

Alma Mater Studiorum Università di Bologna
Archivio istituzionale della ricerca

Continual Adaptation for Deep Stereo

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Poggi, M., Tonioni, A., Tosi, F., Mattoccia, S., Di Stefano, L. (2022). Continual Adaptation for Deep Stereo. IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE, 44(9), 4713-4729 [10.1109/TPAMI.2021.3075815].

Availability:

This version is available at: <https://hdl.handle.net/11585/819561> since: 2021-04-30

Published:

DOI: <http://doi.org/10.1109/TPAMI.2021.3075815>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

This is the final peer-reviewed accepted manuscript of:

M. Poggi, A. Tonioni, F. Tosi, S. Mattoccia and L. D. Stefano, "Continual Adaptation for Deep Stereo," in IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 44, no. 9, pp. 4713-4729, 1 Sept. 2022

The final published version is available online at
<https://dx.doi.org/10.1109/TPAMI.2021.3075815>

Rights / License:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>)

When citing, please refer to the published version.

Continual Adaptation for Deep Stereo

Matteo Poggi, *Member, IEEE*, Alessio Tonioni, *Member, IEEE*, Fabio Tosi, *Student Member, IEEE*, Stefano Mattoccia, *Member, IEEE*, and Luigi Di Stefano, *Member, IEEE*

Abstract—Depth estimation from stereo images is carried out with unmatched results by convolutional neural networks trained end-to-end to regress dense disparities. Like for most tasks, this is possible if large amounts of labelled samples are available for training, possibly covering the whole data distribution encountered at deployment time. Being such an assumption systematically unmet in real applications, the capacity of *adapting* to any unseen setting becomes of paramount importance. Purposely, we propose a continual adaptation paradigm for deep stereo networks designed to deal with challenging and ever-changing environments. We design a lightweight and modular architecture, *Modularly ADaptive Network (MADNet)*, and formulate *Modular ADaptation* algorithms (*MAD*, *MAD++*) which permit efficient optimization of independent sub-portions of the entire network. In our paradigm, the learning signals needed to continuously adapt models online can be sourced from self-supervision via right-to-left image warping or from traditional stereo algorithms. With both sources, no other data than the input images being gathered at deployment time are needed. Thus, our network architecture and adaptation algorithms realize the first real-time self-adaptive deep stereo system and pave the way for a new paradigm that can facilitate practical deployment of end-to-end architectures for dense disparity regression.

Index Terms—Stereo Matching, Deep Learning, Self-supervision, Real-time Adaptation, Continual Learning

1 INTRODUCTION

Estimating dense and accurate depth maps is a key perception step to pursue scene comprehension tasks dealing with navigation and interaction with the environment. Passive, image-based techniques aimed at depth perception compare favourably to active sensors in terms of cost, bulkiness as well as - more often than not- working range and flexibility. Among such techniques, stereo vision [1] is usually the preferred choice, requiring just a pair of synchronized and calibrated cameras to measure depth by triangulation between matching pixels.

Akin to most computer vision problems, in the last years deep learning has entered into solutions for stereo matching, at first replacing certain specific steps of the pipeline by neural networks (*e.g.*, matching cost computation [2]) then rapidly converging toward end-to-end architectures [3], [4]. Although end-to-end networks have established the new state-of-the-art in challenging benchmarks such as KITTI [5], [6], they require a large amount of images labelled with ground-truth disparities, in order to effectively learn how to tackle stereo. As obtaining ground-truth disparities, *i.e.* depths, for real images is particularly challenging and expensive, computer graphics has become a popular alternative to gather thousands of *synthetic* images endowed with depth labels for free [3]. Although highly realistic, these images can hardly encompass all the nuisances occurring in the real world, such as, *e.g.*, sensor noise, reflective surfaces and challenging illumination conditions. Thus, due to the *domain shift* between the training and testing environments [7], deep networks trained by computer-generated imagery suffer from a large loss in accuracy when deployed in the real world. A partial solution to this issue consists in fine-tuning the stereo network on few labelled samples from the real domain. Yet, to obtain such ground-truth labels,




Fig. 1. **Continual adaptation on real images.** We show the reference image of a stereo pair from DrivingStereo [9] (a) and the disparity maps computed by *MADNet* when trained on synthetic data only (b) or adapted online by either *MAD* (c) or *MAD++* (d).

costly active sensors (*e.g.*, LIDAR) and manual intervention or post-processing are required [8]. Even more importantly, despite fine-tuning by a few real-images may address the *synthetic-to-real* domain shift, it cannot take into account the countless diverse environmental conditions that a stereo network meant to be deployed *in-the-wild* may encounter, such as, in autonomous driving scenarios, urban and countryside roads, tunnels, varying weather and sudden changes of the surroundings.

In our vision, the only viable approach to practical end-to-end deep stereo deals with departing from the traditional *training-validation-testing* workflow towards a *continual adaptation* paradigm, so as to realize neural networks able to *adapt* promptly to new situations and environments. Of course, this novel paradigm cannot leverage standard supervised learning approaches since ground-truth labels would not be available for any new setting faced *in-the-wild*.

In this paper we extend our preliminary work on *continual adaptation* for deep stereo [10], which proposed the first-ever real-time, self-adapting, deep stereo network by relying on self-supervision obtained from the input pair of frames

- M. Poggi, F. Tosi, S. Mattoccia and L. Di Stefano are with the Department Computer Science and Engineering, University of Bologna, Italy.
- A. Tonioni is with Google Zurich.

via a photometric loss [11], [12], [13]. Given a sequence of stereo pairs, a straightforward continual adaptation scheme might be realized through the following steps: 1) output the prediction and compute the loss for the current pair of frames; 2) update the whole network by back-propagation; 3) move forward to the next pair of frames, with enriched knowledge encoded in the updated weights. However, due to the computational overhead associated with on-line back-propagation, most state-of-the-art stereo architectures would not operate in real-time under a *continual adaptation* paradigm mandating full update of the network via back-propagation. Thus, we designed a Modularly ADaptive Network (MADNet) architecture that is fast, accurate and features small memory requirements. Moreover, we developed a Modular ADaptation (MAD) algorithm that, in each on-line back-propagation step, enables to select and update only a portion of the whole *MADNet*, thereby vastly reducing the overhead required by online model updating and permitting prediction alongside self-adaptation in real-time without a large loss in accuracy compared to back-propagating errors into the full set of weights of network.

We extend and improve the *MADNet-MAD* framework by leveraging on *proxy supervision* obtained from traditional stereo algorithms [7], [14]. Indeed, although deep stereo networks are unmatched when trained and tested on similar domains, traditional methods, like [15], are largely *domain-agnostic*. In fact, they behave similarly and suffer from the same nuisances (e.g., low-textured regions, occlusions, repetitive patterns) with both synthetic and real images as well as across diverse environmental settings. This suggests that traditional algorithms may be exploited to supervise online deep stereo networks if matching errors, i.e. outliers, are properly detected and filtered out. We show that this strategy yields a much stronger adaptation ability and that this results consistently in a significant shrink of the performance gap between modular and full adaptation of *MADNet*, with the former often turning out even more beneficial than the latter. As a matter of fact, Figure 1 shows a qualitative comparison between the *MAD* formulation described in [10] and the novel approach proposed in this paper, referred to hereinafter as *MAD++*. Moreover, we dig into our *continual adaptation* paradigm and thoroughly explore its behaviour across very different datasets, showing, in particular, that our proposal is not affected by *catastrophic forgetting* while, on the contrary, continually adapting the stereo network is beneficial to performance in case of domain changes.

The main contributions of this unabridged paper on our work on continual adaptation for deep stereo can be summarized as follows.

- We introduce *MAD++* which ameliorates our *MAD* framework by leveraging on proxy supervision provided by traditional stereo algorithms. This novel approach outperforms the original proposal and yields often better results than updating all network weights while running twice faster.
- We extensively evaluate both *MAD* and *MAD++* on the raw KITTI dataset already considered in [10]. Besides, we include experiments on two additional datasets, i.e. DrivingStereo [9] and WeanHall [16], so

as to provide stronger evidence on the effectiveness of our proposed framework in a broader variety of target domains.

- We test our methods across datasets to highlight how continual learning is robust to domain shifts. We find no evidence of catastrophic forgetting in any experiments. On the contrary, we show that models adapted elsewhere feature better adaptation ability when facing new domains.

2 RELATED WORK

In this section, we briefly review the literature relevant to our work.

Machine learning for stereo. The first attempts to use machine learning for stereo matching dealt with estimating confidence measures [17], either by random forest [18], [19], [20], [21] or CNNs [22], [23], [24], [25], and often with the purpose of improving the final accuracy when combined with traditional algorithms. Regarding stereo algorithms, the first works proposed matching cost functions realized by patch-based CNNs [2], [26], [27] and allowed to achieve state-of-the-art performance by replacing handcrafted cost functions [28] within the SGM pipeline [15]. Later, Batsos *et al.* [29] combined traditional matching functions within a random forest framework to obtain better generalization across domains. Then, Shaked and Wolf [30] proposed to rely on deep learning for matching cost computation, disparity selection and confidence prediction, while Gidaris and Komodakis [31] focused on disparity refinement.

A true paradigm shift did occur with the first end-to-end stereo network, DispNetC, was proposed alongside large synthetic training datasets [3]. In [3] a custom correlation layer was designed to encode the similarities between pixels as features. Kendall *et al.* [4] designed GC-Net, switching towards 3D networks that build a cost volume by means of features concatenation. These two architectures started the development of two main families of networks, referred to as 2D and 3D, respectively. Proposals belonging to the former class use typically a single or multiple correlation layers [10], [32], [33], [34], [35], [36], [37], [38], while 3D networks build 4D volumes by means of concatenation [39], [40], [41], [42], [43], [44], [45], [46], [47], features difference [48] or group-wise correlations [49], both combined with active sensors such as LIDAR in [50]. Although most works focus on accuracy, others deploy lightweight architectures [10], [42], [48] aimed at real-time performance, sometimes combining stereo with semantic segmentation [51] or pursuing scene flow [52], [53]. Unfortunately, however, all end-to-end stereo networks are prone to domain shift, as performance decay dramatically when the model is run in environments different from those observed at training time, as shown in [7], [10], [14], [54], [55].

Self-supervision from photometric losses. View synthesis has been recently used to train depth estimation networks in a self-supervised manner by photometric losses [12], [56]. For monocular depth estimation, multiple images are deployed at training time in order to replace ground-truth labels by warping the different views, coming either from stereo pairs or image sequences, according to the predicted depth and minimizing the photometric error between

Algorithm 1 Full Adaptation (*FULL*)

```

1: Require: Stereo model  $\mathcal{N}$  parametrized by  $\Theta$ 
2:  $t = 0$ 
3: while not stop do
4:    $x_t \leftarrow \text{ReadFrames}(t)$ 
5:    $y_t \leftarrow \text{ForwardPass}(\mathcal{N}, \Theta_t, x_t)$ 
6:    $\mathcal{L}_t \leftarrow \text{Loss}(x_t, y_t)$ 
7:    $\Theta_{t+1} \leftarrow \text{UpdateWeights}(\mathcal{L}_t, \Theta_t)$ 
8:    $t \leftarrow t + 1$ 
9: end while

```

real and warped images [12], [56], [57], [58], [59]. Other recent works follow a similar approach for deep stereo matching [10], [11], [13], [60], [61], [62]. Unlike monocular ones, however, in stereo setups the input images used to compute the photometric loss are available at both training and testing time, which renders this self-supervised learning protocol amenable to *continual adaptation*.

Proxy-supervision from distilled labels. A further approach consists in sourcing *pseudo* ground-truth annotations, namely *proxy labels*, accurate enough to allow for effective supervision during training. The process to obtain these annotations is usually referred to as *distillation*. In the field of depth estimation, the work by Tonioni *et al.* [7], [14] was the first to use traditional stereo algorithms filtered out by means of confidence measures for offline adaptation of deep stereo networks. Pang *et al.* [54] used iterative optimization over proxy labels of the network itself obtained at higher resolutions. Recently, these approaches have been applied to monocular depth estimation, sourcing proxies either from traditional stereo algorithms [63], a teacher architecture [64] or the network itself trained in two stages [65].

3 CONTINUAL ADAPTATION

State-of-the-art deep stereo networks are severely challenged by *Out-of-Distribution* generalization, frequently exhibiting large accuracy drops when deployed across different environments. This issue is typically alleviated by fine-tuning the network on additional labelled samples from the target distribution. We argue that this is definitely unpractical as 1) it requires collecting data and tuning the model before real deployment for any target environment and 2) ground-truth labels need to be acquired together with images.

Self-supervision and *distillation* allow to circumvent the need for ground-truth labels, making offline fine-tuning the main obstacle towards seamless deployment in-the-wild. This can be addressed by moving from a traditional *train-validation-test* procedure to a *continual adaptation* paradigm, whereby the distinction between offline training and online testing is relaxed due to both being performed online and at once. In Algo 1 we provide a description of a continual adaptation process referred to as *Full Adaptation*. Given a stereo network \mathcal{N} parametrized by a set of weights Θ , at any given time frame t we read a new stereo pair x_t , made out of a left and right image (l_t, r_t), and predict a disparity map y_t for image l_t based on the current set of parameters θ_t . Then, a suitable loss function \mathcal{L}_t is computed from x_t, y_t

and used to update the network weights Θ before reading the next stereo pair x_{t+1} .

