

# Detection of GAN-synthesized street videos

Omran Alamayreh\*, Mauro Barni†

Department of Information Engineering and Mathematics, University of Siena

Via Roma 56, 53100 - Siena, ITALY

Email: \*omran@diism.unisi.it, †barni@dii.unisi.it

**Abstract**—Research on the detection of AI-generated videos has focused almost exclusively on face videos, usually referred to as deepfakes. Manipulations like face swapping, face re-enactment and expression manipulation have been the subject of an intense research with the development of a number of efficient tools to distinguish artificial videos from genuine ones. Much less attention has been paid to the detection of artificial non-facial videos. Yet, new tools for the generation of such kind of videos are being developed at a fast pace and will soon reach the quality level of deepfake videos. The goal of this paper is to investigate the detectability of a new kind of AI-generated videos framing driving street sequences (here referred to as DeepStreets videos), which, by their nature, can not be analysed with the same tools used for facial deepfakes. Specifically, we present a simple frame-based detector, achieving very good performance on state-of-the-art DeepStreets videos generated by the Vid2vid architecture. Noticeably, the detector retains very good performance on compressed videos, even when the compression level used during training does not match that used for the test videos.

**Index Terms**—DeepStreets, DeepFake, GANs, XceptionNet, Vid2vid, Video Forensics.

## I. INTRODUCTION

The recent development of Artificial Intelligence (AI) tools, enabling even non-expert users to generate fake videos of extraordinary quality, is raising increasing concerns about the credibility of digital media accessible on the internet and diffused by information and social networks. The consequences of such a loss of credibility are devastating and impact all domains of our lives, ranging from everyday life, through journalism, criminal justice, and national security. The necessity of distinguishing between fake and original media has triggered the birth of a new discipline, known as Multimedia Forensics MMF [1], which has attracted the interest of researchers and research funding agencies all around the world [2].

With regard to the detection of fake videos generated by means of AI techniques based on Deep Learning (DL), most of the attention has been devoted to the detection of face facial videos, known as Deepfakes. As a consequence, a wide range of tools have been developed to detect several kinds of deepfake manipulations, including face swapping [3] and face re-enactment [4]. The interest towards this kind of videos is justified by the fact that information conveyed by speeches and facial expressions has an essential role in human social interactions. Besides, human face manipulation techniques have received an increasing attention in the last years, reaching a maturity level making them ready to be exploited in several real-life applications (including malicious ones) [1].

In comparison, the detection of non-facial artificial video sequences has received much less attention. Yet, new AI tools for the generation of non-facial videos are being developed at a fast pace and will soon reach the quality level of deepfake videos [5]. To move a first step to fill this gap, in this paper we focus on the detection of a new class of fake videos, hereafter referred to as DeepStreets, framing driving street sequences, like those described in [6], [7], [5].

To start with, we observe that the forgery detection methods designed to work on facial videos can not be directly applied to DeepStreets. The main reason for such a difficulty is that deepfake detectors rely strongly on the facial features of the videos, all the more that they usually crop the face region before inputting the video frames to the detectors. Besides, some detectors are based on the analysis of a set of geometric and semantic features that are directly related to human faces; like eye color [8] or facial pose [9]. No such features obviously exist in DeepStreets videos, which are characterised by a greater diversity and for which it is not possible to identify a single attention region to focus the analysis on.

To address the above limitations, we propose an end-to-end detector based on CNN. The detector works on a frame by frame basis, and does not focus on any specific region of the input frames. To train (and test) the detector, we generated 600 videos of fake driving scenes by means of the Vid2vid architecture presented in [5] and briefly summarized in Section II-B. To evaluate the generalization capability of the detector and its ability to work in non-ideal conditions, we considered both raw sequences and videos compressed at different quality levels. The experiments we carried out prove the validity of the detector, which exhibited an extremely high accuracy, demonstrating that even DeepStreets kind of videos can reliably be distinguished from real videos (at least in the highly controlled setting adopted in this paper). A remarkable property of our detector is its immunity to video compression, given that the detection accuracy remains good even when the compression setting used during training is different than that used for the test videos.

