Learning on the Edge: Investigating Boundary Filters in CNNs

Carlo Innamorati[1] · Tobias Ritschel[1] · Tim Weyrich[1] · Niloy J. Mitra[1]

Abstract

Convolutional neural networks (CNNs) handle the case where filters extend beyond the image boundary using several heuristics, such as `zero`, `repeat` or `mean` padding. These schemes are applied in an ad-hoc fashion and, being weakly related to the image content and oblivious of the target task, result in low output quality at the boundary. In this paper, we propose a simple and effective improvement that *learns* the boundary handling itself. At training-time, the network is provided with a separate set of `explicit` boundary filters. At testing-time, we use these filters which have learned to extrapolate features at the boundary in an optimal way for the specific task. Our extensive evaluation, over a wide range of architectural changes (variations of layers, feature channels, or both), shows how the `explicit` filters result in improved boundary handling. Furthermore, we investigate the efficacy of variations of such boundary filters with respect to convergence speed and accuracy. Finally, we demonstrate an improvement of 5–20% across the board of typical CNN applications (colorization, de-Bayering, optical flow, disparity estimation, and super-resolution). Supplementary material and code can be downloaded from the project page: http://geometry.cs.ucl.ac.uk/projects/2019/bounce-neural-resim/http://geometry.cs.ucl.ac.uk/projects/2019/investigating-edge/.

1 Introduction

When performing convolutions on a finite domain, boundary rules are required as the kernel's support extends beyond the edge. For convolutional neural networks (CNNs), many discrete filter kernels "slide" over a 2D image and typically boundary rules including `zero`, `reflect`, `mean`, `clamp` are used to extrapolate values outside the image.

Considering a simple detection filter (Fig. 1a) applied to a diagonal feature (Fig. 1b), we see that no boundary rule is ever ideal: `zero` will create a black boundary halo (Fig. 1c), using the `mean` color will reduce but not remove the issue (Fig. 1d), `reflect` and `clamp` (Fig. 1e, f) will create different kinks in a diagonal edge where the ground-truth continuation would be straight. In Fig. 1 we visualize this as the error between the ideal response and the response we would observe at a location if a feature was present. In practical feature channels, these will manifest as false pos-

itive and negative images. These deteriorate overall feature quality, not only on the boundary but also inside. Another, equally unsatisfying, solution is to execute the CNN only on a "valid" interior part of the input image (crop), or to execute it multiple times and merge the outcome slide. Working in lower or multiple resolutions, the problem is even stronger, as low-resolution images have a higher percentage of boundary pixels. In a typical modern encoder–decoder (Ronneberger et al. 2015), *all* will eventually become boundary pixels at some step.

Having a second thought on what a 2D image actually is, we see, that the ideal boundary rule would be the one that extends the content exactly to the values an image taken with a larger sensor would have contained. Such a rule appears elusively hard to come by as it relies on information not observed. We cannot decide with certainty from observing the yellow part inside the image in Fig. 1b how the part outside the image continues—what if the yellow structure really stopped?—and therefore it is unknown what the filter response should be. However, neural networks have the ability to extrapolate information from a context, for example in inpainting tasks (Ren et al. 2015). Here, this context is the image part inside the boundary. Given this observation,

---

Communicated by Ling Shao, Hubert P. H. Shum, Timothy Hospedales.

✉ Carlo Innamorati
  c.innamorati@cs.ucl.ac.uk

[1] University College London, London, UK

Fig. 1 Applying a feature detection-like filter (**a**) to an image with different boundary rules (**b**)–(**f**). We show the error as the ratio of the ideal and the observed response. A bright value means a low error due to a ratio of 1 i.e., the response is similar to the ideal condition. Darker values indicate a deterioration

not every extension is equally likely. Most human observers would follow the Gestalt assumption of continuity and predict the yellow bar to continue at constant slope outside the image.

Addressing the boundary challenge, and making use of a CNN's extrapolating power, we propose the use of a novel `explicit` boundary rule in CNNs. As such rules will have to depend on the image content and the spatial location of that content, we advocate to model them as a set of learned *boundary filters* that simply replace the non-boundary filters when executed on the boundary. Every boundary configuration (upper edge, lower left corner, etc.) has a different filter. This implies, that they incur no time or space overhead at runtime. At training-time, boundary and non-boundary filters are jointly optimized and no additional steps are required. Additionally, in this extended version of the original work (Innamorati et al. 2018), we study the efficiency and performance of different variants of boundary-filters and report quantitative results and convergence rate, averaged over multiple runs. It seems, that introducing more degrees of freedom increases the optimization challenge. However, introducing the right degrees of freedom, can actually turn an unsolvable problem into separate tasks that have simple independent solutions, as we conclude from a reduction of error both at the interior and at the edges, when using our method.

After reviewing previous work and introducing our formalism, we demonstrate how using `explicit` boundary conditions can improve the quality across a wide range of possible architectures (Sect. 4). We next show improvement in performance for tasks such as de-noising and de-bayering (Gharbi et al. 2016), colorization (Zhang et al. 2016) as well as disparity, scene flow (Dosovitskiy et al. 2015) and super-resolution (Dong et al. 2016), in Sect. 5. Finally, in Sect. 6, we investigate the benefits of different implementation strategies of `explicit` boundary conditions.