Since a train iteration is performed on-the-fly on each incoming stereo pair, the network always learns and -potentially- improves by gathering knowledge from the sensed environment. It is worth pointing out that this approach is different from standard practice dealing with updating a network on a batch of images to optimize it over a set of variegated samples that better approximate the entire training set and lead to more stable gradients. In our setting, however, the update steps occur on images that are acquired closely in time and are, therefore, very similar. Thus, the gradients provided by the individual samples are quite similar to the average gradient that would be computed in a batch. In the supplementary material we show that updating the network on a batch of samples does not provide significant benefits.

As shown in [10], our straightforward formulation, though intuitive and effective, introduces a non-negligible computational overhead that does increase the network latency dramatically. To address this drawback, we introduce a modular neural network architecture and a learning algorithm, which are designed to work in synergy to achieve effective continual adaptation with a limited computational overhead.

As pointed out, updating the whole network to achieve continual adaptation is time consuming and may hinder applicability to real-world applications calling for tight low-latency requirements. Due to the time required by back-propagation being proportional to the number of network layers to be traversed, we may speed-up the computation by having fewer layers, *i.e.* fewer weights, to update. Intuitively this is similar to accelerating forward inference by early-stopping the network processing in order to calculate only a subset of the total number of operations [42], [58]. Our work leverages on a similar intuition to speed up online back-propagation, *i.e.* the main computational overhead introduced by continual adaptation.

We start from the general encoder-decoder architecture illustrated in Figure 2, which can be thought of as an abstraction of the majority of deep networks proposed for stereo matching in the past years. The layers in network \mathcal{N} can be organized into non-overlapping, but inter-connected, modules $[\Theta^1, \dots, \Theta^p]$ according to any arbitrary grouping policy, *e.g.*, the resolution at which they operate, as shown in different colors in the figure. Then, we might think of performing back-propagation on only one module at a time to speed it up. However, standard network architectures provide only a prediction amenable to computing the loss at the very last layer. Therefore to back-propagate the training signal to each module we would still need to go through the deeper ones in the architecture. To overcome this limitation, we introduce shortcut back-propagation routes for each module. Our network predicts as many outputs as the modules, $[y^1, \dots, y^p]$, and includes at least one back-propagation route from each output to all the layers belonging to the associated module, *i.e.* (y^i, Θ^i) with $i \in [1, \dots, p]$. Hence, by computing a loss \mathcal{L}^i for each output, y^i we can directly back-propagate into the corresponding module Θ^i avoiding the remaining ones by means, for instance, of skip connections (depicted as arcs in Figure 2). The gradients computed with

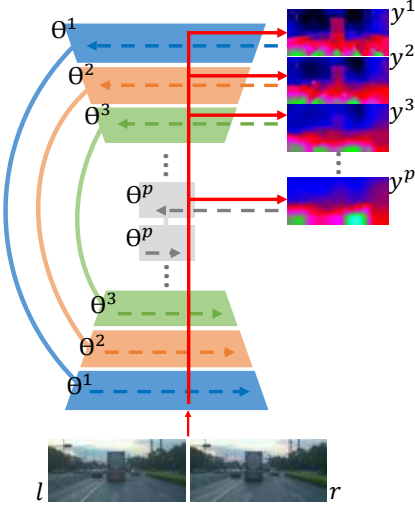


Fig. 2. **Generic design of a modular adaptive network.** The network \mathcal{N} is organized as a set of non-overlapping modules $[\Theta^1, \dots, \Theta^p]$ and is trained to estimate a set of corresponding outputs $[y^1, \dots, y^p]$. During adaptation, a full forward pass (red line) is performed to obtain the outputs, on which losses $[\mathcal{L}^1, \dots, \mathcal{L}^p]$ are computed. By selecting a single \mathcal{L}^i , only one of the back-propagation routes (dashed lines) is followed so to update a single module Θ^i .

this strategy are an approximation of the trues ones, but they do provide a good training signal, as it will be shown experimentally.

An example of such a design is detailed in subsection 3.3 and used in the experiments reported in this paper. In this architecture each y^i is a disparity prediction at a different spatial resolution, with y^1 denoting the highest resolution disparity map delivered as output, while each Θ^i includes all the network layers processing features at that resolution, *i.e.*, both in the encoder and the decoder. Due to all the layers in Θ^i being connected through at least a direct back-propagation path, we can approximate the gradients for all layers by back-propagating only through the direct connection and skipping all the other back-propagation routes. This paradigm *approximates* back-propagation into the whole network by updating layers through time, *e.g.* in p steps should the modules be sequentially updated, while providing a fast inference time, as required by many practical applications.

3.1 Modular ADaptation – MAD

To pursue the modular adaptation approach described in the previous section as effectively as possible, we have developed a selection strategy aimed at choosing the module Θ^i to be updated at each time step. Purposely, we have devised the reward/punishment algorithm outlined in Algo 2. At bootstrap (2), a histogram \mathcal{H} consisting of p bins (one per module) is initialized to zero. Then, at each time step t the disparity maps $[y^1, \dots, y^p]_t$ are predicted (6) and the corresponding losses $[\mathcal{L}^1, \dots, \mathcal{L}^p]_t$ computed (7). Then, we select a network module Θ^{ϕ_t} by sampling an index ϕ_t from the probability distribution associated with \mathcal{H} (8) and perform back-propagation into Θ^{ϕ_t} only (9). At this point, our network has been updated and it is ready to process the next stereo pair. Before moving on, we also update \mathcal{H}

Algorithm 2 Modular ADaptation (MAD, MAD++)

```

1: Require: Stereo model  $\mathcal{N}$  parametrized by  $[\Theta^1, \dots, \Theta^p]$ 
2:  $\mathcal{H} = [h^1, \dots, h^p] \leftarrow 0$ 
3:  $t = 0$ 
4: while not stop do
5:    $x_t \leftarrow \text{ReadFrames}(t)$ 
6:    $[y^1, \dots, y^p]_t \leftarrow \text{ForwardPass}(\mathcal{N}, \Theta_t, x_t)$ 
7:    $[\mathcal{L}^1, \dots, \mathcal{L}^p]_t \leftarrow \text{Loss}(x_t, [y^1, \dots, y^p]_t)$ 
8:    $\phi_t \leftarrow \text{Sample}(\text{softmax}(\mathcal{H}))$ 
9:    $\Theta_{t+1}^{\phi_t} \leftarrow \text{UpdateWeights}(\mathcal{L}_t^{\phi_t}, \Theta_t^{\phi_t})$ 
10:  if  $t == 0$  then
11:     $\tilde{\mathcal{L}}_{t-2}^1 \leftarrow \mathcal{L}_t^1$ ,
     $\mathcal{L}_{t-1}^1 \leftarrow \mathcal{L}_t^1$ 
12:  end if
13:   $\tilde{\mathcal{L}}_t^1 \leftarrow 2 \cdot \mathcal{L}_{t-1}^1 - \mathcal{L}_{t-2}^1$ 
14:   $\gamma \leftarrow \tilde{\mathcal{L}}_t^1 - \mathcal{L}_t^1$ 
15:   $\mathcal{H} \leftarrow \delta \cdot \mathcal{H}$ 
16:   $\mathcal{H}[\phi_{t-1}] \leftarrow \mathcal{H}[\phi_{t-1}] + \lambda \cdot \gamma$ 
17:   $\mathcal{L}_{t-2}^1 \leftarrow \mathcal{L}_{t-1}^1$ ,  $\mathcal{L}_{t-1}^1 \leftarrow \mathcal{L}_t^1$ ,  $\phi_{t-1} \leftarrow \phi_t$ 
18:   $t \leftarrow t + 1$ 
19: end while

```

in order to reward or punish the module updated in the previous time step, namely $\Theta^{\phi_{t-1}}$, depending on whether this has proven to be effective or not. To do so, we linearly extrapolate the expected value for the highest resolution loss (*i.e.*, that computed on the final output of the network), $\tilde{\mathcal{L}}^1$, from the previous ones at times $(t-1)$ and $(t-2)$ (13). Then, we compute the difference, γ , between the expected and computed losses (14). In case of a positive/negative difference we deem the update step on the module selected at time $(t-1)$ to have been effective/ineffective as at time t the highest resolution loss turns out smaller/larger than the value we would have expected had module $\Theta^{\phi_{t-1}}$ not been updated. Accordingly, we gradually apply a decay factor δ to all entries in \mathcal{H} (15), then we reward/punish module ϕ_{t-1} by adding a contribution γ into histogram bin $\mathcal{H}[\phi_{t-1}]$ (16). In our experiments, we set δ and λ to 0.99 and 0.01, respectively, and in the supplementary material we show how the sensitivity of the algorithm to these parameters is moderate.

In [10] we realized Algo 2 using as $\mathcal{L}^i, i \in [1, \dots, p]$, the self-supervised loss provided by the photometric error between the left image l_t and the right image \tilde{r}_t warped according to the estimated disparity y_t^i . In particular, according to a popular choice in literature, we compute this photometric error as

$$\mathcal{L}_t^i = \alpha \cdot \frac{1 - \text{SSIM}(l_t, \tilde{r}_t)}{2} + (1 - \alpha)|l_t - \tilde{r}_t| \quad (1)$$

with α set to 0.85 [12]. Thus the approach refereed to as MAD) in this paper performs real-time continual adaptation by deploying a popular self-supervised loss within Algo 2. Although fast and effective, MAD consists in diluting over time the network optimization process, thereby requiring more frames (*i.e.*, update steps) than the straightforward full adaptation approach (Algo 1) to acquire the knowledge needed to adapt a model to a novel environment. In the next section we describe how to leverage on a different kind of

loss which deploys a stronger source of supervision while still being amenable to continual adaptation. In section 4 we will show how this novel formulation can effectively accelerate a network optimization process distributed over time according to Algo 2 and reduce the performance gap with respect to Algo 1 dramatically.

3.2 Proxy-Supervised Modular Adaptation – *MAD++*

To speed-up the model adaptation process, we move toward a stronger form of supervision. In particular, we propose to rely on *proxy supervision* by leveraging on a reliable external source of disparities used as proxies for ground-truth labels. For instance, the use of active sensors, like LIDARs, have been proposed to supervise a depth prediction network [66]. Yet, a cheaper - and far more practical- source of proxy labels is described in recent works concerning both stereo [7], [14] and monocular [63] depth estimation. Accordingly, the noisy disparities computed by traditional stereo algorithms are filtered by a confidence estimator and deployed as proxy ground-truth labels to either adapt or train from scratch a depth prediction model. Since the procedure described in Algo 2 is agnostic to the actual loss function, we can extend our modular adaptation approach so as to rely on proxy supervision by simply specifying a suitable loss $\mathcal{L}^i, i \in [1, \dots, p]$. This novel formulation of Algo 2 will be referred to hereinafter as *MAD++*.

Given a generic stereo matching pipeline \mathcal{M} , we can obtain a noisy disparity map z_t by processing an input stereo pair $x_t = (l_t, r_t)$. However, as discussed in [7], [14], to effectively supervise a depth estimation network it is crucial to filter out most of the noisy disparities. This can be achieved by estimating a confidence map, c_t , encoding the reliability of each pixel in z_t [7], [14]. Then, supervision for any estimated y_t^i can be obtained from z_t by the following loss function:

$$\mathcal{L}_t^i = \eta_t \cdot |y_t^i - z_t| \quad (2)$$

where η_t denotes an indicator function that selects the measurements in z_t characterized by a sufficiently high confidence, *e.g.* a threshold operator applied to each pixel p according to the estimated confidence $c_t(p)$:

$$\eta_t(p) = \begin{cases} 1 & \text{if } c_t(p) \geq \varepsilon \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

Stereo Matching. Proxy labels can be obtained from an external stereo algorithm with a negligible overhead compared to the computational complexity of a deep stereo network. Indeed, a number of stereo cameras endowed with on-board processing hardware designed to deliver disparity maps at 50+ FPS are available nowadays [67], [68], [69], [70], [71], [72], [73]. As these cameras do not offload the stereo matching computation to the host device, they are amenable to distilling knowledge, *i.e.* proxy ground-truth disparities, to a deep stereo network so as to run continual adaptation without slowing down the process. Figure 3 illustrates how a stereo camera equipped with on-board processing can be deployed to support our real-time continual adaptation framework for deep stereo.




Fig. 3. **Deployment of on-camera disparity computation within *MAD++*.** During the forward pass (green arrows) the acquired frames are processed by *MADNet* to predict a disparity map as well as, in parallel, by a dedicated platform on-board the camera (*e.g.*, an FPGA) to compute proxy disparity labels. During the backward pass (red arrows), the network is updated so as to minimize the loss given by the discrepancy between the predicted and proxy disparities.