The rest of this paper is structured as follows; In Section II, we briefly review the state of the art of deepfake generation methods, video-to-video synthesis and deepfake detection methods. In Section III, we present the dataset we used in our experiments and the detection pipeline. The results of the experiments we carried out to validate the proposed detector are described in Section IV. In Section V, we conclude the paper with some final remarks and hints for future research.

## II. RELATED WORK

### A. Synthesis of deepfake videos

With the rise of deep learning techniques, automatic generation of falsified media has reached an unprecedented level. Several manipulation approaches based on AI have been proposed in the literature for fake videos generation. These methods generally rely on two deep learning architectures; autoencoders (AE [10]) and generative adversarial networks (GANs [11]). A popular approach to create deepfake videos is face replacement, where two autoencoders, with a shared encoder, are trained in parallel on two datasets. The datasets include face images of person *A* and *B*, respectively. The idea is to allow the shared encoder to learn common features for both persons *A* and *B* while keeping the corresponding decoders detached. To swap the faces between person *A* and *B*, a video framing person *A* is fed to the common encoder. Then, the resulting features in the latent space are input to the decoder trained on *B*. This principle is applied in several manipulation tools, like DeepFaceLab [12].

In addition to autoencoders, generative adversarial networks (GANs [11]) are often used to generate falsified media. GANs consist of two networks facing each other in a zero-sum game. In the training phase, the generator network aims at producing falsified images that are undistinguishable from those contained in the training dataset of real images. Meanwhile, the discriminator network is trained to detect the images produced by the generator. After training, the generator is used to create realistic fake images that are hardly identified as such by the discriminator. GANs and autoencoders can also be used together to improve the quality of the synthesised videos, as in Faceswap-GAN [13].

### B. Video-to-Video synthesis

Most deepfake manipulation methods have been developed to create fake videos that contain human faces. Their counterpart for non-facial fake videos has been less explored. However, general-purpose video synthesis techniques are present in several forms, including video prediction [14], unconditional video synthesis [15] and conditional video synthesis [5]. The first methods produce future video frames based on past frames. The second kind of methods take in random variables and generate synthesised videos. The latter methods convert input semantic sequences into photorealistic videos.

In our work, we used a conditional video synthesis method, known as Vid2vid [5], to generate the training and testing datasets. Vid2vid is a video-to-video synthesis architecture operating under the generative adversarial learning framework. The goal of Vid2vid is to convert an input semantic video, for example, a sequence of semantic segmentation masks, to an output photorealistic video that represents the theme of the source video. Vid2vid is considered the counterpart of image-to-image translation methods such as Pix2pix [16] and Covst [17]. The advantage of Vid2vid with respect to the application of image-based architectures, is that working on a frame-by-frame basis, image-based techniques do not

guarantee the temporal coherence of the synthetic output video. In contrast, Vid2vid is designed in such a way to preserve the temporal dynamics of the source video.

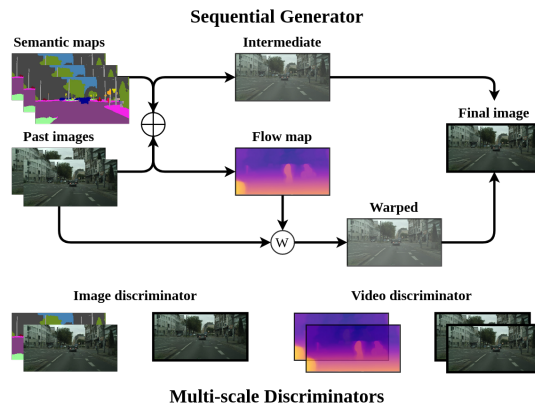


Fig. 1: Vid2vid network architecture.