2 Previous Work

Our work extends deep convolutional neural networks (Goodfellow et al. 2016) (CNNs). To our knowledge, the

immediate effect of boundary handling has not been looked into explicitly. CNNs owe a part of their effectiveness to weight-sharing or -invariance property: only a single convolution needs to be optimized that is applied to the entire image (Fukushima and Miyake 1982). Doing so, inevitably, the filter kernel will touch upon the image boundary at some point. Classic CNNs use zero padding (Ciregan et al. 2012), i.e., they enlarge the image by the filter kernel size they use, or directly crop, i.e., run only on a subset (Krizhevsky et al. 2012) and discard the boundary. Another simple solution is to perform filtering with an arbitrary boundary handling and crop the part of the image that remains unaffected: if the filter is centered and 3 pixels wide, a 100×100 pixel image is cropped to 98×98 pixels. This works in a single resolution, but multiple layers, in particular at multiple resolutions, grow the region affected by the boundary linearly or even exponentially. For example, the seminal U-net (Ronneberger et al. 2015) employs a complicated sliding scheme to produce central patches from a context that is affected by the boundary, effectively computing a large fraction of values that are never used. We show how exactly such a U-net-like architecture can be combined with `explicit` boundaries to realize a better efficacy with lower implementation and runtime overhead. Other work has extended the notion of invariance to flips (Cohen and Welling 2016) and rotations (Worrall et al. 2017). Our extension could be seen as attempting to add invariance under boundary conditions. For some tasks like inpainting, however, invariance is not desired, and translation-variant convolutions are used (Ren et al. 2015). This paper shares the idea to use different convolutions in different spatial locations. Uhrig et al. (2017) have weighted convolutions to skip pixels undefined at test time. In our setting, the undefined pixels are known at train time to always fall on the boundary. By making this explicit to the learning, it can capitalize on knowing how the image extends.

3 Explicit Boundary Rules

In this section we will define convolutions that can account for `explicit` boundary rules, before discussing the loss and implementation options.

Fig. 2 Example domain decomposition for a 5×5 input image and 3×3 filters. Edges and corners of the input are highlighted and color coded. Kernels $g_{2,...n}$ are color coded to match the portion of the input that they will read from. White areas denote entries outside of the input domain. As shown, each kernel is centered in the corresponding location of the input that it is responsible for (Color figure online)

Convolution Key to `explicit` boundary handling is a domain decomposition. Intuitively, in our approach, instead of running the same filter for every pixel, different filters are run at the boundary. This is done independently for every convolution kernel in the network. For simplicity, we will here explain the idea for a single kernel that computes a single feature. The extension to many kernels and features is straightforward. Again, for simplicity, we describe the procedure for a 2D convolution, mapping scalar input to scalar output. The 3D convolution, mapping higher-dimensional input to scalar output is derived similarly.

A common `zero` boundary handling convolution $*_0$ of an input image $f^{(in)}$ with the kernel g is defined as

$$f^{(out)}[\mathbf{x}] *_0 g = \sum_{\mathbf{y} \in \mathcal{K}} \begin{cases} f^{(in)}[\mathbf{x} + \mathbf{y}] \cdot g[\mathbf{y}] & \text{if } \mathbf{x} + \mathbf{y} \in \mathcal{D} \\ 0 & \text{otherwise,} \end{cases}$$

$$(1)$$

where $\mathcal{K}$ is the kernel domain, such as $\{-1, 0, 1\}^2$ and $\mathcal{D}$ is the image domain in pixel coordinates from zero to image width and height, respectively. We extend this to `explicit` boundary handling $*_e$ using a family of kernels $g_{1,...n}$ as

$$f^{(out)}[\mathbf{x}] *_e g_{1,...,n} = \sum_{\mathbf{y} \in \mathcal{K}} \begin{cases} f^{(in)}[\mathbf{x} + \mathbf{y}] \cdot g_{s[\mathbf{x}]}[\mathbf{y}] & \text{if } \mathbf{x} + \mathbf{y} \in \mathcal{D} \\ 0 & \text{otherwise,} \end{cases}$$

$$(2)$$

where $s[\mathbf{x}]$ is a selection function that returns the index from 1 to n of the filter to be used at position $\mathbf{x}$ (Fig. 2). The number of filters n depends on the size of the receptive field: For a 3×3 filter it is 9 cases, for larger fields it is more.

Loss The loss is defined on multiple filter kernel values $g_{1,...n}$ instead of a single kernel. As this construction comprises of linear operations only (the selection function can be written as nine multiplications of nine convolution results with nine masks that are 0 or 1 and a final addition), it is back-propagatable.

Implementation A few things are worth noting for the implementation. First, applying multiple kernels in this fashion has the same complexity as applying a single kernel. Convolution in the Fourier domain, where costs would differ, is typically not done for kernels of this size. Second, the memory requirement is the same as when running with common boundary conditions. All kernels jointly output one single feature image. The boundary filters are never run in the interior part of the input and no result is stored for it. The only overhead is in storing the filter masks. In practice however, implementation constants might differ between implementations, in particular for parallel machines (GPUs).

The first practical option for implementation is the most *compatible* one that just performs all nine convolutions on the entire image and later composes the nine images into a single image. This indeed has compute and memory cost linear in the number of filters, i.e., nine times more expensive, both for training and deployment.

To avoid the overhead, without having to access the low level code of the framework in use, the additional kernels can be trained on the specific sub-parts of the input that they act on and then composited back to form the output.

4 Analysis

We will now analyze the effect of border handling for a simplified task and different networks: learning how to perform a Gaussian blur of a fixed size. This task is suitable for our experiments as it provides a rotational invariant filtering, in a setting where exactly this is violated when going over the boundary. Despite the apparent simplicity, we will see, how many different variants of a state-of-the-art U-net-like (Ronneberger et al. 2015) architecture all suffer from similar

boundary handling problems. This indicates, that the deteriorating effect of unsuccessful boundary handling cannot be overcome by adapting the network structure, but needs the fundamentally different domain decomposition we suggest.