Due to existing hardware platforms relying mainly on the Semi-Global Matching (SGM) [15] or the basic Block Matching stereo algorithms, we will consider these two options in order to distill proxy disparities within *MAD++*.

Proxy Filtering. As traditional stereo matchers deliver noisy disparity maps, an effective criterion, $\eta(p)$ in Equation 3, is necessary to filter out outliers and provide reliable supervision to the continual adaptation process. In [7], [14] confidence estimation came from a large neural network, which, in our framework, would add a substantial computational overhead and prevent continual adaptation in real-time. Thus, in *MAD++* we pursue a different approach and rely on computationally efficient strategies geared toward the adopted stereo algorithm. Following [74], we deploy a simple left-right check for SGM and a combination of six confidence measures proposed in [75] for Block Matching. From now on, the two settings will be dubbed as SGM and WILD.

In Figure 4 we show qualitative examples of proxies obtained with the aforementioned pipelines on the DrivingStereo dataset [9]. We point out that both succeed in providing reliable proxy labels only, as required by the framework set forth in [7], [14], with density depending on the accuracy of the actual stereo matcher.

3.3 Modularly Adaptive Network – *MADNet*

In this section we detail the network architecture designed following the abstraction depicted in Figure 2 and deployed throughout the experimental evaluation presented in section 4. Besides the requirements set forth by modular adaptation (Algo 2), we have developed *Modularly Adaptive Network (MADNet)* to achieve a good balance between speed and accuracy. Indeed, *MADNet* has a smaller memory footprint and delivers disparity maps much more rapidly with a small loss in accuracy compared to complex networks such as [4], [33], [39]. To design (*MADNet*) we took inspiration from recent architectures for optical flow [76], [77] and conceived a novel light-weight model for stereo depth estimation.




Fig. 4. **Proxy labels by the considered stereo pipelines.** The first row depicts a reference image from DrivingStereo (a) alongside the available ground-truth disparities (b). The next two rows report the raw disparities and proxy labels (*i.e.* filtered disparities) obtained by the WILD (c),(d) and SGM (e),(f) pipelines.

Layer	Kernel	Stride	Dilate	Output Channels	Input/Output Resolution	Input
Encoder						
f_conv1_a	3	2	1	16	1/2	input
f_conv1_b	3	1	1	16	1/2	f_conv1_a
f_conv2_a	3	2	1	32	1/4	f_conv1_b
f_conv2_b	3	1	1	32	1/4	f_conv2_a
f_conv3_a	3	2	1	64	1/8	f_conv2_b
f_conv3_b	3	1	1	64	1/8	f_conv3_a
f_conv4_a	3	2	1	96	1/16	f_conv3_b
f_conv4_b	3	1	1	96	1/16	f_conv4_a
f_conv5_a	3	2	1	128	1/32	f_conv4_b
f_conv5_b	3	1	1	128	1/32	f_conv5_a
f_conv6_a	3	2	1	192	1/64	f_conv5_b
f_conv6_b	3	1	1	192	1/64	f_conv6_a
Decoder						
d_conv1	3	1	1	128	1	input
d_conv2	3	1	1	128	1	d_conv1
d_conv3	3	1	1	96	1	d_conv2
d_conv4	3	1	1	64	1	d_conv3
d_conv5	3	1	1	32	1	d_conv4
y	3	1	1	1	1	d_conv5
Refinement						
r_conv1	3	1	1	128	1	input
r_conv2	3	1	2	128	1	r_conv1
r_conv3	3	1	4	128	1	r_conv2
r_conv4	3	1	8	64	1	r_conv3
r_conv5	3	1	16	32	1	r_conv4
r_conv6	3	1	1	1	1	r_conv5

TABLE 1

MADNet architecture. We detail the layers of the pyramidal features extractor (Encoder), disparity estimators (Decoder) and refinement module (Refinement).

We pursue dense disparity regression by a pyramidal approach amenable to modularize the architecture and conducive to fast processing. Following Figure 2, our network estimates disparities y^i , with $i \in [1, \dots, 5]$, ranging from $\frac{1}{4}$ to $\frac{1}{64}$ resolution, respectively. As for the encoding section, we design two pyramidal features extractors with shared weights that process the left and right image through a cascade of blocks, as detailed in Table 1 (**Encoder**). Based on the lowest resolution feature maps, a correlation layer [3] computes raw matching costs between the left and right images. Then, we deploy a disparity decoder, detailed in Table 1 (**Decoder**), predicting the disparity prediction at the lowest resolution, *i.e.* y^5 .

The lowest resolution disparity map is up-sampled by bilinear interpolation to $2\times$ the resolution and used to warp the right features towards the left ones, with both feature maps then forwarded to a further correlation layer and

another disparity Decoder in order to estimate y^4 . Again, this map is up-sampled to the next resolution level and the same computation as in the previous level is carried in order to come up with disparity estimate y^3 . This scheme is repeated until estimate y^1 is reached. The warping mechanism is instrumental to use a small search range at any resolution level, *i.e.* $[-2, 2]$, that is equivalent, at the lowest scale, to a search range of 128, to which a further 2^i range is added at each up-sampling. The highest resolution disparity estimate, y^1 , is further processed by a refinement module [77], as detailed in Table 1 (**Refinement**). Finally, the refined y^1 is up-sampled from $\frac{1}{4}$ to full resolution by bilinear interpolation.

With reference to Figure 2, in *MADNet* a generic module Θ^i consists of a block from the encoding section together with the corresponding disparity decoder, with the arcs linking together the encoder and decoder within a module realized by the warping and correlation layers. Yet, due to y^1 being at quarter resolution, Θ^1 is slightly different: it includes the first two of the six encoding blocks alongside both the disparity decoder and the refinement network.

4 EXPERIMENTAL RESULTS

In this section, we wish to evaluate thoroughly the effectiveness of our continual adaptation framework. Purposely, we run a set of experiments on a variegated family of datasets.

4.1 Datasets

Here, we provide a description of the datasets used for the experiments.

FlyingThings3D. A collection of synthetic images, made out of about 22k training stereo pairs with dense ground-truth labels, part of the SceneFlow synthetic dataset [3]. This dataset has been used to pre-train *MADNet* before deployment on real images, according to the standard practice in recent deep stereo literature outlined in [10]. Specifically, we perform 1.2M training iterations using Adam Optimizer and a learning rate of 10^{-4} , halved after 400K steps and further every 200K until convergence. As loss function, we compute the L1 difference between disparity maps estimated at each resolution and downsampled ground-truth labels. The final loss is a weighted sum of the contributions from the different resolutions, with weights set to 0.005, 0.01, 0.02, 0.08, 0.32 from y^2 to y^6 according to [77].

KITTI 2015 train set. A collection 200 stereo pairs with sparse ground-truth maps, obtained from post-processed LIDAR measurements and 3D CAD objects [6]. This dataset has been used to fine-tune *MADNet* for 500K steps with learning rate 10^{-4} by computing the loss only on the full-resolution disparity map and using 0.001 as weight, as described in [10]. From now on, a *MADNet* model trained on FlyingThings3D and fine-tuned on KITTI 2015 with explicit supervision will be referred to as *MADNet-GT*.

Raw KITTI. A large dataset featuring 61 stereo sequences, for a total of about 43k pairs with different image resolution. We use a constant resolution of 320×1226 pixels by taking central crops of the original frames [10]. As depth ground-truths, we use filtered LIDAR measurements [8] converted to disparities through known calibration parameters [10]. According to the classification reported in

	LEAStereo [46]	CSPN [45]	GANet [43]	AcfNet [44]	HD ³ [37]	DeepPruner [38]	GWCNet [49]	Bi3D [78]	...	AANet [47]	DispNetC [3]	MADNet [10]	StereoNet [48]
D1-all	1.65	1.74	1.81	1.89	2.02	2.15	2.11	2.21		2.55	4.34	4.66	4.83
Time†	0.30	1.00	1.80	0.18	0.48	0.14	0.32	0.48		0.06	0.06	0.02	0.02

TABLE 2

Comparison between stereo architectures on the KITTI 2015 test set without adaptation. Detailed results available in the KITTI online leader-board. † Times (measured on different GPUs) taken from the online benchmark.

Starting Model	Adapt. Mode	Proxy src.	City (8027 frames)		Residential (28067 frames)		Campus (1149×2 frames)		Road (5674 frames)	
			D1-all(%)	EPE	D1-all(%)	EPE	D1-all(%)	EPE	D1-all(%)	EPE
MADNet	No	×	37.42	9.96	37.04	11.34	51.98	11.94	47.45	15.71
MADNet	FULL	×	3.35	1.11	2.38	0.94	10.62	1.78	2.72	1.04
MADNet	MAD	×	7.51 (+4.16)	1.63 (+0.52)	4.37 (+1.99)	1.32 (+0.37)	22.27 (+11.65)	3.66 (+1.88)	9.38 (+6.66)	2.04 (+1.00)
MADNet	FULL++	SGM [15]	3.51	1.12	2.27	0.94	9.69	1.63	3.18	1.05
MADNet	MAD++	SGM [15]	4.12 (+0.62)	1.18 (+0.06)	3.31 (+1.04)	1.04 (+0.10)	11.24 (+1.55)	1.76 (+0.13)	5.32 (+2.14)	1.22 (+0.17)
MADNet	FULL++	WILD [75]	5.11	1.23	2.82	0.99	11.79	1.89	4.28	1.11
MADNet	MAD++	WILD [75]	5.75 (+0.65)	1.30 (+0.07)	2.88 (+0.06)	0.99 (-0.01)	13.93 (+2.13)	2.04 (+0.15)	5.39 (+1.11)	1.24 (+0.13)

TABLE 3

Online adaptation within a single domain. Results on the *City*, *Residential*, *Campus* and *Road* sequences from KITTI [79].

the official website, we group sequences into four main categories: *Road*, *Residential*, *Campus* and *City*. Then, we concatenate the sequences belonging to the same category so as to obtain new, longer sequences of 5674, 28067, 1149×2¹ and 8027 frames for the above mentioned categories, respectively. In this manner, we simulate four macro environments characterised by different peculiarities, *i.e.* *City* and *Residential* mostly show roads surrounded by buildings, while *Road* images are collected while driving in highways and country roads, thus mainly depicting cars and vegetation. The dataset provides also raw LIDAR measurements.

DrivingStereo. A recent dataset [9] collecting about 170k stereo images grouped in 38 sequences with average resolution of 384×832 pixels. ground-truth is obtained by iterative filtering of LIDAR labels by means of a stereo CNN. We select three challenging sequences, namely *2018-08-17-09-45*, *2018-10-11-17-08* and *2018-10-15-11-43*, consisting of 1667, 1119 and 4950 frames, respectively. We rename the above mentioned sequences as *Rainy*, *Cloudy* and *Country*, respectively, according to their main peculiarities. We selected these sequences to 1) evaluate short-term adaptation (*i.e.*, after few hundreds frames) in challenging conditions (*e.g.*, rainy) and 2) assess the impact of prior continual adaptation (*e.g.*, on KITTI) when moving to a new environment.

WeanHall. An indoor dataset [16] which includes 6510 stereo pairs. As the working environment is very different from the autonomous driving scenarios addressed by previous datasets, we deem it worth evaluating performance also when continually adapting *MADNet* in so diverse settings. Unfortunately, no ground-truth disparities/depths are provided in WeanHall, neither are we aware of any other indoor stereo dataset providing sequences of real images alongside with the corresponding ground-truth labels. Therefore, we rely on the photometric error (Equation 1) to provide quantitative performance figures on WeanHall.

1. About *Campus*, it represents the most challenging environment characterized by low-textured buildings, yet the shortest sequence. For this reason, we loop twice over the sequence as in [10].

4.2 Experimental protocol

To assess the performance of our continual adaptation schemes, we run disparity prediction on the stereo pairs of a given sequence according to their order, *i.e.* as if they were acquired online in the field. On KITTI and DrivingStereo we measure the D1-all error rate as the percentage of pixels having absolute disparity error larger than 3 and relative error larger than 5%, as well as the End-Point-Error (EPE), whilst on WeanHall we measure the photometric error upon reprojection, as detailed in Equation 1. In case of baseline performance dealing with prediction without adaptation, we simply compute error metrics for each stereo frame and average them across the entire sequence. In case of adaptation, we process stereo pair x_t to predict a disparity map y_t and compute the error metrics on it, then we update the network according to either Algo 1 (FULL) or Algo 2 (*i.e.* MAD or MAD++). As a consequence, the impact of the continual adaptation step at time t will affect the error metrics from time $(t + 1)$. Akin to baseline performance, per-sequence metrics are computed by averaging across frames those dealing with the per-frame predictions y_t . During adaptation, we use a momentum optimizer, with a constant learning rate of 10^{-4} and a momentum of 0.9. Different optimizers or learning rate schedules resulted in minor fluctuations in our experiments. All experiments have been carried out using the source code and trained models available at github.com/CVLAB-Unibo/Real-time-self-adaptive-deep-stereo.