The architecture of Vid2vid consists of two networks; a sequential generator and a set of multiscale discriminators. The generator takes as input a sequence of segmentation semantic maps and the frames generated previously and produces an intermediate frame and a flow map. Afterwards, the flow map is used to warp the previous frame. Then, the warped frame is combined with the intermediate frame to output the final frame, which is then used to generate the next frame and so on. The multiscale discriminators contain one discriminator for images and one for videos. The image discriminator analyses the input semantic map and the output images to ensure that each output frame resembles a real image. Simultaneously, the video discriminator takes in the flow maps and neighbouring frames to guarantee the temporal coherence of the generated videos. The architecture of Vid2vid is depicted in figure 1.

Recently, two versions of Vid2vid have been released: Few-shot video-to-video synthesis [7] and World-Consistent Video-to-Video Synthesis [6]. The first work increases the generalisation capability of Vid2vid by using a novel network weight generation module based on attention mechanism. The latter improves the generated videos' temporal consistency by exploiting, so-called, *guidance images*, which are a consolidation of the *3D world* rendered in the previous frames. In our experiments, we used Vid2vid and Wc-vid2vid architectures to generate DeepStreets dataset, as described in section III.

### C. Deepfake detection methods

The techniques developed so far to distinguish real and deepfake videos belong to two main classes: those based on handcrafted features and those relying on end-to-end completely automatic detectors based on DL. On the one hand, handcrafted methods exploit a set of predefined traces that result from the production pipeline of deepfakes, like, visual artifacts [18], semantic inconsistencies [8]. On the other hand, DL-based methods, specifically those based on convolutional neural networks (CNN), deliver end-to-end solution



(a) Semantic segmentation masks.

(b) Sequence used to generate the mask.

(c) Generated fake video.

Fig. 2: Example of frames generated by Vid2vid architecture.

to the forgery detection problem. Approaches like MesoInception [19], and XceptionNet [20], provide superior performance compared to the traditional handcrafted methods [20]. CNN based methods still suffer from significant shortcomings like the difficulty of interpreting the output they provide, the vulnerability to adversarial attacks, and the difficulty to maintain good performance when the training and test phases work in mismatched conditions. Moreover, the performance of DL-based methods drop in the presence of video compression [1].

Here we propose a forgery detector that relies on CNN. As demonstrated in section IV, the detector shows impressive performance, and a good degree of robustness against video coding, while its generalization capability strongly depends on the experimental conditions.

### III. METHODOLOGY

#### A. DeepStreets Dataset

To generate the dataset we used in our experiments, we started from the pre-trained models of Vid2vid and Wc-vid2vid on the Cityscapes dataset [21]. Both models require a sequence of semantic segmentation masks as input to generate a video. Additionally, Wc-vid2vid requires a sequence of guidance images to improve the output video’s temporal consistency [6]. The models’ output is a photorealistic video of street scenes that render the content of the input sequence of the semantic segmentation masks. Figure 2 shows an example of the generated videos.

We used Cityscapes [21] and Kitti [22] datasets to produce the input segmentation maps. The cityscapes dataset consists of diverse urban street videos from several German cities. The videos have  $2048 \times 1024$  resolution captured using a pair of stereo cameras at varying times of the year. The Kitti dataset contains driving videos that have been captured near Karlsruhe, Germany, using colour and grayscale camera images. In our experiments, we used only the colour videos with  $1392 \times 512$  resolution.

Since not all the images of Cityscapes are labelled with segmentation masks, and to be consistent with Vid2vid pipeline, we have annotated the images of Cityscapes and Kitti by using the network from Zhu et al. [23]. We also followed Wc-vid2vid pipeline to generate the guidance images, and we performed *structure from motion* SfM on the video sequences using OpenSfM [24].

We divided the dataset into three subsets; cityvid, Citywcvid and Kittivid. We generated Cityvid by inputting the segmentation masks of Cityscapes into Vid2vid, and Citywcvid by feeding Wc-vid2vid with segmentation masks and guidance images of Cityscapes. For Kittivid, we used the segmentation masks of Kitti dataset as input to Vid2vid. Each sub-dataset contains 400 videos, half of them are fake, and the other half are real (the real videos are the original sequences from Cityscapes and Kitti datasets). The videos resolution is  $512 \times 1024$  and their duration is 3 seconds, with 30 frames. Table I summarises the characteristics of the sub-datasets.

| sub-dataset | network    | segmentation map |
|-------------|------------|------------------|
| Cityvid     | vid2vid    | cityscapes       |
| Citywcvid   | wc-vid2vid | cityscapes       |
| Kittivid    | vid2vid    | kitti            |

TABLE I: A summary of DeepStreets dataset.