4.1 Methods

Task The tasks is to learn the effect of a Gaussian filter of size 13×13 to 128×128 images, obtained from the dataset used for the ILSVRC (Russakovsky et al. 2015) competition, comprising of over one million images selected from ImageNet (Deng et al. 2009). The ground truths were computed over $(128 + 12) \times (128 + 12)$ images, which were then cropped to 128×128.

Metrics We compare to the reference by means of the MSE metric, which was also used as the loss function. The models were selected by comparing the loss values over validation set, while the reported loss values were separately computed over a test set comprising of 10 k examples.

Architecture We use a family of architectures to cover both breadth and width of the network. The breadth is controlled by the number of feature channels and the depth by the number of layers. More specifically, the architecture comprises of n_1 layers. Each layer performs a convolution to produce n_f feature channels, followed by a ReLU non-linearity. We

choose such an architecture, to show that the effect of boundary issues is not limited to a special setting but remains fundamental.

Boundary Handling We include our `explicit` handling, as well as the classic `zero` strategy that assumes the image to be 0 outside the domain and `reflect` padding, that reflects the image coordinate around the edge or corner.

4.2 Experiments

Here, we study how different architecture parameters affect boundary quality for each type of boundary handling.

Varying Depth When varying depth n_1 from a single up to 7 layers (Fig. 3a) we find, that our `explicit` boundary handling performs best on all levels, followed by `reflect` boundary handling and `zero`. The feature channel count is held fixed at $n_f = 3$.

Varying Feature Count When varying feature channel count n_f, it can be seen that `explicit` leads the board, followed by `reflect` and `zero` (Fig. 3b). The depth is held fixed at $n_d = 2$.

Varying Feature Count and Depth When varying both depth n_1 and feature channel count n_f, seen in Fig. 3, c we find, that again no architectural choice can compensate for
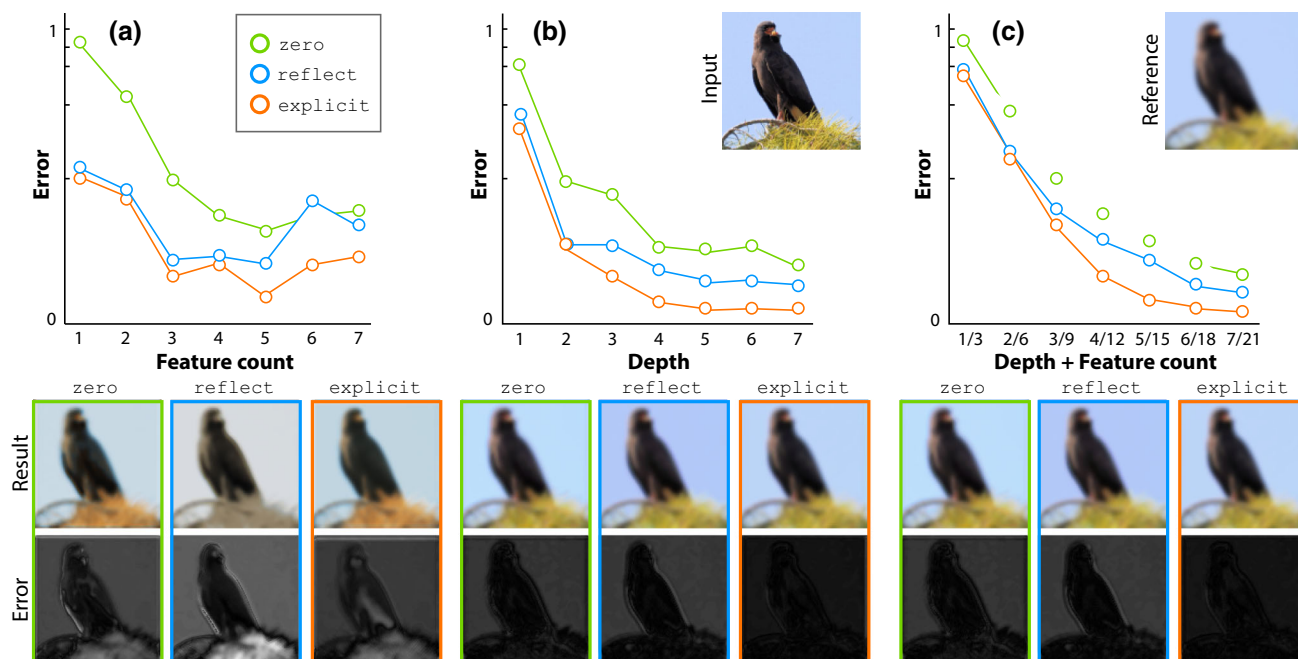

Fig. 3 Analysis of different architectural choices using different boundary handling (colors). First, **a** we increase feature channel count (first plot and columns of insets). The vertical axis shows log error for the MSE (our loss) and the horizontal axis different operational points. Second **b**, depth of the network is increased (second plot and first 4 columns of insets). Third, **c** both are increased jointly. The second row of insets shows the best (most similar to a reference) result for each boundary method (column) across the variation of one architectural parameter for a random image patch (input an reference result seen in corner) (Color figure online)