4.3 Comparison with state-of-the-art stereo networks

We start by comparing *MADNet* to the current state-of-the-art networks on the standard KITTI 2015 benchmark setting, *i.e.* without performing online adaptation. Table 2 collects performance figures retrieved from the online benchmark (to which we refer the reader for a complete overview) for both accuracy and runtime, although the latter is measured on GPUs belonging to different hardware generations. We can notice that the most accurate architectures trade precision for speed, whereas *MADNet* can operate in real-time with a modest increase of the error rate (about 3% with

Starting Model	Adapt. Mode	Proxy src.	D1-all(%)	EPE
<i>MADNet</i>	No	✗	38.84	11.68
<i>MADNet</i>	FULL	✗	2.43	0.95
<i>MADNet</i>	MAD	✗	4.09	1.19
			(+1.66)	(+0.24)
<i>MADNet</i>	FULL++	SGM [15]	2.28	0.95
<i>MADNet</i>	MAD++	SGM [15]	2.46	0.98
			(+0.18)	(+0.03)
<i>MADNet</i>	FULL++	WILD [75]	2.64	0.98
<i>MADNet</i>	MAD++	WILD [75]	2.44	0.96
			(-0.20)	(-0.02)

TABLE 4

Online adaptation across different domains. Results on the sequence *Campus* → *City* → *Residential* → *Road* (~43k frames, the whole KITTI dataset)




Fig. 5. **Adaptation speed on *Campus*.** *MAD++* adapts much faster than *MAD*, rapidly converging to the same error level as *FULL* and *FULL++* (blue and red solid lines, almost completely overlapped).

respect to the top-performing network). The focus of our work concern deploying *MADNet* alongside suitable online adaptation strategies to maximize accuracy while running in real-time in unknown environments.

4.4 Evaluation on KITTI

We begin our evaluation by studying different aspects of continual adaptation on the KITTI dataset. First, we address short-term adaptation within a domain by considering the KITTI sequences belonging to the same category independently. Then, we tackle a setup dealing with long-term adaptation across domains by concatenating together the sequences belonging to the different categories. We report results obtained by both the continual adaptation schemes discussed in section 3, namely full adaptation (Algo 1) and the more efficient - though approximated- modular adaptation approach (Algo 2). We also assess upon steering both adaptation schemes by either self-supervision of proxy supervision. As for the latter source of supervision, we consider proxy labels yielded by the previously described SGM and WILD pipelines. As in KITTI raw LIDAR measurements are available alongside stereo pairs, we also investigate on the effectiveness of this form of proxy supervision. Finally, we dig deeper into our framework by analysing the distribution of the update steps across the modules in modular adaptation (Algo 2 and investigating on the computational savings that may be achieved by employing slower adaptation rates.

Online adaptation *in-the-wild*. We address here the reference scenario concerning practical deployment *in-the-wild*: a stereo network pre-trained on synthetic data is run in a wholly unknown environment. In this setting, the experimental results reported in Table 3 and Table 4 deals with adaptation within a single domain and across different domains, respectively, both tables following the same organization. The first row highlights the baseline




Fig. 6. **Comparison between different proxy labels.** We show a reference image (a) from the *City* domain and proxy labels sourced by SGM (b) and WILD (c).

Starting Model	Adapt. Mode	D1-all(%)	EPE
<i>MADNet-Synthia</i>	FULL (baseline)	2.76	1.03
<i>MADNet-Synthia-FOL2A</i> [55]	FULL	2.60	1.01
<i>MADNet-Synthia</i>	FULL++ (ours)	2.23	0.95

TABLE 5

Comparison of online adaptation strategies across different domains. Results on the sequence *Campus* → *City* → *Residential* → *Road*. Comparison between *MADNet* trained on FlyingThings3D and finetuned on Synthia, with (FOL2A) or without meta-learning optimization.

performance measured by pre-training *MADNet* on FlyingThings3D and then running the model without any kind of online adaptation. Then, we report the results achieved by the two online adaptation schemes, *i.e.* FULL (Algo 1) and *MAD* (Algo 2), realized by self-supervision through the photometric error loss, as formulated in our previous work [10]. The next rows concern the novel formulation described in this paper, which exploits different sources of proxy supervision, namely SGM and WILD. Again, *FULL++* and *MAD++* refer to Algo 1 and Algo 2, respectively. For each source of supervision, we report in blue the difference in terms of D1-all and EPE between the *MAD* and *FULL* adaptation schemes, so as to highlight the gap between modular and full adaptation when relying on the same form of supervision. The difference between the two Tables deals with the KITTI sequences belonging to the same category having been processed individually in Table 3, after concatenation according to the order *Campus* → *City* → *Residential* → *Road*, in Table 4.

Firstly, the comparison between the baseline performance reported in the first row of both Tables and the figures in the successive ones highlights the dramatic error drops yielded by all considered methods and vouches for the utmost effectiveness of adapting online a stereo model pre-trained on synthetic data and run in a wholly unknown environment. As expected, between the two schemes, full adaptation (*FULL*/*FULL++*) consistently outperforms modular adaptation (*MAD*/*MAD++*) when steered by the same kind of supervision. Then, as for the former scheme, self-supervision (*FULL*) and proxy supervision (*FULL++*) seem, overall, rather equivalent options, one or the other performing slightly better in some experiments: *e.g.*, in Table 3 *FULL++* provides the lowest D1-all error in *City* and *Road*, *FULL++*(SGM) in *Residential* and *Campus*, whereas in Table 4 *FULL++*(SGM) yields smaller errors than *FULL* which, in turn, outperforms *FULL++*(WILD).

Starting Model	Adapt. Mode	Proxy src.	City (8027 frames)		Residential (28067 frames)		Campus (1149×2 frames)		Road (5674 frames)	
			D1-all(%)	EPE	D1-all(%)	EPE	D1-all(%)	EPE	D1-all(%)	EPE
MADNet-GT	No	✗	2.08	0.80	2.55	0.91	6.51	1.31	1.63	0.83
MADNet-GT	FULL	✗	1.60	0.89	1.87	0.86	4.70	1.31	1.14	0.81
MADNet-GT	MAD	✗	1.62	0.90	1.54	0.85	4.81	1.28	1.28	0.84
			(+0.02)	(+0.01)	(-0.33)	(-0.01)	(+0.10)	(-0.03)	(+0.14)	(+0.03)
MADNet-GT	FULL++	SGM [15]	1.59	0.92	1.49	0.84	3.50	1.05	1.24	0.85
MADNet-GT	MAD++	SGM [15]	1.59	0.92	1.52	0.86	3.73	1.14	1.35	0.86
			(-0.01)	(0.00)	(+0.03)	(+0.02)	(+0.22)	(+0.09)	(+0.11)	(+0.01)
MADNet-GT	FULL++	WILD [75]	1.58	0.90	1.50	0.85	4.19	1.13	1.24	0.83
MADNet-GT	MAD++	WILD [75]	1.57	0.91	1.79	0.87	4.26	1.23	1.30	0.85
			(-0.01)	(+0.01)	(+0.29)	(+0.02)	(+0.07)	(+0.10)	(+0.06)	(+0.02)

TABLE 6

Online adaptation within a single domain after fine-tuning. Results on the *City*, *Residential*, *Campus* and *Road* sequences from KITTI [79]. -GT denotes fine-tuning by ground-truths on the KITTI training set after pre-training on synthetic imagery.

Starting Model	Adapt. Mode	Proxy src.	D1-all(%)	EPE
MADNet-GT	No	✗	2.45	0.89
MADNet-GT	FULL	✗	1.83	0.88
MADNet-GT	MAD	✗	1.94	0.86
			(+0.11)	(-0.01)
MADNet-GT	FULL++	SGM [15]	1.46	0.85
MADNet-GT	MAD++	SGM [15]	1.76	0.89
			(+0.30)	(+0.03)
MADNet-GT	FULL++	WILD [75]	1.48	0.85
MADNet-GT	MAD++	WILD [75]	1.64	0.86
			(+0.16)	(+0.01)

TABLE 7

Online adaptation across different domains after fine-tuning. Results on the sequence *Campus* → *City* → *Residential* → *Road*. -GT denotes fine-tuning by ground-truths on the KITTI training set after pre-training on synthetic imagery.

However, when dealing with modular adaptation, proxy supervision (*MAD++*) consistently outperforms self-supervision (*MAD*), often by a very large margin. Indeed, unlike self-supervision, proxy supervision allows for reducing the performance gap between full and modular adaptation dramatically, as highlighted by the figures reported in blue in Table 3. We also point out that this is particularly evident in short sequences, such as *Campus* in Table 3, for which the gap is reduced from about 11.83% to less than 3%. Indeed, as shown in Figure 5, *MAD++* is much faster than *MAD* in reaching the same accuracy level as *FULL*/*FULL++*, which suggests performance differences measured by error metrics averaged along a sequence likely showing up more evidently in shorter ones. Similar to Table 3, proxy supervision turns out particularly beneficial to improve modular adaptation with respect to the formulation based on self-supervision in the long-term, cross-domain adaptation experiments considered in Table 4, with *MAD++* turning out almost as effective as *FULL++* while observing a substantial gap between *MAD* and *FULL*.

When it comes to reasoning on the different proxy sources adopted with either *FULL++* or *MAD++*, we can notice that, more often than not, SGM delivers a more effective supervision than *WILD*. In fact, the latter provides better performance only with *MAD++* in the Residential domain (Table 3) and, though rather slightly, in case of cross-domain adaptation (Table 4). We ascribe this to the much higher density of proxy labels featured by the SGM pipeline compared to *WILD*, as illustrated qualitatively in Figure 6.

Proxy supervision vs meta-learning. We compare proxy supervision to the meta-learning framework proposed in [55]. Purposely, according to the setting suggested in [55], we take the *MADNet* model trained on FlyingThings3D

and 1) fine-tune it on Synthia in a traditional, supervised manner or 2) using the First Order approximation variant of L2A (FOL2A). Thus, for this experiment, 1) provides the baseline with respect to which we evaluate the adaptation performance provided by FOL2A and our proposed approach dealing with proxy labels gathered by SGM. In both cases, we use a learning rate of 10^{-5} and a batch size of 16 for 10K steps. For meta-learning optimization we choose FOL2A [55], since all other strategies resulted unstable when coupled with *MADNet*, and use 3 consecutive frames. Table 5 collects the outcome of this experiment. We can notice how using L2A during pre-training can improve the performance with respect to the baseline. However, our proposal turns out significantly more effective, *i.e.* proxy supervision is more effective than L2A pre-training. A thorough comparison dealing with all the individual sequences is reported in the supplementary material.

Online adaptation after fine-tuning. As proposed in [10], we also investigate on the effectiveness of the different on-line adaptation schemes in case the pre-trained model may be fine-tuned by real stereo pairs with ground-truths before running inference. Indeed, this is the case of some research datasets like KITTI.

Thus, in Table 6 and Table 7 we report the results dealing with adaptation on each of the four KITTI domains and across domains, the only difference with respect to Table 3, Table 4 being that now *MADNet* has been pre-trained on FlyingThings3D and then fine-tuned on the KITTI 2015 training set before undergoing online adaptation by the considered schemes and forms of supervision.

The first row in Table 6 and Table 7 show the baseline performance yielded by running the pre-trained and fine-tuned model with online adaptation turned off. As now the network has been fine-tuned by samples endowed with ground-truth disparities, the baseline model performs considerably better than in Table 3 and Table 4. Nevertheless, both full adaptation (*FULL*/*FULL++*) as well as modular adaptation (*MAD*/*MAD++*) allow to further improve over this strong baseline in all the considered experiments, typically yielding substantial relative performance gains. For instance, in Table 6, *FULL++*(SGM) and *MAD++*(SGM) can provide a relative D1-all error reduction of about 46% and 42% on Campus (the shortest sequence), whereas *FULL* can decrease such error by about 30% in Road and *MAD++*(WILD) by about 24% in City. Similarly, in Table 7, the relative D1-all error reduction ranges from about 20% (*MAD*) to as much as 40% (*FULL++*(SGM)).

Starting Model	Adapt. Mode	Proxy src.	City (8027 frames)		Residential (28067 frames)		Campus (1149×2 frames)		Road (5674 frames)	
			D1-all(%)	EPE	D1-all(%)	EPE	D1-all(%)	EPE	D1-all(%)	EPE
MADNet	No	✗	37.42	9.96	37.04	11.34	51.98	11.94	47.45	15.71
MADNet	FULL++	LIDAR	3.66	0.97	2.94	0.89	9.10	1.54	3.24	0.93
MADNet	MAD++	LIDAR	4.63	1.12	3.99	1.06	19.33	2.32	4.74	1.12
			(+0.97)	(+0.15)	(+1.05)	(+0.18)	(+10.23)	(+0.78)	(+1.50)	(+0.20)
MADNet-GT	No	✗	2.08	0.80	2.55	0.91	6.51	1.31	1.63	0.83
MADNet-GT	FULL++	LIDAR	2.00	0.67	2.17	0.75	4.29	0.96	1.59	0.66
MADNet-GT	MAD++	LIDAR	3.16	0.88	2.86	0.92	4.96	1.20	1.89	0.79
			(+1.15)	(+0.21)	(+0.69)	(+0.17)	(+0.67)	(+0.23)	(+0.30)	(+0.14)

TABLE 8

Online adaptation within a single domain with proxy supervision from raw LIDAR. Results on the *City*, *Residential*, *Campus* and *Road* sequences from KITTI [79]. -GT denotes fine-tuning by ground-truths on the KITTI training set after pre-training on synthetic imagery.