In addition to the raw videos produced by the Vid2vid models, we have compressed the videos with two quality levels; *HQ*, and *LQ*. In particular, we compressed the raw videos by using the H.264 codec with a constant rate quantisation parameter equal to 23 and 40, respectively. The compressed videos as well as the raw videos are available on our website; <http://clem.dii.unisi.it/~vipp/datasets.html>.

#### B. The Proposed Detector

To distinguish genuine videos from DeepStreets ones, we propose to use a data-driven forgery detector based on CNN. Since the proposed dataset has no clear predefined traces, it is difficult to rely on handcrafted features to differentiate between fake and genuine videos. In addition, the existing deepfake detectors usually detect and crop the faces in each frame of the videos. In our case, we need to input the whole frame to the detector, thus making a CNN-like architecture a good candidate for our task. Having said that, we used Xception network [25] as the backbone for our detector.

Xception is a CNN-like architecture based on depthwise separable convolution layers. It uses 36 convolutional layers for the feature extraction part. The layers are organised into 14 modules; each module has linear residual connections, excluding the first and last modules. The convolutional part is followed by an optional fully-connected layer and a logistic

regression layer. For our detector, we added a fully-connected layer with two outputs on top of the feature extraction part.

We formalise the detection task as a *per-frame binary classification* problem. We start by extracting the frames from each video. Then, to limit the computational burden, we resized each frame to  $256 \times 512$  resolution. Afterwards, the frames are sequentially fed to the Xception network. The detection pipeline is shown in figure 3.




Fig. 3: Detection pipeline.

To ensure the reproducibility of our results, we report the setting that we have used to conduct the experiments. We split the dataset into training, testing, and validation sets, with ratios of 60%, 25%, and 15%, respectively. All the results reported hereafter regard the detection accuracy on the test set. During the training, we used Adam optimizer with learning rate ( $\alpha$ ) = 0.0001 and the default values for the first and second-order moments ( $\beta_1 = 0.9, \beta_2 = 0.999$ ), and an epsilon value ( $\epsilon$ ) =  $10^{-8}$ . Finally, the batch size was set to 8. Training was stopped whenever validation loss did not decrease for 10 consecutive epochs. All the experiments were performed using PyTorch framework on a workstation with one intel core i9 and four NVIDIA GeForce RTX 2080 Ti GPUs.

#### IV. EXPERIMENTAL RESULTS

We have conducted several experiments to assess the performance of the proposed forgery detector on DeepStreets dataset at various compression levels, including matched and mismatched conditions. Furthermore, we have investigated its generalization capability by carrying out cross-dataset analysis between all subsets of DeepStreets dataset.

| Dataset\Quality | RAW    | HQ     | LQ     |
|-----------------|--------|--------|--------|
| Cityvid         | 100.00 | 100.00 | 97.93  |
| Citywcvd        | 99.76  | 99.62  | 99.96  |
| Kittivid        | 100.00 | 100.00 | 100.00 |
| DeepStreets     | 99.86  | 100.00 | 99.19  |

TABLE II: Detection accuracy (the detector is asked to distinguish the original sequences used to create the segmentation masks and the synthetic videos).

In our first experiments, we trained the detector to distinguish the original sequences used to generate the segmentation masks and the sequences produced by the synthetic video generators. Table II shows the detection accuracy of the detector on all subsets of DeepStreets dataset. The results refer to matched training and testing, that is when training and testing is carried out on the same subsets. As it can be seen, the classification accuracy is almost perfect for all the subsets even when training and testing is carried out on the whole dataset.