Table 1 Quantitative results

Task	Src	Absolute error				Error ratio		
		Other metric				MSE ratio		
		Unit	`refl`	`zero`	Ours	`refl`(%)	`zero`(%)	Ours (%)
Gaussian filtering		DSSIM	.0018	.0022	**.0016**	79	83	**100**
De-noise/Bayer	Gharbi et al. (2016)	PSNR	31.46	31.50	**31.94**	90	89	**100**
Colorization	Zhang et al. (2016)	DSSIM	.1593	.1604	**.1577**	99	98	**100**
Disparity	Dosovitskiy et al. (2015)	px	1.538	1.511	**1.403**	84	88	**100**
Scene flow	Dosovitskiy et al. (2015)	px	1.380	1.183	**1.096**	56	73	**100**
Super-resolution	Dong et al. (2016)	DSSIM	7.3e–4	8.8e–4	**5.3e–4**	78	86	**100**

Different rows are different tasks. Different columns express different measures and different methods. *Absolute error* is measured using different metrics (eventually not identical to the loss), while the *Error ratio* is expressed as the ratio of the loss of our method over the opposing boundary rules. Best is bold. The architecture used for these tests is described in Sect. 4

the boundary effects. Each of the seven steps increase feature count by 3 and depth by 1.

Statistical Analysis A two-sided t test ($N = 10{,}000$) rejects the hypothesis that our method is the same as any other method for any task with $p < .001$.

5 Applications

Now, we compare different boundary handling methods in several typical applications.

5.1 Methods

Architecture We use an encoder–decoder network with skip connections (Ronneberger et al. 2015) optimized for the MSE loss using the ADAM optimizer.

Details are shown in our supplemental materials. The architecture is different from the simplified one in the previous section where it was important to systematically explore many possible variants. The encoding proceeds in 3×3 convolution steps 1 to 7, increasing the number of feature channels from 1 to 256. There is a flat 1×1 convolution at the most abstract representation at stage 8. Decoding happens on stages 9 to 14. This step resizes the image, convolves with stride 1 and outputs the stated number of features, followed by a concatenate convolution by the stated skip ID and finally a convolution with stride 1 that outputs the stated number of features (`ResConv`).

Note, that boundary handling is required at all stages except 8. For the down-branch 1–7 it can be less relevant as strides do not produce all edge cases we handle. This is because, with an even resolution scheme and a stride of 2, the last row and column are skipped. Consequentially, three of the four corners and two of the four edges are skipped, too.

Measure We apply different task-specific measures: Gaussian filtering and Colorization produce images for human observers and consequently are quantified using DSSIM. De-Bayering, as a de-noising task, is measured using the PSNR metric while disparity and scene flow are image correspondence problems with results in pixel units.

Additionally, we propose to measure the success as the *loss ratio* between the test loss of our architecture with and the test loss of an architecture without explicit boundary handling, using the MSE metric. We suggest to use the ratio as it abstracts away from the unit and the absolute loss value that depends on the task, allowing to compare effectiveness across tasks (Table 1).

5.2 Results

Gaussian Blur Gaussian filtering is a simple baseline task with little relevance to any practical application as we know the solution (Sect. 4). It is relevant to our exposition, as we know that, if the network had seen the entire world (and not just the image content) it would be able to solve the task. It is remarkable, that despite the apparent simplicity of the task—it is a single linear filter after all—the absolute loss is significant enough to be visible for classic boundary handling. It is even more surprising, that the inability to learn a simple Gaussian filter does not only result in artifacts along the boundaries, but also in the interior. This is to be attributed to the inability of a linear filter to handle the boundary. In other words, it is surprisingly complex for a network without explicit boundary handling to learn a task as easy as blurring an image. We will see that this observation can also be made for more complex tasks in the following sections.

De-noising and De-bayering In this application we learn a mapping from noisy images with a Bayer pattern to clean images using the training data of Gharbi et al. (2016). The measure is the PSNR, peak signal-to-noise ratio (more is better). We achieve the best PSNR at 31.94, while the only

change is the boundary handling. In relative terms, traditional boundary handling can achieve only up to 90% of MSE.

Colorization Here we learn the mapping from grey images to color images using data from Zhang et al. (2016). The metric again is DSSIM. We again perform slightly better in both absolute and relative terms.

Disparity and Scene Flow Here we learn the mapping from RGB images to disparity and scene flow using the data from Dosovitskiy et al. (2015). We measure error in pixel distances (less is better). Again, adding our boundary handling improves both absolute and relative error. In particular, the error of `reflect` and `zero` is much higher for scene flow.

Super-Resolution In this application, we learn a mapping from low to high resolution images. In order to keep the same architecture and maintain consistency, images of equal size had to be used. For this reason, low resolution images are mapped to high resolution sub-patches of the original image using a 1:4 ratio. This definition of the problem allows us to keep both the input and output size consistent with the other tasks, while still performing the super resolution task, albeit only on a sub-portion of the input image. Super-resolution (Dong et al. 2016) is a classic Deep Learning task that has been gaining traction in recent years. We used the DSSIM metric to measure our success, with the `explicit` strategy over-performing both in relative and absolute error.

5.3 Discussion

We now will discuss the benefit and challenges of `explicit` boundary handling.

Overhead Here we study four implementation alternatives for Sect. 3. They were implemented as a combination of OpenGL geometry and fragment shaders. The test was ran on a Nvidia Geforce 480, on a 3 mega-pixel image and a 3×3 receptive field.

The *first* method uses a simple `zero`-padding provided by OpenGL's `sampler2D`, invoking the GS once to cover the entire domain and applying the same convolution everywhere. This requires 2.5 ms. This is an upper bound for any convolution code.

The *second* implementation executes nine different convolutions, requiring 22.5 ms. This invokes the GS nine times, each invoking all pixels.

The *third* variant invokes the GS once and a conditional statement for all pixels selects the kernel weights per-pixel in the domain. This requires 11.2 ms.