Starting Model	Adapt. Mode	Proxy src.	D1-all(%)	EPE
MADNet	No	✗	38.84	11.68
MADNet	FULL++	LIDAR	2.87	0.87
MADNet	MAD++	LIDAR	3.86	1.02
			(+0.99)	(+0.15)
MADNet-GT	No	✗	2.45	0.89
MADNet-GT	FULL++	LIDAR	2.06	0.73
MADNet-GT	MAD++	LIDAR	2.86	0.89
			(+0.81)	(+0.16)

TABLE 9

Online adaptation across different domains with proxy supervision from raw LIDAR. Results on the sequence *Campus* → *City* → *Residential* → *Road*. -GT denotes fine-tuning by ground-truths on the KITTI training set after pre-training on synthetic imagery.




Fig. 7. **Raw LIDAR for proxy supervision.** We show a reference image (a) from the *City* domain alongside proxy labels sourced by raw Lidar (b), with the latter exhibiting wrong measurements at depth boundaries.

Besides, due to the base model undergoing adaptation being stronger, full and modular adaptation tend to exhibit a much smaller gap when driven by the same form of supervision (figures in blue). Moreover, unlike the previous experiment, in most cases proxy supervision provides better performance than self-supervision non only with modular adaptation (*MAD++* vs. *MAD*) but also with full adaptation (*FULL++* vs *FULL*).

As for the two kinds of proxy labels, the WILD pipeline seems now competitive with respect to SGM, as it can provide better or equivalent performance also in case of full adaptation (*City* and *Road* in Table 6) and turns out generally more effective when the model undergoes modular adaptation, in particular in the long-term, cross-domain experiment (Table 7). We would be led to ascribe this finding to the fact that, although fewer in number, the proxies extracted by WILD are more accurate and thus more amenable to refine the already good disparities predicted by a strong base model fine-tuned by real images equipped with ground-truths. Conversely, as observed in the previous experiment, denser proxies seem instrumental to break down the gross errors spread throughout the image delivered by a baseline prone to the synthetic-to-real domain shift.

Proxy supervision by LIDAR. We also inquire about the effectiveness of continual adaptation in case proxy supervision may be obtained from raw measurements provided by a LIDAR sensor, as it is the case of the KITTI dataset. Akin to previous experiments, we consider both a baseline *MADNet* pre-trained on FlyingThings3D as well as a model further fine-tuned by labelled stereo pairs from the KITTI 2015 training set (*MADNet-GT*). Table 8 and Table 9 collect the results dealing with online adaptation on each of the four KITTI domains and across them, respectively. We can notice a trend similar to previous experiments relatively

to several key findings. Indeed, in the *in-the-wild* scenario, online adaptation does matter a lot as in both Tables we observe a dramatic reduction of errors compared to the baseline model pre-trained on synthetic imagery. Besides, *FULL++* consistently outperform *MAD++*, the margin turning out generally higher than with proxy supervision by SGM and WILD (Table 3 and Table 4). In case stereo pairs with ground-truth are available to fine-tune the pre-trained model, online adaptation by LIDAR proxies and full adaptation (*FULL++*) is still beneficial, whilst modular adaptation (*MAD++*) tend to perform worse than the baseline. Hence, we are lead to conclude that modular adaptation (Algo 2) with supervision by LIDAR is less effective than in case the proxy labels are gathered by the SGM and WILD pipeline.

Moreover, by comparing Table 8 and Table 9 to Tables 1-4, we can notice that, in general, both *MADNet* and *MADNet-GT* exhibit a lower EPE when steered by LIDAR supervision thanks to the higher depth resolution provided by the sensor. Yet, the D1-all error tends to be higher. We ascribe this to raw LIDAR measurements featuring a larger number of outliers compared to SGM and WILD because both the latter pipelines include a filtering step to disregard low-confidence disparities. In particular, it is worth observing how raw LIDAR measurements often yield gross errors near depth discontinuities, as illustrated in Figure 7.

Complementing low-density supervision. As shown in both Figure 6 and Figure 7, we can notice how different the density of the labels is by changing the source, with WILD, in particular, resulting the method providing sparser supervision. To obtain denser supervision, we might combine WILD with self-supervision [10] or with LIDAR measurements, that often results complementary to it. We show this experiment in the supplementary material.

Distribution of update steps in modular adaptation. We dig deeper into the behavior of Algo 2 in order to study the distribution of the update steps across the modules




Fig. 8. **Percentage of update steps across MADNet’s modules over time.** Update frequency (%) for $\Theta^i, i \in [2, \dots, 6]$ MAD (top) and MAD++ with SGM proxy labels (bottom) as function of the processed frames. Experiment dealing with continual adaptation from synthetic pre-training on *Campus* \rightarrow *City* \rightarrow *Residential* \rightarrow *Road*.

of MADNet. Figure 8 plots the distribution of update steps (expressed as a percentage of the processed frames) occurring for each module during continual adaptation *in-the-wild* across the four KITTI domains (same setting as in Table 4): in the top and bottom figures we show the results achieved by running MAD and MAD++ with SGM labels, respectively. We can notice how, with the original MAD strategy based on self-supervision by the photometric error, most of the steps concern Θ_5 , which shows consistently a higher percentage with respect to the other modules, while Θ_6 results the module less frequently updated. Conversely, MAD++ steered by SGM yields a more balanced distribution of updates across the modules, with almost no gap between the percentages observable after 20K frames. We argue that the more even distribution of the update steps featured by MAD++ leads to a better approximation of continual adaptation by back-propagation into the whole network (Algo 1), which, indeed, updates all modules in each step. A round-robin strategy would allow for uniform sampling of the modules alike, but such a fixed sampling schedule proved to be less effective than MAD [10]. MAD++, seems to provide a balanced update distribution without being bound to a fixed schedule, which, in fact, turns out beneficial to the continual adaptation process. As such, Figure 8 may help explaining why, as observed in Table 3 and Table 4, MAD++ is more effective than MAD in filling the performance gap between modular adaptation and full adaptation.

Experiments with different adaptation rates. To pursue further computational efficiency in continual adaptation, we investigate on updating the model on only a subset of the incoming frames, so as to reduce the computational overhead due to continuously gathering knowledge about the sensed environment. Purposely, we consider a more general adaptation approach, whereby the model is updated every K frames, and evaluate the speed-performance trade-off yielded by Algo 1 and Algo 2 while varying K . Figure 9 plots the D1-all error as a function of the frame-rate² achieved with $K = 1, 2, 5, 10$ by considering self-supervision (FULL and MAD) and proxy supervision




Fig. 9. **Performance vs. speed with different adaptation rates.** FPS vs. D1-all measurements for Algo 1 (blue) and Algo 2 (red) when adapting every $K=1, 2, 5, 10$ frames. Solid lines represent linear interpolations between measurements. Results dealing with continual adaptation from synthetic pre-training on *Campus* \rightarrow *City* \rightarrow *Residential* \rightarrow *Road*.

(FULL++ and MAD++ with SGM proxies).

Performing adaptation at every frame ($K=1$) allows MADNet to run at about 14 FPS with FULL and FULL++, while the network can achieve more than 25 FPS by MAD and MAD++. By increasing the interval between subsequent adaptations, we observe a corresponding increase of both speed and disparity prediction errors for all the considered methods.

We can appreciate the better performance-speed trade-off provided by MAD++ by fixing a target frame-rate and comparing the accuracy yielded by the different methods, or vice-versa. For instance, considering 30 FPS as the target speed requirement, we can notice that this is met by Algo 2 and Algo 1 with $K=2$ and $k=10$, respectively. In these settings, MAD++ achieves the lowest D1-all error, with FULL and FULL++ running faster but yielding higher errors and the original MAD formulation resulting both slower and less accurate than both variants of Algo 1. Should the application demand a lower real-time requirement, *i.e.* 25 FPS, MAD++ would turn out again the most accurate method to meet the given target speed. On the other hand, we may compare methods achieving close D1-all scores, *e.g.* MAD++ and FULL++, both with $K=1$, and highlight how MAD++ can run much faster with equivalent accuracy. Likewise, by considering Algo 1 with $K=5$ and Algo 2 with $K=2$, we can observe how MAD would turn out significantly less accurate than FULL and FULL++, whilst MAD++ can provide equivalent accuracy and faster speed.

Thus, we are lead to conclude that MAD++ achieves a more favourable performance-speed trade-off with respect to the other continual adaptation methods, as also suggested in Figure 9 by the trend of the interpolating curves (the lower the better).

4.5 Evaluation on DrivingStereo

Here we move to a different dataset dealing with autonomous driving scenarios in order to assess the performance of continual adaptation in diverse weather conditions. Purposely, we select three sequences from the DrivingStereo dataset [9]: *Rainy* and *Dusky* are short videos featuring less than 2k frames, while *Cloudy* counts about 5k frames. Moreover, *Rainy* depicts imagery acquired in presence of rain and wet road surface, making it particularly hard for traditional stereo algorithms too. These three sequences have been selected to study: 1) how continual

2. Ideal performance measured on a 1080Ti GPU.

Starting Model	Adapt. Mode	Proxy src.	Rainy (1667 frames)		Dusky (1119 frames)		Cloudy (4950 frames)	
			D1-all(%)	EPE	D1-all(%)	EPE	D1-all(%)	EPE
<i>MADNet</i>	No	✗	31.40	4.46	39.04	5.77	25.37	3.13
LEAStereo [46]	No	✗	22.14	3.90	22.99	3.45	8.64	1.61
<i>MADNet</i>	FULL	✗	19.64	2.71	20.40	2.48	9.54	1.57
<i>MADNet</i>	MAD	✗	26.11 (+6.47)	3.67 (+0.97)	33.23 (+12.84)	5.43 (+2.96)	16.03 (+6.49)	2.83 (+1.25)
<i>MADNet</i> -K-FULL	FULL	✗	15.75	2.46	9.21	1.43	8.09	1.48
<i>MADNet</i> -K-MAD	MAD	✗	12.22 (-3.53)	1.82 (-0.63)	9.81 (+0.60)	1.43 (0.00)	7.32 (-0.77)	1.41 (-0.06)
<i>MADNet</i>	FULL++	SGM [15]	17.28	2.62	12.86	1.73	6.63	1.33
<i>MADNet</i>	MAD++	SGM [15]	17.70 (+0.42)	2.48 (-0.15)	13.61 (+0.76)	1.84 (+0.11)	7.76 (+1.13)	1.45 (+0.12)
<i>MADNet</i> -K-FULL++	FULL++	SGM [15]	12.97	2.50	6.99	1.47	6.61	1.66
<i>MADNet</i> -K-MAD++	MAD++	SGM [15]	12.65 (-0.33)	2.32 (-0.18)	5.93 (-1.06)	1.40 (-0.07)	6.26 (-0.35)	1.72 (+0.06)
SGM [15]	✗	✗	17.88 (4.59) (66.90% density)	6.55 (1.25)	11.83 (4.49) (77.35% density)	2.60 (1.01)	7.31 (2.10) (79.53% density)	2.33 (0.86)
<i>MADNet</i>	FULL++	WILD [75]	17.93	2.60	19.04	2.32	8.31	1.49
<i>MADNet</i>	MAD++	WILD [75]	17.71 (-0.22)	2.40 (-0.20)	19.24 (+0.20)	2.48 (+0.16)	8.38 (+0.07)	1.52 (+0.03)
<i>MADNet</i> -K-FULL++	FULL++	WILD [75]	14.01	2.24	7.94	1.30	6.21	1.30
<i>MADNet</i> -K-MAD++	MAD++	WILD [75]	13.70 (-0.31)	2.09 (-0.15)	7.50 (-0.44)	1.27 (-0.03)	5.96 (-0.25)	1.29 (-0.01)
WILD [75]	✗	✗	36.55 (2.45) (23.37% density)	17.85 (1.03)	33.55 (1.92) (25.12% density)	13.83 (0.87)	22.89 (1.13) (28.21% density)	9.77 (0.83)

TABLE 10

Online adaptation on DrivingStereo. Results on the *Rainy*, *Dusky* and *Cloudy* sequences. -K denotes prior adaption on KITTI (*Campus* → *City* → *Residential* → *Road*) before further adaptation on DrivingStereo.

adaptation performs in presence of challenging weather conditions (e.g., rain), 2) how prior continual adaptation does affect performance when facing a new environment. Short sequences allow to better investigate on the latter issue, as long ones would hide the effects of prior adaptation on the initial frames due to performance figures being averaged across all frames. As for experiments concerning prior continual adaptation, we use the *MADNet* instance trained on synthetic images and adapted on *Campus* → *City* → *Residential* → *Road*, dubbed as *MADNet*-K-A, with A being the specific adaptation mode adopted.