With regard to the impact of video compression on detection accuracy, we carried out two sets of experiments. In Table II, the detector is trained and tested in matched compression conditions, while in Table III different training and testing conditions are used. In both cases, compression has a minor impact on the accuracy. This marks a difference with respect to deepfakes videos, where the presence of compression leads to a significant performance drop (see for instance the results reported in [20] with regard to low-quality videos).

| Training\Testing | RAW    | HQ     | LQ    |
|------------------|--------|--------|-------|
| RAW              | 99.89  | 99.90  | 95.41 |
| HQ               | 100.00 | 100.00 | 95.72 |
| LQ               | 99.70  | 99.63  | 99.19 |

TABLE III: Compression mismatch experiments. The detector has been trained (and tested) on the entire DeepStreet dataset.

To investigate the generalisation capability of the proposed forgery detector, we performed a cross-dataset analysis between all the subsets of DeepStreets dataset. As illustrated in Table IV, in some cases, dataset mismatch causes a dramatic drop of the performance, while in other cases the performance loss is less significant. Such behaviour can be explain by the different production pipeline used for the three sub-datasets. As mentioned in section III-A, we have generated each subset by adopting different settings, including the source of the semantic maps and the specific Vid2vid architecture that generated the videos.

| Training\Testing | Cityvid | Citywcvd | Kittivid |
|------------------|---------|----------|----------|
| Cityvid          | 100.00  | 71.50    | 88.16    |
| Citywcvd         | 98.76   | 99.76    | 50.00    |
| Kittivid         | 50.03   | 50.00    | 100.00   |

TABLE IV: Cross-dataset analysis between subsets of DeepStreets dataset (RAW videos).

For instance, we can see that performance significantly deteriorates if we perform cross-dataset analysis between Kittivid and Citywcvd sub-datasets, since these sub-datasets have a different source for the semantic sequences. We used Kitti [22] and Cityscapes [21] dataset to generate Kittivid and Citywcvd sub-datasets, respectively. Besides, the Vid2vid architectures used to generate them are different (see table I). In contrast, the performance is quite adequate when we perform cross-dataset analysis between Cityvid and Citywcvd sub-datasets since both of them have the same source of semantic sequences (Cityscapes [21] in this case) and differ only for the Vid2vid architecture used for the generation.

We also conducted a second set of experiments, wherein the detector is asked to distinguish the synthetic sequences from the real sequences used to train the generators (namely the sequences from real images of the Cityvid sub-dataset). The reason why we carried out this additional experiment is that the generators tend to produce scenes with colours

and light conditions similar to those of the sequences used to train them (that is Cityvid real images), so it may be relatively easy for the discriminator to distinguish original images belonging to a different dataset, e.g. Kitti sequences, since they are characterised by different environmental conditions (light and sun conditions, for instance). This is not the case, when the real images are taken from the same dataset used to train the GANs. In the new setting, the real videos are driving scenes from the Cityscapes dataset, and the fake videos are synthesised videos generated by using the semantic sequences from the Kitti dataset. The detector trained in this way achieved 100% accuracy on the test set of the new sub-dataset. We also tested the ability of the detector to distinguish between real Cityvid videos and fake videos generated starting from segmentation maps of Cityvid (the Cityvid sub-datasets in Table I). Interestingly, the detector achieved a 80.80% accuracy, showing a certain capability to detect examples of fake videos that are not represented in the training set.

## V. CONCLUSIONS

We have demonstrated the possibility of detecting non-facial fake videos generated by video-to-video translations tools by using data-driven methods, even when training and testing are carried out in compression mismatched conditions. We chose the street views as an example of non-facial fake videos. Nevertheless, the generation and detection pipeline can be extended to other contents. Future research will be devoted to create a dataset containing street videos and other non-facial fake videos, such as F1 videos, tennis videos etc. (despite the computational burden required to train new video-to-video models [5]). Another direction for future research is to investigate the generalization capability of the proposed forgery detector to a class of facial and non-facial synthetic videos that is not included in the training set. Finally, it would also be interesting to assess the robustness of the proposed detector against adversarial attacks.

## ACKNOWLEDGMENT

This work has been supported by the PREMIER project under contract PRIN 2017 2017Z595XS-001, funded by the Italian Ministry of University and Research and by the Defense Advanced Research Projects Agency (DARPA) and the Air Force Research Laboratory (AFRL) under agreement number FA8750-20-2-1004. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA and AFRL or the U.S. Government.