The *fourth* variant, a domain decomposition, invokes the GS nine times to draw nine quads that cover the respective interior and all boundary cases as seen in Fig. 2. Even after averaging a high number of samples, we could not find evidence for this to be slower than the baseline method i.e.,


Fig. 4 Convergence rate with different types of boundary

2.5 ms. This is not unexpected, as the running time for a few boundary pixels is below the variance of the millions of interior pixels.

In practice, the learning is limited by other factors such as disk-IO. Our current implementation in Keras (Chollet 2015), offers a simple form of domain decomposition. We tested the performance loss over epochs with an average duration of 64 seconds. Our method results in a 0.2% average performance loss over the classic `zero` rule.

Scalability in Receptive Field Size For small filters, the number of cases is small, but grows for larger filters. Fortunately, the trend is to rather cascade many small filters in deeper network, instead of shallower networks with large filters.

Convergence Convergence of both our approach and traditional `zero` boundary handling is seen in Fig. 4. We find, that our method is not only resulting in a smaller loss, but also does so at the same number of epochs. Before we have established that the duration of epoch are the same for both methods. We conclude there is no relevant training overhead for our method.

Structure Here we seek to understand where spatially in the image the differences are strongest. While our approach changes the processing on edges, does it also affect the interior? We compute the per-pixel MAE and average this over all images in the corpus. The resulting error images are seen in Fig. 5. We found the new method to consistently improve results in the interior regions. It looks as if the new boundary rules effectively "shield" the inner regions from spurious boundary influences. The results at the boundaries are very competitive too, often better than `zero` and `reflect` boundary handling. Note, that it is not expected for any method, also not ours, to have a zero error at the boundary: this would imply we were able to perfectly predict unobserved data outside of the image.

Practical Alternatives There are simpler alternatives to handle boundaries in an image of n_p pixels. We will consider a
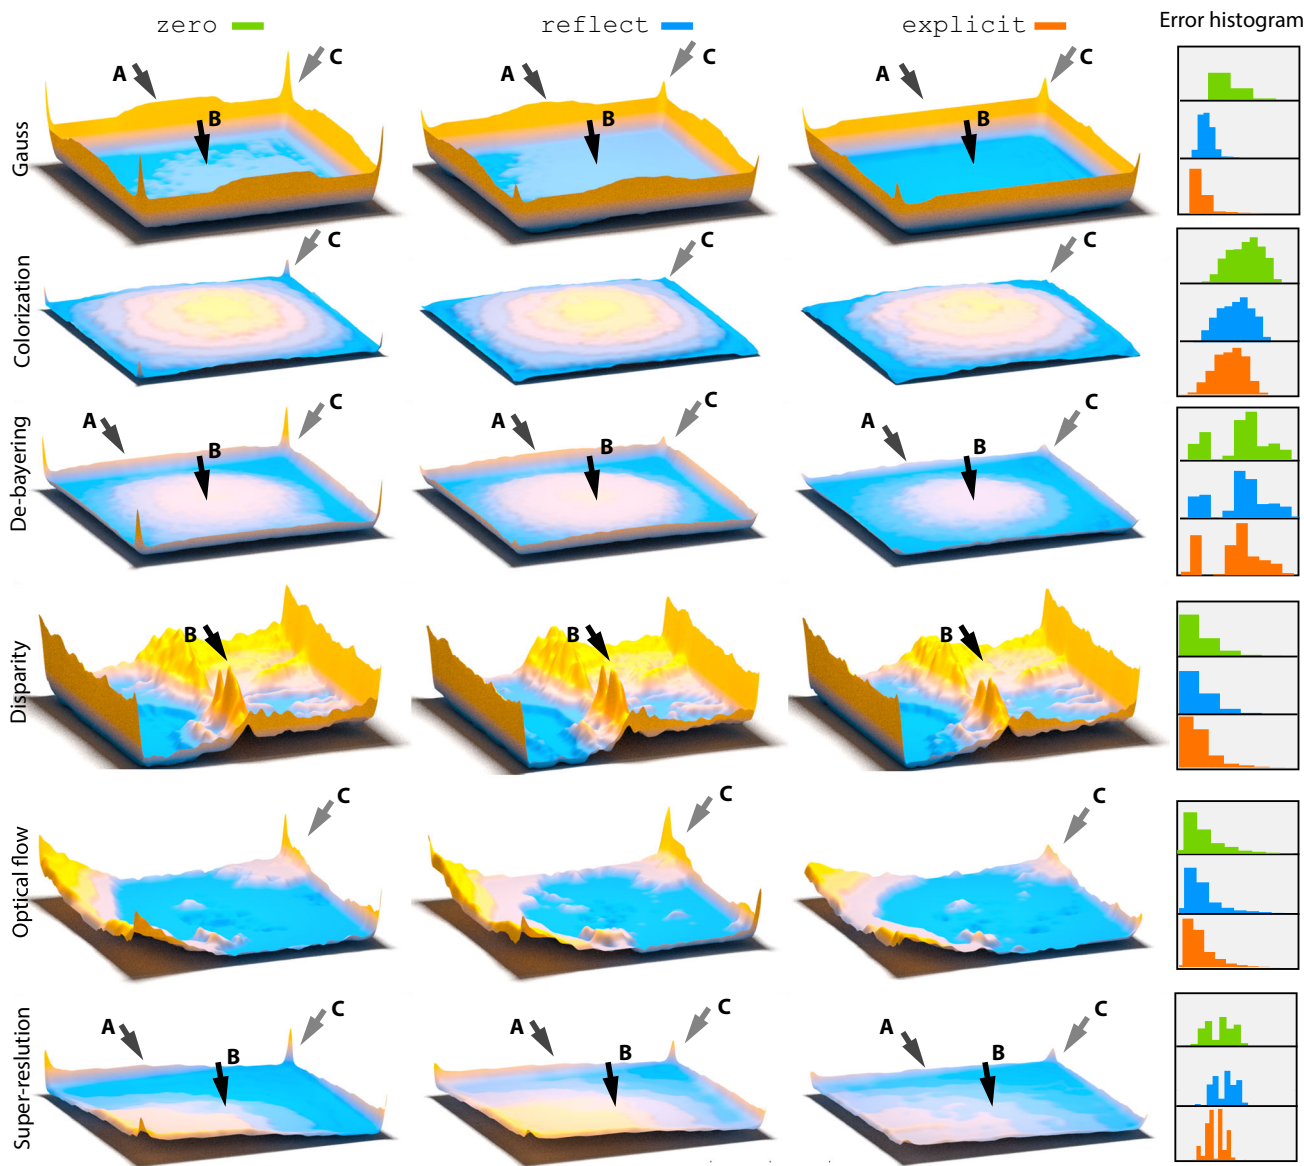
Fig. 5 Mean errors across the corpus visualized as height fields for different tasks and different methods. Each row corresponds to one task each column to one way of handling the boundary. Arrow A marks the edge that differs (ours has no bump on the edge). Arrow B mark the interior that differs (ours is flat and blue, others is non-zero, indicting we improve also inside). Arrow C shows corners, that are consistently lower for us