Keep adapting! Table 10 collects the experimental results on DrivingStereo. In the first row we report the baseline performance achieved by running *MADNet* after pre-training on synthetic imagery and without any further adaptation, followed by the results achieved by the current top-1 method on the KITTI online benchmark pre-trained on the same synthetic data, i.e. LEAStereo [46], that highlights how current state-of-the-art methods are prone to the domain-shift issue. The three successive sub-tables, i.e. rows 2-7, 8-16 and 17-25, collect results dealing with continual adaptation realized through self-supervision by the photometric error loss, proxy supervision by SGM and proxy supervision by WILD, respectively. For the two sources of proxy labels we also show the errors yielded by the pipeline providing the supervision, both before and after the outlier removal step (rows 14-15 and 23-24), alongside the resulting label density (rows 16 and 25). We do not consider supervision by LIDAR as such measurements are not provided in DrivingStereo.

Considering self-supervision, we first adapt on a sequence by FULL and MAD (rows 2 and 3). As observed on KITTI, both strategies are effective, though a large margin between the two exists. In rows 5 and 6 we report the results achieved by FULL (*MADNet*-K-FULL) and MAD (*MADNet*-K-MAD) with model instances obtained through

prior continual adaptation on KITTI (i.e., the models saved after the experiments in Table 4). Hence, by keep adapting to the current domain we achieve much better results, as highlighted by the comparison between rows 5 and 2 as well as 6 and 3. In particular, prior adaptation is particularly effective in the short-term, as shown by the *Dusky* sequence where the error rate is roughly halved with FULL and brought down by about 70% with MAD, while the benefit tends to be smaller in longer sequences, such as *Cloudy*, in particular with FULL. These experimental findings show that continual adaptation realized through both FULL and MAD is always beneficial. Interestingly, we point out how prior adaptation turns out much more effective with MAD, allowing it to even outperform FULL in several sequences (as shown by row 7).

Moving to the experiments dealing with proxy supervision by SGM, we can notice, in general, a substantial performance improvement. In particular, when adapting *MADNet* starting from pre-training on synthetic imagery, SGM proxies (rows 8 and 9) consistently outperform self-supervision by the photometric error loss (rows 2 and 3) and allow for breaking down the margin between Algo 1 and Algo 2 (row 10 vs row 4). By keep adapting from KITTI (rows 11 and 12), results get much better on all sequences and SGM neatly outperforms adoption by the photometric loss in the same training protocol (rows 5 and 6). In this case, moreover, MAD++ outperforms FULL++ on all sequences (as highlighted in row 13).

Finally, by analyzing the experimental results obtained by WILD proxies, we can observe a trend similar to that already discussed for SGM. In particular, keep adapting on DrivingStereo following prior adaptation on KITTI is highly beneficial and MAD++ turns out more effective than FULL++ in this setting (row 22).

Proxy labels comparison. We use DrivingStereo se-

Starting Model	Adapt. Mode	Proxy src.	Photometric Error
<i>MADNet</i>	No	\times	0.124
<i>MADNet</i>	FULL	\times	0.084
<i>MADNet</i>	MAD	\times	0.094 (+0.010)
<i>MADNet-K-FULL</i>	FULL	\times	0.080
<i>MADNet-K-MAD</i>	MAD	\times	0.083 (+0.003)
<i>MADNet</i>	FULL++	SGM [15]	0.086
<i>MADNet</i>	MAD++	SGM [15]	0.088 (+0.002)
<i>MADNet-K-FULL++</i>	FULL++	SGM [15]	0.082
<i>MADNet-K-MAD++</i>	MAD++	SGM [15]	0.082 (0.000)
<i>MADNet</i>	FULL++	WILD [75]	0.087
<i>MADNet</i>	MAD++	WILD [75]	0.088 (+0.001)
<i>MADNet-K-FULL++</i>	FULL++	WILD [75]	0.082
<i>MADNet-K-MAD++</i>	MAD++	WILD [75]	0.082 (0.000)

TABLE 11

Online adaptation on WeanHall. Results on the *Rainy*, *Dusky* and *Cloudy* sequences. -K denotes prior adaption on KITTI (*Campus* \rightarrow *City* \rightarrow *Residential* \rightarrow *Road*) before further adaptation on WeanHall.

quences also to further evaluate proxy labels. Rows 14-16 and 23-25 report the accuracy and density of the labels produced by SGM and WILD. In particular, we show first the error rates achieved without filtering out the outliers, then, in brackets, those computed only on the final labels used to provide the supervision and finally the density of the filtered labels. We can notice that the SGM pipeline extracts much more labels at the cost of a lower accuracy. This confirms the finding already discussed in subsection 4.4 about the different traits of the two pipelines deployed to attain proxy supervision.

Eventually, we point out that *MADNet* models adapted by high-confidence (*i.e.* filtered) SGM and WILD proxies tend to consistently outperform the dense (*i.e.*, without filtering) pipeline providing the supervision, more often than not by a large margin. Indeed, only with *Dusky* (the shortest sequence) the SGM pipeline (row 14) slightly outperforms *MADNet* adapted from synthetic pre-training (rows 8 and 9), though when keep adapting the model following prior adaptation on KITTI the latter turns out about 40-50% more accurate (rows 11 and 12).

However, the accuracy of the filtered points remains higher, yet much sparser and thus less amenable for practical applications. Although approaches to restore density exists [74] and might represent an alternative to our framework, in the supplementary material we show that, in practise, using the sparse proxies to adapt *MADNet* often leads to better results.

4.6 Evaluation on WeanHall

Finally, we make a further step into evaluating the effectiveness of continual adaptation across very different domains by moving to an indoor environment, like that featured by the WeanHall dataset. As previously pointed out, given the absence of ground-truth labels for this set of images, we measure the photometric error according to Equation 1.

Table 11 collects the outcome of the experiments carried out on WeanHall following the same protocol adopted in subsection 4.5 for DrivingStereo. Although the margins in terms of photometric error are smaller compared to the D1-all and EPE metrics, we observe findings consistent to

Adapt. Mode	Proxy src.	Before WeanHall		After WeanHall		After WeanHall (adapting)	
		D1-all(%)	EPE	D1-all(%)	EPE	D1-all(%)	EPE
FULL	\times	2.43	0.95	2.03	0.93	1.49	0.87
MAD	\times	4.09	1.19	1.34	0.77	2.43	0.95
		(+1.66)	(+0.24)	(-0.69)	(-0.16)	(+0.94)	(+0.08)
FULL++	SGM	2.28	0.95	2.45	0.94	1.51	0.88
MAD++	SGM	2.46	0.98	2.52	0.95	1.94	0.91
		(+0.18)	(+0.03)	(+0.07)	(+0.01)	(+0.43)	(+0.03)

TABLE 12

Experimental results concerning catastrophic forgetting. Results on the *City* \rightarrow *Residential* \rightarrow *Campus* \rightarrow *Road* sequences from KITTI. From left to right, results from first adaptation on KITTI (Table 4) and from a second run on KITTI after having adapted on WeanHall, respectively without and with adaptation enabled during the second run. Starting model: *MADNet*.

previous experiments. Indeed, continual adaptation yields a significant improvement with respect to the baseline, with the adoption of proxy labels rather than self-supervision shrinking the gap between Algo 1 and Algo 2, and *MAD++* always outperforming *MAD*. More importantly, starting from models previously adapted on KITTI always lead to better results, despite the successive adaptation being conducted in a totally different environment (*i.e.* indoor vs. outdoor).

4.7 Catastrophic forgetting experiment

Finally, to highlight how our framework is not affected by catastrophic forgetting during continuous adaptation across very different domains, we perform a second adaptation run on the entire KITTI raw dataset, after having adapted the model first on KITTI and then on WeanHall. Should catastrophic forgetting occur, the accuracy of *MADNet* during this second run would dramatically drop. Table 12 collects the outcome of this experiment, showing the results achieved by *MADNet* with different adaptation strategies. In particular, we report, from left to right, the results concerning 1) the first run over KITTI (*i.e.*, those shown in Table 4, 2) the second run over KITTI after having adaptation on WeanHall, without further adaption on KITTI and 3) the second run over KITTI while keeping adapting. For FULL++/MAD++, we show the results achieved when using SGM proxy labels. We can notice how, even without adapting a second time to KITTI, the accuracy is comparable to that achieved during the first run, highlighting that catastrophic forgetting has not occurred, *i.e.* upon adaptation on WeanHall the network has not forgotten how to predict disparities on KITTI. Indeed, with FULL/MAD the results are improved, while with FULL++/MAD++ we observe a negligible drop. By turning adaptation on in the second run on KITTI we observe how performance can further improve, except for MAD.

5 QUALITATIVE RESULTS

To conclude, we refer the reader to a video available at <https://www.youtube.com/watch?v=YnPGbQE2dLQ> which shows the disparity maps and associated error curves for the methods and datasets considered throughout the paper. Starting from KITTI, we point out that, after 20-30 frames of instability due to the domain shift (*i.e.*, less than 3 seconds in the video), *MAD++* rapidly adapts to




Fig. 10. **Qualitative results for different continual adaptation strategies.** We show, across time, the reference image of a stereo pair from DrivingStereo *Dusky* sequence and different adaptation strategies among those reported in Table 10 (MAD++ uses SGM labels).

the new environment, outperforming *MAD*. When moving across KITTI sequences, we observe steps in the error curve upon scene changes, with *MAD++* consistently yielding better performance, *i.e.* faster adaptation to the new scene and smaller error. When moving to DrivingStereo and WeanHall, we point out how keeping adapting from KITTI (*K-MAD* and *K-MAD++*) turns out more effective than starting the adaptation process from scratch (*MAD* and *MAD++*). On DrivingStereo, where the adopted metric is the disparity error with respect to the ground-truth, we observe how *K-MAD++* provides better performance than *K-MAD*. A qualitative example is shown in Figure 10. On WeanHall, *K-MAD++* behaves quite equivalently to *K-MAD*, although it is worth pointing out that, due to the lack of ground-truth disparities, the metric adopted to assess performance is exactly the photometric error minimized by the latter to pursue continual adaptation. A qualitative example is shown in Figure 11.

6 CONCLUSION

We have presented a novel continual adaptation paradigm for deep stereo networks conceived to deal with challenging and ever-changing environments. By coupling a Modularly Adaptive Network with a Modular ADaptation strategy leveraging on either proxy labels sourced from traditional algorithms or self-supervision via the photometric error, our framework realizes the first-ever real-time and self-adaptive deep stereo network. Experimental results on a variety of datasets support the effectiveness of our proposal, highlighting in particular how proxy supervision is more

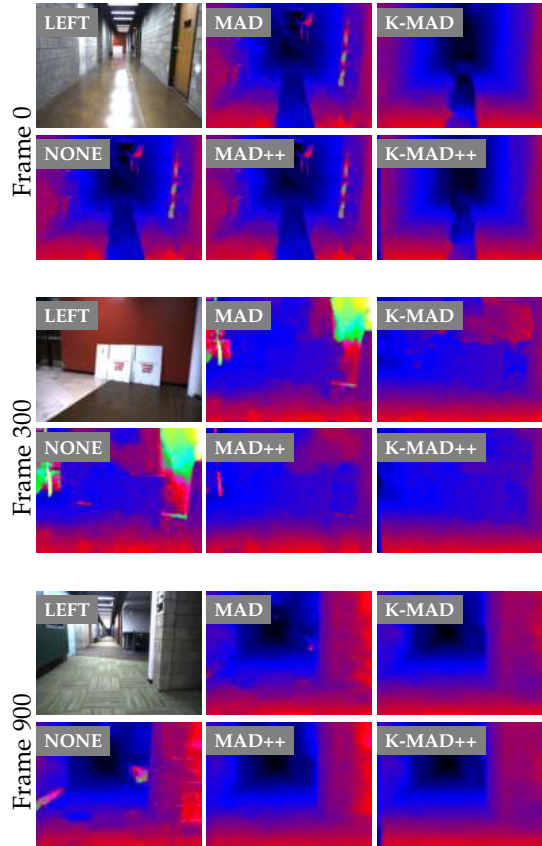


Fig. 11. **Qualitative results for different continual adaptation strategies.** We show, across time, the reference image of a stereo pair from WeanHall sequence and different adaptation strategies among those reported in Table 11 (MAD++ uses SGM labels).

beneficial than self-supervision and that continual adaptation holds the potential to address the unavoidable domain shifts that would occur when deploying deep stereo in many practical applications. The experimental findings also provide evidence on the ability of our paradigm to learn knowledge that can transfer well across domains, avoiding, in particular, catastrophic forgetting.

Acknowledgments. We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan Xp GPU used for this research.