## REFERENCES

- [1] Luisa Verdoliva. Media forensics and deepfakes: an overview. *IEEE Journal of Selected Topics in Signal Processing*, 14(5):910–932, 2020.
- [2] Semantic forensics. <https://www.darpa.mil/program/semantic-forensics>.
- [3] Yuval Nirkin, Yosi Keller, and Tal Hassner. Fsgan: Subject agnostic face swapping and reenactment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [4] Justus Thies, Michael Zollhofer, Marc Stamminger, Christian Theobalt, and Matthias Nießner. Face2face: Real-time face capture and reenactment of rgb videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2387–2395, 2016.
- [5] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Guilin Liu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. Video-to-video synthesis. *arXiv preprint arXiv:1808.06601*, 2018.
- [6] Arun Mallya, Ting-Chun Wang, Karan Sapra, and Ming-Yu Liu. World-consistent video-to-video synthesis. *arXiv preprint arXiv:2007.08509*, 2020.
- [7] Ting-Chun Wang, Ming-Yu Liu, Andrew Tao, Guilin Liu, Jan Kautz, and Bryan Catanzaro. Few-shot video-to-video synthesis. *arXiv preprint arXiv:1910.12713*, 2019.
- [8] Falko Matern, Christian Riess, and Marc Stamminger. Exploiting visual artifacts to expose deepfakes and face manipulations. In *2019 IEEE Winter Applications of Computer Vision Workshops (WACVW)*, pages 83–92. IEEE, 2019.
- [9] Xin Yang, Yuezun Li, and Siwei Lyu. Exposing deep fakes using inconsistent head poses. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8261–8265. IEEE, 2019.
- [10] Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on Machine learning*, pages 1096–1103, 2008.
- [11] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *arXiv preprint arXiv:1406.2661*, 2014.
- [12] Ivan Petrov, Daiheng Gao, Nikolay Chervoni, Kunlin Liu, Sugasa Marangonda, Chris Umé, Jian Jiang, Luis RP, Sheng Zhang, Pingyu Wu, et al. Deepfacelab: A simple, flexible and extensible face swapping framework. *arXiv preprint arXiv:2005.05535*, 2020.
- [13] GitHub. Faceswap-gan. <https://github.com/shaoanlu/faceswap-GAN>.
- [14] Tianfan Xue, Jiajun Wu, Katherine L Bouman, and William T Freeman. Visual dynamics: Probabilistic future frame synthesis via cross convolutional networks. *arXiv preprint arXiv:1607.02586*, 2016.
- [15] Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Generating videos with scene dynamics. *arXiv preprint arXiv:1609.02612*, 2016.
- [16] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8798–8807, 2018.
- [17] Dongdong Chen, Jing Liao, Lu Yuan, Nenghai Yu, and Gang Hua. Coherent online video style transfer. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1105–1114, 2017.
- [18] Yuezun Li and Siwei Lyu. Exposing deepfake videos by detecting face warping artifacts. *arXiv preprint arXiv:1811.00656*, 2018.
- [19] Darius Afchar, Vincent Nozick, Junichi Yamagishi, and Isao Echizen. Mesonet: a compact facial video forgery detection network. In *2018 IEEE International Workshop on Information Forensics and Security (WIFS)*, pages 1–7. IEEE, 2018.
- [20] Andreas Rossler, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias Nießner. Faceforensics++: Learning to detect manipulated facial images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1–11, 2019.
- [21] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016.
- [22] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *International Journal of Robotics Research (IJRR)*, 2013.
- [23] Yi Zhu, Karan Sapra, Fitsum A Reda, Kevin J Shih, Shawn Newsam, Andrew Tao, and Bryan Catanzaro. Improving semantic segmentation via video propagation and label relaxation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8856–8865, 2019.
- [24] GitHub. Opensfm. <https://github.com/mapillary/OpenSfm>.
- [25] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1251–1258, 2017.