1D domain as an example here. The first is to crop n_c pixels on each side and compute only $n_p - 2n_c$ output pixels. The cropping n_c is to be made sufficiently large, such that no result is affected by a boundary pixel and n_c depends on the network structure. In a single-resolution network of depth n_d with a receptive field size of $2n_r+1$, we see, that $n_c = n_d \times n_r$. In a multi-resolution network however, the growth is exponential, so $n_c = n_r^{n_d}$, and for a typical encoder–decoder that proceeds to a resolution of 1×1, every pixel is affected. This leaves two options: either the minimal resolution is capped and the CNN is applied in a sliding window fashion (Ronneberger et al. 2015), computing always only the unaffected result part, incurring a large waste of resources, or the network simply has to use its own resources to make do with the inconsistent input it receives.

6 Refined Boundary Filters

Our results indicate that learning additional filters can reduce the error that classic convolutions introduce when reading the input at its boundaries. We also show our implementation to
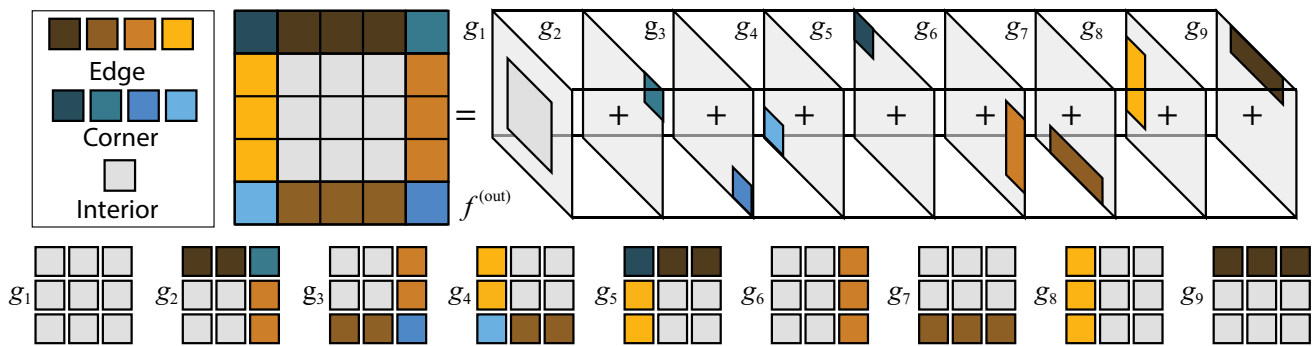
Fig. 6 Example domain decomposition for a 5×5 input image and 3×3 `explicit-shift` boundary filters. Edges and corners of the input are highlighted and color coded. Kernels $g_{2,...,n}$ are color coded to match the portion of the input that they will read from. As opposed to Fig. 2, where each kernel is centered in the corresponding location of the input that it is responsible for, each filter is offset by one row and/or column to avoid reading values that are outside of the input domain. Consequently, during training, the filters will be required to learn to *shift* back the information that they access (Color figure online)

have limited impact on performance. In this section, we investigate the effect of different variations of boundary filters. Our basic implementation of the boundary filters, referred to in the paper as `explicit` boundary filters, is what was used for the results in all previous sections of the paper. In addition, we now detail two alternatives: `explicit-shift` boundary filters and `explicit-randomized` boundary filters.

Explicit For `explicit` boundary filters, as shown in Fig. 2, filters $g_{2,...,n}$ are reading values outside of the input's boundary. Said values, for `explicit` filters, are filled with zeros. In this mode, while the network can learn to treat the boundary values of the input differently, it still has *less* information to solve the task at the boundaries than elsewhere.

We now introduce a new strategy with the objective of minimizing the aforementioned limitations. For `explicit-shift` boundary filters, we avoid any padding altogether by only reading values from the input. This is achieved by shifting the filter centers away from the boundary, allowing the filters to read the same values, but additionally seeing further data from the input, instead of (padded) zeros. Figure 6 re-codes filters $g_{2,...,n}$ from Fig. 2 to match the behaviour described here. Using the `explicit-shift` strategy, the filters will have further information to solve the task at hand, but will additionally have to learn about the shift.