REFERENCES

- [1] D. Scharstein and R. Szeliski, “A taxonomy and evaluation of dense two-frame stereo correspondence algorithms,” *International journal of computer vision*, vol. 47, no. 1-3, pp. 7–42, 2002.
- [2] J. Zbontar and Y. LeCun, “Stereo matching by training a convolutional neural network to compare image patches,” *Journal of Machine Learning Research*, vol. 17, no. 1-32, p. 2, 2016.
- [3] N. Mayer, E. Ilg, P. Hausser, P. Fischer, D. Cremers, A. Dosovitskiy, and T. Brox, “A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [4] A. Kendall, H. Martirosyan, S. Dasgupta, P. Henry, R. Kennedy, A. Bachrach, and A. Bry, “End-to-end learning of geometry and context for deep stereo regression,” in *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
- [5] A. Geiger, P. Lenz, and R. Urtasun, “Are we ready for autonomous driving? the kitti vision benchmark suite,” in *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*. IEEE, 2012, pp. 3354–3361.

- [6] M. Menze and A. Geiger, "Object scene flow for autonomous vehicles," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [7] A. Tonioni, M. Poggi, S. Mattoccia, and L. Di Stefano, "Unsupervised adaptation for deep stereo," in *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
- [8] J. Uhrig, N. Schneider, L. Schneider, U. Franke, T. Brox, and A. Geiger, "Sparsity invariant cnns," in *International Conference on 3D Vision (3DV)*, 2017.
- [9] G. Yang, X. Song, C. Huang, Z. Deng, J. Shi, and B. Zhou, "DrivingStereo: A large-scale dataset for stereo matching in autonomous driving scenarios," in *CVPR*, 2019.
- [10] A. Tonioni, F. Tosi, M. Poggi, S. Mattoccia, and L. Di Stefano, "Real-time self-adaptive deep stereo," in *CVPR*, June 2019.
- [11] C. Zhou, H. Zhang, H. Shen, and J. Jia, "Unsupervised learning of stereo matching," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 1567–1575.
- [12] C. Godard, O. Mac Aodha, and G. J. Brostow, "Unsupervised monocular depth estimation with left-right consistency," in *CVPR*, vol. 2, no. 6, 2017, p. 7.
- [13] Y. Zhang, S. Khamis, C. Rhemann, J. Valentin, A. Kowdle, V. Tankovich, M. Schoenberg, S. Izadi, T. Funkhouser, and S. Fanello, "Activestereonet: End-to-end self-supervised learning for active stereo systems," in *15th European Conference on Computer Vision (ECCV 2018)*, 2018.
- [14] A. Tonioni, M. Poggi, S. Mattoccia, and L. Di Stefano, "Unsupervised domain adaptation for depth prediction from images," 2019.
- [15] H. Hirschmüller, "Accurate and efficient stereo processing by semi-global matching and mutual information," in *Computer Vision and Pattern Recognition*, 2005. *CVPR 2005. IEEE Computer Society Conference on*, vol. 2. IEEE, 2005, pp. 807–814.
- [16] H. Alismail, B. Browning, and M. B. Dias, "Evaluating pose estimation methods for stereo visual odometry on robots," in *the 11th International Conference on Intelligent Autonomous Systems (IAS-11)*, 2011.
- [17] M. Poggi, F. Tosi, and S. Mattoccia, "Quantitative evaluation of confidence measures in a machine learning world," in *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
- [18] R. Haeusler, R. Nair, and D. Kondermann, "Ensemble learning for confidence measures in stereo vision," in *CVPR. Proceedings*, 2013, pp. 305–312, 1.
- [19] A. Spyropoulos, N. Komodakis, and P. Mordohai, "Learning to detect ground control points for improving the accuracy of stereo matching," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2014, pp. 1621–1628.
- [20] M. G. Park and K. J. Yoon, "Leveraging stereo matching with learning-based confidence measures," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.
- [21] M. Poggi and S. Mattoccia, "Learning a general-purpose confidence measure based on o(1) features and a smarter aggregation strategy for semi global matching," in *Proceedings of the 4th International Conference on 3D Vision, 3DV*, 2016.
- [22] —, "Learning from scratch a confidence measure," in *Proceedings of the 27th British Conference on Machine Vision, BMVC*, 2016.
- [23] A. Seki and M. Pollefeys, "Patch based confidence prediction for dense disparity map," in *British Machine Vision Conference (BMVC)*, 2016.
- [24] F. Tosi, M. Poggi, A. Benincasa, and S. Mattoccia, "Beyond local reasoning for stereo confidence estimation with deep learning," in *15th European Conference on Computer Vision (ECCV)*, September 2018.
- [25] S. Kim, S. Kim, D. Min, and K. Sohn, "LAF-Net: Locally adaptive fusion networks for stereo confidence estimation," in *CVPR*, 2019.
- [26] Z. Chen, X. Sun, L. Wang, Y. Yu, and C. Huang, "A deep visual correspondence embedding model for stereo matching costs," in *The IEEE International Conference on Computer Vision (ICCV)*, December 2015.
- [27] W. Luo, A. G. Schwing, and R. Urtasun, "Efficient deep learning for stereo matching," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 5695–5703.
- [28] H. Hirschmüller and D. Scharstein, "Evaluation of stereo matching costs on images with radiometric differences," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, pp. 1582–1599, 08 2008.
- [29] K. Batsos, C. Cai, and P. Mordohai, "CbmV: A coalesced bidirectional matching volume for disparity estimation," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [30] A. Shaked and L. Wolf, "Improved stereo matching with constant highway networks and reflective confidence learning," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [31] S. Gidaris and N. Komodakis, "Detect, replace, refine: Deep structured prediction for pixel wise labeling," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [32] J. Pang, W. Sun, J. S. Ren, C. Yang, and Q. Yan, "Cascade residual learning: A two-stage convolutional neural network for stereo matching," in *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
- [33] Z. Liang, Y. Feng, Y. G. H. L. W. Chen, and L. Q. L. Z. J. Zhang, "Learning for disparity estimation through feature constancy," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [34] E. Ilg, T. Saikia, M. Keuper, and T. Brox, "Occlusions, motion and depth boundaries with a generic network for disparity, optical flow or scene flow estimation," in *The European Conference on Computer Vision (ECCV)*, September 2018.
- [35] X. Song, X. Zhao, H. Hu, and L. Fang, "Edgestereo: A context integrated residual pyramid network for stereo matching," in *ACCV*, 2018.
- [36] G. Yang, H. Zhao, J. Shi, Z. Deng, and J. Jia, "Segstereo: Exploiting semantic information for disparity estimation," in *ECCV*, 2018, pp. 636–651.
- [37] Z. Yin, T. Darrell, and F. Yu, "Hierarchical discrete distribution decomposition for match density estimation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 6044–6053.
- [38] S. Duggal, S. Wang, W.-C. Ma, R. Hu, and R. Urtasun, "Deep-pruner: Learning efficient stereo matching via differentiable patch-match," in *ICCV*, 2019.
- [39] J.-R. Chang and Y.-S. Chen, "Pyramid stereo matching network," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [40] G.-Y. Nie, M.-M. Cheng, Y. Liu, Z. Liang, D.-P. Fan, Y. Liu, and Y. Wang, "Multi-level context ultra-aggregation for stereo matching," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [41] S. Tulyakov, A. Ivanov, and F. Fleuret, "Practical deep stereo (pds): Toward applications-friendly deep stereo matching," in *Advances in Neural Information Processing Systems*, 2018, pp. 5871–5881.
- [42] Y. Wang, Z. Lai, G. Huang, B. H. Wang, L. van der Maaten, M. Campbell, and K. Q. Weinberger, "Anytime stereo image depth estimation on mobile devices," in *ICRA*, 2019.
- [43] F. Zhang, V. Prisacariu, R. Yang, and P. H. Torr, "Ga-net: Guided aggregation net for end-to-end stereo matching," in *CVPR*, 2019, pp. 185–194.
- [44] Y. Zhang, Y. Chen, X. Bai, S. Yu, K. Yu, Z. Li, and K. Yang, "Adaptive unimodal cost volume filtering for deep stereo matching," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 07, 2020, pp. 12 926–12 934.
- [45] X. Cheng, P. Wang, and R. Yang, "Learning depth with convolutional spatial propagation network," *IEEE transactions on pattern analysis and machine intelligence*, vol. 42, no. 10, pp. 2361–2379, 2019.
- [46] X. Cheng, Y. Zhong, M. Harandi, Y. Dai, X. Chang, H. Li, T. Drummond, and Z. Ge, "Hierarchical neural architecture search for deep stereo matching," *Advances in Neural Information Processing Systems*, vol. 33, 2020.
- [47] H. Xu and J. Zhang, "Aanet: Adaptive aggregation network for efficient stereo matching," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 1959–1968.
- [48] S. Khamis, S. Fanello, C. Rhemann, A. Kowdle, J. Valentin, and S. Izadi, "Stereonet: Guided hierarchical refinement for real-time edge-aware depth prediction," in *15th European Conference on Computer Vision (ECCV 2018)*, 2018.
- [49] X. Guo, K. Yang, W. Yang, X. Wang, and H. Li, "Group-wise correlation stereo network," in *CVPR*, 2019.
- [50] M. Poggi, D. Pallotti, F. Tosi, and S. Mattoccia, "Guided stereo matching," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [51] P. L. Dovesi, M. Poggi, L. Andraghetti, M. Martí, H. Kjellström, A. Pieropan, and S. Mattoccia, "Real-time semantic stereo matching," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020.
- [52] H. Jiang, D. Sun, V. Jampani, Z. Lv, E. Learned-Miller, and J. Kautz, "Sense: A shared encoder network for scene-flow estimation,"

- in *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.
- [53] F. Aleotti, M. Poggi, F. Tosi, and S. Mattoccia, "Learning end-to-end scene flow by distilling single tasks knowledge," in *Thirty-Fourth AAAI Conference on Artificial Intelligence*, 2020.
 - [54] J. Pang, W. Sun, C. Yang, J. Ren, R. Xiao, J. Zeng, and L. Lin, "Zoom and learn: Generalizing deep stereo matching to novel domains," *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
 - [55] A. Tonioni, O. Rahnama, T. Joy, L. Di Stefano, A. Thalaiyasingam, and P. Torr, "Learning to adapt for stereo," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
 - [56] T. Zhou, M. Brown, N. Snavely, and D. G. Lowe, "Unsupervised learning of depth and ego-motion from video," in *CVPR*, vol. 2, no. 6, 2017, p. 7.
 - [57] R. Garg, V. K. BG, G. Carneiro, and I. Reid, "Unsupervised cnn for single view depth estimation: Geometry to the rescue," in *European Conference on Computer Vision*. Springer, 2016, pp. 740–756.
 - [58] M. Poggi, F. Aleotti, F. Tosi, and S. Mattoccia, "Towards real-time unsupervised monocular depth estimation on cpu," in *IEEE/IROS Conference on Intelligent Robots and Systems (IROS)*, 2018.
 - [59] M. Poggi, F. Tosi, and S. Mattoccia, "Learning monocular depth estimation with unsupervised trinocular assumptions," in *6th International Conference on 3D Vision (3DV)*, 2018.
 - [60] Y. Zhong, H. Li, and Y. Dai, "Open-world stereo video matching with deep rnn," in *15th European Conference on Computer Vision (ECCV 2018)*, 2018.
 - [61] Y. Wang, P. Wang, Z. Yang, C. Luo, Y. Yang, and W. Xu, "Unos: Unified unsupervised optical-flow and stereo-depth estimation by watching videos," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 8071–8081.
 - [62] H.-Y. Lai, Y.-H. Tsai, and W.-C. Chiu, "Bridging stereo matching and optical flow via spatiotemporal correspondence," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
 - [63] F. Tosi, F. Aleotti, M. Poggi, and S. Mattoccia, "Learning monocular depth estimation infusing traditional stereo knowledge," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
 - [64] A. Pilzer, S. Lathuiliere, N. Sebe, and E. Ricci, "Refine and distill: Exploiting cycle-inconsistency and knowledge distillation for unsupervised monocular depth estimation," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
 - [65] M. Poggi, F. Aleotti, F. Tosi, and S. Mattoccia, "On the uncertainty of self-supervised monocular depth estimation," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
 - [66] Y. Kuznetsov, J. Stuckler, and B. Leibe, "Semi-supervised deep learning for monocular depth map prediction," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
 - [67] C. Banz, S. Hesselbarth, H. Flatt, H. Blume, and P. Pirsch, "Real-time stereo vision system using semi-global matching disparity estimation: Architecture and fpga-implementation," in *ICSAMOS*, 2010, pp. 93–101.
 - [68] S. K. Gehrig, F. Eberli, and T. Meyer, "A real-time low-power stereo vision engine using semi-global matching," in *ICVS*, 2009, pp. 134–143.
 - [69] K. Schmid and H. Hirschmuller, "Stereo vision and imu based real-time ego-motion and depth image computation on a handheld device," in *ICRA*, 2013.
 - [70] S. Mattoccia and M. Poggi, "A passive rgb-d sensor for accurate and real-time depth sensing self-contained into an fpga," in *9th ICDSC*, 2015.
 - [71] O. Rahnama, T. Cavallari, S. Golodetz, A. Tonioni, T. Joy, L. Di Stefano, S. Walker, and P. H. Torr, "Real-time highly accurate dense depth on a power budget using an fpga-cpu hybrid soc," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 66, no. 5, pp. 773–777, 2019.
 - [72] O. Rahnama, T. Cavalleri, S. Golodetz, S. Walker, and P. Torr, "R3sgm: Real-time raster-respecting semi-global matching for power-constrained systems," in *2018 International Conference on Field-Programmable Technology (FPT)*. IEEE, 2018, pp. 102–109.
 - [73] D. Honegger, H. Oleynikova, and M. Pollefeys, "Real-time and low latency embedded computer vision hardware based on a combination of fpga and mobile cpu," in *IROS*, 2014.
 - [74] F. Aleotti, F. Tosi, L. Zhang, M. Poggi, and S. Mattoccia, "Reversing the cycle: self-supervised deep stereo through enhanced monocular distillation," in *16th European Conference on Computer Vision (ECCV)*. Springer, 2020.
 - [75] F. Tosi, M. Poggi, A. Tonioni, L. Di Stefano, and S. Mattoccia, "Learning confidence measures in the wild," in *28th British Machine Vision Conference (BMVC 2017)*, September 2017.
 - [76] A. Ranjan and M. J. Black, "Optical flow estimation using a spatial pyramid network," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
 - [77] D. Sun, X. Yang, M.-Y. Liu, and J. Kautz, "Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
 - [78] A. Badki, A. Troccoli, K. Kim, J. Kautz, P. Sen, and O. Gallo, "Bi3d: Stereo depth estimation via binary classifications," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 1600–1608.
 - [79] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The kitti dataset," *International Journal of Robotics Research (IJRR)*, 2013.

