V\|_2^2 \leq \tau_s$ . Each flow is run for  $N_i = 1000$  iterations, and trajectories shorter than  $\tau_{len}$  are discarded, e.g. trajectories of type  $C_s$ . We used  $\sigma_{Pf} = 5$  for blurring out the noise before computing the fields. The implemented pointflow algorithm for finding contours in our circle images is presented in Algorithm 1.

After computing the list of contours  $\mathcal{C}$  in the image  $I$ , we estimate the radius using the average curve length  $\hat{r} = \frac{1}{2\pi} \sum_{i=1}^{|\mathcal{C}|} \text{length}(\mathcal{C}_i)$ . Since the average of

---

**Algorithm 1** Contour integration with Pointflow on image  $I$

---

```

Compute  $I_b = g_{\sigma_{Pf}} * I$ 
Compute  $V_a = \nabla \|\nabla I_b\|_2$  and  $V_r = \nabla I_b^\top$ 
Compute  $V_+ = \frac{1}{2}(V_a + V_r)$  and  $V_- = \frac{1}{2}(V_a - V_r)$ 
Choose  $N_{Pf}$  random points independently and uniformly in the image domain
 $[-\frac{D-1}{2}, \frac{D-1}{2}]^2$ 
Let  $\mathcal{C} = []$  be an empty list of computed contours
for  $i = 1 \dots N_{Pf}$  do
  Let  $(traj_+, C)$  be the flow along  $V_+$  starting from the  $i$ -th point
  if  $C = C_l$  and  $\text{length}(traj_+) \geq \tau_{len}$  then
    Let  $(traj_l, C_l)$  be the reflow along  $V_+$  starting from the endpoint of  $traj_+$ 
     $C.append(traj_l)$ 
  else if  $C = C_o$  and  $\text{length}(traj_+) \geq \tau_{len}$  then
    Let  $(traj_-, C_-)$  be the reflow along  $V_-$  starting from the endpoint of  $traj_+$ 
    if  $C_- = C_l$  and  $\text{length}(traj_-) \geq \tau_{len}$  then
      Let  $(traj_l, C)$  be the reflow along  $V_-$  starting from the endpoint of  $traj_-$ 
       $C.append(traj_l)$ 
    end if
  end if
end for
Return  $\mathcal{C}$ 

```

---

the points did not yield the best estimation of the circle centre, we estimate it instead using least squares. The equation of a circle is naturally given by  $(x - c_x)^2 + (y - c_y)^2 = r^2$ , which can be written as  $\theta_1 x + \theta_2 y + \theta_3 = x^2 + y^2$ , where  $\theta_1 = 2c_x$ ,  $\theta_2 = 2c_y$ , and  $\theta_3 = r^2 - c_x^2 - c_y^2$ . We can thus estimate for each contour  $\theta = (\theta_1, \theta_2, \theta_3)^\top$  by least squares as  $\hat{\theta} = A^\top (AA^\top)^{-1} B$ , with  $A_{i,:} = (x_i, y_i, 1)$  and  $B_i = x_i^2 + y_i^2$  and  $i$  ranging in the number of computed points on the contour. From  $\hat{\theta}$  we can estimate  $\hat{c} = (\frac{\theta_1}{2}, \frac{\theta_2}{2})$ . The final centre estimation is then given by the average of this estimation over all contours. Note that we can also estimate  $r$  using  $\theta_3$  but we found that it did not outperform the length strategy so we do not use it.

**Acknowledgements** This work is in part supported by the French government under management of Agence Nationale de la Recherche as part of the "Investissements d'avenir" program, reference ANR-19-P3IA-0001 (PRAIRIE 3IA Institute).

## References

1. Alwosheel, A., van Cranenburgh, S., Chorus, C.G.: Is your dataset big enough? sample size requirements when using artificial neural networks for discrete choice analysis. *Journal of choice modelling* **28**, 167–182 (2018)
2. Anthony, M., Bartlett, P.L.: *Neural Network Learning: Theoretical Foundations*. Cambridge University Press (1999)

3. Bai, B., Yang, F., Chai, L.: Point flow edge detection method based on phase congruency. In: 2019 Chinese Automation Congress (CAC). pp. 5853–5858. IEEE (2019)
4. Bengio, Y., Lecun, Y., Hinton, G.: Deep learning for AI. *Communications of the ACM* **64**(7), 58–65 (2021)
5. Butcher, J.: *Numerical Methods for Ordinary Differential Equations*. Wiley (2008), <https://books.google.co.il/books?id=opd2NkBmMxsC>
6. Dagès, T., Lindenbaum, M., Bruckstein, A.M.: From compass and ruler to convolution and nonlinearity: On the surprising difficulty of understanding a simple cnn solving a simple geometric estimation task. *arXiv preprint arXiv:2303.06638* (2023)
7. Elad, M.: *Sparse and redundant representations: from theory to applications in signal and image processing*. Springer (2010)
8. Geirhos, R., Narayanappa, K., Mitzkus, B., Thieringer, T., Bethge, M., Wichmann, F.A., et al.: Partial success in closing the gap between human and machine vision. *Advances in Neural Information Processing Systems* **34**, 23885–23899 (2021)
9. Geirhos, R., Temme, C.R., Rauber, J., Schütt, H.H., Bethge, M., Wichmann, F.A.: Generalisation in humans and deep neural networks. *Advances in neural information processing systems* **31** (2018)
10. Goodfellow, I., Bengio, Y., Courville, A.: *Deep learning*. MIT press (2016)
11. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016)
12. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. *Communications of the ACM* **60**(6), 84–90 (2017)
13. Lakshmanan, V., Robinson, S., Munn, M.: *Machine learning design patterns*. O’Reilly Media (2020)
14. LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. *Nature* **521**(7553), 436–444 (2015)
15. LeCun, Y., Boser, B., Denker, J., Henderson, D., Howard, R., Hubbard, W., et al.: Handwritten digit recognition with a back-propagation network. *Advances in neural information processing systems* **2** (1989)
16. Mohamed, A.r., Dahl, G., Hinton, G.: Deep belief networks for phone recognition. In: *Nips workshop on deep learning for speech recognition and related applications*. vol. 1, p. 39 (2009)
17. Novikoff, A.B.: On convergence proofs for perceptrons. Tech. rep., Stanford Research Institute, Menlo Park, CA, USA (1963)
18. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., et al.: Imagenet large scale visual recognition challenge. *International journal of computer vision* **115**(3), 211–252 (2015)
19. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014)
20. Vapnik, V.N., Chervonenkis, A.Y.: On the uniform convergence of relative frequencies of events to their probabilities. In: *Measures of complexity*, pp. 11–30. Springer (2015)
21. Wiener, N.: *Extrapolation, interpolation, and smoothing of stationary time series: with engineering applications*, vol. 113. MIT press Cambridge, MA (1949)
22. Yang, F., Cohen, L.D., Bruckstein, A.M.: A model for automatically tracing object boundaries. In: *2017 IEEE International Conference on Image Processing (ICIP)*. pp. 2692–2696. IEEE (2017)