Randomized As boundary filters see substantially less data than interior filters at training time, we experimented whether allowing the boundary and interior filters to see an equal amount of data would benefit the training procedure. To this end, we randomized, at training time, the filter that would take the interior part of the image. The implementation follows the `explicit-shift` paradigm, meaning that the interior patch being read is dependent on the randomly selected filter. For example, if the `left` edge is selected as the filter for the interior, the input patch will comprise the rightmost column, but not the leftmost.

Evaluation We evaluate the aforementioned implementation choices both by comparing the quantitative results obtained with the different filter strategies and by observing the loss behaviour of each choice at training time. As the differences between the methods is intuitively smaller when compared to previous results, we opted to only test out tasks that could be run on the ImageNet dataset, for statistical consistency.

Figure 7 displays the loss behaviour of the different boundary filter implementations. The graphs were obtained by averaging the trends over three differently seeded runs, for each task. This was done to further decrease variance. Interestingly, while we notice an improvement for the `explicit-shift` strategy for the De-Bayering and Super-resolution tasks, this is not the case for Gaussian filtering and Colorization. This suggests the improvement to be possibly task-dependent. Comparing the plots in Fig. 5, we notice how for both Gaussian filtering and Colorization tasks, the bulk of the improvement did not reside at the corners and edges themselves, but rather in the mitigation of the overall bias, which might explain why the different boundary strategies appear the perform similarly. With regard to the `explicit-randomized` strategy, the values at training time are expected to be lower, as the boundary filters cannot be as effective as the inner filters they randomly substitute.

Further, we quantitatively evaluate over a test set comprising of 10k examples, similar to Sec. 5.1. Table 2 displays the results of such analysis. The results confirm the expectations from Fig. 7. Namely, that the `explicit-shift` paradigm appears to be performing better, on average, than other boundary filter strategies, in addition of not requiring any padding. On the other hand, the `explicit-randomized` boundary filter strategy came out as the worst among the three variants. We conclude that the advantages of further training
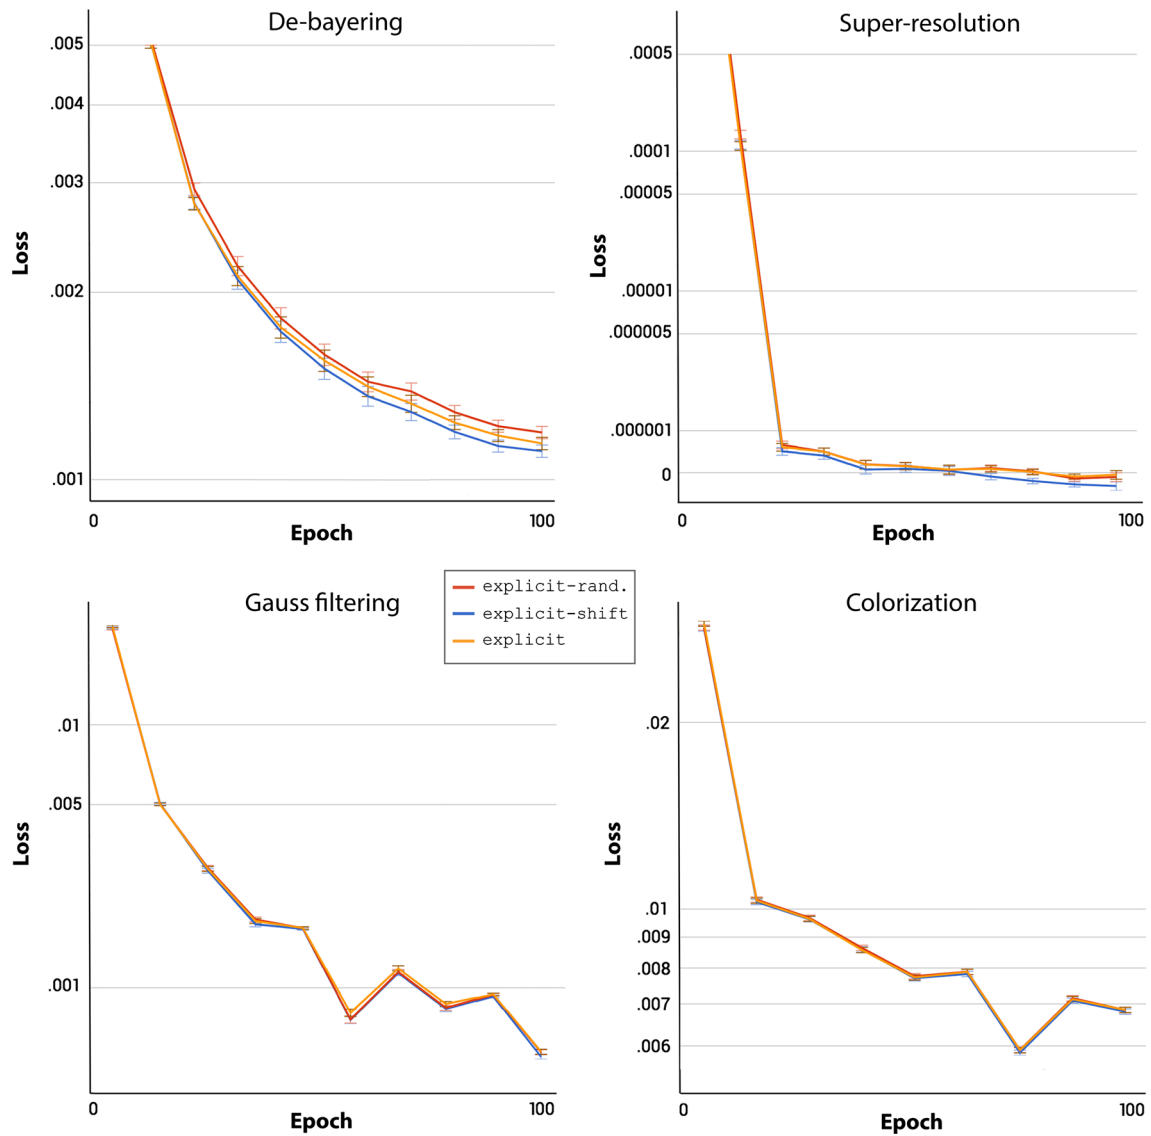
Fig. 7 Convergence rate for different tasks and boundary filter strategies. The plots are constructed by averaging the trends over ten differently seeded runs, for each task, in order to minimize the impact of variance. Following the same reasoning, we only test on tasks that can be trained using the ImageNet dataset. For each plot, we show error bars to display the standard deviation across the runs