Continual Adaptation for Deep Stereo – Supplementary material

Matteo Poggi¹ Alessio Tonioni² Fabio Tosi¹
Stefano Mattoccia¹ Luigi Di Stefano¹

¹University of Bologna, Italy ²Google Inc.

In this document, we provide additional experiments concerning “Continual Adaptation for Deep Stereo”, complementing those reported in the main paper.

1. Batched adaptation

As discussed in the main paper, updating the network using batches of images instead of individual ones might improve performance. However, this usually happens because building a batch allows to optimize the network over a set of variegated samples, that should better approximate the entire training set and thus lead to more stable gradients compared to those computed on a single sample. In our setting, however, the update steps occur on images that are acquired closely in terms of time and are, therefore, very similar. Thus, the gradients provided by the individual samples are quite similar to the average gradient that would be computed in a batch. We show an experiment to investigate on whether updating the network based on a batch of 10 samples, which slows down the inference rate by a factor of about 4, may bring in significant improvements. The results are shown in [Table 1](#): by comparing the first and second column pairs, one can observe how batch-updating (second column pair) provides just a very small performance improvement wrt our approach (first pair). On the other hand, one might recover full-speed by updating the network every 10 frames based on a batch of 10 samples, but this deteriorates accuracy significantly, as shown in the third column pair of [Table 1](#).

		step=1		step=1		step=10	
		batch=1		batch=10		batch=10	
Adapt. Mode	Proxy src.	D1	EPE	D1	EPE	D1	EPE
FULL	X	2.43	0.95	2.35	0.95	4.43	1.21
FULL++	SGM	2.28	0.95	2.25	0.95	4.44	1.18
MAD	X	4.09	1.19	4.03	1.18	11.37	2.31
MAD++	SGM	2.46	0.98	2.43	0.97	7.36	1.41

Table 1. **Adaptation on image batches.** Experiment dealing with continual adaptation from synthetic pretraining on *Campus* → *City* → *Residential* → *Road*. We perform updates with single images (batch=1) or stacking the latest 10 (batch=10), at any time a new sample is available (step=1) or after collecting 10 new samples (step=10). Pretraining on synthetic images.

2. Sensitivity of parameters δ and λ of Algorithm 2

The two parameters δ and λ , respectively a decay and an update weight, are introduced in Section 3.1 to gradually decay the probability distribution over time and balance the impact of a successful update, respectively. We have tuned them to 0.99 and 0.01 after an analysis carried out on Campus (the shortest and hardest sequence) and starting from *MADNet* weights trained on synthetic data. Table 2 shows D1 error for different decay-weight configurations, highlighting better results with 0.99, 0.01. However, as also noticeable in Table 2, we found that these hyper-parameters are not critical for making adaptation work.

		decay			
		1	0.99	0.75	0.50
weight	1	22.86	22.35	22.59	22.72
	0.1	22.59	22.31	22.55	22.62
	0.01	22.35	22.27	22.37	22.36
	0.001	22.40	22.40	22.46	22.68
MAD					

		decay			
		1	0.99	0.75	0.50
weight	1	12.35	12.78	12.70	14.31
	0.1	11.70	12.69	13.35	11.53
	0.01	13.65	11.50	13.24	13.00
	0.001	12.38	12.02	11.95	11.80
MAD++					

Table 2. **Hyper-parameters tuning** on Campus sequence (1149×2 frames), with *MADNet* pretrained on synthetic data.

3. Proxy supervision vs meta-learning.

Following the L2A paper, we take a *MADNet* model trained on FlyingThings3D and 1) fine-tune it on Synthia in a traditional, supervised manner or 2) using the First Order approximation variant of L2A (FOL2A). We tried also to use the full meta learning formulations without approximation (L2A and L2A-Wad), but unfortunately, the training of the network was very unstable and always lead to diverging models. Thus, for this experiment, 1) provides the baseline with respect to which we evaluate the adaptation performance provided by FOL2A and our proposed approach dealing with proxy labels gathered by SGM. Table 3 collects the outcome of this experiment on the KITTI raw dataset. We can notice how using L2A during pre-training can improve the performance with respect to the baseline. However, our proposal turns out significantly more effective, *i.e.* proxy supervision is more effective than L2A pre-training. For the sake of space, we report in the main paper only the numbers obtained over the whole dataset (All entry in Table 3), while Table 3

Starting Model	Adapt. Mode	City (8027 frames)		Residential (28067 frames)		Campus (1149×2 frames)		Road (5674 frames)		All	
		D1-all(%)	EPE	D1-all(%)	EPE	D1-all(%)	EPE	D1-all(%)	EPE	D1-all(%)	EPE
<i>MADNet</i> -Synthia	FULL (baseline)	3.99	1.26	3.14	1.10	15.09	2.55	3.89	1.19	2.76	1.03
<i>MADNet</i> -Synthia-FOL2A [1]	FULL	3.90	1.25	2.93	1.06	13.98	2.49	3.76	1.21	2.60	1.01
<i>MADNet</i> -Synthia	FULL++ (ours)	3.58	1.14	2.29	0.95	10.62	1.79	3.04	1.05	2.23	0.95

Table 3. **Comparison of online adaptation strategies across different domains.** Results on the *City*, *Residential*, *Campus* and *Road* sequences from KITTI. Comparison between *MADNet* trained on FlyingThings3D and finetuned on Synthia, with (FOL2A) or without meta-learning optimization.

4. Complementing low-density supervision.

We have carried out additional experiments by combining WILD labels with LIDAR as well with self-supervision, so as to study whether one might effectively complement the low density supervision provided by WILD. As for the former setting, we use LIDAR measurements where WILD proxy labels are missing due to low confidence. This allows to consider LIDAR points only where WILD does not provide labels (e.g., untextured regions), while using WILD where LIDAR usually fail, *i.e.* close to object boundaries. In the latter, we just add the two losses, as also proposed in previous works [2, 3]. Results are collected in Table 4. We can notice that complementing WILD with LIDAR supervision consistently improves accuracy. As for self-supervision, this tend to be the case mostly for FULL, while MAD seems to benefit when adapted by the stronger supervision provided by a regression loss, based on either proxy labels, LIDAR measurements or both.

Adapt. Mode	Source	City (8027 frames)		Residential (28067 frames)		Campus (1149×2 frames)		Road (5674 frames)		All	
		D1-all(%)	EPE	D1-all(%)	EPE	D1-all(%)	EPE	D1-all(%)	EPE	D1-all(%)	EPE
FULL++	WILD+Self-Sup.	3.08	1.06	2.13	0.91	8.04	1.49	2.34	0.96	2.04	0.92
FULL++	WILD+LIDAR	3.58	1.02	2.37	0.91	9.09	1.47	3.17	0.92	2.18	0.89
MAD++	WILD+Self-Sup.	5.56	1.29	4.30	1.13	12.13	1.87	5.16	1.20	3.09	1.00
MAD++	WILD+LIDAR	4.10	1.08	2.66	0.93	11.28	1.76	4.71	1.15	2.43	0.91

Pretraining on **synthetic** data


FULL++	WILD+Self-Sup.	1.51	0.88	1.75	0.85	3.80	1.11	1.08	0.80	1.66	0.85
FULL++	WILD+LIDAR	1.48	0.79	1.28	0.76	3.26	0.95	1.19	0.75	1.25	0.75
MAD++	WILD+Self-Sup.	1.65	0.91	2.97	0.95	4.22	1.25	1.34	0.85	2.20	0.91
MAD++	WILD+LIDAR	1.43	0.80	1.67	0.80	3.68	1.10	1.19	0.76	1.34	0.77

Pretraining on **synthetic** data, fine-tuning on KITTI 2015 GT.

Table 4. **Online adaptation within a single domain (WILD+Lidar labels).** Results on the *City*, *Residential*, *Campus* and *Road* sequences from KITTI.

5. Comparison with sparse (and filled) SGM

As we have observed in the experiments on the DrivingStereo sequences, the sparse proxy labels generated by SGM and WILD are very accurate, although they have low density. However, we point out that lower density might be rather undesirable in real applications. For instance, let us consider an autonomous driving scenario, as depicted in Figure 1 below from the Rainy sequence used in our experiments. Although selecting the proxy algorithm with higher density, the sparse SGM map lacks measurements in critical areas close to the vehicle, whilst the smooth and dense map perceived by our method seems quite more amenable to support high-level tasks dealing with planning and navigation.



Reference image SGM (before filtering) SGM (after filtering) *MADNet-K (MAD++)*

Figure 1. **Qualitative results achieved by SGM and *MADNet*.** Example taken from *Rainy* sequence.

Given the good accuracy provided by SGM sparse proxies, one might try to densify these sparse depth measurements in order to use them in practical applications. Thus, we have performed depth completion on the sparse SGM maps by a standard hand-crafted interpolation method (*i.e.*, the hole-filling algorithm used to densify disparity maps when submitting to the KITTI online benchmark) and a recent learning-based solution (the network used in [4]) trained on the same synthetic data as *MADNet*. We report the results of this experiment in Table 5. We notice that hole-filling can outperform our proposal in the Cloudy sequence only, *i.e.* the one where SGM achieves the lowest error among the three sequences, as shown in the main paper. However, more challenging conditions (as observed in the Rainy and Dusky sequences) are better tackled by *MADNet* thanks to online adaptation. We would also like to highlight, that, as shown in Figure 2, standard hand-crafted hole filling is particularly prone to yield gross artifacts that do not affect the error metrics due to the sparsity of the ground-truth provided with DrivingStereo (filtered LIDAR). Yet, more advanced learning-based depth densification methods seem significantly affected by the domain shift, as vouched by the second row in Table 5.

Completion method	Rainy (1667 frames)		Dusky (1119 frames)		Cloudy (4950 frames)	
	D1-all(%)	EPE	D1-all(%)	EPE	D1-all(%)	EPE
SGM + Hole-filling [†]	12.68	3.05	7.88	1.30	4.50	1.19
SGM + ECCV 2020	17.90	3.91	14.53	1.91	9.74	1.83
<i>MADNet</i> -K-MAD++ (SGM)	12.65	2.32	5.93	1.40	6.26	1.72

Table 5. **Comparison between online adaptation and SGM densification on DrivingStereo.** Results on the *Rainy*, *Dusky* and *Cloudy* sequences. Learned methods are trained on synthetic data. [†]means non-learned algorithm.




Figure 2. **Comparison between SGM densification and online adaptation.** Example taken from *Rainy* sequence.

References

- [1] A. Tonioni, O. Rahnama, T. Joy, L. Di Stefano, A. Thalaiyasingam, and P. Torr, “Learning to adapt for stereo,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2
- [2] A. Tonioni, M. Poggi, S. Mattoccia, and L. Di Stefano, “Unsupervised domain adaptation for depth prediction from images,” 2019. 3
- [3] F. Tosi, F. Aleotti, M. Poggi, and S. Mattoccia, “Learning monocular depth estimation infusing traditional stereo knowledge,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 3
- [4] F. Aleotti, F. Tosi, L. Zhang, M. Poggi, and S. Mattoccia, “Reversing the cycle: self-supervised deep stereo through enhanced monocular distillation,” in *16th European Conference on Computer Vision (ECCV)*. Springer, 2020. 3