Table 2 Quantitative results

Task	Unit	Other metric			MSE		
		expl.	rand.	shift	expl.	rand.	shift
Gaussian filtering	DSSIM	.0016	.0016	**.0016**	.0007	.0007	**.0006**
De-noise/Bayer Gharbi et al. (2016)	PSNR	31.94	31.42	**32.13**	.0014	.0016	**.0011**
Colorization Zhang et al. (2016)	DSSIM	**.1577**	.1581	.1579	0059	.0063	**.0057**
Super-resolution Dong et al. (2016)	DSSIM	5.3e–4	5.1e–4	**4.4e–4**	2.1e–7	2.1e–7	**1.8e–7**

Different rows are different tasks. Columns express different measures and methods. The architecture used for these tests is described in Sect. 4

the boundary filters is not justified by comparing the additional training time against marginal improvements.

7 Conclusion

In traditional image processing, the choice of boundary rule was never fully satisfying. In this work, we provide evidence, that CNNs offer the inherent opportunity to jointly extract features and handle the boundary as if the image continues naturally. We do this by learning filters that are executed on the boundary along with traditional filters executed inside the image. We further investigate the benefit of boundary filters by introducing different implementation strategies for the additional filters. We show that the `explicit-shift` paradigm performs better, on average, than other boundary filter strategies, while also lifting the need to pad its input. Incurring little learning and no execution overhead, the concept is simple to integrate into an existing architecture, which we demonstrate by increased result fidelity for a typical encoder–decoder architecture on practical CNN tasks. We therefore encourage testing our method for any CNN application where the visual quality of the output carries reasonable weight. Future work could investigate whether the statistical bias of humans to favor framing certain objects and patterns over others at the boundary, could affect the benefit of boundary rules and how it could be leveraged. In a similar direction, future research could analyze differences in the learned weights of individual boundary filters and draw conclusions from potential recurring patterns.

References

Chollet, F., et al. (2015). Keras. Retrieved July, 2019, from https://keras.io

Ciregan, D., Meier, U., & Schmidhuber, J. (2012). Multi-column deep neural networks for image classification. In *CVPR* (pp. 3642–49).

Cohen, T., & Welling, M. (2016). Group equivariant convolutional networks. In *ICML* (pp. 2990–2999).

Deng, J., Dong, W., Socher, R., Jia Li, L., Li, K., & Fei-fei, L. (2009). Imagenet: A large-scale hierarchical image database. In *CVPR*.

Dong, C., Loy, C. C., He, K., & Tang, X. (2016). Image super-resolution using deep convolutional networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *38*(2), 295–307. https://doi.org/10.1109/TPAMI.2015.2439281.

Dosovitskiy, A., Fischer, P., Ilg, E., Häusser, P., Hazırbaş, C., Golkov, V., v.d. Smagt, P., Cremers, D., & Brox, T. (2015). Flownet: Learning optical flow with convolutional networks. In *ICCV*.

Fukushima, K., & Miyake, S. (1982). Neocognitron: A self-organizing neural network model for a mechanism of visual pattern recognition. In Competition and cooperation in neural nets (pp. 267–85). Berlin: Springer.

Gharbi, M., Chaurasia, G., Paris, S., & Durand, F. (2016). Deep joint demosaicking and denoising. *ACM Transactions on Graphics*, *35*(6), 191.

Goodfellow, I., Bengio, Y., Courville, A., & Bengio, Y. (2016). *Deep learning* (Vol. 1). Cambridge: MIT Press.

Innamorati, C., Ritschel, T., Weyrich, T., & Mitra, N. J. (2018). Learning on the edge: Explicit boundary handling in CNNs. In *Proceedings of the british machine vision conference (BMVC)*. BMVA Press. Selected for oral presentation

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In *NIPS* (pp. 1097–1105).

Ren, J.S., Xu, L., Yan, Q., & Sun, W. (2015). Shepard convolutional neural networks. In *NIPS* (pp. 901–09).

Ronneberger, O., Fischer, P., & Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. In *International conference on medical image computing and computer-assisted intervention* (pp. 234–41).

Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., et al. (2015). ImageNet large scale visual recognition challenge. *International Journal of Computer Vision (IJCV)*, *115*(3), 211–252. https://doi.org/10.1007/s11263-015-0816-y.

Uhrig, J., Schneider, N., Schneider, L., Franke, U., Brox, T., & Geiger, A. (2017). Sparsity invariant CNNs. arXiv:1708.06500.

Worrall, D. E., Garbin, S. J., Turmukhambetov, D., & Brostow, G. J. (2017). Harmonic networks: Deep translation and rotation equivariance. In *CVPR*.

Zhang, R., Isola, P., & Efros, A. A. (2016). Colorful image colorization. In *ECCV*.